Imperial College London
Department of Earth Science and Engineering

Reservoir Modeling and Inversion using Generative Adversarial Network Priors

Lukas J. Mosser

August 2019

Submitted in part fulfilment of the requirements for the degree of
Doctor of Philosophy of Imperial College London

Declaration of Originality

I herewith certify that all material in this dissertation which is not my own work has been properly acknowledged.

Lukas J. Mosser

Abstract

Determining the spatial distribution of geological heterogeneities and their petrophysical properties is key to successful hydrocarbon production and carbon capture and storage. Due to the sparse nature of direct observations of the earth's interior from borehole data, most inferences about the interior structure of the earth and its properties have to be made by indirect observation such as seismic reflection or dynamic data. Determining these property distributions from indirect observations requires solving an ill-posed inverse problem which can be defined as a Bayesian inference problem where we seek to obtain the posterior distribution of the subsurface properties given the observed data.

Recently, deep generative modeling has enabled multi-modal probability distributions of large three-dimensional natural images to be represented. Generative Adversarial Networks (GANs) are deep generative models that learn a representation of the probability distribution implicitly defined by a set of training images using two competing neural networks.

This thesis introduces GANs as probabilistic models of geological features and petrophysical properties at the reservoir scale and images of porous media at the pore-scale. A GAN can be trained to represent pore-scale micro-CT images of segmented and grayscale porous media. After training, the GAN generator is used to sample large high-fidelity realizations that follow the same statistical and physical properties as represented in the training images. Using GANs as a probabilistic generative model allows them to be incorporated in a Bayesian inversion workflow. Based on a synthetic test-case, two inverse problems were considered: inversion of acoustic properties from seismic observations and reservoir history matching of a two-phase flow problem at the reservoir-scale. In both cases, the posterior distribution of the petrophysical property distributions was obtained using approximate Bayesian inference over the latent variables. The samples obtained from the posterior match the observed seismic or production data, and can be conditioned to direct observations at wells.

This approach of deep stochastic inversion based on deep generative models such as GANs opens new opportunities for geological modeling and solving ill-posed inverse problems.

List of Publications

Articles Published and in Print

Mosser, L., O. Dubrule, and M. J. Blunt (2017). "Reconstruction of three-dimensional porous media using generative adversarial neural networks". In: *Physical Review E* 96.4, p. 043309.

— (2018a). "Conditioning of generative adversarial networks for pore and reservoir scale models". In: *80th EAGE Conference and Exhibition 2018*.

— (2018b). "Stochastic reconstruction of an oolitic limestone by generative adversarial networks". In: *Transport in Porous Media* 125.1, pp. 81–103.

Articles Submitted for Peer-Review

Submitted to the *Journal of Mathematical Geosciences* and accepted with minor revisions

Mosser, L., O. Dubrule, and M. J. Blunt (2018). "Stochastic seismic waveform inversion using generative adversarial networks as a geological prior". In: *ArXiv e-prints*. arXiv: 1806.03720.

Submitted to the *Journal of Computational Physics-X*

Mosser, L., O. Dubrule, and M. J. Blunt (2019). "DeepFlow: History matching in the space of deep generative models". In: *ArXiv e-prints*. arXiv: 1905.05749.

Available Code and Data

Mosser, L., O. Dubrule, and M. J. Blunt (2017). "Reconstruction of three-dimensional porous media using generative adversarial neural networks". In: *GitHub*. DOI: 10.5281/zenodo.3358586. URL: https://github.com/LukasMosser/PorousMediaGAN.

— (2018a). "Conditioning of generative adversarial networks for pore and reservoir scale models". In: *GitHub*. DOI: 10.5281/zenodo.3358588. URL: https://github.com/LukasMosser/geogan.

— (2018b). "Stochastic seismic waveform inversion using generative adversarial networks as a geological prior". In: *GitHub*. DOI: 10.5281/zenodo.3358582. URL: https://github.com/LukasMosser/stochastic_seismic_waveform_inversion.

— (2019). "DeepFlow: History matching in the space of deep generative models". In: *GitHub*. DOI: 10.5281/zenodo.3358579. URL: https://github.com/LukasMosser/DeepFlow.

Copyright Declaration

Acknowledgments

First, I would like to thank my supervisors Olivier Dubrule and Martin J. Blunt for their unwavering support and guidance throughout my PhD.

Olivier, thank you for the countless conversations and discussions we've had these three years which have been a source of great inspiration for me. This thesis only represents a small part of the many joyful collaborations and projects I've had the chance to work on with you, many of which shape my academic career beyond my time as a student.

Thank you, Martin, for being a role-model of academic diligence for me - every form of my research, whether as part of this thesis or beyond, has been improved by your supervision. Your passion for imparting knowledge has been deeply inspiring.

Thank you, Olivier and Martin, for being not only my supervisors but also my mentors. I am deeply grateful for having had the chance to learn from both of you.

I would like to thank my family, my parents Ursula and Jakob, and my sister Sophie, for their unconditional love, their support all throughout my education, and regardless of where life took me, for providing a place I could turn to and call *home*.

I want to thank my partner, Andrea, for her love and support, and for walking this path far from home with me.

I would also like to thank Alfredo de la Fuente Briceño, for the many supportive conversations and discussions about the struggles of an academic career as well as sharing his passion for science and the world of machine learning with me.

Furthermore, I want to thank my dear friends Dominik, Fabian, Jakob, and Manuel, for their endless friendship.

I would also like to thank my friends and colleagues at the Department of Earth Sciences and Engineering, for providing good times and an open ear over the last years.

Finally, I am indebted to the maintainers and contributors to the open-source software, which I have used in my projects that have contributed to this work. Without their efforts, this thesis would not exist.

Dedication

To my grandparents - Erich and Elisabeth, Jakob and Anna

Don't cut corners - Grandpa Erich on a drive to pick mushrooms in the forest.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

1.1 Motivation

Modeling at the Pore-scale

The main motivation for this thesis is to find probabilistic representations of geological heterogeneities at the pore- and reservoir-scales. At the pore-scale, modern micro-CT scanning allows three-dimensional images of porous media to be acquired at high-resolution (Blunt, 2017). Petrophysical properties can be estimated from these three-dimensional images using numerical solutions of flow and transport equations. This allows critical flow related properties such as absolute and relative permeability to be determined. Understanding the effect of pore structure is a key element to understanding flow and transport at the pore-scale. Stochastic modeling of porous media at the pore-scale allows flow and transport to be simulated using synthetic realizations based on statistical properties of porous micro-structures, the sedimentation and diagenetic history, or when possible from images of the pore and grain space. Stochastic models of porous media at the pore-scale have been constructed using object-based models which generate a grain arrangement using characteristic geometric shapes such as ellipsoids or spheres. Simulated annealing methods (Yeong and Torquato, 1998) reconstruct three-dimensional images of porous media from two-point or higher-order statistical information at high computational cost (Pant, 2016). Process-based models aim to generate three-dimensional realizations of porous media by emulating sedimentation and diagenesis processes (Øren and Bakke, 2002; Øren and Bakke, 2003). Multiple-point statistical (MPS) methods use a training-image of a porous medium to deduce multi-scale statistical properties from which new realizations are generated in a sequential manner (Mariethoz and Caers, 2014). MPS methods have been adapted to pore-scale stochastic reconstruction allowing three-dimensional stochastic modeling from orthogonal cross-sections only (Okabe and Blunt, 2004; Okabe and Blunt, 2007). Patch-based methods rearrange extracted subdomains of a training image to form new stochastic realizations (Tahmasebi and Sahimi, 2016). Many of these approaches suffer from high computational cost for generating a single stochastic realization and when a large number of realizations are required e.g. for uncertainty quantification, most of these approaches become prohibitively computationally expensive because the computational cost is proportional to the number of realizations.

Modeling at the Reservoir-scale

Modeling the spatial distribution of petrophysical properties within the earth's subsurface is a key step in reservoir characterization. Traditionally, geostatistical methods have been used to create continuous petrophysical property models and categorical distributions for lithofacies modeling (Dubrule, 2003; Chiles and Delfiner, 2009). At the reservoir-scale, geostatistical methods for reservoir modeling can be categorized into continuous and categorical random function modeling approaches. Continuous properties can be represented by calculating variograms from well data and fitting a variogram model to model their spatial distribution. Indicator variables which are commonly used to represent rock-types or facies distributions have been modeled using pixel-based approaches, such as sequential indicator simulation (SIS) (Bierkens and Burrough, 1993), pluri-Gaussian simulation (PGS) (Armstrong et al., 2011), multiple-point statistics (MPS)(Guardiano and Srivastava, 1993), and object-based approaches, where discrete volumetric representations of geological features such as fluvial channels are randomly placed within a reservoir model. While object-based models allow characteristic geological features to be modeled discretely, conditioning them to well data is computationally challenging when there are many wells.

The Bayesian Framework

Limited direct observations of the earth's interior from borehole data and uncertainties affecting indirect observations such as seismic or dynamic data motivate a Bayesian framework to quantify subsurface structures and petrophysical properties. Furthermore, inferring subsurface structures and their petrophysical properties means solving an ill-posed inverse problem i.e. there exist many solutions to the inverse problem that match the observed data. Tarantola, (2005) set the geophysical inverse problem in a Bayesian setting. The Bayesian framework requires a definition of a likelihood that characterizes the relationship between the underlying model parameters and the observations. Furthermore a prior distribution of the model parameters that quantifies the *a priori* knowledge about the model parameters and their uncertainty needs to be defined. In many geophysical inverse problems the relationship between the model parameters and the observed data is non-linear. Finding the posterior distribution may therefore require Monte-Carlo methods where many evaluations of the forward problem have to be performed. Simplifying the forward problem by linearization can make Bayesian inversion feasible (Buland and Omre, 2003; Hansen et al., 2006). Conditioning a reservoir model with historical dynamic production data for the sake of production forecasting and reservoir management, so called history matching, is a common non-linear inverse problem in petroleum engineering. The changes to the petrophysical parameters obtained by manual and automated history matching methods often cause non-geological changes which violate prior geological and geophysical assumptions (Hoffman et al., 2005). Bissell et al., (1997) introduced a method based on the pilot-point method (Marsily, 1978) that allows for history matching under geostatistical constraints. Bayesian approaches to solve the reservoir history matching problem, such as the Ensemble Kalman Filter (EnKF) (Evensen, 2003; Oliver and Chen, 2011) have been applied in numerous history matching case studies (Lorentzen et al., 2005).

Advances in Deep Learning and Generative Modeling

Recent algorithmic advances in automatic differentiation such as efficient back-propagation (Rumelhart et al., 1988; LeCun et al., 1998) and optimization of high-dimensional parametric functions using stochastic gradient-based optimization techniques (Robbins and Monro, 1951; Kiefer and Wolfowitz, 1952; Bottou, 2010), combined with increased computational resources using graphics processing units (GPUs) (Krizhevsky et al., 2012) and the availability of very large datasets (Deng et al., 2009) have led to numerous advances in the field of machine learning (Bishop, 2006; Goodfellow et al., 2016). Specifically, training very deep neural networks has been made feasible by these advances, which are finding successful applications in many industrial sectors such as healthcare (Wang et al., 2017; Gómez-Bombarelli et al., 2016). While many successful applications of machine learning rely on supervised training i.e. learning from pairs of data and corresponding labels, unsupervised learning aims to classify data or reduce their dimensionality without any label information. Estimating the probability distribution associated with a given dataset and possibly learning a model of how these data can be synthesized is referred to as generative modeling. Using deep neural networks to parameterize generative models has enabled modeling of multi-modal probability distributions of large natural images (Bengio et al., 2014; Goodfellow et al., 2014) and making inferences about their underlying factors (Kingma and Welling, 2013).

Thesis Aims

The aims of this thesis are therefore two-fold: first, I investigate the use of recent advances in deep learning and specifically deep generative modeling to create probabilistic representations of pore-scale images and reservoir-scale models which allow fast sampling of new high-fidelity realizations. Second, I aim to leverage the probabilistic definition of deep generative models to define a new set of prior distributions over petrophysical model parameters that can be used to solve high-dimensional subsurface inversion problems in a Bayesian context.

1.2 Contributions

The papers published during this PhD and presented in this thesis are under the names of myself, Lukas J. Mosser, and my supervisors Olivier Dubrule and Martin J. Blunt. As the first author I have performed all research and code development myself. All implementations in code, data and results, have been implemented in Python and have been published as software packages under open-source licenses. (Github: `https://www.github.com/LukasMosser` - references to the code repositories are in the individual chapters). The contributions presented in this thesis and published as journal papers have been written by myself first and were then proof-read by my supervisors who helped me to improve my writing in style and structure.

Four papers have been submitted for peer-review in the course of this thesis which form the main parts of this thesis. Two papers have been published in print, the first in *Physical Review E* (Mosser et al., (2017), Chapter 3) and the second in *Transport in Porous Media* (Mosser et al., (2018c), Chapter 4). Chapter 6 (Mosser et al., (2018d) has been accepted for publication in *Mathematical Geosciences* with minor revisions, and has been resubmitted for publication. The final chapter of this thesis (Chapter 7, Mosser et al., (2019)) has been submitted to the *Journal of Computational Physics* and is currently under peer-review. All publications have been published as a pre-print on *ArXiv* at time of submission and all journal publications have been published under open licenses. This also includes the contribution presented in Chapter 5 (Mosser et al., 2018a) which has been presented at the 80th EAGE Annual Conference (2018).

During my PhD I have also presented my work at a number of conferences including oral presentations at *Interpore* (Rotterdam 2017), the *80th and 81st EAGE Annual Conferences*(Copenhagen 2018, and London 2019), the *PESGB/EAGE Workshop on Machine Learning* (London 2018), the *IAMG Conference* (Olumuc 2018), the *IPTC Conference* (Beijing 2019) and the *Seminar for Deep Learning in Inverse Problems 2019* held at KTH Stockholm . I was also fortunate to have received a stipend by the *Simons Foundation* to attend the *Math+X Seminar on Data Science in Geophysics*. Furthermore I was accepted to present a poster at the *Gordon Conference and Research Seminar on Flow and Transport in Porous Media*. Finally, I have also had the opportunity to present my work to colleagues in seminars at the Department of Earth Sciences and Engineering and in industry at Total Paris, Saclay and Houston.

1.3 Statement of Originality

The work contained in this thesis has not been previously submitted for a degree or diploma at any other higher education institution. To the best of my knowledge and belief, the thesis contains no material previously published or written by another person except where due references are made.

1.4 Thesis Outline

In Chapter 2, I present an overview of existing methods for inversion, geostatistical simulation and deep generative modeling which serves as a basis to the methods and concepts that are used in later chapters of this thesis. The end of Chapter 2 outlines a number of research questions that motivate the presented research.

Chapter 3 introduces GANs as a method to create stochastic three-dimensional samples of porous media at the pore-scale using a number of binary training images.

The method is then extended in Chapter 4 to grayscale images of porous media and introduces an evaluation of generated samples based on morphological image descriptors.

I then show in Chapter 5 that the same approach can be applied to three-dimensional reservoir-scale training images and how generated realizations can be conditioned to hard data such as wells.

By generalizing GANs as a probabilistic method to model distributions of images I incorporate trained GAN generators as prior distributions in a Bayesian inversion framework and in Chapter 6 apply the proposed method to a synthetic two-dimensional acoustic full-waveform inversion example.

In Chapter 7 I apply the proposed GAN-based inversion to infer spatial distributions of rock properties in the frame of reservoir history matching.

Chapter 8 is a discussion of the contributions presented in the individual chapters and publications, highlighting some of the challenges and opportunities of using GANs for pore and reservoir-scale modeling, and places the methods into a larger context of using deep generative models for solving inverse problems.

Final conclusions of the thesis are presented in Chapter 9.

Chapter 2

Bayesian Inversion and Modeling

2.1 Inversion

2.1.1 Maximum Likelihood

Determining the structural and petrophysical properties of the earth's interior is a key step in hydrocarbon exploration and production, seismology and carbon capture and sequestration (CCS). While observations of the geological features can be made in analog outcrops on the surface, hard data from boreholes are expensive to acquire and represent a sparse set of measurements. The geological properties measured from borehole data are at a very different scale compared to larger geological heterogeneities such as reservoirs bodies or fault complexes. Direct observation of subsurface geological features and their petrophysical properties is impossible and geoscientists must therefore rely on indirect observations such as seismic or dynamic data. Determining the underlying petrophysical properties that govern these indirect measurements is called inversion. To formalize this, it is commonly assumed that the observed data $\mathbf{d}_{obs}$ are related to a three-dimensional earth model $\mathbf{m}(\mathbf{x})$ by

$$\mathbf{d}_{obs} = \mathcal{F}(\mathbf{m}(\mathbf{x})) \tag{2.1}$$

where $\mathcal{F}$ is the forward operator (Tarantola, 2005). Due to the heterogeneous nature of the earth's interior a common assumption is that the model parameters vary in space $\mathbf{m}(\mathbf{x})$ and when dynamic forward operators are considered these parameters may also be time-dependent $\mathbf{m}(\mathbf{x}, t)$.

Indirect geophysical observations are often accompanied by large uncertainties due to data acquisition and processing, as well as error regarding the conceptual interpretation of the observed data and the forward model that is assumed to generate them. Geological and petrophysical properties are often correlated spatially. Knowledge about spatial correlation represents valuable information to deduce the spatial distribution of geological heterogeneities given observed measurements under uncertainty.

Due to the sparsity of hard data and the resolution of the observations $\mathbf{d}_{obs}$ it is possible to find numerous realizations of the model $\mathbf{m}(\mathbf{x})$ that can explain the observed data $\mathbf{d}_{obs}$ and hence inverse problems related to geoscience applications are often ill-posed.

The simplest approach to finding a set of model parameters that match the observed data is to perform *maximum likelihood estimation* (Tarantola, 2005). For

a given set of observations $\mathbf{d}_{obs,i}$ we seek to maximize the likelihood and assuming conditional independence of the observations

$$\mathbf{m}^{MLE} = \underset{\mathbf{m}}{\operatorname{argmax}} \prod_i p(\mathbf{d}_{obs,i}|\mathbf{m}) = \underset{\mathbf{m}}{\operatorname{argmax}} \sum_i \log p(\mathbf{d}_{obs,i}|\mathbf{m}) \tag{2.2}$$

where in practise due to numerical precision issues we prefer to maximize the log-likelihood (Eq. 2.2) (Murphy, 2012). Assuming the observed data are a sum of the forward model and Gaussian uncorrelated noise with variance σ_ϵ^2 our representation of the forward model becomes

$$\mathbf{d}_{obs} = \mathcal{F}(\mathbf{m}(\mathbf{x})) + \varepsilon,\ \varepsilon \sim \mathcal{N}(\mathbf{0}, \sigma_\epsilon^2 \mathbb{I}) \tag{2.3}$$

and hence to find the maximum likelihood estimator of the model parameters we maximize the Gaussian log-likelihood of the observations given the parameters

$$\mathbf{m}^{MLE} = \underset{\mathbf{m}}{\operatorname{argmax}} \sum_i \log p(\mathbf{d}_{obs,i}|\mathbf{m}) \tag{2.4a}$$

$$= \underset{\mathbf{m}}{\operatorname{argmin}} \sum_i \frac{\|\mathbf{d}_{obs,i} - \mathcal{F}_i(\mathbf{m})\|_2^2}{\sigma_{\epsilon_i}^2} \tag{2.4b}$$

When ideal measurements without noise are assumed the maximum likelihood approach reduces to the least-squares method

$$\mathbf{m}^{LS} = \underset{\mathbf{m}}{\operatorname{argmin}} \sum_i \|\mathbf{d}_{obs,i} - \mathcal{F}_i(\mathbf{m})\|_2^2 \tag{2.5}$$

The maximum likelihood approach does not take into account any prior knowledge of the distribution of model parameters and only provides a point estimate that maximizes the match with the observations. Therefore, only maximizing the likelihood may lead to models $\mathbf{m}(\mathbf{x})$ that are geologically unrealistic.

2.1.2 Posterior Distribution

Setting the inversion problem into a Bayesian context (Tarantola, 1987) allows prior knowledge and uncertainties in observed and modeling parameters to be incorporated. By treating the observations $\mathbf{d}_{obs}$ and model parameters as random variables and hence properties that follow an underlying stochastic process we can define the joint distribution $p(\mathbf{d}_{obs}, \mathbf{m})$. Factorizing the joint distribution leads to

$$p(\mathbf{d}_{obs}, \mathbf{m}) = p(\mathbf{d}_{obs}|\mathbf{m})p(\mathbf{m}) = p(\mathbf{m}|\mathbf{d}_{obs})p(\mathbf{d}_{obs}) \tag{2.6}$$

which allows the conditional distribution of model parameters on the observed data $p(\mathbf{m}|\mathbf{d}_{obs})$ - also called the posterior distribution - to be expressed by rearranging Eq. 2.6

$$p(\mathbf{m}|\mathbf{d}_{obs}) = \frac{p(\mathbf{d}_{obs}|\mathbf{m})p(\mathbf{m})}{p(\mathbf{d}_{obs})} = \frac{p(\mathbf{d}_{obs}|\mathbf{m})p(\mathbf{m})}{\int p(\mathbf{d}_{obs}|\mathbf{m})p(\mathbf{m})\,\mathrm{d}\mathbf{m}} \tag{2.7}$$

which is Bayes' law. To fully define Bayes' law requires definition of the marginal or prior distributions of the observations $p(\mathbf{d}_{obs})$ and model parameters $p(\mathbf{m})$

$$\mathbf{m} \sim p(\mathbf{m})\ (\textit{prior probability}) \tag{2.8a}$$

$$\mathbf{d}_{obs} \sim p(\mathbf{d}_{obs})\ (\textit{evidence}) \tag{2.8b}$$

The use of conjugate probability distributions such as multi-Gaussian distributions allow us to evaluate the posterior analytically, which in the case of a multi-Gaussian prior and linear likelihood is also a multi-Gaussian distribution (Bishop, 2006).

In many applications the prior distribution of model parameters may be very high-dimensional; a grid representation of the spatial distribution of lithological facies, porosity and permeability in a hydrocarbon reservoir can easily exceed the order of millions of individual grid blocks. While in many cases spatial correlations are assumed to represent the prior distribution of model parameters $p(\mathbf{m})$ exact numerical integration of the marginal distribution of the observations $p(\mathbf{d}_{obs})$ quickly becomes intractable as the number of parameters in $\mathbf{m}$ grows. Therefore other approaches must be considered in the multivariate case to perform Bayesian inference i.e. finding the posterior distribution of model parameters, without the need for analytical integration of the model evidence $p(\mathbf{d}_{obs})$.

The model evidence in Eq. 2.7 is a constant and therefore we can determine the posterior up to a scaling factor

$$p(\mathbf{m}|\mathbf{d}_{obs}) \propto p(\mathbf{d}_{obs}|\mathbf{m})p(\mathbf{m}) \tag{2.9}$$

This simplification allows us to find the maximum of the posterior distribution (MAP)

$$\mathbf{m}^{\mathrm{MAP}} = \underset{\mathbf{m}}{\operatorname{argmax}}\, p(\mathbf{d}_{obs}|\mathbf{m})p(\mathbf{m}) \tag{2.10a}$$

due to the fact that maxima of functions remain maxima under rescaling by a positive constant. Therefore if we are only interested in finding a set of parameters that maximize the posterior probability $\mathbf{m}^{MAP}$ i.e. finding the mode of the posterior, it is sufficient to maximize the product of the prior and the likelihood.

When the prior and likelihood distributions are conjugate we can find the maximum of the posterior analytically. Where gradients of the posterior can be obtained, gradient-based optimization methods can be used to obtain the MAP estimate. When the posterior is multi-modal finding the global MAP using gradient-based methods can become challenging and even if the maximum can be found it may not provide a good representation of the posterior distribution. This is further compounded in the high-dimensional case where the mode i.e. the area of highest probability density is far from the area of highest probability mass due to the curse of dimensionality.

2.1.3 Sampling from the Posterior

Posterior sampling methods, also called *simulating from the posterior*, allows samples from the posterior to be obtained by first drawing samples from the prior and using additional criteria to determine whether a sample should be accepted as being part of the posterior distribution or rejected. The *rejection sampling algorithm* draws samples from a known density $g(\mathbf{m})$, the proposal or envelope distribution and accepts samples with probability

$$p_i = \frac{p(\mathbf{d}_{obs}|\mathbf{m})p(\mathbf{m})}{Mg(\mathbf{m})} \tag{2.11}$$

where M is a scaling constant so that $g(\mathbf{m})$ fully covers the posterior distribution $p(\mathbf{m}|\mathbf{d}_{obs})$. For high-dimensional distributions or complex posterior distributions

with many modes this approach can be very inefficient as a large number of samples have to be simulated for one sample to be accepted.

Importance sampling uses the ratio defined in Eq. 2.11 as a weight that determines a sample's contribution towards an estimate of the posterior mean. As such every sample contributes, although possibly very small weight towards the posterior mean, whereas rejection sampling considers samples in a purely binary fashion.

Approximate Bayesian Computation (ABC) relaxes the requirement of defining a likelihood to perform sampling from the posterior (Beaumont et al., 2002; Csilléry et al., 2010). This is useful when the exact likelihood is unknown. We first define a scalar measure of mismatch $\rho(\cdot,\cdot)$ or distance between the observed and modeled data $\rho(\mathbf{d}_{\text{obs}}, \mathcal{F}(\mathbf{m}))$ for a sample drawn from the prior $\mathbf{m} \sim p(\mathbf{m})$. Samples from the approximate posterior $p(\mathbf{m}|\rho(\mathbf{d}_{\text{obs}}, \mathcal{F}(\mathbf{m})) \leq \epsilon)$ can be drawn by accepting samples with mismatch less or equal to a threshold ϵ. Two limiting cases need to be considered: first, when the threshold $\epsilon \to 0$ the true posterior is recovered and second as $\epsilon \to \infty$ obtains samples from the prior. While ABC is a pragmatic way of performing approximate Bayesian inference, choosing different values if ϵ allows us to trade-off efficiently between accuracy of the approximate posterior and the necessary computational effort. Small threshold values ε may require excessive computational resources. When the compared observations are high-dimensional using a lower-dimensional measure of mismatch ρ such as the cross-correlation or mean squared difference should be preferred to accept and reject simulated realizations.

Rejection and importance sampling (Gelman et al., 2013) assumes independence between individual simulated samples accepted as part of the posterior. To explore posterior distributions more effectively and hence be able to sample from high-dimensional multi-modal posterior distributions, *Markov-Chain Monte-Carlo* (MCMC) approaches assume a dependence between individual proposed samples (Gilks et al., 1995). A Markov-Chain is a series of parameter states, where in its simplest definition, the next state is only dependent on the current state. By performing numerous and often millions or even billions of these correlated steps the series of samples tends towards an equilibrium distribution corresponding to the posterior distribution.

The Metropolis-Hastings (M-H) algorithm defines the most basic form of such a Markov-Chain (Metropolis et al., 1953; Hastings, 1970). If the target density is the unnormalized posterior distribution $p(\mathbf{m}|\mathbf{d}_{obs})$ and $q(\mathbf{m}_i|\mathbf{m}_{i-1})$ defines a transition probability that determines how the current state of parameters $\mathbf{m}_{i-1}$ transitions to the next state $\mathbf{m}_i$, acceptance of a new proposed state $\mathbf{m}_{i+1}$ is given by the probability

$$\alpha_i = min\left\{\frac{p(\mathbf{d}_{obs}|\mathbf{m}_i)p(\mathbf{m}_i)q(\mathbf{m}_i|\mathbf{m}_{i-1})}{p(\mathbf{d}_{obs}|\mathbf{m}_{i-1})p(\mathbf{m}_{i-1})q(\mathbf{m}_{i-1}|\mathbf{m}_i)}, 1\right\} \tag{2.12}$$

If the proposed state is rejected a new proposal $\mathbf{m}_i \sim p(\mathbf{m}_i|\mathbf{m}_{i-1})$ is drawn. Random-walk Metropolis assumes the proposal distribution

$$\mathbf{m}_i = \mathbf{m}_{i-1} + \varepsilon_m,\ \varepsilon_m \sim \mathcal{N}(0, \Sigma_m) \tag{2.13}$$

where ε_m is sampled from a multi-Gaussian with variance-covariance matrix Σ_m. The challenge with MCMC methods is to find good proposal strategies that allow exploration of the target posterior while maintaining high acceptance rates. In the case of the random-walk Metropolis algorithm the acceptance rate is high, but the method suffers from poor exploration of the posterior and is equivalent to a random

walk with a guide provided by the M-H criterion. The Gibbs-Sampling strategy (Tanner and Wong, 1987; Gelfand and Smith, 1990) tries to improve exploration of the posterior by alternating between multiple correlated Markov-Chains that propose changes to a subset of the parameters at a time.

When gradients of the target posterior can be computed efficiently Hamiltonian-Markov-Chain methods (Duane et al., 1987; Neal, 2011; Betancourt, 2017) are able to leverage the underlying structure of the posterior distribution and in combination with automatic step-size selection algorithms provided by methods such as the No-U-Turn-Sampler (NUTS) (Hoffman and Gelman, 2014), this family of gradient-based MCMC algorithms allows efficient sampling of high-dimensional posterior distributions. In Chapter 6 I will use an approximation to a gradient-based MCMC method, the so-called approximate Metropolis-adjusted Langevin Algorithm (Roberts and Tweedie, 1996; Roberts and Rosenthal, 1998; Nguyen et al., 2016) to obtain samples of the posterior distribution for a seismic full-waveform inversion problem where gradients of the likelihood and prior are obtained by the adjoint-state method and traditional neural network backpropagation.

2.1.4 Latent Variable Models

A common family of probabilistic models are latent variable models. These probabilistic models are defined by a set of observations $\mathbf{d}_{obs}$ and unobserved parameters $\mathbf{z}$. Together this defines a joint distribution $p(\mathbf{d}_{obs}, \mathbf{z})$. We may wish to determine the posterior distribution of the latent variables given the observations $p(\mathbf{z}|\mathbf{d}_{obs})$. Gaussian mixture models (GMM) are a member of the family of latent variable models. A GMM can be represented as a combination of a discrete latent variables $\mathbf{z}$ and continuous Gaussian distributions associated with the observations $\mathbf{d}_{obs}$. The posterior distribution $p(\mathbf{z}|\mathbf{d}_{obs})$ corresponds to finding the distribution of the initially unknown parameters of a set of Gaussian distributions from a set of observed data. This allows us to learn something about the hidden structure of the data e.g. forming clusters under the GMM. In modern statistical learning (Bishop, 2006; Goodfellow et al., 2016) probabilistic models (Koller and Friedman, 2009) can become very complex and, with the availability of enormous datasets, using traditional sampling-based approaches can quickly become intractable and we therefore need to expand our toolkit of Bayesian inference methods to techniques that can overcome these hurdles.

Variational inference (VI) reformulates the problem of finding an approximation to the true posterior distribution as one of optimization (Jordan et al., 1999; Wainwright et al., 2008; Blei et al., 2017). To perform variational inference we choose a parametric distribution $q_\theta(\mathbf{z})$ and minimize a measure of similarity between two distributions, the posterior $p(\mathbf{z}|\mathbf{d}_{obs})$ and our parametric proposal distribution $q_\theta(\mathbf{z})$. By optimizing the parameters θ of our proposal distribution we seek to obtain an approximate representation of the posterior. A common choice as a measure of dissimilarity between distributions is the Kulback-Leibler divergence (KL) (Gelman et al., 2013)

$$\theta^{\mathrm{VI}} = \underset{\theta}{\operatorname{argmin}}\, \mathrm{KL}(q_\theta(\mathbf{z}) \| p(\mathbf{z}|\mathbf{d}_{obs})) \tag{2.14}$$

Computing the KL-divergence is intractable and therefore it is not possible to formulate a direct minimization problem where we optimize the parameters of the proposal distribution. Instead, we optimize a lower bound on the model evidence

(Evidence Lower BOund - ELBO) which can be shown to be equivalent to minimizing the KL-divergence.

For the mean-field family (Blei et al., 2017) where the proposal distribution is factorized into independent distributions

$$q(\mathbf{z}) = \prod_i q_i(\mathbf{z}_i) \tag{2.15}$$

we can then perform coordinate descent to maximize the ELBO objective. I will make use of variational inference methods later during this chapter to introduce a very flexible family of methods to represent high-dimensional distributions that are parameterized by neural networks.

2.2 Representing the Prior Distribution

In the first part of this chapter I have introduced how subsurface properties can be inferred from measured data in a Bayesian setting. Bayes' Law tells us how we can derive knowledge about the subsurface from the observed data by combining our prior understanding (formulated in a prior distribution) with the likelihood function. I have introduced a number of deterministic and stochastic techniques that allow us to derive knowledge about the posterior distribution of the subsurface properties. I have previously assumed that a prior distribution has been formulated from which we can obtain random samples. In the following I will give an overview of the most important approaches that have been developed to create prior probabilistic models of properties at the pore and reservoir-scale.

The first distinction I will make is based on the type of random function being modeled; continuous or categorical. At the reservoir-scale continuous properties are often associated to petrophysical parameters such as porosity and permeability or the P-wave velocity of the subsurface lithologies, while categorical properties are associated with lithologies or geological facies. In many cases we can obtain sparse direct measurements of the subsurface properties at boreholes. These observations should be honored by any random sample that we obtain from our prior distribution and in the following I refer to these as *hard data*. *Geostatistical conditional simulation* (GCS) allows us to draw realizations of multi-dimensional property distributions that honor observed hard data while being statistically representative of the subsurface heterogeneities. There are a number of differences between GCS and the simpler spatial interpolation or *Kriging* approach. Kriging is a deterministic method that while honoring hard data leads to very smooth representations of the properties away from the observed data. Kriging represents the mean of all possible realizations that match the data. This representation of subsurface properties is not reconcilable with actual property distributions - where there is strong heterogeneity in a reservoir this should be the case at the wells and between the wells. In contrast to Kriging, samples created by GCS are stochastic i.e. we can draw random samples from a distribution where each realization honors observed data while also statistically representing heterogeneities.

2.2.1 Continuous Random Functions

In the continuous random variable case, a number of approaches have been developed to create realizations of stationary Gaussian random variables. *Sequential Gaussian simulation* (SGS) (Deutsch and Journel, 1992) creates realizations of property distributions by sequentially populating the simulation domain following a random path with random samples. Each random value that is used to populate a new location is drawn from a Gaussian distribution with a mean and standard deviation obtained by kriging in a moving window neighborhood. Observed data and previously simulated values act as conditioning locations through kriging and are left unchanged in the simulation process. *LU decomposition* of the covariance matrix as introduced by (Davis, 1987; Alabert, 1987) lends itself to simulation of Gaussian random variables but can become computationally challenging for large grid sizes. Two-step iterative methods create conditional realizations of Gaussian random variables by first generating an unconditional realization using methods such as the *Fourier Integral method* (Yao, 1999) or the moving average technique (Oliver, 1995; Doyen, 1988) and then correct for the mismatch at the conditioning data (Chiles and Delfiner, 2009). Non-Gaussian variables can be simulated either by first transforming non-Gaussian distributions to a Gaussian distribution using the normal score transform or by *direct simulation* i.e. for data that show close to Gaussian behavior performing SGS from a local distribution (Deutsch and Journel, 1992) obtained by kriging. For the case where data have been transformed using the normal score transform, these Gaussian realizations usually obtained by SGS have to be transformed back to the non-Gaussian case after simulation.

2.2.2 Pixel-based Approaches for Random Functions

Geological facies or lithologies at the reservoir-scale, as well as the pores and grains of a pore-scale image after segmentation, can be represented by indicator variables. *Pixel-based* methods consider each pixel value within an image or grid block in the case of 3D reservoir models as a random variable and these methods generate realizations on a per-pixel basis. *Sequential indicator simulation* (SIS) (Bierkens and Burrough, 1993) represents a generalization of SGS to categorical random variables. In the case where multiple indicator variables are present (e.g. 0-sand, 1-shale, 2-limestone) these are first transformed to separate binary indicator variables and experimental variograms are fit by a model. Realizations are obtained analogous to the SGS approach in a sequential manner along a random path where kriging in the SIS case is used to obtain the probability of an indicator variable within a local neighborhood. *Truncated pluri-Gaussian simulation* uses conditional simulations of Gaussian variables which are thresholded using a truncation rule that allows spatial dependence and chronological ordering of indicator variables to be modeled (Rudkiewicz et al., 1990; Emery, 2007; Armstrong et al., 2011).The simulated Gaussian variables are often called latent variables.

At the pore-scale, pixel-based approaches such as *simulated annealing* (Yeong and Torquato, 1998; Jiao et al., 2008) allow high-quality three-dimensional reconstruction and incorporation of numerous statistical descriptors of porous media. Pant, (2016) introduced a multi-scale simulated annealing algorithm allowing simulation of three-dimensional porous media achieving a lower computational cost and higher reconstruction quality than previous methods.

Many pixel-based approaches use variograms as a measure of spatial variability, whereas recent methods in geostatistics use training images as a basis for pixel-based geostatistical simulation (Caers and Zhang, 2004; Mariethoz et al., 2010; Meerschman et al., 2013). Training images are usually assumed to exhibit stationarity of the probability distribution of the properties of interest and provide higher order *multiple-point statistics* (MPS) to reconstruct stochastic random media. With MPS, the probability distributions are represented by training images and are sampled using a limited multi-scale neighborhood that captures the variation on a large scale, as well as fine structural details on smaller scales (Tahmasebi et al., 2014). MPS-based methods have been successfully applied in generating realistic three-dimensional realizations of geological features (Guardiano and Srivastava, 1993; Caers, 2001; Mariethoz and Caers, 2014). The computational complexity of these methods is highly dependent on individual algorithms as well as the size of the domains used to sample from the training images (Mariethoz and Caers, 2014). Parallelized versions have been developed, reducing the computational time required to perform reconstruction using MPS (Straubhaar et al., 2011; Huang et al., 2013). At the pore-scale, micron-resolution X-ray tomography (micro-CT imaging) (Flannery et al., 1987) provides training images for MPS-based simulation techniques that enable stochastic reconstruction of three-dimensional porous media (Okabe and Blunt, 2004; Okabe and Blunt, 2005; Okabe and Blunt, 2007).

Tahmasebi et al., (2012) have introduced a patch-based approach where sub-domains are simulated along a predefined path and populated based on a cross-correlation distance criterion (CCSIM) (Tahmasebi and Sahimi, 2012; Tahmasebi and Sahimi, 2013). This approach is similar to the image quilting algorithm by Efros and Freeman, (2001) but corrects mismatching patches in overlapping or neighboring domains.

2.2.3 Object-based Methods

Object-based methods create random spatial arrangements of geological bodies that represent the building blocks of the observed geology or porous media. Many sedimentary rocks consist of granular siliciclastic or carbonate materials. Boolean models use this fundamental characteristic of natural granular materials to emulate the shape of the arising pore space, due to an underlying random process that controls the distribution of the individual grains (Matheron, 1975; Serra, 1980). While for the classical Boolean model, the centers of the grains are uniformly distributed in space and grains can arbitrarily overlap, more complicated models with rigid hard sphere grains and more complex grain interaction functions have been developed (Matheron, 1971; Arns et al., 2009; Rikvold and Stell, 1985; Torquato, 2013). The framework of Boolean models also allows extension beyond spherical particles and enables derivation of the properties of material models as a function of the parameters of the underlying random process (Bretheau and Jeulin, 1989; Jeulin, 2000; Lin and Cohen, 1982; Yeong and Torquato, 1998). At the reservoir-scale, object-based models have been successfully applied as representations of many geological structures such as fluvial channels and conditioned to available seismic data in a Bayesian framework (Holden et al., 1998) . However, their conditioning can be challenging when the data spacing is smaller than the average size of the modeled objects (Dubrule, 2003).

While many of the presented approaches to represent continuous and categorical

random functions rely on two-point or higher-order statistical information to obtain samples from a prior distribution of spatial properties, more complex models may allow more sophisticated models of random variable functions. In the following I will introduce the family of neural network type functions and how we can use these to build sophisticated models of probability distributions based on a set of training data or images, from which we can sample new realizations.

2.3 Generative Modeling using Neural Networks

I will now address modeling distributions of data in a more general way. In the following we will assume that we are modeling a distribution of images $\mathbf{m} \sim p_\theta(\mathbf{m})$ - geological training images at reservoir-scale or CT-images at pore-scale. As introduced in Section 2.2 we represent the distribution of properties $\mathbf{m}$ that follow an unknown distribution $p_\theta(\mathbf{m})$

$$\mathbf{m} \sim p_\theta(\mathbf{m}) \tag{2.16}$$

If we assume that our training data can be represented by a multi-Gaussian distribution $\mathcal{N}(\mu, \Sigma)$ then the set of model parameters θ correspond to the mean and variance-covariance matrix. Without loss of generality the training data we are modeling can be of any form, but for the sake of simplicity we consider here the training data to be a set of geological training images. For a set of two-dimensional or three-dimensional images $\mathbf{m}$ we may seek to find values of the parameters θ that allow us to create a parametric probability density function $p_\theta(\mathbf{m})$ that approximates the probability distribution of the training images. This process is commonly referred to as learning. Natural data often follow complex multi-modal and non-Gaussian distributions and hence our models of the probability distributions need to be sufficiently powerful to represent the training data.

2.3.1 Using Neural Networks for Supervised Learning

In recent years, neural networks have become a popular choice as parametric function representations. One common application for neural networks today is for representing conditional distributions such as in the case of classification or regression problems (Bishop, 2006). For example, in binary classification problems the class probabilities of a Bernoulli distribution given the data can be parameterized by a neural network. Neural networks have recently earned enormous attention thanks to significant advances such as the development of efficient training algorithms using backpropagation (Rumelhart et al., 1988). Backpropagation allows gradients of a loss function with respect to the parameters of a neural network to be efficiently computed. The availability of very large datasets and the ability of stochastic gradient descent techniques (Bottou, 2010) to optimize the parameters of neural networks by using noisy gradient estimates from small so-called batches of data have also played a key role in the success of neural networks. Modern graphics processing units (GPUs) are a specialized type of computational processors that allow fast and parallel matrix-vector operations to be performed (Krizhevsky et al., 2012). Efficient implementations of the backpropagation algorithm (LeCun et al., 1998) use matrix-vector algorithms to evaluate the forward calculation which in the case of a classification problem corresponds to predicting the class probabilities given the data,

and to evaluate gradients of the mismatch from the true labels for a small batch of training data with respect to all the parameters of the neural network.

The development of convolutional neural networks has sparked many successful applications in the domain of computer vision, such as image recognition and segmentation (LeCun et al., 1995). Convolutional neural networks extend the idea of traditional kernel-based computer vision, where an image is convolved by a small kernel of parameters with fixed weights, to detect specific features such as edges. The parameters of a convolutional neural network are the weights of the kernel that is convolved with a given image. Convolutional neural networks allow the structured representation of images to be exploited and make an implicit assumption of translational invariance to learn feature representations with a much smaller number of parameters. Finally, the combination and stacking of many of these parametric functional layers in so-called *deep* neural networks drive most of today's successful applications of deep learning (Goodfellow et al., 2016). Affine transformations of the input data (also called fully connected layers) and convolutional operations are linear operations and are combined with non-linear activation functions. Vanishing gradients due to the use of sigmoid-type activation functions in deep neural networks make learning the parameters very time consuming and inefficient. Rectified Linear Unit (ReLU) activation functions have a unit gradient regardless of the magnitude of the input feature and when used in very deep neural networks allow gradients of sufficient magnitude to be backpropagated in very deep neural networks so that learning the parameters using gradient descent can be performed within reasonable time and computational effort (Krizhevsky et al., 2012). In this thesis, I use deep neural networks to represent and parameterize families of probability functions of very high-dimensional data such as geological models or three-dimensional images of porous media.

2.3.2 Using Neural Networks to represent Probability Distributions

Maximum likelihood estimation allows us to learn the parameters θ of a neural network used to parameterize the density $p_\theta(\mathbf{m})$ for a given training dataset of images $\mathcal{M} = \{\mathbf{m}_1, ..., \mathbf{m}_i\}$ assumed to be independent and identically distributed (iid)

$$\max_\theta \{\log p_\theta(\mathcal{M})\} = \max_\theta \{\sum_{\mathbf{m}_i \in \mathcal{M}} \log p_\theta(\mathbf{m}_i)\} \tag{2.17}$$

Due to the differentiable nature of neural networks we can find a set of parameters θ that maximizes the log-likelihood over the elements $\mathbf{m}_i$ of the training dataset $\mathcal{M}$ using gradient descent

$$\theta_{t+1} = \theta_t - \eta \nabla_\theta \log p_\theta(\mathbf{m}) \tag{2.18}$$

where t denotes the optimization step, η is the step size or learning rate and gradients are computed with respect to model parameters θ.

While for fully-observed distributions $p_\theta(\mathbf{m})$ we can use maximum likelihood to learn the parameters θ for a training dataset $\mathcal{M}$, this is not the case for deep latent variable models i.e. latent variable models where deep neural networks are used to parameterize the density function

$$p_\theta(\mathbf{m}) = \int p_\theta(\mathbf{m}, \mathbf{z}) \, \mathrm{d}\mathbf{z} \tag{2.19}$$

The additional challenge comes from the integral expression in Eq. 2.19. Before being able to evaluate the likelihood it is necessary to marginalize over the latent variables which requires high-dimensional integration and makes the maximum likelihood approach intractable. Approximate inference techniques such as variational inference (Section 2.1.4) (Blei et al., 2017) approximate the probability distributions involved in deep latent variable models and find such representations without the need for expensive high-dimensional integration techniques such as Markov-Chain-Monte-Carlo approaches. I will now review the family of generative models whose parameters θ that can be efficiently estimated using the maximum likelihood approach (Goodfellow, 2016).

2.3.3 Explicit Density

Generative models that represent the density explicitly allow to draw new samples $\mathbf{m} \sim p_\theta(\mathbf{m})$. We can further distinguish between two types of *explicit density* representations; *tractable* and *approximate* density methods.

Tractable Density

Tractable density methods represent the density by tractable computations. *Autoregressive* methods use a factorization of the density into a sequence of conditional distributions where each random variable is conditional on all previously occurring random variables

$$p(\mathbf{x}) = \prod_{i=1}^{n} p(x_i | x_1, ..., x_{i-1}) \tag{2.20}$$

Autoregressive techniques have been developed for sampling generative models of continuous and categorical data such as images and audio. Recurrent neural networks (PixelRNN Oord et al., (2016b)) and convolutional neural networks (Oord et al., 2016a; Oord et al., 2016c) lend themselves specifically to represent the sequential ordering associated with the conditional distributions in autoregressive techniques. While autoregressive techniques have the advantage to provide explicit values of the density and all conditional distributions, their sequential ordering which requires sequential generation makes them comparatively slow to learn the parameters of their neural network representations as well as at creating new realizations.

In the case of deep latent variable models where the dimension of the latent space is the same as that of the image space it is possible to formulate a generative model that allows a tractable explicit representation of the density through a *change of variables*

$$p(\mathbf{m}) = p(\mathbf{z}) \left| \det \left(\frac{\partial g_\theta(\mathbf{z})}{\partial z^T} \right) \right|^{-1} \tag{2.21}$$

where g_θ is a function that maps from the latent to the image space and $\frac{\partial g_\theta(\mathbf{z})}{\partial z^T}$ is the Jacobian of g_θ with respect to $\mathbf{z}$. If the transformation g_θ is a bijective function we can learn the parameters θ using a maximum likelihood approach (Dinh et al., 2016). For high-dimensional distributions such as images computing the Jacobian would be very computational demanding and therefore specialized bijective transformations g_θ are required that allow efficient computation of the transformation and the Jacobian. A number of these bijective transformations can be coupled together forming a *self-normalizing flow* (Rezende and Mohamed, 2015). Dinh et al., (2016) introduce

specialized coupling layers that preserve invertibility of the transformation $g(\cdot)$ allowing for complex and high-dimensional distributions such as natural images to be represented by transformation from a simple multi-Gaussian prior distribution $p(\mathbf{z})$. Kingma and Dhariwal, (2018) have shown that self-normalizing flows can be trained using a maximum likelihood approach and introduce invertible 1×1 convolution operations to learn density representations of high-dimensional datasets such as high-resolution images.

Approximate Density

I have introduced a number of approaches that allow explicit representation learning using a maximum likelihood approach. A different approach to the methods such as autoregressive techniques or self-normalizing flows of finding a tractable representation of the density is to find approximate methods to obtain the parameters of density models where the evaluation of the likelihood is intractable as is the case for deep latent variable models (Eq. 2.19). Kingma and Welling, (2013) use the factorization of the joint distribution of a deep latent variable model

$$p_\theta(\mathbf{m}, \mathbf{z}) = p_\theta(\mathbf{m}|\mathbf{z})p_\theta(\mathbf{z}) \tag{2.22}$$

with a simple multi-Gaussian prior distribution $p_\theta(\mathbf{z})$ of the latent variables. They define an *encoder* distribution $q_\phi(\mathbf{z}|\mathbf{m})$ parameterized by a neural network with parameters ϕ that approximates the intractable true posterior

$$q_\phi(\mathbf{z}|\mathbf{m}) \approx p_\theta(\mathbf{z}|\mathbf{m}) \tag{2.23}$$

and a *decoder* network $p_\theta(\mathbf{m}|\mathbf{z})$ that represents the conditional distribution of the images given the latent variables. They decompose the log-likelihood or model evidence $\log p_\theta(\mathbf{m})$ as

$$\log p_\theta(\mathbf{m}) = \mathcal{L}_{\theta,\phi}(\mathbf{m}) + D_{KL}(q_\phi(\mathbf{z}|\mathbf{m})\|p_\theta(\mathbf{z}|\mathbf{m})) \tag{2.24}$$

where $\mathcal{L}_{\theta,\phi}(\mathbf{m})$ is the evidence lower bound (ELBO)

$$\mathcal{L}_{\theta,\phi}(\mathbf{m}) = \mathrm{E}_{q_\phi(\mathbf{z}|\mathbf{m})}\left[p_\theta(\mathbf{m}, \mathbf{z}) - q_\phi(\mathbf{z}|\mathbf{m})\right] \tag{2.25}$$

The ELBO is a lower bound due to the KL-Divergence (Eq. 2.14) in Eq. 2.24 being positive by definition. By leveraging a reparameterization of the prior distribution of the latent variables the so-called *Variational Autoencoder* (VAE) by Kingma and Welling, (2013) model allows end-to-end learning of a generative and approximate inference network for intractable deep latent variable models. Due to the fully-differentiable nature of the VAE the parameters of the encoder and decoder network can be learned by leveraging stochastic gradient descent techniques on noisy estimates of the gradient of the ELBO using mini-batches of data, making the VAE a scalable approach suited for approximate inference for very large datasets. After training, the marginal likelihood of new data can be estimated using an importance sampling technique (Rezende and Mohamed, 2015).

2.3.4 Implicit Density

I will now introduce deep generative models that do not require an explicit representation of the density $p_\theta(\mathbf{m})$. Implicit density methods aim to learn a representation of the underlying probability distribution defined implicitly by a set of training examples.

Generative Adversarial Networks

Recently, a new approach that casts learning probability distributions of data as a two-player game called Generative Adversarial Networks (GANs) has been proposed (Goodfellow et al., 2014). Formally, GANs are deep latent variable models that learn an implicit representation of the probability distribution defined by a training dataset. Let us assume here that the training dataset is composed of images.

GANs consist of two parametric functions: a discriminator $D_\omega(\mathbf{m})$ and a generator $G_\theta(\mathbf{z})$. The discriminator evaluates samples from the set of training examples $\mathbf{m} \sim p_{data}(\mathbf{m})$ and realizations created by the generator by applying the parametric generator function $G_\theta(\mathbf{z})$ to a sample drawn from a prior distribution of the latent variables $\mathbf{z} \sim p(\mathbf{z})$. Common choices for the prior distribution of the latent variables $p(\mathbf{z})$ are multi-Gaussian or uniform distributions. The generator $G_\theta(\mathbf{z})$ maps random samples from the latent space into the space of images. The discriminator's role is to assign a probability that a random sample $\mathbf{m}$ is from the data distribution $p_{data}(\mathbf{m})$. The discriminator tries to label each sample correctly, while the generator tries to "fool" the discriminator into labeling the fake images as part of the true data distribution and therefore achieving $D_\omega(G_\theta(\mathbf{z}))$ close to one. The objectives of the generator and discriminator counteract each other and each player acts as an adversary to the other's objective - hence the name Generative *Adversarial* Networks.

More formally we can define the loss i.e. the cost function for GANs as a minimization-maximization problem

$$\min_\theta \left[\max_\omega \{ \mathbb{E}_{\mathbf{m}\sim p_{data}(\mathbf{m})}[\log(D_\omega(\mathbf{m}))] + \mathbb{E}_{\mathbf{z}\sim p_\mathbf{z}(\mathbf{z})}[log(1 - D_\omega(G_\theta(\mathbf{z})))] \} \right] \tag{2.26}$$

Solutions to this optimization problem have been shown to be Nash equilibria, where each player achieves a local optimum of their loss function with respect to their parameters (Goodfellow, 2016). In practice we represent G_θ and D_ω by convolutional neural networks where parameters θ and ω are optimized using stochastic gradient descent. Training is performed in an iterative alternating manner: first the discriminator is trained to maximize

$$J^{(D_\omega)} = \mathbb{E}_{\mathbf{m}\sim p_{data}(\mathbf{m})}[\log(D_\omega(\mathbf{m}))] + \mathbb{E}_{\mathbf{z}\sim p_\mathbf{z}(\mathbf{z})}[log(1 - D_\omega(G_\theta(\mathbf{z})))] \tag{2.27}$$

while the parameters of the generator are kept constant. This improves the ability of the discriminator to distinguish between real images from the training set and those created by the generator.

In a subsequent step we generate synthetic samples by applying the generator $G_\theta(\mathbf{z})$ to samples from the prior $\mathbf{z} \sim p(\mathbf{z})$ and train the generator to minimize

$$J^{(G)} = \mathbb{E}_{\mathbf{z}\sim p_\mathbf{z}}[log(1 - D_\omega(G_\theta(\mathbf{z})))] \tag{2.28}$$

while keeping the parameters of the discriminator fixed. In the ideal case convergence is reached when the value of the discriminator is $\frac{1}{2}$ as this means that the discriminator cannot distinguish between data from the training set and samples created by the generator and therefore $p_g(\mathbf{m}) = p_{data}(\mathbf{m})$.

The two-player game defined by the training procedure for GANs shows highly unstable behavior making training of GANs for general high-dimensional datasets such as images difficult in practice, where large number of trial and error runs are required to find an optimal set of hyperparameters to allow stable training.

2.4 Research Questions

At the beginning of this chapter I have reviewed a number of inversion techniques and methods to generate realizations of geological features at the pore- and reservoir-scale. I have shown that for Bayesian inversion techniques it is necessary to find good representations of the prior distribution of the model parameters $p(\mathbf{m})$ that allow us to obtain representations that honor observed data while capturing the relevant features of the prior distribution. At the pore-scale having statistical or generative models of the data allows us to quantify the distribution of flow and transport properties under changing structural features and boundary conditions allowing us to build representative Digital Rock Physics (Blunt, 2017) models calibrated by laboratory measurements. I have then introduced parametric generative models where neural networks are used to represent the probability distribution of a given dataset of observed examples and I have shown how the parameters of these neural networks can be learned using a maximum likelihood approach. Generative models based on explicit density representations allow us to draw new samples and compute the likelihood of a dataset given the model, while implicit representations allow us only to generate new high-fidelity realizations. Given the recent advances in deep generative modeling we hypothesize that it may be possible to find new representations of prior distributions for Bayesian inversion in subsurface applications and to create data-driven stochastic reconstruction methods for digital rock applications. From this we formulate the following research questions and contributions.

Research Question 1: *How can we build generative models of porous media at the pore-scale that allow fast sampling while reproducing statistical and physical properties relevant for flow?*

In Chapter 3 I propose using generative adversarial networks to learn a probabilistic model of micro-CT image data that allow fast sampling of large three-dimensional stochastic realizations of porous media at the pore-scale. I represent the generative model by a deep three-dimensional convolutional neural network trained on subdomains or patches of large segmented three-dimensional micro-CT images of synthetic materials and porous rocks. Using a fully-convolutional network allows me to train on a large number of small training images while being able to generate large three-dimensional realizations after training. I show that the obtained samples honor the statistical properties of the porous media and the flow-properties such as effective permeability.

In Chapter 4 I extend the analysis of GANs for stochastic image reconstruction to three-dimensional greyscale micro-CT image of the oolitic Ketton limestone. Due to the diverse nature of the features in the unsegmented micro-CT image this represents a challenging dataset. Furthermore, I propose an evaluation of the generated images based on a so-called Minkowski functionals - a set of morphological properties that are inherently linked to key flow and transport properties of porous media such as the Euler characteristic which is closely linked to the connectivity of porous media. Finally, I investigate the representations learned in the deep convolutional neural network and highlight some of the challenges inherent in training GANs to represent high-dimensional datasets such as their ability to represent multi-modal data distributions, and I address the issue of training instabilities.

Research Question 2: *Can we condition deep generative models trained to represent three-dimensional geological features to generate stochastic realizations that honor lower-dimensional data such as two-dimensional thin-sections or one-dimensional borehole data?*

In Chapter 5 I propose conditioning of stochastic samples created by a trained GAN generator using a gradient-based search to find three-dimensional pore-scale images and reservoir-scale models in the latent space of the GAN that honor the available lower-dimensional data. This approach builds upon the idea of image inpainting (Yeh et al., 2016) where I constrain the output of the generator function to minimize the mismatch between existing lower-dimensional data and the output of the generator function. By leveraging the discriminator function I additionally regularize the generator to create samples that show features of the underlying data distribution implicitly defined by the set of training images used to train the GAN. I demonstrate that starting from various initial samples from the prior distribution of the latent variables we obtain stochastic realizations of three-dimensional porous media and reservoir models.

Due to the uncertainty and sparsity of observed geophysical data, many inverse problems in the geosciences can be formulated in a Bayesian inference setting. Given a prior distribution of model parameters the goal is to find the posterior distribution i.e. the conditional distribution of the model parameters given observed data. The probabilistic interpretation of GANs as models of high-dimensional data distributions leads to the final research question:

Research Question 3: *Can GANs be used in a Bayesian setting to represent prior distributions in ill-posed large-scale subsurface inverse problems?*

In Chapter 6 I propose the use of GANs as probabilistic models to represent the prior distribution of possible geological and petrophysical models. I obtain solutions to the ill-posed non-linear full-waveform inversion problem of acoustic seismic data by performing Bayesian inference on the latent variables that govern the output of the GAN generator using an approximate Markov-Chain-Monte-Carlo algorithm. This allows samples from the posterior distribution of reservoir models to be obtained by conditioning to observed seismic data. I use an approximate Metropolis-Adjusted Langevin Algorithm (MALA) that leverages gradients of the posterior to perform sampling of this posterior. Gradients of the posterior are obtained by combining the GAN generator with a finite-difference solution to the wave equation that allows gradients of the likelihood to be obtained using the adjoint-state method. Gradients at grid-block-scale are then propagated to the latent variables by neural network backpropagation and use of the chain-rule.

In Chapter 7 I use GANs to represent the space of possible geological models for the reservoir history matching problem. An ensemble of history-matched reservoir models is obtained by starting from a number of random initial sets of latent variables and performing maximum a posteriori inference on each of them given observed two-phase flow rate and pressure data.

Chapter 3

Reconstruction of Three-Dimensional Porous Media using Generative Adversarial Neural Networks

3.1 Abstract

To evaluate the variability of multiphase flow properties of porous media at the pore-scale, it is necessary to acquire a number of representative samples of the void-solid structure. While modern X-ray computer tomography has made it possible to extract three-dimensional images of the pore-space, assessment of the variability in the inherent material properties is often experimentally not feasible. We present a novel method to reconstruct the solid-void structure of porous media by applying a generative neural network that allows an implicit description of the probability distribution represented by three-dimensional image datasets. We show, by using an adversarial learning approach for neural networks, that this method of unsupervised learning is able to generate representative samples of porous media that honor their statistics. We successfully compare measures of pore morphology, such as the Euler characteristic, two-point statistics and directional single-phase permeability of synthetic realizations with the calculated properties of a bead pack, Berea sandstone, and Ketton limestone. Results show that GANs can be used to reconstruct high-resolution three-dimensional images of porous media at different scales that are representative of the morphology of the images used to train the neural network. The fully convolutional nature of the trained neural network allows the generation of large samples while maintaining computational efficiency. Compared to classical stochastic methods of image reconstruction, the implicit representation of the learned data distribution can be stored and reused to generate multiple realizations of the pore structure very rapidly.

3.2 Introduction

3.2.1 Image Reconstruction

The reconstruction and the evaluation of the material properties of porous media plays a key role across many engineering disciplines. Many physical processes such as

the movement of multiple phases of fluids through sedimentary rocks are controlled by individual pores at the micron and sub-micron scale (Blunt, 2017).

In carbon capture and sequestration (CCS), the long term storage behavior is controlled by the physical and chemical interaction of super-critical CO_2 with the reservoir brine, as well as the spatial distribution and connectivity of minerals in the pore-space (Juanes et al., 2006; Kang et al., 2010). The variability of the controlling properties such as the permeability of the host rock is determined by repeated experiments or numerical modeling of these processes.

Using modern computer tomographic methods, it is possible to observe porous materials and evaluate their material properties at the micrometer scale (micro-CT) under static and transient conditions at high pressures and temperatures in near real time. Performing micro-CT imaging of porous media requires specialized, expensive equipment and in the case of CCS, only a single image of the investigated rock type is typically acquired.

To evaluate the variability associated with the geometrical and mineralogical morphology of the pore-space, numerous physical experiments using the same rock type would have to be performed to obtain a distribution over larger volumes. Due to time and cost limitations inherent with the experimental acquisition of high-resolution images, this is often deemed unfeasible. Material properties governing the single and multiphase flow behavior of porous media can be estimated from numerical solution of partial differential equations at a scale larger than that of a representative elementary volume (REV) and verified by experimental results (Mostaghimi et al., 2013).

Many sedimentary rocks consist of granular siliciclastic or carbonate materials. Boolean models use this fundamental characteristic of natural granular materials to emulate the shape of the arising pore-space, due to an underlying random process that controls the distribution of the individual grains (Matheron, 1975; Serra, 1980). While for the classical Boolean model, the centers of the grains are uniformly distributed in space and grains can arbitrarily overlap, more complicated models with rigid hard sphere grains and more complex grain interaction functions have been developed (Matheron, 1971; Arns et al., 2009; Rikvold and Stell, 1985; Torquato, 2013). The framework of Boolean models also allows extension beyond spherical particles and enables derivation of the properties of material models as a function of the parameters of the underlying random process (Bretheau and Jeulin, 1989; Jeulin, 2000; Lin and Cohen, 1982; Yeong and Torquato, 1998).

In clastic rocks, the arrangement of individual grains occurs due to the transport of material from a high energy source to a low energy sink. Process models, where depositional mechanisms are simulated, have been shown to reproduce realistic granular reconstructions capturing the pore-space morphology of granular sedimentary rocks (Øren and Bakke, 2003).

Spatial probabilistic models such as truncated Gaussian processes or sequential indicator simulation have been widely applied in the geosciences to model the spatial distribution of materials (Pyrcz and Deutsch, 2014). Many of these methods rely on two-point probability functions as a measure of spatial variability, whereas recent methods in geostatistics use training images as a basis for sample reconstruction (Caers and Zhang, 2004; Mariethoz et al., 2010; Meerschman et al., 2013). These images are usually assumed to exhibit stationarity of the probability distribution of the properties of interest and rely on higher order multiple-point statistics (MPS) to

reconstruct stochastic random media.

With MPS, the probability distributions are represented by training images and are sampled using a limited multi-scale neighborhood that captures the variation on a large scale, as well as fine structural details on smaller scales (Tahmasebi et al., 2014). MPS-based methods have been used in two and three-dimensional conditional simulation of spatial properties in reservoir-scale earth modeling applications (Comunian et al., 2011). The computational complexity of these methods is highly dependent on individual algorithms as well as the size of the domains used to sample from the training images (Mariethoz and Caers, 2014). Parallelized versions have been developed, reducing the computational time required to perform reconstruction using MPS (Straubhaar et al., 2011; Huang et al., 2013).

Three-dimensional porous media have been reconstructed using a modified multiple-point statistics approach based on two-dimensional images of porous media (Okabe and Blunt, 2004; Okabe and Blunt, 2005; Okabe and Blunt, 2007). Stochastic methods based on simulated annealing allow the incorporation of arbitrary cost functions of statistical and morphological properties used in unconditional three-dimensional image reconstruction (Smith et al., 1983; Svergun, 1999). Recent advances have reduced the computational runtime of simulated-annealing-based methods for reconstruction of porous media, to the order of tens of hours per realization at the scale of 300^3 voxels (Pant, 2016).

In the following section, we introduce a recently developed class of unsupervised machine learning methods called generative adversarial networks (GAN), as discussed in Chapter 2.3.4 and applied here to generate pore-space images, that allow simulation of probability distributions given a set of training data (Goodfellow et al., 2014). Volumetric generative adversarial networks have previously been applied to low-resolution three-dimensional CAD model synthesis, and practical applications of 3D-GANs are few compared to their two-dimensional counterparts (Wu et al., 2016). Integration of multi-resolution datasets incorporating image data across a number of length scales is possible in the GAN framework by using a Laplacian pyramid approach such as LapGAN (Denton et al., 2015).

We investigate the applicability of GANs to model three-dimensional textures of rocks based on three-dimensional binary representations of porous media acquired at the micrometer scale. We compare statistical, morphological and transport properties of the simulated images with those of the training images. We evaluate the single-phase directional permeability to show that the synthetic realizations sampled from the learned representation of the input data can capture single-phase flow properties of sedimentary rocks.

Training of these neural networks involves finding a set of hyperparameters that lead to stable training (Goodfellow, 2016). While this training can take on the order of tens of hours, the sampling of large volumetric domains occurs on the order of seconds on the current generation of graphical processing units (GPU). We show that in favorable cases convolutional neural networks incorporated in the GAN framework allow the generation of synthetic reconstructions of porous media that exceed the dimensions of their training images. Contrary to most existing simulation techniques the set of parameters used to generate synthetic realizations can be stored once trained, allowing rapid generation of new samples to assess the variability of material properties.

While we apply GANs to a set of micro-CT images of porous media, the method

can readily be applied to volumetric images of porous media obtained from other three-dimensional microscopy instruments such as nano or medical-CT instruments. We discuss the challenges involved in training GANs for stochastic image reconstruction of porous media, as compared to other stochastic image reconstruction methods and we evaluate the computational efficiency of GAN-based image reconstruction. Finally, we provide empirical guidelines on the requirements of the input dataset to allow successful training of GANs on large three-dimensional voxel representations of natural porous media.

All data used in this study, as is the case for the rest of this thesis, are available in the public domain and we have made the code used for training, as well as example pre-trained models, available as additional supporting material [1]. A public dataset of high-resolution micro-CT images made available by the Imperial College Pore-Scale Modelling Group [2], of a spherical beadpack, Berea sandstone, and oolitic Ketton limestone will serve as benchmark cases to study the application of GANs to three-dimensional stochastic image reconstruction.

3.3 Generative Adversarial Networks

In the following section, we present generative adversarial networks (GAN) for three-dimensional image generation developing the discussion presented in Chapter 2 to focus on the generation of three-dimensional pore-space images. Generative neural networks have been developed in the context of deep learning by Goodfellow et al., (2014) as a methodology to learn a representation of a high-dimensional probability distribution from a given dataset. In the context of image reconstruction, we refer to this dataset as a set of training images that are representative samples of the probability distribution underlying the image space.

GANs learn an implicit representation of the probability density as opposed to explicit density models. The main drawback of explicit density models is their computational cost which grows with the dimensionality of the samples and requires sequential simulation of each voxel. For high-dimensional samples such as volumetric image data, the computational cost is $O(N)$ where N represents the number of voxels in the domain of interest and can easily exceed 10^9 voxels for modern high-resolution micro-CT image data. Using any of these methods would make it intractable to generate a large number of very large samples. GANs have been designed to perform fast sampling from the learned density representation and allow full parallel generation, making them an ideal candidate to generate large volumetric images (Goodfellow, 2016).

GANs consist of two differentiable functions: a discriminator D and a generator G. The discriminator receives samples of the "real" dataset (Label 1) $x \sim p_{data}$ and "fake" samples $G(\mathbf{z})$ (Label 0) created by the generator from the hidden latent space Z (Fig. 3.1). The latent space Z is composed of independent real random variables, typically normally or uniformly distributed, that represent the random input to the generator G. The generator G maps random variables from the latent space into the space of images. The discriminator's role is to assign a probability that a

[1] `https://github.com/LukasMosser/PorousMediaGan`

[2] `http://www.imperial.ac.uk/earth-science/research/research-groups/perm/research/pore-scale-modelling/micro-ct-images-and-networks/`

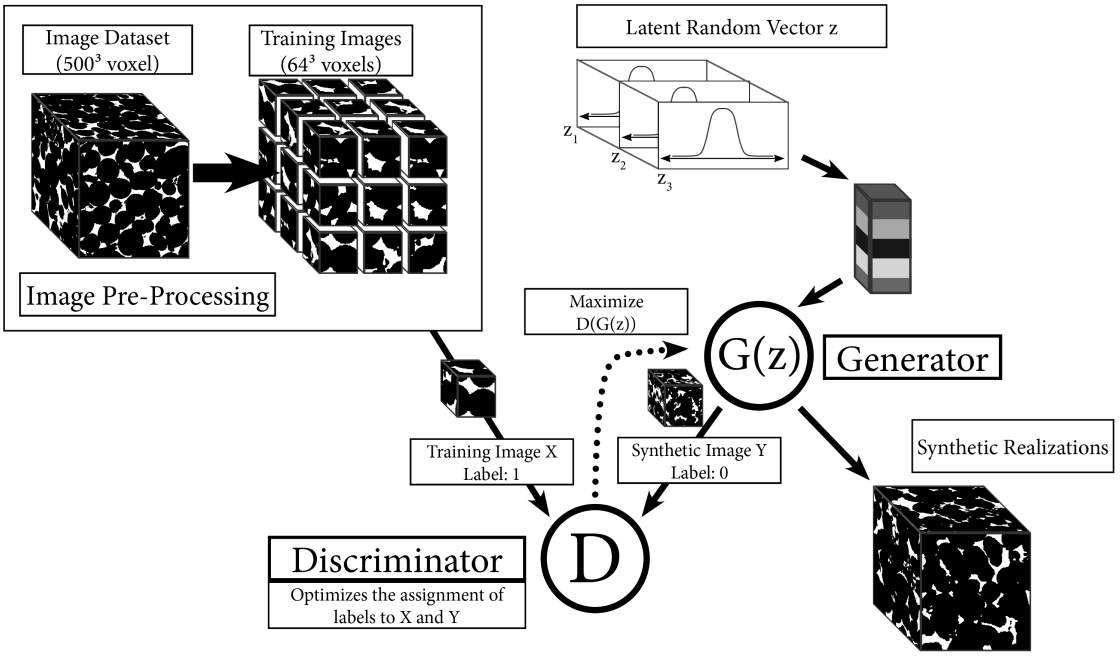

Figure 3.1: Overview of the GAN training process. Segmented volumetric images are split into 64^3 or 128^3 voxel training images. The generator G is a function that is applied to a sample from a latent random space Z and creates a synthetic realization. We assume that samples drawn from the hidden latent space Z are normally distributed (Section 3.3). The discriminator's role is to determine whether a sample is part of the training image dataset (Label 1) or from the generator (Label 0). The misclassification error is computed as a binary cross-entropy criterion and the error back-propagated to improve the discriminator's ability to distinguish real from "fake" images. Then the generator is updated to improve the quality of the produced samples and "fool" the discriminator. When sufficient image quality is obtained, training is stopped, and the discriminator may be discarded. The generator can now be used to create new samples. By providing larger latent vectors than used initially for training, larger output images can be produced.

random sample is from the "real" data distribution p_{data}. The discriminator tries to label each sample correctly, while the generator tries to "fool" the discriminator into labeling the fake images as part of the true data distribution and therefore achieving $D(G(\mathbf{z}))$ close to one.

More formally we can define the loss i.e. the cost function for GANs as a minimization-maximization problem

$$\begin{aligned} \min_G \max_D \{ & \mathbb{E}_{\mathbf{x}\sim p_{data}(\mathbf{x})}[log(D(\mathbf{x}))] \\ & + \mathbb{E}_{\mathbf{z}\sim p_{\mathbf{z}}(\mathbf{z})}[log(1 - D(G(\mathbf{z})))]\} \end{aligned} \tag{3.1}$$

Solutions to this optimization problem have been shown to be Nash equilibria, where each player achieves a local minimum of their loss function with respect to their parameters (Goodfellow, 2016).

In practice we represent G and D by convolutional neural networks that are trained by a gradient descent-based optimization method. Training is performed in

two steps: First the discriminator is trained to maximize

$$J^{(D)} = \mathbb{E}_{\mathbf{x}\sim p_{data}(\mathbf{x})}[log(D(\mathbf{x}))] + \mathbb{E}_{\mathbf{z}\sim p_{\mathbf{z}}(\mathbf{z})}[log(1 - D(G(\mathbf{z})))] \tag{3.2}$$

while the parameters of the generator are fixed. This improves the ability of the discriminator to distinguish between real and fake images. In a subsequent step we generate synthetic samples $G(\mathbf{z})$ by drawing samples $\mathbf{z}$ from an N-dimensional normal distributed latent space and train the generator to minimize

$$J^{(G)} = \mathbb{E}_{\mathbf{z}\sim p_{\mathbf{z}}}[log(1 - D(G(\mathbf{z})))] \tag{3.3}$$

while keeping the discriminator fixed.

By minimizing Eq. 3.3 the generator tries to "fool" the discriminator into believing that the samples $G(\mathbf{z})$ are real data samples. In this way the generator learns to represent a distribution $p_g(\mathbf{x})$ that is as close as possible to the real data distribution $p_{data}(\mathbf{x})$. When convergence is reached $p_g(\mathbf{x}) = p_{data}(\mathbf{x})$ and the value of the discriminator becomes $\frac{1}{2}$ as it cannot distinguish between the two anymore.

Initially, the discriminator D outperforms the generator significantly making the gradient used to train the generator close to zero. Therefore, instead of minimizing $log(1-D(G(\mathbf{z}))$ for the generator, it is helpful to maximize $log(D(G(\mathbf{z}))$ (Goodfellow, 2016). GANs show highly unstable behavior during training and a large number of trial and error runs are required to find an optimal set of hyperparameters that allow stable training. A number of heuristics have been published which have been shown to stabilize GAN training, such as one-sided label smoothing and adding white noise to the input layer of the discriminator (Salimans et al., 2016; Sønderby et al., 2016). We provide a more detailed overview of the neural networks used in this study in Section 3.4.2 and later provide suggestions on how to facilitate efficient training (Section 3.7) for volumetric image datasets of porous media.

3.4 Methodology

In the following section we outline the criteria used to evaluate the quality of simulations based on the training image datasets. We treat all images under the assumption of stationarity and the existence of a representative elementary volume.

3.4.1 Evaluation Criteria

Two-Point Statistics

We characterize the second order structure of the porous media by calculating the two-point probability function of the pore phase. By assuming stationarity, this function is equivalent to the non-centered covariance (Matheron, 1971):

$$S_2(\mathbf{r}) = \mathbf{P}(\mathbf{x} \in P, \mathbf{x} + \mathbf{r} \in P) \; for \; \mathbf{x}, \mathbf{r} \in \mathbb{R}^d \tag{3.4}$$

which is the probability $\mathbf{P}$ that two points $\mathbf{x}$ and $\mathbf{x} + \mathbf{r}$, separated by the lag vector $\mathbf{r}$, are located in the pore phase P. At the origin, $S_2(0)$ is equal to the porosity ϕ. S_2 stabilizes around ϕ^2 as $r \to \infty$ (Fig. 3.2). Due to the anisotropic nature of many

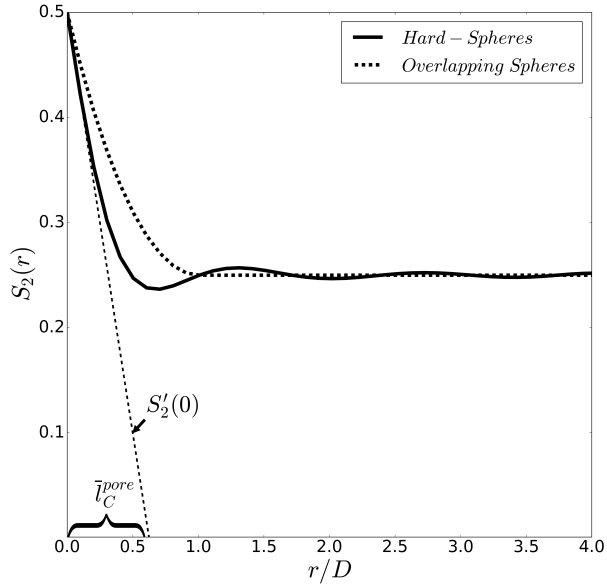

Figure 3.2: Comparison of $S_2(r)$ for a Boolean model and a packing of hard spheres at a porosity $\phi = 0.5$. S_2 exhibits exponential decay for the Boolean model, whereas a dampened oscillation is characteristic for packings of spheres. The mean chord length can be found at the intersection of the slope of S_2 at the origin with the x-axis (Eq. 3.6).

porous media, we compute $S_2(r)$ along the three Cartesian directions, as well as the radial average of $S_2(r)$.

It is a well known result that the specific surface area S_V of a porous medium can be expressed as a function of S_2 (Debye et al., 1957). In the case of an isotropic porous medium and in three-dimensions S_V is related to S_2 by

$$S_V = -4S_2'(0) \tag{3.5}$$

where $S_2'(0)$ is the derivative of $S_2(r)$ at the origin.

Furthermore, the average chord length within the pore and the grain phases are (Torquato, 2013)

$$\bar{l}_c^{pore} = \frac{\phi}{S_2'(0)} \tag{3.6a}$$

$$\bar{l}_c^{grain} = \frac{1-\phi}{S_2'(0)} \tag{3.6b}$$

which for the pore phase can be readily found from the intersection of the slope of $S_2(r)$ with the x-axis (Fig. 3.2).

In favorable cases, it is possible to find analytical expressions of $S_2(r)$ from the spatial distribution and geometry of the grains. A Boolean model of overlapping spherical grains of uniform spatial distribution exhibits an exponential decay of the covariance until the lag distance is equal to the diameter of the grains where it becomes zero (Matheron, 1971). For porous media that can be well described by a Boolean model, we can estimate the size of the elementary Boolean grain from the decay of S_2. Semi-analytical expressions for more complex models such as for a packing of hard spheres have been developed (Torquato and Lado, 1985). Models

of $S_2(r)$ for spherical packings exhibit a dampened oscillation. The shape of the estimated covariance, therefore, allows us to obtain information on the structure of the porous medium (Fig. 3.2). The covariance $S_2(r)$ was estimated for the training images and the stochastic reconstructions generated by the trained GAN model. For each GAN model, we evaluate the non-centered covariance S_2 as well as the specific surface area S_V (Eq. 3.5) and compare these to the values obtained from the original training images. In our discussion on the required training image sizes (Section 3.7), we will use the average chord length and the specific surface area as possible indicators of the necessary training image size.

Morphological Measures

It has been shown that flow properties at the pore-scale can be related to morphological characteristics of the void-solid interface of a porous medium (Scholz et al., 2015). Hadwiger's theorem states that any continuous rigid motion invariant valuation on compact-convex subsets of $\mathbb{R}^d$ can be described by a linear combination of $d+1$ independent parameters characterizing the body. In three dimensions we can, therefore, define four so-called Minkowski functionals that characterize the topology of a three-dimensional object. We compute estimates of three Minkowski functionals; the porosity ϕ, the specific surface area S_V and the Euler characteristic χ_V corresponding to the zero, first and 3$^{\text{rd}}$ order functionals. We compute the densities of the Minkowski functionals by dividing by the volume V.

The Minkowski functional of order zero is the porosity, defined as the ratio of volume of the void space to the bulk volume of the sample

$$\phi = \frac{V_{pore}}{V} \tag{3.7}$$

and is, therefore, a measure of the ability of a porous medium to store fluids.

The Minkowski functional of rank one is the specific surface area S_V.

$$S_V = \frac{1}{V} \int dS \tag{3.8}$$

where integration occurs over the void-solid interface S. The specific surface area S_V has dimensions of $\frac{1}{length}$ and its inverse allows us to define a characteristic pore size.

The specific Euler characteristic is closely related to the order three Minkowski functional and represents a dimensionless quantity defined as

$$\chi_v = \frac{1}{4\pi V} \int \frac{1}{r_1 r_2} dS \tag{3.9}$$

where r_1 and r_2 are the principal radii of curvature of the void-solid interface. To compute χ_V we do not directly evaluate the integral in Eq. 3.9 but instead make use of a relationship for the Euler characteristic of arbitrary polyhedra,

$$\chi = V - E + F - O \tag{3.10}$$

where V is the number of vertices, E the number of edges, F the number of faces and O the number of objects (Blunt, 2017). This expression is the basis for efficient algorithms to compute Minkowski functionals of arbitrary geometric bodies represented as volumetric voxelized domains (Lang et al., 2001). To compute these

Table 3.1: Neural network configurations and hyperparameters used to train on voxelized image subsets.

	Training Image Dataset		
	Beadpack	Berea	Ketton
Training Image Size	128^3 voxels	64^3 voxels	64^3 voxels
Latent Space $\mathbf{z}$ Dimension	100	512	100
Generator Filters N_G	64	64	64
Discriminator Filters N_D	8	16	16
Optimizer	Generator + Discriminator: Adam		
Learning Rate / Momentum	2×10^{-4} / 0.5	2×10^{-4} / 0.5	2×10^{-4} / 0.5
Stabilization	White Noise ($\sigma = 0.1$)	Label Smoothing ($\epsilon = 0.1$)	White Noise ($\sigma = 0.1$)

three Minkowski functionals we have used the open-source image morphological software library MorphoLibJ (Legland et al., 2016). While the porosity expresses the ability to store fluids in a porous medium, adsorption and dissolution processes are controlled by the specific surface area. The Euler characteristic allows the connectivity of the porous medium to be characterized, which is a critical component in the ability of fluids to flow. Reconstructions of porous media should therefore closely match the observed Minkowski functionals to represent the behavior of relevant physical processes at the pore-scale. The direct computation of the specific surface area S_V and porosity ϕ from images allows us to perform a comparison with the values obtained from estimates obtained by computing the empirical non-centered covariance $S_2(r)$ (Eq. 3.5).

Single-Phase Permeability

To evaluate the single-phase permeability of the porous media and their generated synthetic reconstructions we solve the Stokes equations for slow, incompressible flow assuming small inertial forces.

$$\nabla \cdot \mathbf{v} = 0 \tag{3.11a}$$

$$\mu \nabla^2 \mathbf{v} = \nabla p \tag{3.11b}$$

The Stokes equations are solved on the domain that is connected to the fluid inlet and outlet. This allows us to define an effective porosity where only the fraction of the pore-space that also contributes to flow is considered

$$\phi_{eff} = \frac{V_{flow}}{V} \tag{3.12}$$

A finite difference method to solve Eqs. 3.11a-3.11b on pore-space representations has been implemented as a parallel flow solver, in the free open source numerical framework OpenFOAM (Weller et al., 1998; Mostaghimi et al., 2013).

3.4.2 Neural Network Architecture

The neural network architecture used for the three-dimensional image reconstruction corresponds to a volumetric version of the DCGAN network (Radford et al., 2015). The network consists of two independent fully convolutional neural networks, the generator G and the discriminator D. Upsampling from the input latent vector $\mathbf{z}$ is performed by volumetric transposed convolution, followed by batch normalization

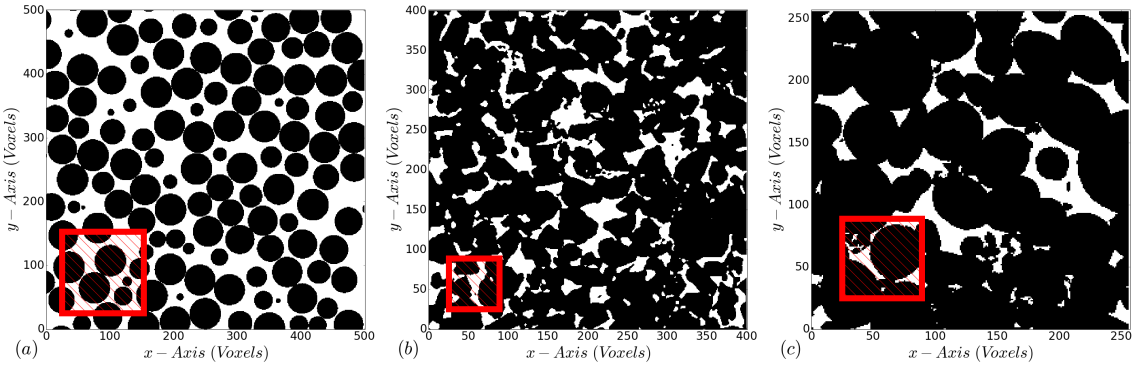

Figure 3.3: Cross-sections of the three image datasets. The bordered regions indicate the size of the training images extracted from the full dataset. The beadpack (a) consists of spheres of equal diameter ($d = 50$ $voxels$). The Berea sandstone (b) is an angular granular sandstone that shows traces of dispersed clay. The oolitic Ketton limestone (c) consists of ellipsoidal grains showing inter and intra-granular porosity. The voxel sizes are 3 μm for the bead pack and Berea sandstone and 15.2 μm for the Ketton sample.

and a rectified linear unit (ReLU) activation function in all layers except the last (Ioffe and Szegedy, 2015; Hinton et al., 2010). The discriminator D receives images sampled from the latent space by the generator $G(\mathbf{z})$ and images from the set of training images representing $p_{data}(\mathbf{x})$. Therefore, the size of the input layer of the discriminator corresponds to the dimensions of the input training images. The discriminator consists of volumetric convolution layers combined with LeakyReLU activation functions (Maas and Ng, 2013). The final convolutional layer of the discriminator is followed by a Tanh activation function. This combination of generator and discriminator neural network architectures has previously been applied to subsets of the Imagenet and CIFAR-10 datasets (Radford et al., 2015). The hyperparameters for the generator to be used in the optimization of the neural network architecture are the number of trainable convolutional filters in each layer of the neural network $N_{G,F}$, $N_{D,F}$ and the size of the latent vector $\mathbf{z}$.

The generator and discriminator are optimized using a gradient descent-based method where the parameters w are changed by taking k steps in the gradient

$$w^{k+1} = w^k - \alpha \nabla f(w^k) \tag{3.13}$$

where α is the learning rate. We have used the gradient descent-based optimizer ADAM for optimization of both the generator and discriminator (Kingma and Ba, 2014). GANs have been shown to exhibit unstable behavior during training. The addition of Gaussian noise to the input of the discriminator is an effective way to stabilize the training process (Sønderby et al., 2016). An additional stabilization measure called one-sided label smoothing, wherein the class label of 1 for real images is replaced by a new value of $1 - \epsilon$ has been empirically shown to improve training of GANs (Salimans et al., 2016). Both label smoothing and white noise addition to the input of the discriminator have been used in this study to stabilize the training based on the volumetric image datasets. Table 3.1 gives an overview of the neural network hyperparameters used for each evaluated sample, the hyperparameters and the stabilization measure used during training. Images generated by the GAN were post-processed using a 3^3 median filter to remove single-pixel noise. The resulting

images are grayscale images with all voxel values close to zero or one. To compare the resulting images to the binary training images, we segment the generated images using Otsu's method (Otsu, 1975).

3.5 Experimental Data

3.5.1 Image Data and Processing

To evaluate the applicability of GANs for reconstruction of natural porous media we use three previously acquired datasets. All images have been segmented into a three-dimensional binary voxel representation of the pore-space (white) and grain (black) (Fig. 3.3). We create a training database of images by extracting sub-volumes from the voxelized binary images. Ideally, these training images should represent independent domains, but due to the limited size of these images, we extract subsets that overlap. Training image sizes were chosen based on an estimate of the average grain size for each sample. To be able to match the covariance $S_2(r)$ (Eq. 3.4) and image morphological characteristics, training images larger than the structuring element were necessary. We discuss this requirement in more detail in the discussion of our results (Section 3.7). Due to computational limitations, training image sizes exceeding 128^3 voxels were not considered.

Beadpack

The beadpack is based on a real packing of equally sized grains in a disordered close packing (Finney, 1970). The image consists of 500^3 voxels with a size of 3 μm. The size of an individual sphere is 50 voxels. 1727 training images were extracted of size 128^3 voxels corresponding to a spacing of 32 voxels between them in the original image.

Berea

Berea sandstone is a fluvial sandstone of medium to fine grain size (Wentworth classification) (Pepper, 1954). The individual grains are bonded by clays. The sample analyzed in this study was acquired from an outcropping of the Berea sandstone in a quarry near Berea, Ohio. De Witt showed that the Berea sandstone was deposited in the early Carboniferous (354-323 Mya) (de Witt Jr, 1970).

The image of Berea sandstone consists of angular grains with no clay presence in the intergranular pore-space. The image has dimensions of 400^3 voxels with a voxel size of 3 μm. To capture the local interaction of grains we have extracted training images at 64^3 voxels which allows a number of grains to be present in one training image (Section 3.7). Due to the small image size of 400^3 voxels, subvolumes were extracted at a spacing of 16 voxels. In all, 10647 training images were used for the image reconstruction.

Ketton

The Ketton sample is an oolitic limestone of Jurassic age (201.3-145 Mya). The sample was acquired from a quarry of Lincolnshire limestone in the North-East of England. The oolites contained in the Lincolnshire formation are mainly non-ferroan

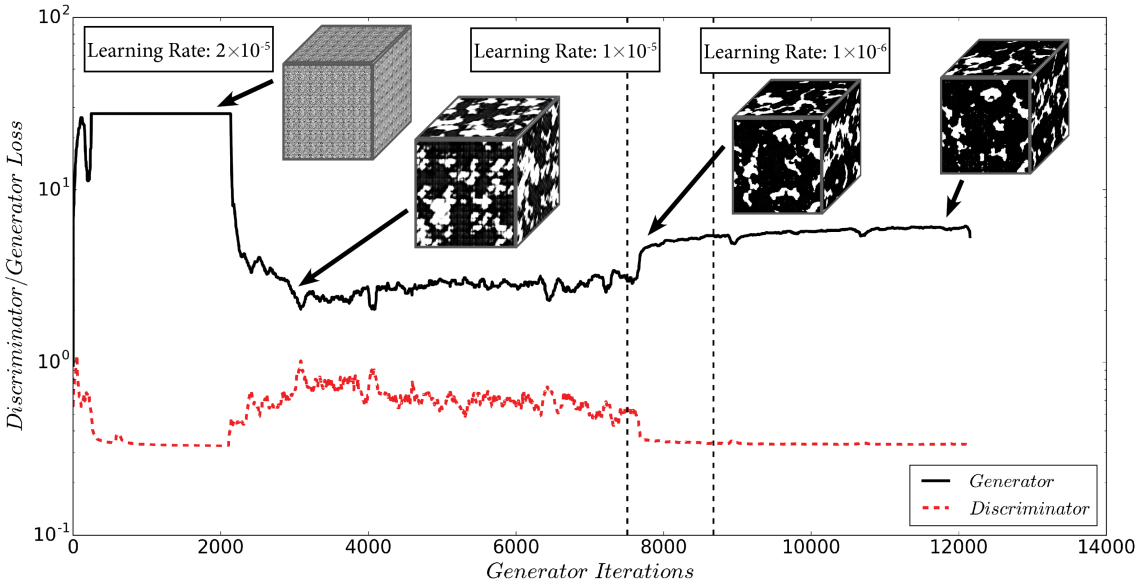

Figure 3.4: Value of the cost function i.e. loss of the discriminator and generator (Eqs. 3.2 - 3.3) for the GAN trained on the Berea sandstone. The samples shown were computed with the same random number seed, showing the evolution of a single realization during training. Initially, image quality is very low and random noise can be observed. After 2000 generator iterations a drop in the generator loss function is observed and coarse structures can be identified in the resulting sample. Loss functions in GAN models do not reflect improvement in image quality which can be observed from samples. Learning rates (Eq. 3.13) were reduced after sufficient image quality was reached and training stopped based on manual inspection of Minkowski functionals and two-point statistics.

calcite grains. The oolitic limestones of the Lincolnshire show a wide variety of cementation, ranging from uncemented oolite sands with no intergranular cement to heavily ferroan spar-cemented oolites with infilled microporosity (Emery, 1988). Microstructures in the pore-space can be observed that lead to a reduction in porosity (Fig. 3.3). The Ketton sample chosen for this study consists of large grains compared to the overall image size. The image used for the following evaluation has been downsampled from a 500^3 voxel representation to an image size of 256^3 voxels. This allows more grains to be resolved per training image extracted from the full volume. The downsampled voxel size is 15.2 μm.

Training images were extracted at a sub-volume size of 64^3 with a spacing of 8 voxels leading to a total of 15624 training images. The small spacing of the training images results from the small CT image size of 256^3 voxels.

3.6 Results

Three GANs were trained based on the network architectures highlighted in Section 3.4.2. The training time for each dataset was 24 hours. Manual inspection of synthetic realizations was performed during training to ensure convergence and intermediate evaluation of the covariance and Minkowski functionals.

Figure 3.4 shows the training curve for the Berea sandstone dataset. Initially the generator loss function (Eq. 3.3) is very high and no structural components can be

Table 3.2: Chord lengths $\bar{l}_C$ for the pore and grain phase (Eq. 3.6) determined from the radial averaged covariance $S_2(r)$ of each training image and corresponding realizations generated by the GAN model. The specific surface area S_V and porosity ϕ were evaluated for each of the samples using direct image morphological computation and derived from the covariance. Close agreement between estimates of the porosity and specific surface area can be observed for values determined by direct image morphological estimation and derived values obtained from the radial averaged covariance.

	Beadpack				Berea				Ketton			
$\bar{l}_C^{pore}$ [voxel]	20				10				9			
$\bar{l}_C^{grain}$ [voxel]	36				41				64			
	Training Image		Synthetic		Training Image		Synthetic		Training Image		Synthetic	
Minkowski Functional	$S_2(r)$	Direct	$S_2(r)$	Direct	$S_2(r)$	Direct	$S_2(r)$	Direct	$S_2(r)$	Direct	$S_2(r)$	Direct
Porosity ϕ	0.363	0.359	0.368	0.366	0.196	0.198	0.199	0.197	0.127	0.119	0.119	0.119
$S_v \times 10^{-2}$ $[\frac{1}{voxel}]$	7.0	7.3	6.9	7.5	7.5	8.2	7.9	8.5	5.2	5.2	4.7	5.2

observed in the samples. After a large reduction in the loss function of the generator, initial structures are observed. Image reconstruction quality significantly improves with the number of generator iterations, but cannot be linked to the loss function of the generator. This can be observed from the increase in generator loss at the end of training while image quality improves significantly. The final GAN models were subsequently evaluated in terms of their directional and radial averaged non-centered covariance $S_2(r)$, Minkowski functionals and the single-phase permeability.

For all datasets, 20 realizations were generated using the trained GAN model. In the following section, we present the results of the evaluation of the properties outlined in Section 3.4.1 and compare these to the properties of the original input training image.

3.6.1 Beadpack

The evaluation of the non-centered covariance $S_2(r)$ for the beadpack (Fig. 3.5) shows a strong hole effect reflecting the spherical nature of the grains. A GAN model was trained for 24 hours on the beadpack training image dataset. The GAN model achieves a small error in the porosity of the generated images with a tendency towards higher porosities (Fig. 3.7). A bias can be observed for the specific surface area and the Euler characteristic of the microstructure (Table 3.2). This bias can be explained by the deviation of the grains from a perfect spherical shape in the synthetic realizations. Due to the smooth nature of the spherical particles in the training image, any deviation from this geometry will lead to an increase in the surface area. This is reflected by a higher specific surface area for the synthetic realizations. In addition we observe a reduction in connectivity, represented by a less negative Euler characteristic. The directional covariance S_2 measured on the generated samples show excellent agreement up to the training image size of 128^3 voxels and stabilizes at ϕ^2 (Fig. 3.8). As expected no directional variation of the covariance is observed and the sample is therefore assumed to be isotropic. Single-phase permeability shows a close agreement in both magnitude and variance between the measured training image and the synthetic realizations (Fig. 3.7). Figure 3.6 shows a crossplot of the effective porosity ϕ_{eff} i.e. the porosity open to flow (Eq. 3.12), and the single-phase permeability exhibiting a similar trend in the distribution of values computed on

training images and synthetic realizations. We provide a comparison of all twenty realizations generated by the GAN model in cross-sections through the x-y plane of the original model and a synthetic realization in Fig. 3.9. Many of the grains show a circular to ellipsoidal shape, which considering the fact that a priori the GAN model does not have any knowledge of the geometry of the grains, learning a representation of a perfect sphere can be considered challenging (Section 3.7). The complex grain-grain interface where individual beads contact at single points can be observed for numerous grain arrangements in the generated realizations.

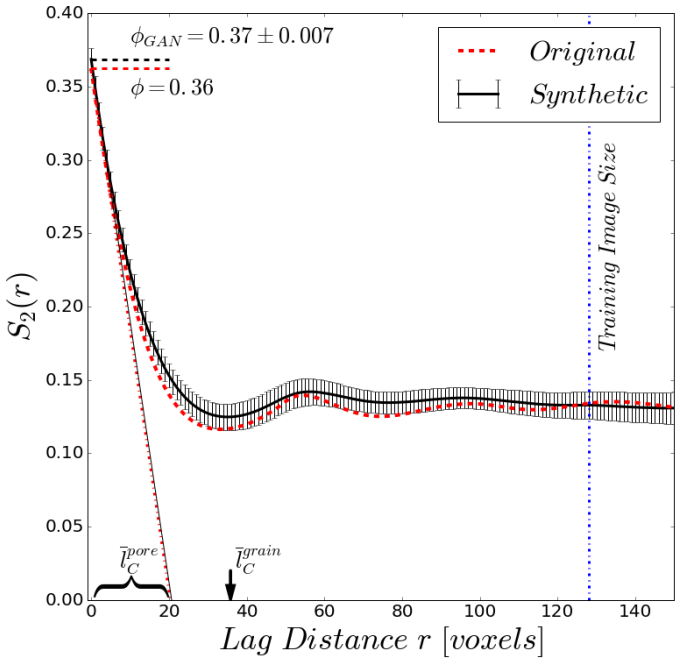

Figure 3.5: Radial averaged covariance $S_2(r)$ for the beadpack sample and 20 synthetic realizations generated by the GAN model.The specific surface area S_V and mean chord lengths $\bar{l}_C$ are derived from the slope of the covariance at the origin (Eqs. 3.5-3.6).

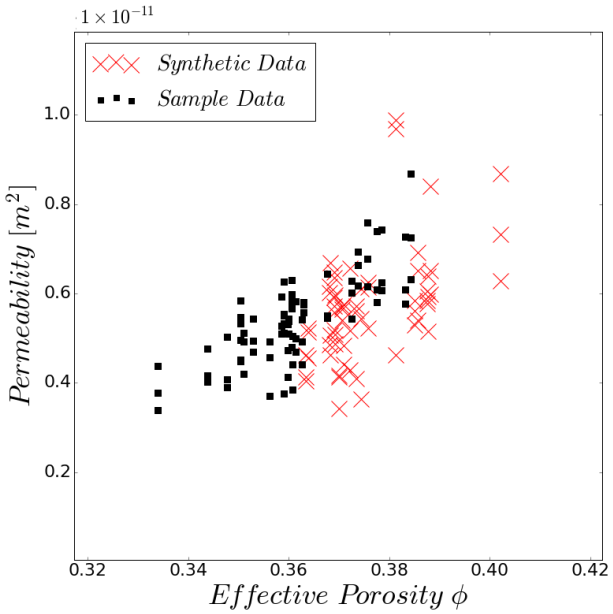

Figure 3.6: A comparison of the numerically estimated single-phase permeability of the beadpack for 128^3 voxel subdomains of the original image and equal sized GAN-based realizations shows a slight overestimation of the effective porosity for the synthetic models. The mean and variance of both permeability distributions are in close agreement (Fig. 3.7).

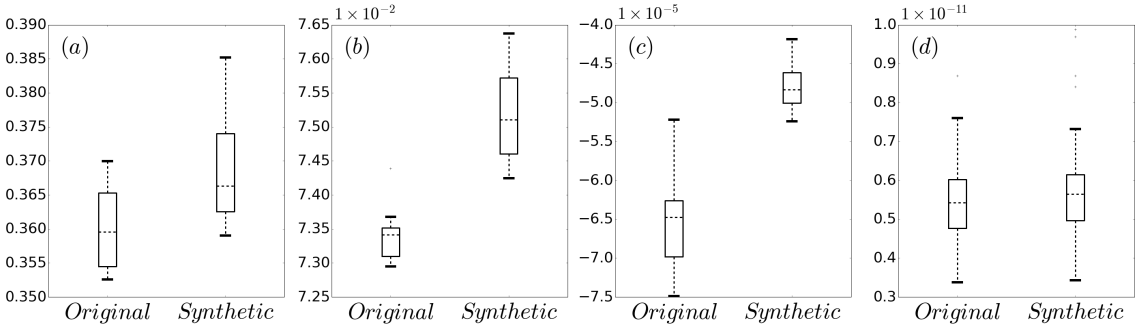

Figure 3.7: Comparison of three Minkowski functionals (porosity (a), specific surface area (b) and Euler characteristic (c)) for the beadpack evaluated on 200^3 voxel subdomains of the original training image and realizations of the GAN model. An error of less than 5% can be observed for the porosity and surface area. A comparison of the estimated values of permeability is shown in (d).

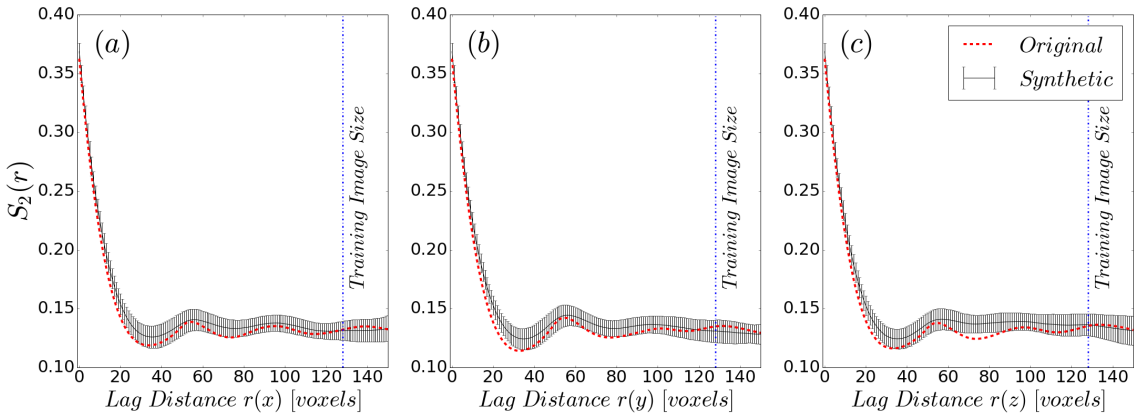

Figure 3.8: Comparison of the directional covariance (x-direction (a), y-direction (b), z-direction(c)) of the beadpack and the average covariance of GAN-based synthetic realizations. A clear hole effect can be observed in the original dataset, which is captured by the GAN model.

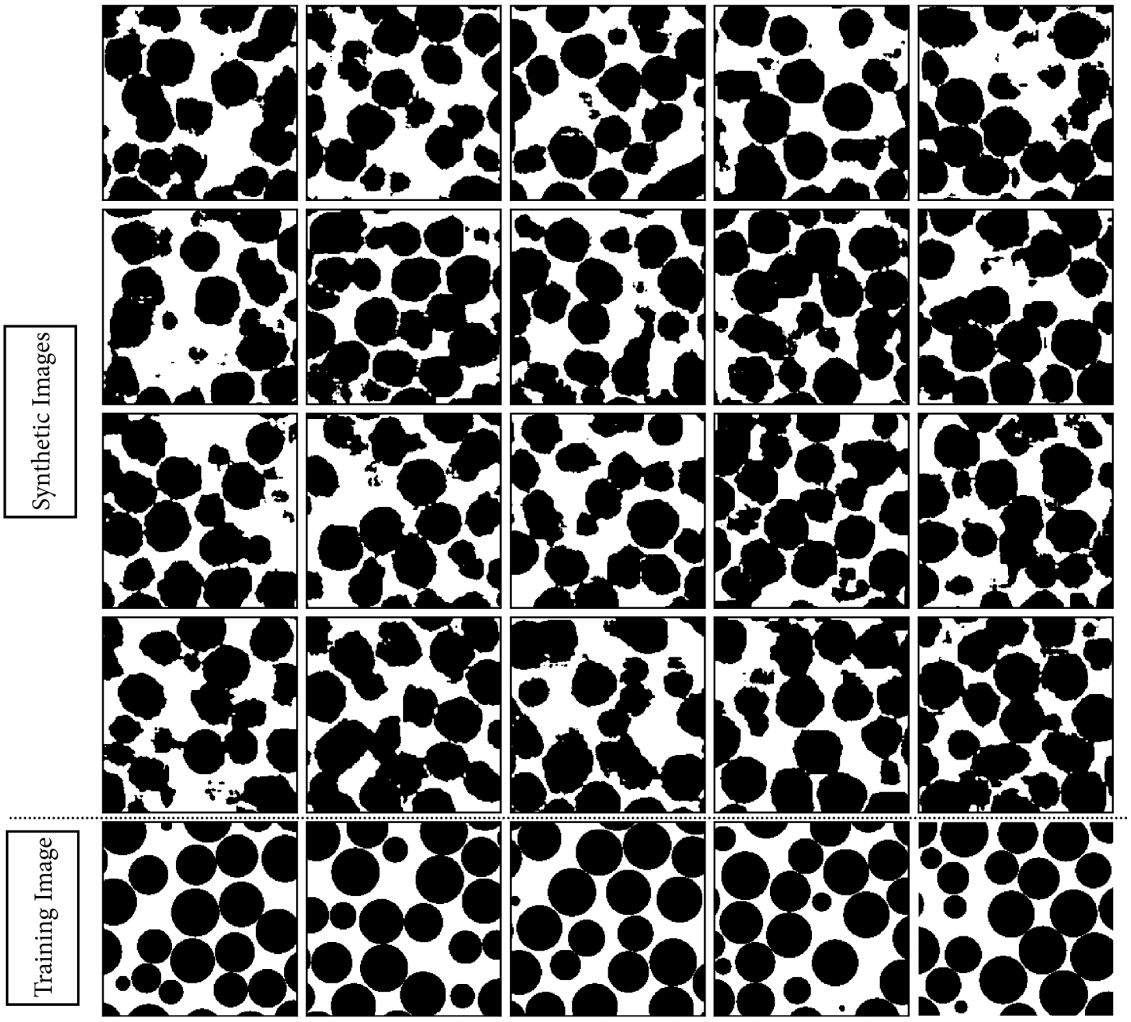

Figure 3.9: Twenty realizations of the spherical beadpack (top) generated to evaluate the statistical, image morphological and transport properties considered in this study. Cross-sectional view of the beadpack training image dataset (bottom).

3.6.2 Berea

The radial averaged covariance $S_2(r)$ in Fig. 3.10, shows a near exponential decay and stabilization occurs at a lag distance of 30 voxels for both covariance functions obtained from the Berea training image and synthetic realizations generated by the GAN model. Additionally, Fig. 3.12 shows that the directional two-point statistics characterized by the directional covariances is captured in the generated images. In all three directions, the GAN model shows excellent agreement and closely follows the trend of the empirical estimates of S_2. The results of the direct computation of the Minkowski functionals is presented in Fig. 3.13 and show comparable distributions for the porosity ϕ, specific surface area S_V and the Euler characteristic χ_V of the training images and the synthetic realizations. A comparison of the specific surface area S_V obtained from the covariance and the direct computation of the Minkowski functional, show nearly equal values (Table 3.2). The estimates of the single-phase permeability show a similar distribution covering the range of effective permeability measured on the training images. Figure 3.11 shows the computed values of permeability and the corresponding effective porosity. The permeability of the synthetic realizations capture the values, variability and trend obtained from the Berea training image dataset. Figure 3.14 shows a comparison of twenty realizations of the GAN model trained on the Berea dataset. A smaller training image size of 64^3 voxels was used, as compared to the beadpack (128^3 voxels). This is due to the smaller size of the structuring elements observed in the training image. A smaller training image size was therefore sufficient to capture the long and short range correlation found in the Berea sample.

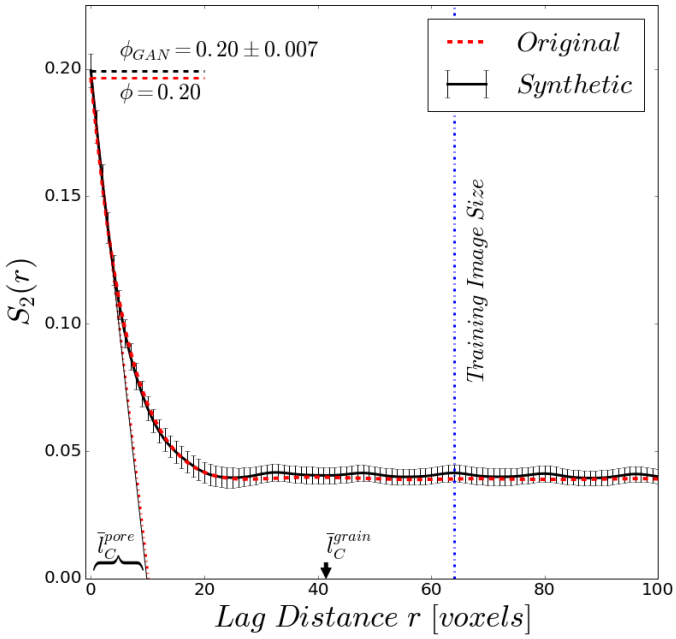

Figure 3.10: Radial averaged covariance $S_2(r)$ for Berea sandstone training images and 20 synthetic realizations generated by the GAN model.

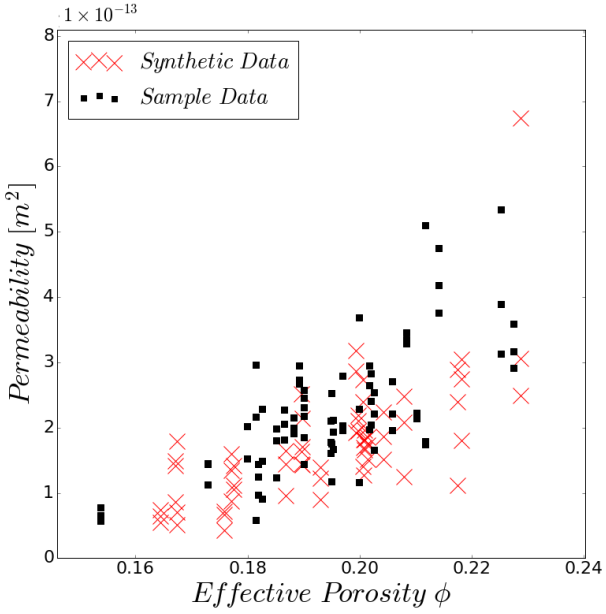

Figure 3.11: The distribution of numerically obtained permeability values on 128^3 voxel subdomains and sampled realizations obtained from a GAN model trained on the Berea sandstone dataset show close agreement in the effective porosity, as well as the evaluated permeability.

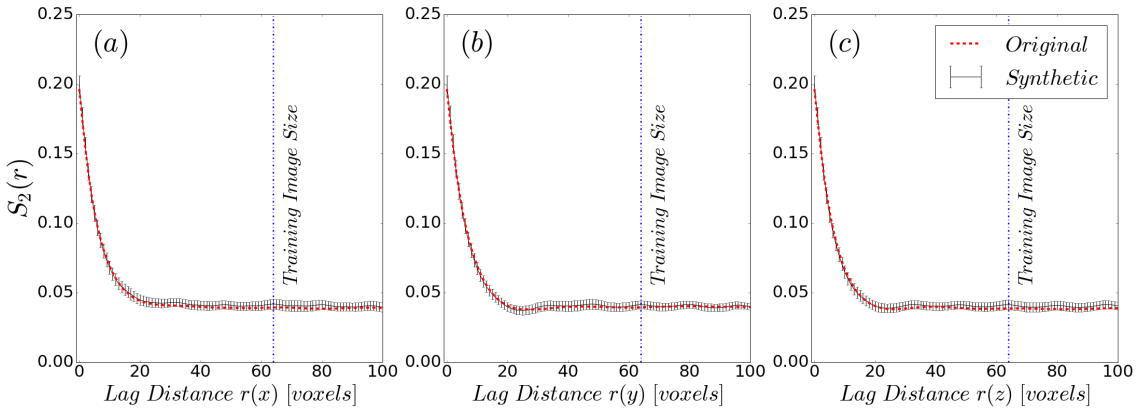

Figure 3.12: Directional non-centered covariance comparison for Berea sandstone. The trained GAN model shows good agreement with the non-centered covariance S_2 of the training image.

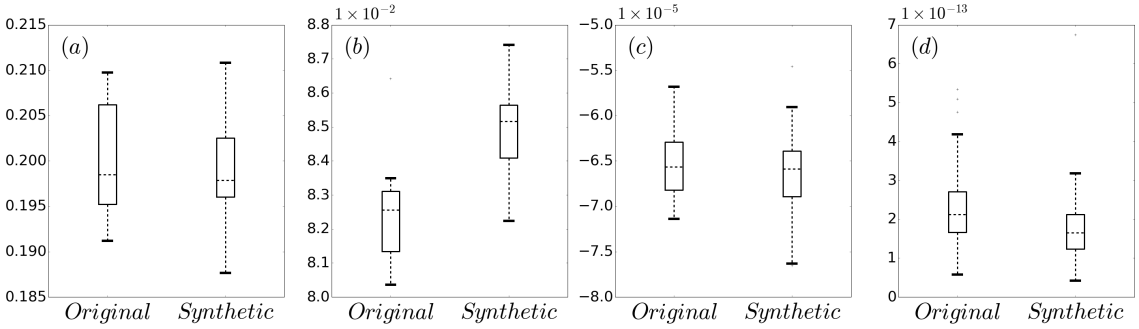

Figure 3.13: Comparison of three Minkowski functionals (porosity (a), specific surface area (b) and Euler characteristic (c)) for Berea sandstone. The porosity, specific surface area and specific Euler characteristic show good agreement between the training image and samples from the trained GAN model. Estimates of permeability for the Berea sandstone and generated samples are presented in (d).

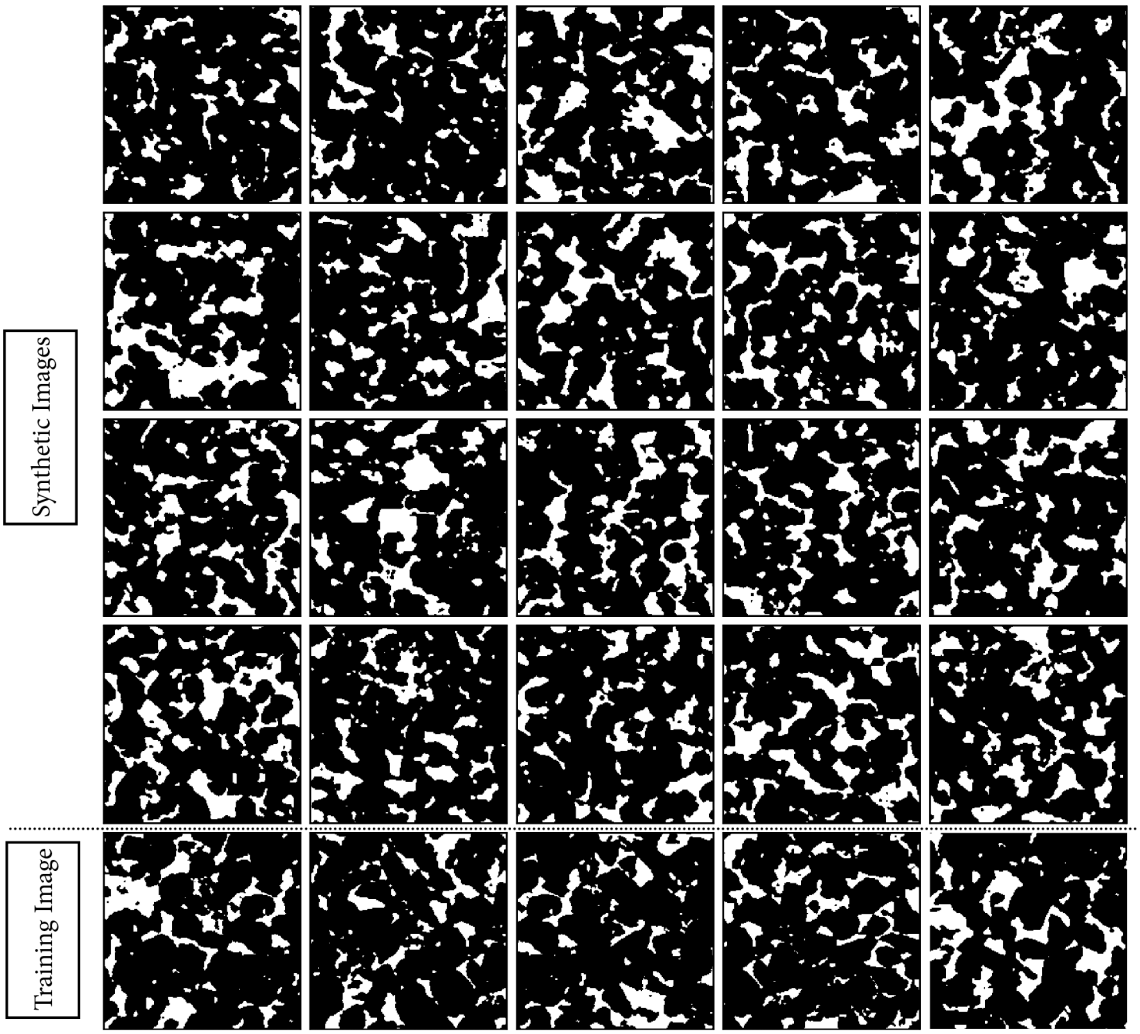

Figure 3.14: Realizations generated by the GAN model (top) compared to training images (bottom) for Berea sandstone.

3.6.3 Ketton

The covariance $S_2(r)$ of the Ketton limestone shown in Fig. 3.15, shows a pronounced hole effect due to the ellipsoidal oolitic grains. Due to the hole effect observed in the radial averaged covariance (Fig. 3.15), we relate the Ketton sample to a hard-sphere model. Figure 3.17 indicates that the images generated by the GAN model trained on the Ketton image, capture the oscillatory and anisotropic behavior of the covariance observed in Ketton. The specific surface area S_V derived from the generated images is in close agreement with the training data. An error of approximately 1% was achieved in the porosity of the GAN generated images compared to the original Ketton dataset (Fig. 3.18). The measured specific surface area of the synthetic images shows a higher variance compared to the original training images. Nevertheless, the average values of the porosity ϕ and specific surface area S_V derived from the non-centered covariance $S_2(r)$ (Eq. 3.5) are in good agreement with values obtained from direct image morphological estimation (Table 3.2). The distribution of single-phase permeability estimates of the synthetic GAN realizations overlies the permeability values of the Ketton training images. The Euler characteristic χ_V and the permeability of the Ketton training dataset are closely matched by the synthetic images and therefore capture the connectivity observed in the oolitic Ketton limestone. We present an overview of the 20 realizations generated by the GAN model trained on the Ketton dataset in Fig. 3.19.

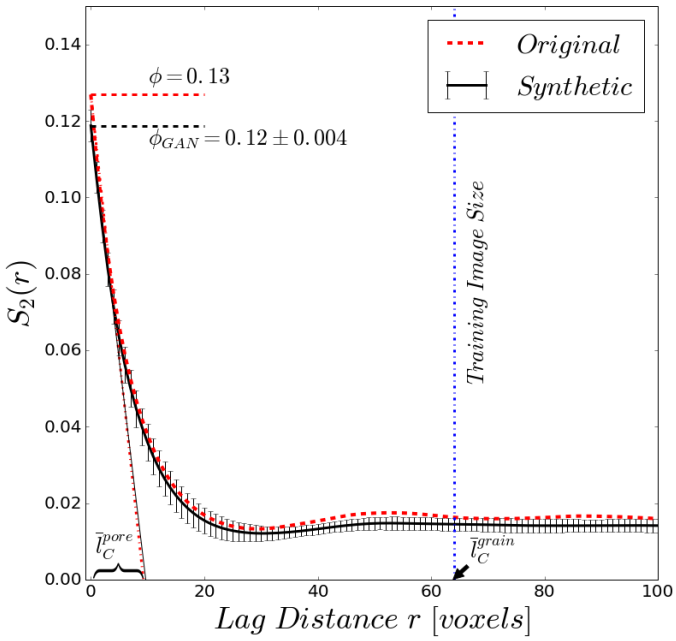

Figure 3.15: Radial averaged covariance $S_2(r)$ for the oolitic Ketton limestone training image and 20 synthetic realizations generated by the GAN model.

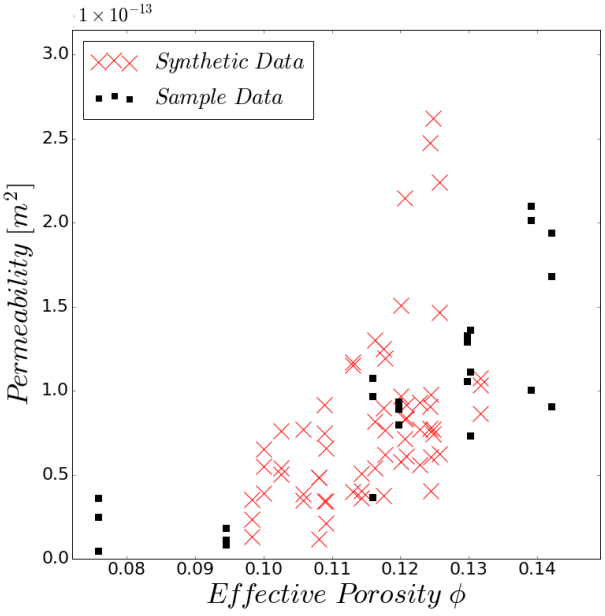

Figure 3.16: Evaluated single-phase permeability for the Ketton training image. The synthetic realizations show similar effective porosity and permeability as the Ketton sample.

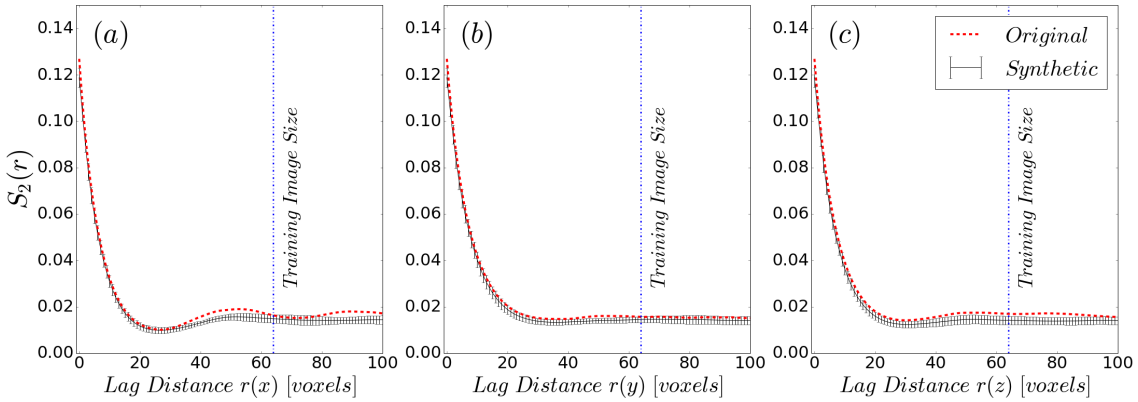

Figure 3.17: The directional covariance of the Ketton sample shows oscillating behavior in the x-direction (a), whereas a nearly exponential decrease can be observed for the y and z directions (b-c). This anisotropy in $S_2(r)$ is also reflected in the covariance of the samples obtained from the GAN model.

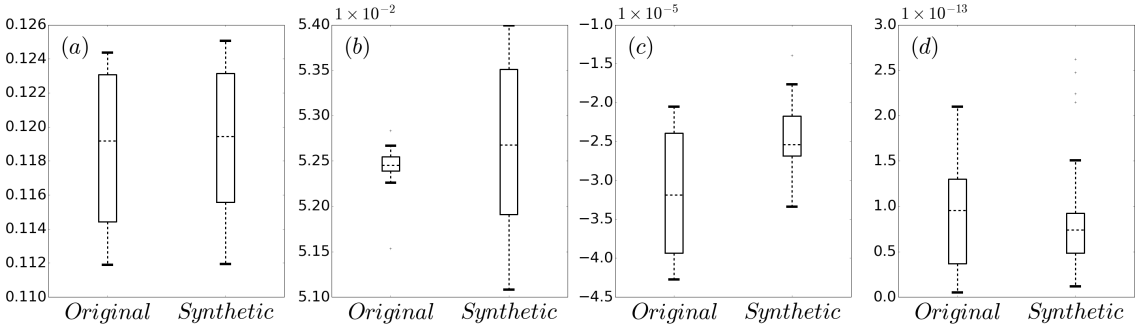

Figure 3.18: Comparison of the Minkowski functionals (porosity (a), specific surface area (b) and Euler characteristic (c)) for the Ketton training image. The three evaluated Minkowski functionals show good agreement. The evaluated Euler characteristic indicates that the sampled synthetic realizations show a similar degree of connectivity as the training image. Estimated permeability values for the Ketton sample are presented in (d).

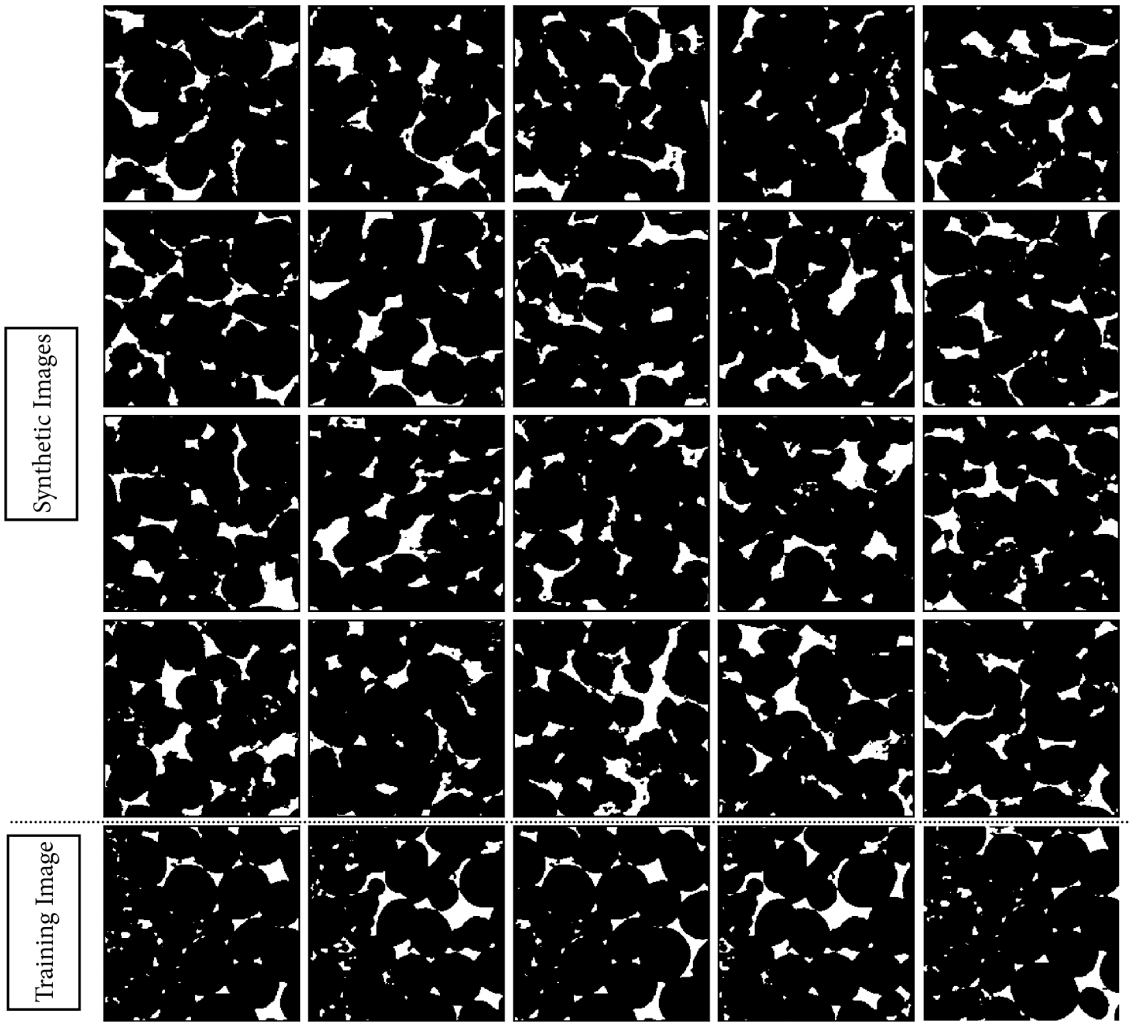

Figure 3.19: Realizations generated by the GAN model (top) compared to training images (bottom) for Ketton limestone.

3.7 Discussion

This chapter presents a novel method for three-dimensional stochastic image reconstruction based on generative adversarial neural networks (GAN) trained on three-dimensional segmented images. To summarize, the objectives of this contribution are threefold. Firstly, generate stochastic reconstructions of porous media such as sedimentary rocks exceeding the size of the acquired image datasets. Secondly, evaluate the ability of GAN models to capture the image morphological and physical properties of micro-scale porous media. Thirdly, establish a method of stochastic image reconstruction that allows a probabilistic treatment of pore-scale properties such as permeability without the need to acquire numerous images of a single rock type.

The first objective stems from technical limitations of micro-CT data acquisition. Images are acquired as a trade-off between sample size i.e. how many representative structures can be captured in one image versus the resolution at which these pore-scale structures are resolved. The generation of large porous domains based on high-resolution images enables this gap in scales to be bridged and micro-scale features to be incorporated in macro-scale models.

Our findings show that GANs can learn an implicit representation of the image space given a limited number of training images subsampled from larger images. These subdomains were extracted based on characteristic length scales (Section 3.4.1) and serve as a training set for the GAN model. For the Ketton limestone, a small spacing of the extracted subdomains was required to increase the size of the training image dataset. While we did not find any evidence of an introduced bias by using correlated subdomains, we believe that these extracted training images should represent independent regions.

We have evaluated the ability to train GANs for a number of training image sizes less than and up to twice the size of the structuring elements. We have found that models trained on images smaller than the average grain size results in artifacts and distorted shapes occurring in the generated micro-structures. For the beadpack, the size of an individual sphere is 50 voxels. A training image of 64^3 voxels would typically only contain parts of an individual grain and only capture the interaction of the particles, but not the geometry of the structuring element. For the beadpack, models trained on 64^3 voxels were successful in learning a representation of the short scale micro-structure but failed to reproduce the long distance correlation. A larger training image of 128^3 voxels, as was used to model the beadpack has a much higher chance to represent the full geometry of the particles and therefore not only learn interactions, but also the shapes of grains.

We, therefore, suggest that training images extracted from large datasets must be larger than the average grain size. For models that are well described by a Boolean model, the size of the structuring element can be readily estimated from stabilization of the covariance $S_2(r)$. For more complex samples a different measure must be used to estimate the size of the required training image.

The chord length is one additional measure that can be obtained to characterize the grain space of porous media. While we have found that the mean chord length of the grain space $\overline{l}_C^{grain}$ is always less than or equal to that of our chosen training image size, $\overline{l}_C^{grain}$ increases with decreasing porosity. This contradicts the need to have the largest training domain for the beadpack sample which also has the highest porosity.

A better estimate may be derived from the representative elementary volume of the specific surface area which by definition is the same for the grain- and pore-space and is, therefore, more representative of the morphology of the porous medium (Bear, 2013). Based on the properties we have evaluated we could not find a measure derived from two-point statistical or image morphological properties that is closely related to the required training image size and we see a theoretical discussion of this as possible future work.

Conceptually the simplest model considered in this study, the spherical beadpack, has proven to be the most challenging as a training image for the GAN model (Section 3.6.1). While we observe spherical and ellipsoidal shapes in the resulting realizations (Fig. 3.9), the shape is exactly defined by the spherical nature of the grains. Any deviation from this shape, which for GANs, is learned implicitly from the data itself, will lead to a misrepresentation of the effective properties. Random hard-sphere models with spherical grains will efficiently capture the nature of this dataset. Therefore we suggest a fit-for-purpose application of GANs, for training images that exhibit variability of grain sizes and shapes, which are not readily captured by a simpler model.

While for many sedimentary granular rocks representative volumetric images can be obtained, this may be more challenging for carbonate samples with complex pore-grain structures. The three training images considered in this study were all treated under the assumption of stationarity i.e. we do not expect a systematic variation in the mean and variance of the averaged properties as a function of location. In theory, GANs are not limited to learning representations of stationary datasets. This is shown by the many successful applications for two-dimensional image and texture synthesis of non-stationary domains, such as learned image representations of human faces (Gauthier, 2014) or galaxies (Schawinski et al., 2017; Ravanbakhsh et al., 2016). Therefore a model that incorporates non-stationarity for a single rock-type would technically be possible in the GAN framework but would require the acquisition of many images of the same porous medium.

A valid representation of the microscale variability and connectivity of the pore-space is critical to assess the single and multiphase flow behavior of porous media. Therefore any stochastic reconstruction method used in the process of deriving or evaluating the variability of micro-scale properties must capture the statistical and image morphological characteristics of the reconstructed porous medium. While we have shown that for the evaluated datasets, the GAN-based image reconstructions capture the variation and characteristics of these porous media, a number of challenges remain in this task that are different from those encountered by classical stochastic methods of image reconstruction.

For porous media, many flow related properties can be related to the porosity. Classical stochastic methods are able to capture the porosity efficiently by defining a specific proportion of the grain and pore domain. The GAN-based model presented in this study initially has no knowledge of the porosity. The porosity, therefore, arises as a feature of the training image data. Matching the porosity distribution of the training image dataset was found to be the main challenge in training a GAN model. An error of three percent in porosity, for instance could lead to a significant mismatch in the permeability of the synthetic images. It is, therefore, necessary to continuously monitor the derived properties such as the Minkowski functionals or estimates of the permeability, in the course of training the neural networks to

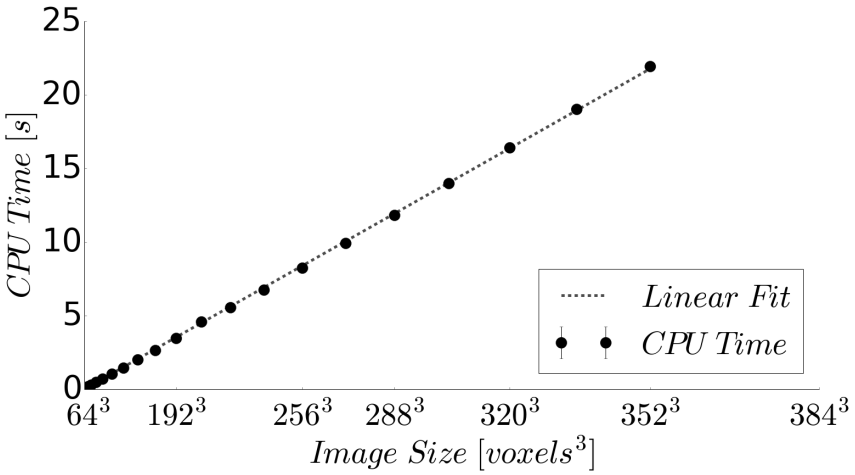

Figure 3.20: Measured CPU time for generating synthetic realizations of Berea sandstone at increasing image size. 100 realizations were computed at each image dimension and CPU time averaged. Computational cost increases linearly with the number of voxels in the generated image.

ensure that synthetic realizations created by the GAN model are able to capture the effective properties of the micro-scale domains.

While this can be considered one of the main challenges in the application of GANs for synthetic image reconstruction, learning an implicit representation of the training data itself can be seen as a strength. Many classic stochastic methods rely on the formulation of an objective function that ensures that statistical properties are captured in the generated realizations e.g. matching $S_2(r)$ and the specific surface area S_V of the stochastic reconstructions to a desired precision. The GAN approach does not require an explicit objective function a priori. The objective function is encoded in the discriminator and adapted in the course of training.

During adversarial training both the generator and discriminator are continuously improved. The discriminator's sole purpose is to be able to distinguish real training data from generated synthetic data. On the other hand, the generator tries to generate synthetic data that the discriminator is not able to distinguish from the training data. Due to the multi-scale representation of the convolutional neural networks, these features must be learned across the full range of length scales present in the training data, leading to a high-resolution image that captures small and large scale features of the image dataset. A number of stacked GAN models can be trained on e.g. low-resolution medical-CT data and high-resolution micro-CT allowing incorporation of spatial information across multiple length scales (Zhang et al., 2016).

Once the GAN model has successfully learned to create physically representative samples of the porous medium, one possible application is to evaluate the variability in the flow properties by evaluating the properties of a large number of samples. This not only requires a physically valid representation of the porous medium but also requires a method that allows fast image reconstruction. In Section 3.6 we have shown that training was performed for approximately 24 hours and may vary due to the need for manual inspection of the generated samples in the training process. Figure 3.20 shows the CPU time required for generation of images at increasing image size. The fully convolutional nature of the GAN architecture allows very large

images, exceeding the size of the original sample to be generated very efficiently and at low computational cost and runtime.

While training requires considerable time and computational resources in the form of modern graphics processors as well as optimized neural network frameworks, image reconstruction requires little computational effort and scales linearly in the total number of voxels of the generated images. This, therefore, enables the generation of ensembles of large domains based on volumetric images acquired from 3D microscopy, that capture the physical behavior of the porous medium. The learned representation of the generator consists of the weights of the convolutional filters learned in the training process and can, therefore, be stored for future use once training has finished.

3.8 Conclusions

We have evaluated the application of generative adversarial neural networks for stochastic image reconstruction of porous media based on previously acquired images of sedimentary rocks. Three image datasets were used as training images: a beadpack, a Berea sandstone, and an oolitic Ketton limestone.

By evaluating two-point statistical measures, image morphological features and computing the single-phase effective permeability we have shown that the synthetic images generated by the GAN model are able to match key characteristic statistical and physical parameters of these porous media. While a large computational effort is required to train the GAN model, the generation of samples from the learned representation is highly efficient and learned models are easily stored for future use.

Future work in the application of GANs to stochastic image reconstruction of porous media will include improving the quality of the image reconstruction by evaluating various generator-discriminator architectures, the use of grayscale and multi-channel training images, as well as the application of large multi-scale domains of porous media to evaluate the ensemble behavior of single and multiphase flow properties in porous media. Recent advances in the understanding of GANs should lead to a more stable and consistent training process (Mao et al., 2016; Arjovsky et al., 2017). This is explored further in the following chapter where we study the reconstruction of grayscale (unsegmented) images of Ketton limestone, as well as associated topological and flow properties.

Chapter 4

Stochastic Reconstruction of an Oolitic Limestone by Generative Adversarial Networks

4.1 Abstract

Stochastic image reconstruction is a key part of modern digital rock physics and materials analysis that aims to create representative samples of micro-structures for upsampling, upscaling, and uncertainty quantification. We present new results of a method of three-dimensional stochastic image reconstruction based on generative adversarial neural networks (GANs). GANs are a family of unsupervised learning methods that require no a priori inference of the probability distribution associated with the training data. Thanks to the use of two convolutional neural networks, the Discriminator and the Generator, in the training phase, and only the Generator in the simulation phase, GANs allow the sampling of large and realistic volumetric images. We apply a GAN-based workflow of training and image generation to an oolitic Ketton limestone micro-CT unsegmented gray-level dataset. Minkowski functionals calculated as a function of the segmentation threshold are compared between simulated and acquired images. Flow simulations are run on the segmented images, and effective permeability as well as velocity distributions of simulated flow are also compared. Results show that GANs allow a fast and accurate reconstruction of the evaluated image dataset. We discuss the performances of GANs in relation to other simulation techniques, and stress the benefits resulting from the use of convolutional neural networks . We address a number of challenges involved in GANs, in particular the representation of the probability distribution associated with the training data.

4.2 Introduction

The micro-structural characteristics of porous media play an important role in the understanding of numerous scientific and engineering applications such as the recovery of hydrocarbons from subsurface reservoirs (Blunt et al., 2013), sequestration of CO_2 (Singh et al., 2017) or the design of new batteries (Siddique et al., 2012). Modern micro-computer tomographic (micro-CT) methods have enabled the acquisition of

high-resolution three-dimensional images at the scale of individual pores. Increased resolution comes at the cost of longer image acquisition time and limited sample size. Individual samples allow numerical and experimental assessment of the effective properties of the porous media, but give no insight into the variance of key microstructural properties. Therefore, an efficient method to generate representative volumetric models of porous media that allow the assessment of the effective properties is required. The generated images serve as an input to a digital rock physics workflow to represent the computational domain for numerical estimation of key physical properties (Berg and Nyström, 2017).

Statistical methods aim at reconstructing porous media based on spatial statistical properties such as two-point pore-grain correlation functions. Quiblier, (1984) has presented an extensive overview of the early literature of porous media reconstruction and provided an extension of the method of Joshi, (1974) by reconstructing three-dimensional porous media based on the empirical covariance function and probability density function obtained from two-dimensional thin-sections. Other statistical methods such as simulated annealing (Yeong and Torquato, 1998; Jiao et al., 2008) allow high-quality three-dimensional reconstruction and incorporation of numerous statistical descriptors of porous media. Pant, (2016) introduced a multi-scale simulated annealing algorithm allowing simulation of three-dimensional porous media at much lower computational cost than previous methods.

Methods to incorporate higher-order multiple-point statistical (MPS) properties of porous media have been developed. These MPS functions are implicitly defined by two- or three-dimensional training images. Simulation algorithms based on multiple-point statistics are therefore considered as training image-based algorithms. MPS simulation was originally developed in the context of generating realistic geological structures (Guardiano and Srivastava, 1993; Caers, 2001; Mariethoz and Caers, 2014). With the advent of micron-resolution X-ray tomography (micro-CT imaging) (Flannery et al., 1987), which provides training images, MPS simulation techniques have been successfully applied to the stochastic reconstruction of three-dimensional porous media (Okabe and Blunt, 2004; Okabe and Blunt, 2005; Okabe and Blunt, 2007).

Tahmasebi et al., (2012) as well as Tahmasebi and Sahimi (Tahmasebi and Sahimi, 2012; Tahmasebi and Sahimi, 2013) have introduced a patch-based approach where subdomains are simulated along a pre-defined path and populated based on a cross-correlation distance criterion (CCSIM). This approach is similar to the image quilting algorithm by Efros and Freeman, (2001) but corrects mismatching patches in overlapping or neighboring domains. Tahmasebi et al., (2017) present a method for fast reconstruction of granular porous media from a single two or three-dimensional training image using a method closely related to CCSIM. They obtain significant speedup in computational time by incorporating a fast-Fourier transform and a multi-scale approach. A graph-based approach is used to resolve non-physical regions at the boundaries of simulated patches of grains.

Object-based methods describe the material domain by locating geometrical bodies of random size at locations provided by a spatial point process. The so-called Boolean model is a particular case where the randomly placed bodies, typically spheres are allowed to overlap (Matheron, 1975; Chiu et al., 2013). Object-based methods may also allow interaction of particles to be incorporated (Torquato, 2013). They have successfully been used to describe complex and heterogeneous materials.

Process models reconstruct the pore and grain structure of materials by mimicking how they were formed. Øren and Bakke, (2003) have created reconstructions of sandstones by reproducing the natural processes of sedimentation, compaction and diagenesis.

This contribution presents a training image-based method of image reconstruction using a class of deep generative methods called generative adversarial networks (GANs) first introduced by Goodfellow et al., (2014). In Chapter 3 and Mosser et al., (2017) we have shown that GANs allow the reconstruction of three-dimensional porous media based on segmented volumetric images. Their study applied GANs to three segmented images of rock samples. They showed that GANs represent a computationally efficient method for the fast generation of large volumetric images that capture the statistical and morphological features, as well as the effective permeability.

We expand on the work of Mosser et al., (2017) and investigate the ability of generative adversarial networks to create stochastic reconstructions of an unsegmented micro-CT scan of a larger oolitic Ketton limestone sample. We evaluate the four Minkowski functionals for the three-dimensional datasets as a function of the gray-level threshold. In addition to the numerical evaluation of permeability as shown by Mosser et al., (2017), we compare velocity distributions of the original porous medium and samples obtained from the GAN. We also provide details of the convolution approach used by GANs. Furthermore we evaluate the reconstruction process within the trained generative function and highlight the parametric and differentiable nature of the obtained generative function. We evaluate the computational cost of GAN-based image simulation with a reported values of computational run time for a variety of other reconstruction methods of equal reconstruction quality. We also investigate how the image representation evolves along the different layers of the GAN network, and discuss the benefits that can be derived from the differentiable nature of the parameterization used by GANs.

4.3 Generative Adversarial Networks

Generative adversarial networks are a deep learning method for generating samples from arbitrary probability distributions (Goodfellow et al., 2014; Goodfellow, 2016). GANs do not impose any a priori model on the probability density function and are therefore also referred to as an implicit method. Without the need to specify an explicit model, GANs provide efficient sampling methods for high dimensional and intractable density functions.

In the case of CT images of porous media we can define an image x to be a sample of a real, unknown probability density function (pdf) of images p_{data} of which we have acquired a number of samples which serve as training images. In our example, the training set is comprised of 64^3 voxel sub-domains of the original micro-CT image. Sub-domains are extracted without any overlap and each training image represents the originally acquired dataset.

GANs consist of two functions: a generator, whose role it is to generate samples of the unknown density $p_{data}(\mathbf{x})$ and a discriminator function that tries to distinguish between samples from the training set and synthetic images created by the generator. The generator G is defined by its parameters θ and performs a mapping from a

random prior $\mathbf{z}$ to the image domain

$$\mathbf{z} \sim \mathcal{N}(0,1)^{d\times 1\times 1\times 1} \tag{4.1}$$

$$G_\theta : \mathbf{z} \rightarrow \mathbb{R}^{1\times 64\times 64\times 64} \tag{4.2}$$

where d is the dimensionality of the random prior.

The discriminator $D_\omega(\mathbf{x})$ assigns a probability to an image x being a sample of the true data distribution p_{data}

$$D_\omega : \mathbb{R}^{1\times 64\times 64\times 64} \rightarrow [0,1] \tag{4.3}$$

where values close to 1 represent a high probability of being a sample of $\mathbf{x} \sim p_{data}(\mathbf{x})$.

We represent both the generator $G_\theta(\mathbf{z})$ and the discriminator $D_\omega(\mathbf{x})$ by differentiable neural networks with parameters θ and ω respectively. This allows us to use backpropagation combined with mini-batch gradient descent to optimize the generator and discriminator according to the functional

$$\min_\theta \max_\omega \{\mathbb{E}_{\mathbf{x}\sim p_{data}}[log\ D_\omega(\mathbf{x})] + \mathbb{E}_{\mathbf{x}\sim p_\mathbf{z}}[log(\ 1 - D_\omega(G_\theta(\mathbf{z})))]\} \tag{4.4}$$

The optimization criterion of the generator and discriminator (Eq. 4.4) is solved sequentially in a two-step procedure. We first train the discriminator to maximize its ability to distinguish real from fake samples. This is done in a supervised manner by training the discriminator on known real samples (Label 1) and samples created by the generator (Label 0). The binary cross-entropy is used as an objective function to compute the miss-classification error

$$H(\mathbf{y}, \mathbf{y}') = -\sum_i (y_i \log(y_i') + (1 - y_i)\log(1 - y_i')) \tag{4.5}$$

where $\mathbf{y}'$ is a vector containing the output probability assigned by the discriminator for each element of a given mini-batch of samples. For each mini-batch of real images we therefore optimize $H(\mathbf{1}, \mathbf{y}')$, and for all fake samples $H(\mathbf{0}, \mathbf{y}')$ (Eq. 4.5). The error is back-propagated while keeping the parameters of the generator constant.

In a second step we train the generator to maximize its ability to "fool" the discriminator into misclassifying the images provided by the generator as real images. This is performed by computing the binary cross-entropy of the output of the discriminator on a mini-batch of sampled from the generator $G_\theta(\mathbf{z})$ and requiring that the created labels be close to one, therefore computing $H(\mathbf{1}, \mathbf{y}')$. The parameters of the generator are then modified to optimize $H(\mathbf{1}, \mathbf{y}')$ by applying stochastic gradient descent while keeping the parameters of the discriminator constant.

Training of these networks is often challenging due to the competing objective functions of the generator and discriminator. Recently, new objective functions and training heuristics have greatly improved the training process of GANs (Arjovsky et al., 2017; Yeh et al., 2016).

GANs follow a different training scheme from other stochastic reconstruction methods (Section 4.2). There are two phases in GAN-based reconstruction: training and generation. Training is expensive, requiring modern graphics processing units (GPU) and for three-dimensional datasets large GPU memory. Parallelization of the training process across numerous GPUs reduces time for training the network.

Nevertheless, finding a set of hyper-parameters, that is a network architecture (number of filters, types, order of layers and activation functions) that leads to the desired quality can require significant trial and error.

The second phase of GAN-based reconstruction, the generation of individual samples, is extremely fast. All operations in the generator network can be represented as matrix-vector operations which are executed efficiently on modern GPU systems and take on the order of seconds for modern GPUs, as shown later in this Chapter.

4.4 Dataset

The sample used in this study is an oolitic limestone of Jurassic age (169—176 million years). The spherical to ellipsoidal grains consist of 99.1% calcite and 0.9% quartz (Menke et al., 2017). Inter and intra-granular porosity can be observed, as well as significant amounts of unresolved sub-resolution micro-porosity. This is characterized by the various shades of gray in individual grains, where the interaction of sub-resolution porosity with X-rays penetrating the sample during imaging leads to an increase in intermediate gray-level values (Fig. 4.1). The sample was imaged using a Zeiss XRM 510 with a voxel size of 27.8 μm. The size of the image domain after resampling to 8bit resolution is 900^3 voxels. We subdivide the original image into a training set of non-overlapping 5832 images at a size of 64^3 voxels. We define a sequential randomized pass over the full training set as an epoch. Evaluation of the effective properties is performed at larger image sizes than the training images to judge whether the GAN is able to generalize to larger domains. To evaluate the reconstruction quality of the GAN model we randomly extract 64 images at a size of 200^3 voxels with no overlap from the original training image (Fig. 4.1) which we refer to as the validation set. A synthetic validation set was created by sampling 64 images at a size of 200^3 voxels from the trained GAN model. To perform numerical computation of the effective permeability as well as measure the two-point correlation function, all images of the synthetic and original Ketton validation set were segmented using Otsu thresholding (Otsu, 1975). Minkowski functionals were evaluated for the unsegmented validation sets.

4.4.1 Neural Network Architecture and Training

Radford et al., (2015) proposed to remove fully connected layers in the input and output of the generator network. They represent the input layer for the latent random vector by a reshaping operation, followed by a stack of strided convolutional layers. Jetchev et al., (2016) introduced the SGAN architecture where the input latent vector has spatial dimension and is immediately followed by a set of convolution operations. This allows images to be generated that are larger than the training images. They also provide evidence that sampling using the SGAN network architecture represents a stationary, ergodic and strongly mixing stochastic process. Our generator architecture represents a fully convolutional network without reshaping operations. The fully convolutional nature of the generator allows us to create images of arbitrary size by providing latent random vectors with larger spatial dimensionality e.g. $\mathbf{z} \sim \mathcal{N}(0,1)^{d\times m\times n\times o}$. During training m, n and o are of size one, which results in an image of 64^3 voxels. For image generation m, n and o may be of any integer size. The main difference to the SGAN architecture of Jetchev et al., (2016) is therefore

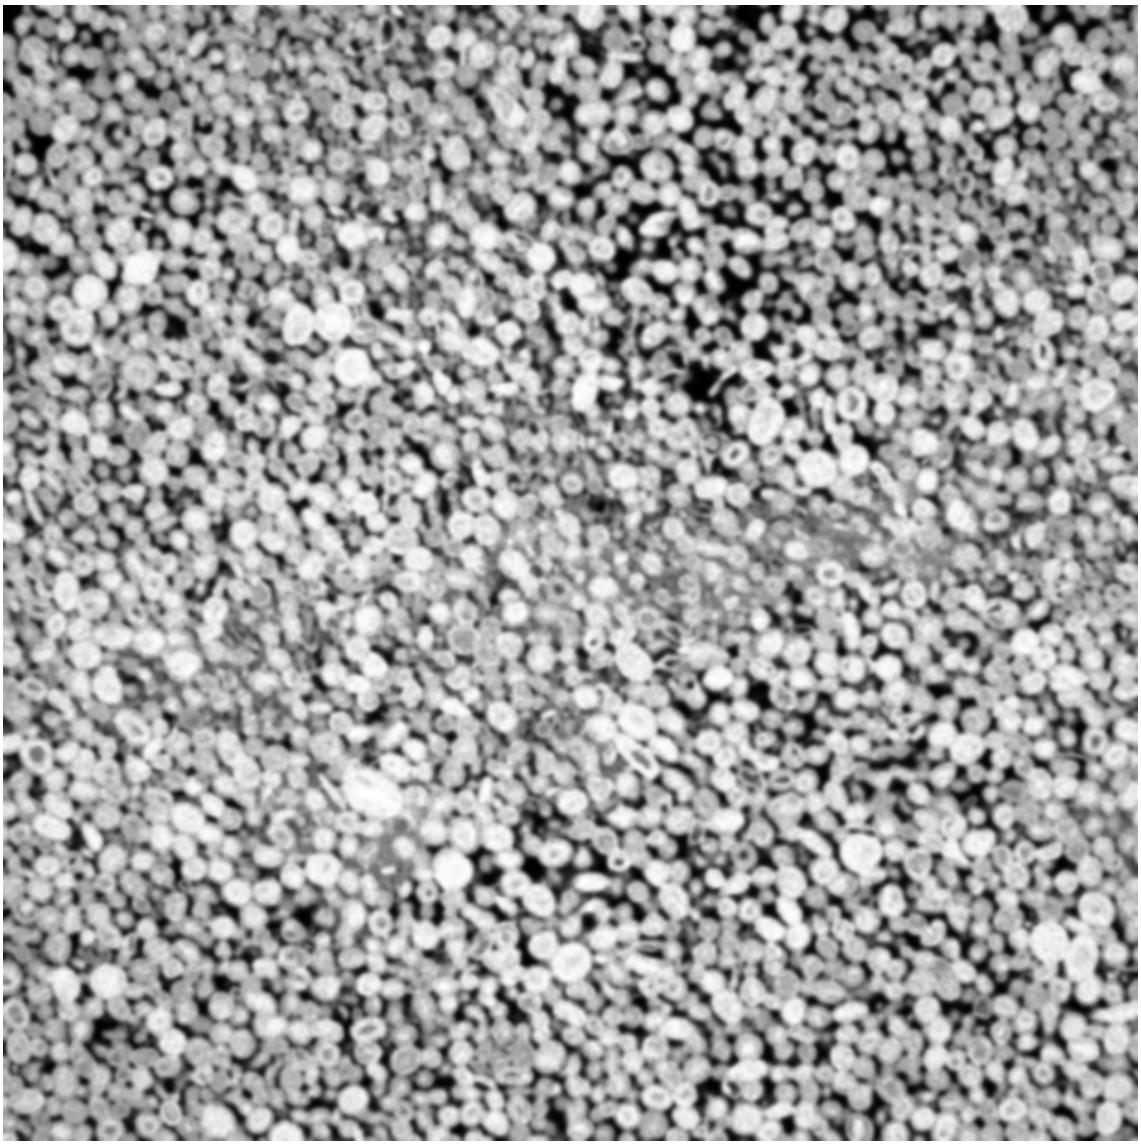

Figure 4.1: Two-dimensional gray-level cross-section of the three-dimensional micro-CT image of the studied oolitic Ketton limestone sample. The image has a size of 900^3 voxels and was acquired with a voxel size of 27.8 μm. Histogram equalization was applied to the image prior to its use as a training image.

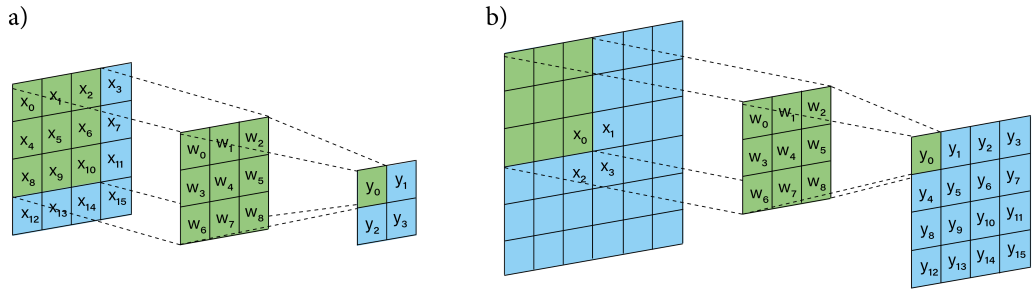

Figure 4.2: Example of a discrete convolution (a) for a 3×3 filter kernel size applied to a 4×4 feature map. (b) The so-called transposed convolution operation for a 3×3 filter kernel size applied to a 2×2 feature map. The active regions to compute the output value are shaded green. Empty cells represent zero-padded values.

that at training time the input random vector has a spatial dimension of one and the output of the discriminator is a single scalar value.

In Fig. 4.2 we show an example of a convolution and transposed convolution operation for the two-dimensional case. The convolution is performed by sliding a filter kernel w_i (Eq. 4.6) over the input feature map x_i (Eq. 4.7) (Dumoulin and Visin, 2016). We rewrite this as an efficient matrix vector operation (Eq. 4.8) by unrolling the discrete convolution

$$\mathbf{W} = \begin{pmatrix} w_0 & w_1 & w_2 & 0 & w_3 & w_4 & w_5 & 0 & w_6 & w_7 & w_8 & 0 & 0 & 0 & 0 & 0 \\ 0 & w_0 & w_1 & w_2 & 0 & w_3 & w_4 & w_5 & 0 & w_6 & w_7 & w_8 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & w_0 & w_1 & w_2 & 0 & w_3 & w_4 & w_5 & 0 & w_6 & w_7 & w_8 & 0 \\ 0 & 0 & 0 & 0 & 0 & w_0 & w_1 & w_2 & 0 & w_3 & w_4 & w_5 & 0 & w_6 & w_7 & w_8 \end{pmatrix} \quad [4\times16] \tag{4.6}$$

The input image $\mathbf{x}$, in this case a single-channel 4×4 image, and the output $\mathbf{y}$ are represented as one dimensional vectors

$$\mathbf{x}\ [16 \times 1],\ \mathbf{y}\ [4 \times 1] \tag{4.7}$$

This allows us to perform the discrete convolution

$$\mathbf{W} * \mathbf{x} = \mathbf{y} \tag{4.8}$$

where $\mathbf{W}$, $\mathbf{x}$, $\mathbf{y}$ are defined according to Eq. 4.7. For each convolutional layer of the network, the input features are convolved with a number of filter kernels $\mathbf{W}$. While the convolution operation (Eq. 4.8) creates an output of smaller size than the input feature map, the so-called transposed convolution operation creates an output that has larger spatial-dimensions. This is shown in Fig. 4.2-b where the input features $\mathbf{x}$ are first padded with zeros and are subsequently convolved with the weight kernel $\mathbf{W}$ to generate a spatially larger output feature map $\mathbf{y}$.

The generator consists of a series of three-dimensional transposed convolutions. In each layer, the number of weight kernels is reduced by a factor of $\frac{1}{2}$. Before the final transposed convolution we add an additional convolutional layer (Fig. 4.3). Each layer in the network except the last is followed by a batch normalization (Ioffe and Szegedy, 2015) and a Leaky Rectified Linear Unit (LeakyReLU) activation function. The final transposed convolution in the generator is followed by a hyperbolic tangent activation function (Tanh) (LeCun et al., 1998). A representation of each activation function used in the network is shown in Fig. 4.4.

We represent the discriminator as a convolutional classification network with binary output using as input the real samples of the 64^3 voxel training set (Label 1) and synthetic realizations of equal size created by the generator (Label 0). Each layer in the network consists of a three-dimensional convolution operation followed by batch normalization and a LeakyReLU activation function. The final convolutional layer outputs a single value between 0 and 1 (Sigmoid activation) which corresponds to the probability that the input image belongs to the original training set or in other words that it is a real image.

We distinguish two sets of parameters for training: the set of weights of a network comprises the adjustable parameters of the filter kernels for convolutional and neurons for linear network layers. The so-called hyper-parameters define the network architecture and training scheme e.g. the number of filters per layer, the number of convolutional layers or learning rates. A chosen set of hyper-parameters

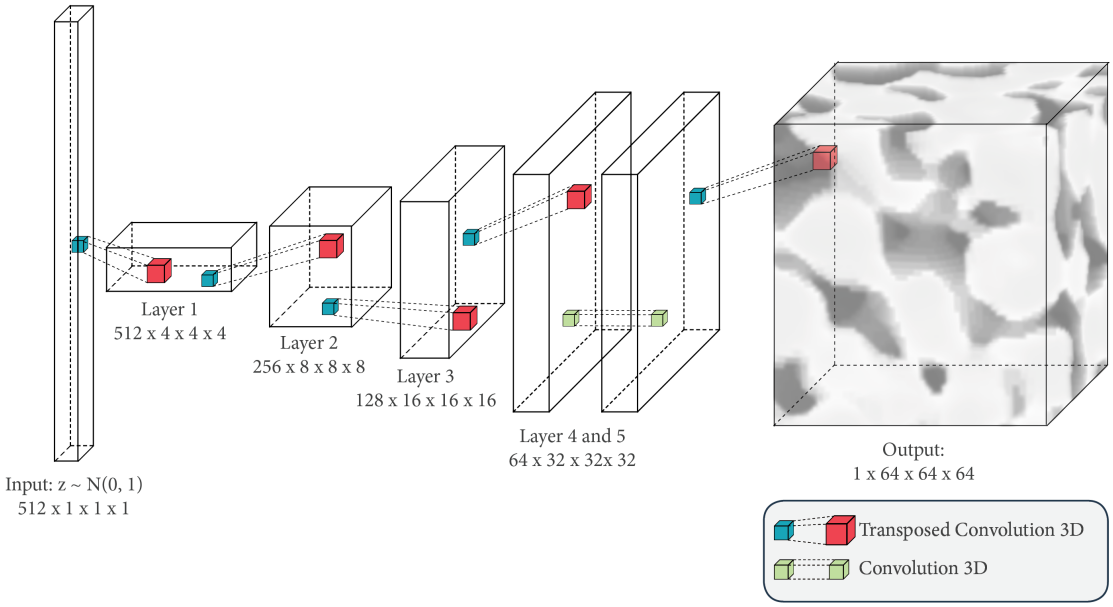

Figure 4.3: Architecture of the neural network used to represent the generator function $G_\theta(\mathbf{z})$. The latent vector $\mathbf{z}$ is passed through a fully convolutional feed-forward neural network. Transposed convolution operations upsample the image in each layer. A single convolutional layer is introduced prior to the final network layer to reduce artifacts due to upsampling using transposed convolution.

defines different networks with their own weights (parameters) which are adapted using a mini-batch gradient descent method at training time.

In total 8 models have been trained on the Ketton image dataset. The main hyper-parameters that were varied for each model are the number of filters in the generator and discriminator, N_{GF} and N_{DF} respectively, as well as the number of convolutional layers before the final transposed convolution in the generator. The dimensionality of the latent random vector $\mathbf{z}$ was kept constant at a size of $512 \times 1 \times 1 \times 1$. Learning was performed by stochastic gradient descent using the ADAM optimizer with momentum constants $\beta_1 = 0.5$, $\beta_2 = 0.999$ and a constant learning rate of 2×10^{-4}. Network training was performed on eight NVIDIA K40 GPUs using a mini-batch size of 64 images and the total run time of each training run is eight hours.

To train the pair of networks $G_\theta(\mathbf{z})$ and $D_\omega(\mathbf{x})$ we make use of two heuristic stabilization methods. First, Gaussian noise ($\mu = 0, \sigma = 0.1$) is added to the input of the discriminator which is annealed linearly over the first 300 epochs of training. A theoretical analysis of why adding Gaussian noise helps to stabilize GAN training was performed by Sønderby et al., (2016). In addition, we make use of a second stabilization method called label switching. Label switching represents a heuristic stabilization method with the aim of weakening the discriminator during the early stages of training. This heuristic stabilization method is performed by training the discriminator every N steps for one step with switched labels of the input real and generator simulated images; a real image is expected to be labeled as false and generated images as real. This corresponds to switching the expected labels of the input image mini-batches in Eq. 4.5.

Among the eight models tested, the network architecture generating realizations

Table 4.1: Architecture of the generator and discriminator networks. The generator is a fully convolutional version of a DCGAN (Radford et al., 2015) with one additional convolution layer prior to the final transposed convolution. LeakyReLU activation functions were used for all layers except the last.

Generator

Layer	Type	Filters	Kernel	Stride	Padding	Batchnorm	Activation
1	ConvTransp3D	512	$4 \times 4 \times 4$	1	0	Yes	LeakyReLU
2	ConvTransp3D	256	$4 \times 4 \times 4$	2	1	Yes	LeakyReLU
3	ConvTransp3D	128	$4 \times 4 \times 4$	2	1	Yes	LeakyReLU
4	ConvTransp3D	64	$4 \times 4 \times 4$	2	1	Yes	LeakyReLU
5	Conv3D	64	$3 \times 3 \times 3$	1	1	Yes	LeakyReLU
6	ConvTransp3D	1	$4 \times 4 \times 4$	2	1	No	Tanh

Discriminator

Layer	Type	Filters	Kernel	Stride	Padding	Batchnorm	Activation
1	Conv3D	64	$4 \times 4 \times 4$	2	1	No	LeakyReLU
2	Conv3D	128	$4 \times 4 \times 4$	2	1	Yes	LeakyReLU
3	Conv3D	256	$4 \times 4 \times 4$	2	1	Yes	LeakyReLU
4	Conv3D	512	$4 \times 4 \times 4$	2	1	Yes	LeakyReLU
5	Conv3D	1	$4 \times 4 \times 4$	1	0	No	Sigmoid

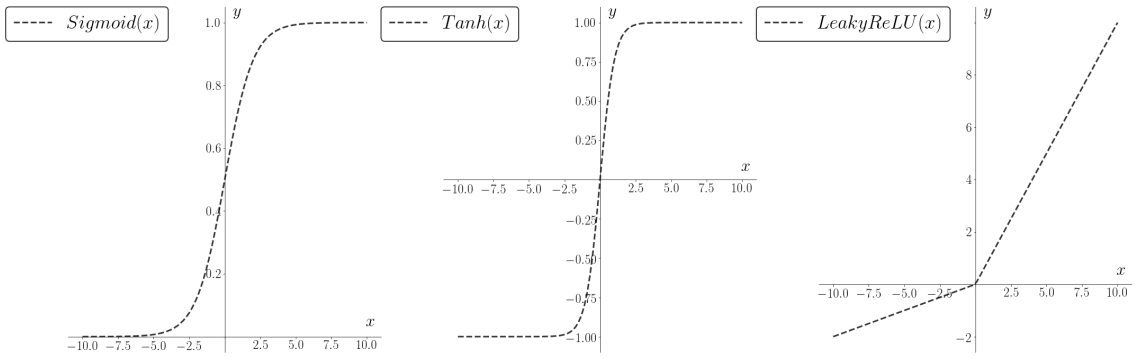

Figure 4.4: Activation functions used in the generator and discriminator networks.

with the smallest mismatch with respect to the evaluated statistical and physical properties is presented in Table 4.1. The presented model has hyper-parameters of $N_{DF} = N_{GF} = 64$. Training was stopped after 170 epochs i.e. full iterations of the training set of images. The generator consists of 27.9×10^6 adjustable parameters and 11.0×10^6 parameters for the discriminator. Visual inspection of the generated images and empirical computation of morphological and statistical properties were used as a measure for reconstruction performance at each iteration.

After training, the generator was used to create 64 reconstructions at a size of 200^3 voxels by sampling from the noise prior $\mathbf{z}$ (Eq. 4.1) and performing the mapping from the latent space to the image space (Eq. 4.2). Figure 4.5 shows slices through 32 non-overlapping sub-domains of the Ketton validation set and slices through 32 synthetic validation samples generated by the GAN model. The samples shown represent a random set of the generator output and were not selected by hand for their visual or statistical quality. The following sections present the a posteriori calculations of statistical, morphological and effective properties of these 64 synthetic

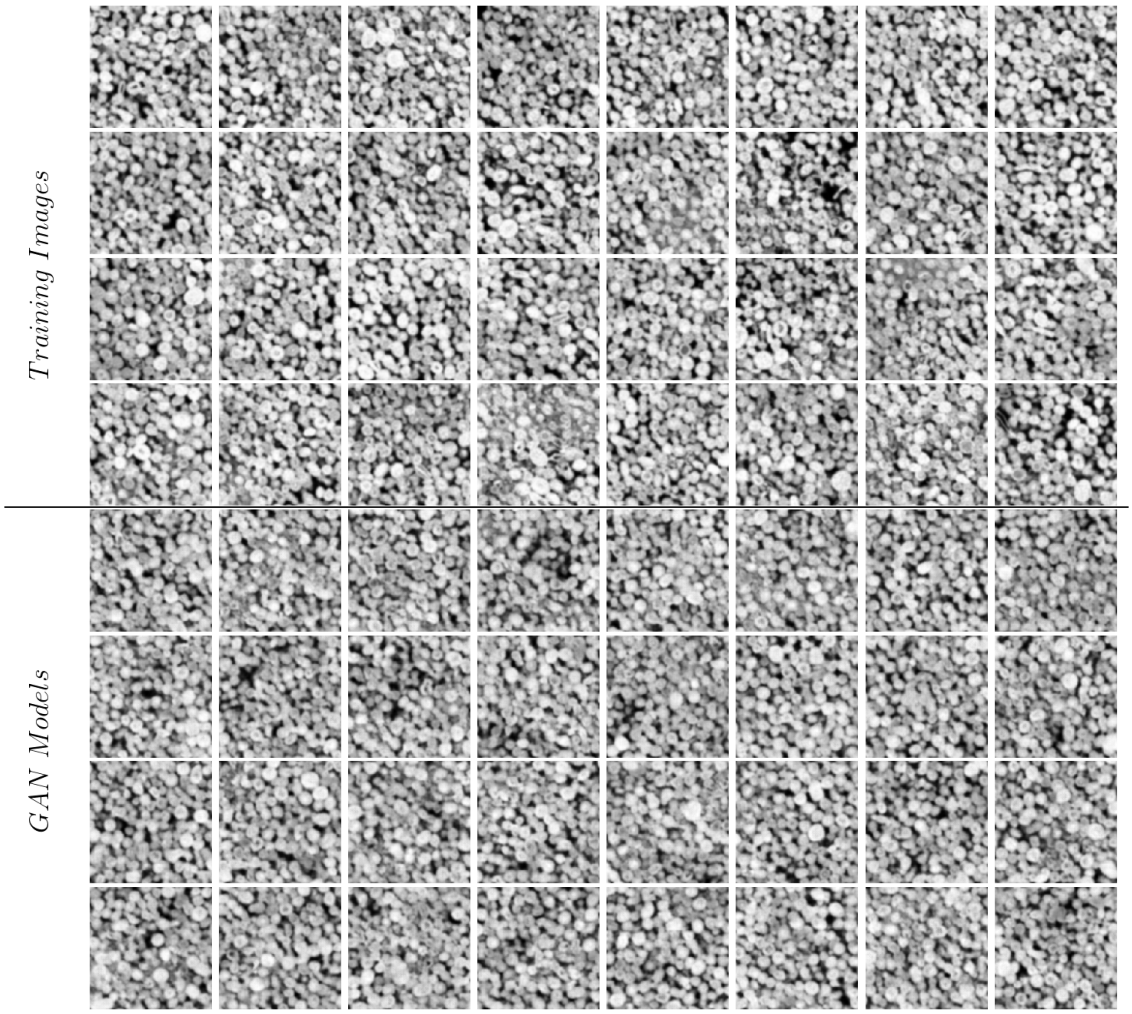

Figure 4.5: Cross-sections of the 200^3 voxel sub-domains of the Ketton micro-CT image (top) and synthetic realizations obtained from the trained generator of the generative network (bottom).

validation images in comparison to the extracted validation set of the original Ketton image (Fig. 4.5).

4.4.2 Two-Point Probability Functions

The two-point probability functions $S_2(\mathbf{r})$ allow the first and second order moments of a micro-structure to be characterized. We define the isotropic non-centered two-point probability function $S_2(\mathbf{r})$ as the probability that two arbitrary points separated by a distance $\|\mathbf{r}\|$ are located in the same phase i.e. grain or void phase of the micro-structure. While $S_2(\mathbf{r})$ may be defined for both phases of a porous medium, we compute the two-point probability function with respect to the pore phase only.

$$S_2(\mathbf{r}) = \mathbf{P}(\mathbf{x} \in P, \mathbf{x} + \mathbf{r} \in P) \; for \; \mathbf{x}, \mathbf{r} \in \mathbb{R}^d \tag{4.9}$$

$S_2(0)$ is equal to the porosity of the porous medium. Stabilization of $S_2(\mathbf{r})$ occurs around a value of ϕ^2 as the distance tends towards infinity. In addition, the specific surface area S_V can be determined from the slope of the two-point probability

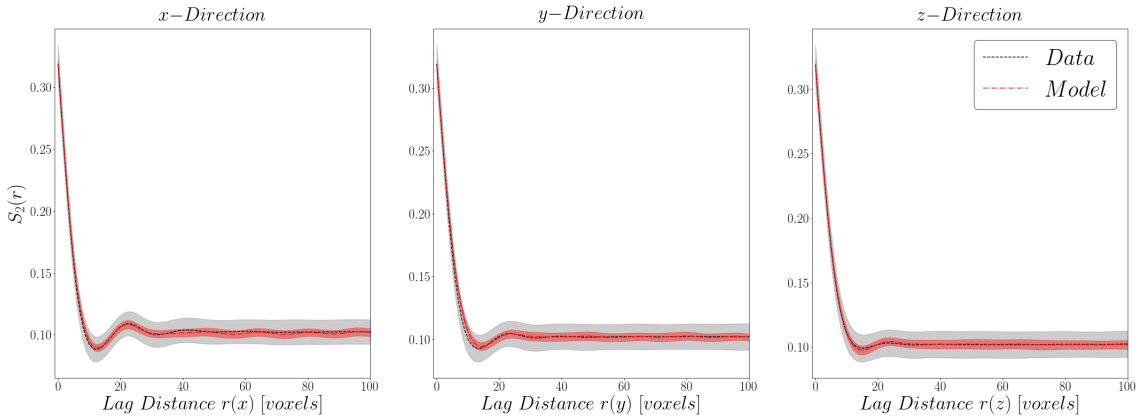

Figure 4.6: Comparison of the two-point probability function $S_2(\mathbf{r})$ measured along the Cartesian axes for Ketton image sub-domains and GAN generated realizations. $S_2(\mathbf{r})$ was measured on images after thresholding using Otsu's method. Gray and red shaded areas respectively show the variation around the average behavior ($\mu \pm \sigma$) of 64 images of the Ketton image and GAN generated validation set.

function at the origin $S_V = -4S_2'(0)$ (Berryman, 1987).

We calculate $S_2(\mathbf{r})$ numerically using the lattice point algorithm described by (Jiao et al., 2008). Figure 4.6 shows the directional two-point probability function for 64 200^3 voxel sub-domains of the original Ketton validation set (gray) and the GAN generated realizations (red). Our findings show that the 64 GAN-generated realizations lie within the standard deviation of the experimental $S_2(\mathbf{r})$ computed for the 64 original Ketton images.

Due to the ellipsoidal nature of the grains found in the Ketton limestone, a significant oscillation can be observed in all three orthogonal directions. This "hole-effect" is characteristic of periodic media (Torquato and Lado, 1985). The hole-effect found in the training image dataset is reproduced by the samples generated by the GAN model, indicating the preservation of periodic features in the pore micro-structure of the synthetic images.

Good agreement between the real and synthetic micro-structures can be observed for the radial averaged two-point probability function (Fig. 4.7). For both the radial averaged and directional estimates of $S_2(\mathbf{r})$ a tight clustering around the mean can be observed, whereas the real porous medium shows a larger degree of variation around the mean.

4.4.3 Minkowski Functionals

To evaluate the ability of the trained GAN model to capture the morphological properties of the studied Ketton limestone, we compute four integral geometric properties that are closely related to the set of Minkowski functionals as a function of the image gray value. This is an extension of the work in the previous Chapter 3 and as presented in Mosser et al., (2017), where each functional had only one value for each segmented image.

For any n-dimensional body we can define n+1 Minkowski functionals to characterize morphological descriptor of the grain-pore bodies structures (Mecke, 2000). The Minkowski functional of zeroth-order is equivalent to the porosity of a porous

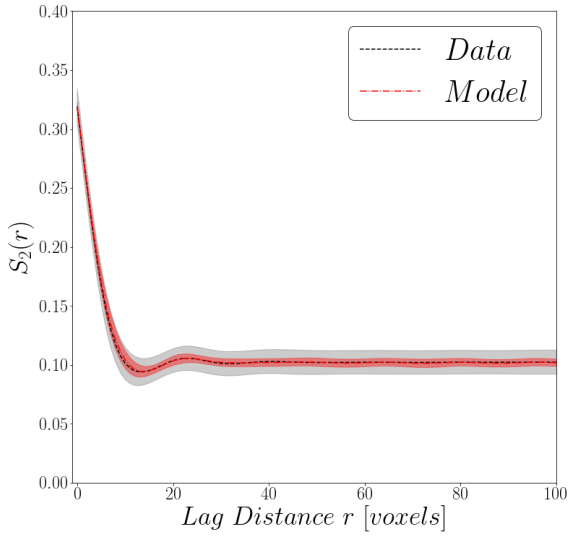

Figure 4.7: Radial average of the average two-point probability function $S_2(\mathbf{r})$) for 64 dataset sub-domains and GAN generated images. Excellent agreement of the average behavior can be observed (dashed line) whereas a lower variation around the mean behavior can be observed for the GAN generated images.

medium and defined as

$$\phi = M_0 = \frac{V_{pore}}{V} \tag{4.10}$$

where V_{pore} corresponds to the pore volume and V to the bulk volume of the porous medium.

We measure the specific surface area S_V defined as an integral geometric relationship

$$S_V = \frac{M_1}{V} = \frac{1}{V}\int dS \tag{4.11}$$

where M_1 is the Minkowski functional of first order. In three dimensions, M_1 corresponds to the surface area of the pore-grain interface. Both S_V and ϕ can be obtained by estimation of the two-point probability function $S_2(\mathbf{r})$ (Section 4.4.2). The specific surface area S_V has dimensions of $\frac{1}{length}$ and its inverse can be used to define a characteristic length scale of the porous medium.

The Minkowski functional of order 2, the integral of mean curvature, M_2, can be related to the shape of the pore-space due to its measure of the curvature of pore-grain interface. We use a bulk volume average of the integral of mean curvature defined as

$$\kappa_V = \frac{M_2}{V} = \frac{1}{2V}\int (\frac{1}{r_1} + \frac{1}{r_2})dS \tag{4.12}$$

where r_1 and r_2 are the principal radii of curvature of the pore-grain interface.

The Euler characteristic, χ_V, is a measure of connectivity that is proportional to the dimensionless third order Minkowski functional M_3

$$\chi_V = \frac{M_3}{4\pi V} = \frac{1}{4\pi V}\int \frac{1}{r_1 r_2} dS \tag{4.13}$$

We evaluate these four image morphologic properties at each of the 256 gray-level values of the 200^3 voxel Ketton image sub-domains and the GAN generated

realizations. This allows us to describe the porous medium as a set of characteristic functions dependent on a global truncation value ρ for each of the four Minkowski functionals (Schmähling, 2006; Vogel et al., 2010). To compute the four properties at each threshold level ρ the publicly available micro-structure analysis software library Quantim was used (Vogel et al., 2010).

Figure 4.8 compares these four estimated properties as a function of the image threshold value for the Ketton image (gray) and the samples generated by the GAN model (red). The shaded regions correspond to the variation around the mean $\mu \pm \sigma$ for both synthetic and real image datasets. The same 64 samples of the validation set used in the evaluation of the two-point probability function have been used for this analysis. Additionally, the vertical dashed lines represent the range of the threshold values obtained by Otsu's method when applied to the individual images. This allows an estimate of the error region that is significant when introducing a thresholding method based on a global truncation value such as Otsu's method.

Our analysis of the GAN-based models shows excellent agreement for the porosity $\phi(\rho)$, specific surface area $S_V(\rho)$ and integral of mean curvature $\kappa_V(\rho)$ as a function of the threshold value ρ. For these three properties a low error is introduced when applying global thresholding. The fourth property, the specific Euler characteristic, $\chi_V(\rho)$, shows an error of 20% in the range of global thresholding values with good agreement outside this range. This implies that care must be taken when segmenting an image - real or generated - to preserve the connectivity of the pore-space. As for the covariances, we also observe that the scatter produced by the GAN simulations is less than the scatter of the training set.

4.4.4 Permeability and Velocity Distributions

To validate GAN-based model generation for uncertainty evaluation and numerical computations it is key that the generated samples capture the relevant physical properties of the porous media that the model was trained on. The permeability and, moreover, the local velocity distributions represent the key properties of the porous medium (Menke et al., 2017).

To evaluate the ability of GAN-based models to capture the permeability and in situ velocity distributions of the Ketton training images, we solve the Stokes equation on a segmented representation of each of the 64 Ketton sub-domains and 64 synthetic pore representations created by the GAN model. The segmented representations used to estimate the two-point probability functions were reused for this evaluation. A finite-difference-based method adapted for binary representations of voxel-based pore representations was used to compute the effective permeability from the derived velocity field (Mostaghimi et al., 2013). The effective permeability was computed in the three Cartesian directions.

$$\nabla \cdot \mathbf{v} = 0 \tag{4.14a}$$

$$\mu \nabla^2 \mathbf{v} = \nabla p \tag{4.14b}$$

We present the resulting distribution of estimated permeability values as a function of the effective porosity

$$\phi_{eff} = \frac{V_{flow}}{V} \tag{4.15}$$

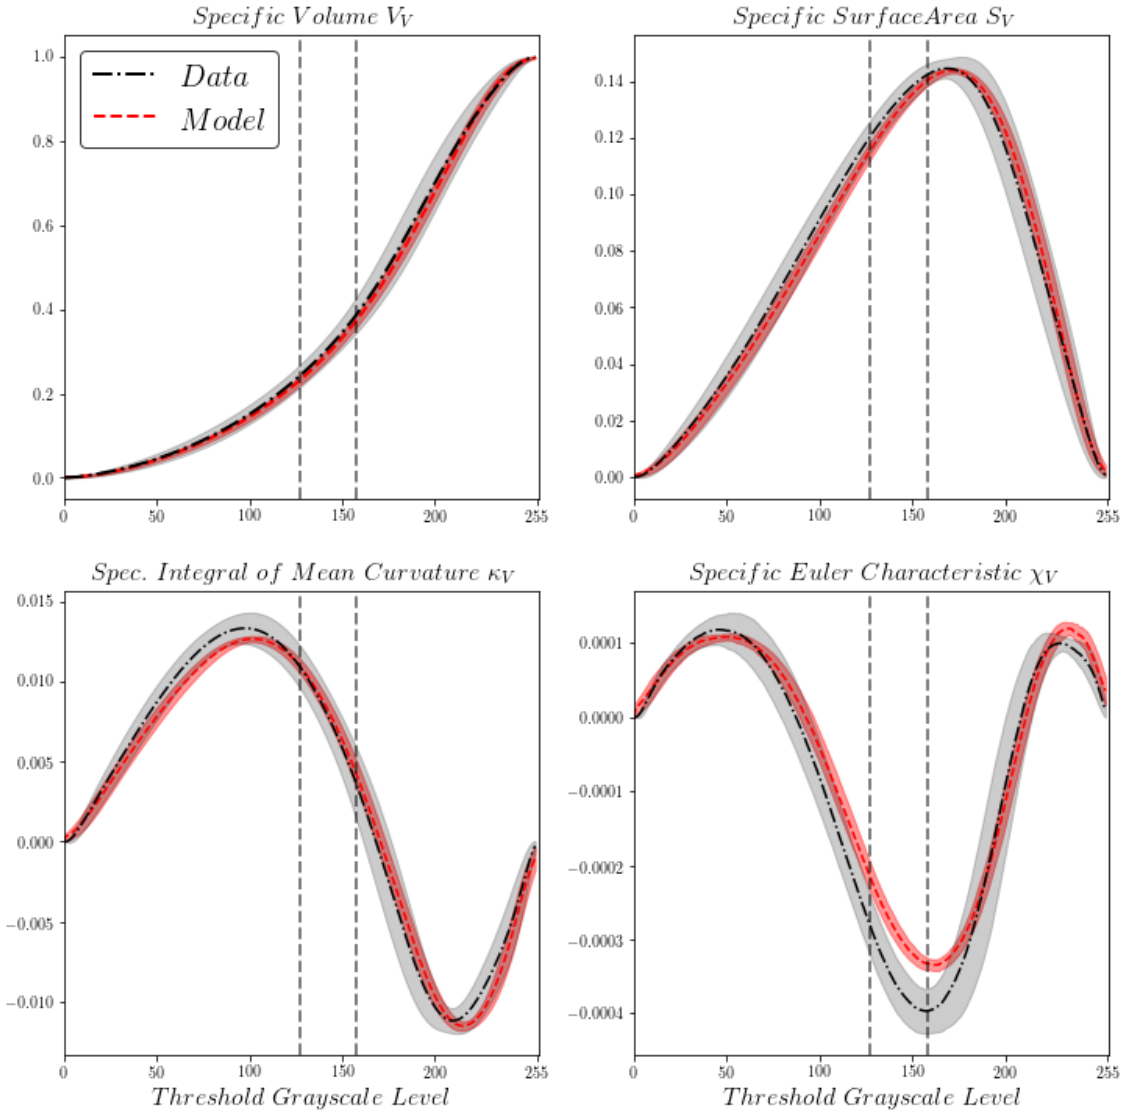

Figure 4.8: Four Minkowski functionals as a function of the segmentation threshold. Shaded regions show variation of the properties around the mean $\mu \pm \sigma$. Vertical dashed lines show the region of segmentation thresholds obtained by applying Otsu's method.

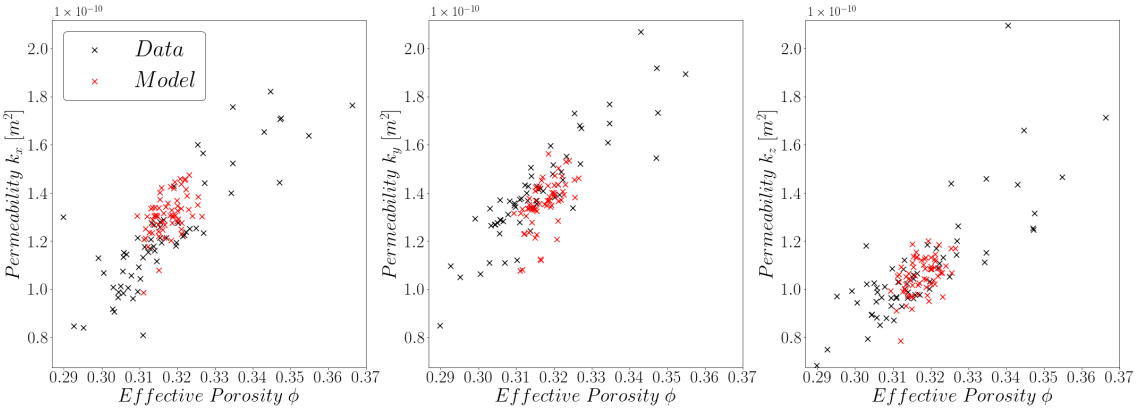

Figure 4.9: Directional permeability computed on the validation dataset (64 images with 200^3 voxels) extracted from the original Ketton limestone micro-CT dataset and realizations obtained from the GAN model. Values of permeability obtained from the synthetic images are tightly clustered around the mean of the original dataset.

where V_{flow} is the volume of the connected porosity.

Our results show (Figs. 4.9 and 4.10) that the GAN model generates stochastic reconstructions that capture the average permeability of the original training image at a scale of 200^3 voxels, with the majority of samples closely centered around the average effective permeability of the Ketton subsets.

The velocity distributions of the numerical simulations performed on the Ketton validation dataset and generated realizations were normalized by the average cell-centered velocity following the approach of Alhashmi et al., (2016) and a histogram with 256 logarithmically-spaced bins in a range from 10^{-4} to 10^2 for each simulation was obtained.

Figure 4.11 shows the per bin arithmetic average of the bin frequencies and a bounding region of one standard deviation $\mu \pm \sigma$ as the shaded area. Due to the high range of velocities spanning six orders of magnitude, the x-axis is represented in logarithmic scaling.

Visually, the distributions of the generated samples and Ketton sub-domains are nearly equivalent with minor deviations in the frequency of the very high and very low velocities. For the GAN model, low velocities are more abundant than in the original image whereas the opposite is true for high velocities.

To evaluate whether the velocity distributions obtained from numerical simulation of flow for the GAN generated images are statistically similar to distributions representative of the original image dataset we perform a two-sample Kolmogorov-Smirnov test. The null hypothesis H_0 states that two samples are of the same underlying distribution. Define $D_{n,m}$ as

$$D_{n,m} = \sup_x |F_{1,n}(x) - F_{2,m}(x)| \tag{4.16}$$

and the null hypothesis H_0 is rejected if

$$D_{n,m} > c(\alpha)\sqrt{\frac{n+m}{nm}} \tag{4.17}$$

where n and m are the sample sizes respectively and $c(\alpha) = \sqrt{-\frac{1}{2}ln(\frac{\alpha}{2})}$.

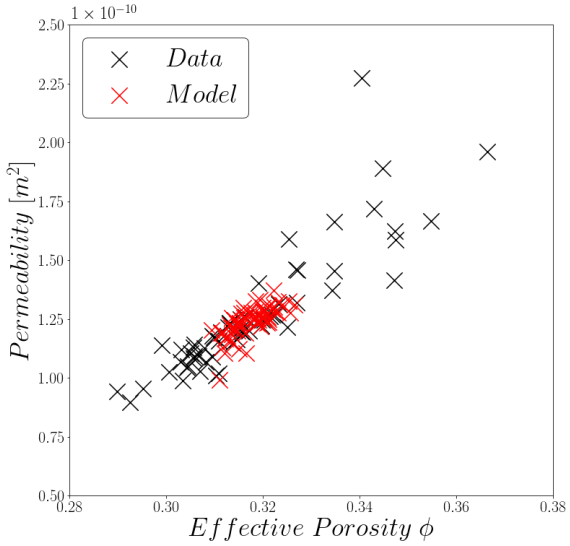

Figure 4.10: Averaged permeability for the original image datasets and synthetic realizations obtained from the GAN model.

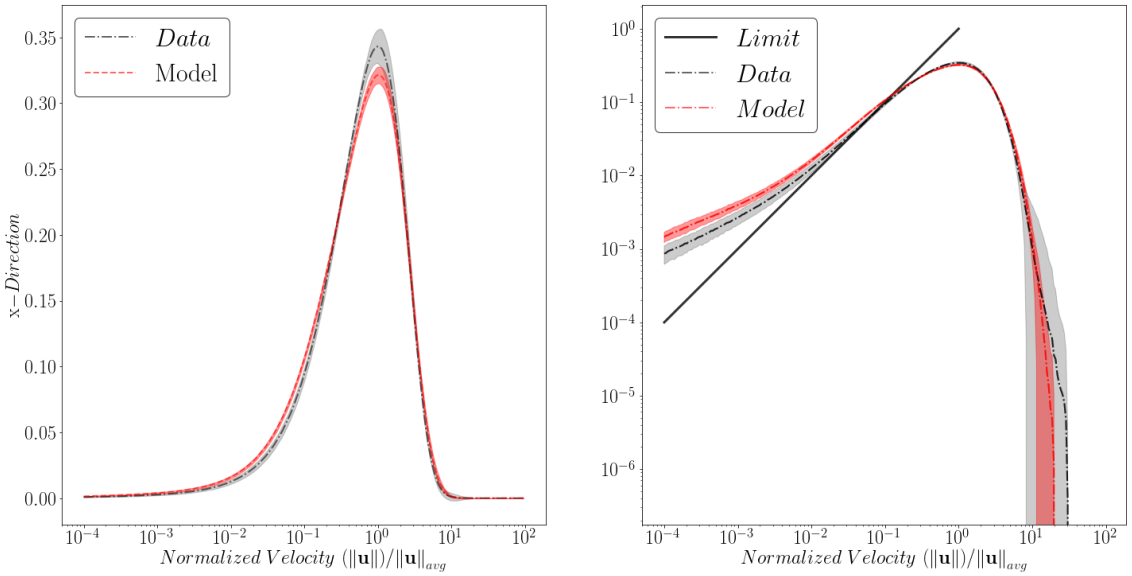

Figure 4.11: Comparison of probability density functions of the magnitude of velocity extracted from the centers of voxels in the pore-space divided by the average flow velocity plotted on semi-logarithmic (left) and double-logarithmic axes (right). The combination of 64 simulations on sub-domains obtained from the original dataset and 64 generated realizations of the GAN model are shown. Shaded regions highlight the variation around the mean of all simulations $\mu \pm \sigma$. The solid line shows the homogeneous limit velocity distribution for a single capillary tube.

Table 4.2: Results of the two sample Kolmogorov-Smirnov test for equality of velocity distributions computed on the image dataset and generated realizations. The null hypothesis of distributional equality is to be accepted at a significance level of $\alpha = 0.05$ for all three directional velocity distributions.

Direction	$D_{n,m}$	$D_{0.05}$
x	0.09	0.12
y	0.09	
z	0.07	

All tests were performed at a significance level of $\alpha = 0.05$ for the per-bin average velocity distributions presented in Fig. 4.11 (dashed curves).

For all three directions the null hypothesis can be accepted at the 5% significance level based on the $D_{0.05}$ statistic, giving evidence to the visual similarity between the velocity distributions of the real Ketton images and their synthetic counterparts (Table 4.2).

4.5 Discussion

We have presented the results of training a generative adversarial network on a micro-CT image of the oolitic Ketton limestone. The image morphological properties were evaluated as a function of the image threshold level and it was shown that the generated images capture the textural features of the original training image. Two-point statistics and effective properties computed on segmented representations of the individual sub-domains have also shown excellent agreement between the realizations generated by the GAN model and subsets of the Ketton image. Nevertheless there remain a number of open questions that need to be addressed.

The predicted statistical and morphological properties have shown a tight bound around the average behavior of the training image. This indicates that there is less variation in the generated samples than in the training samples. This behavior can have a number of origins.

The training images can be regarded as samples of the unknown multivariate pdf $p_{real}(\mathbf{x})$, which is likely to be multi-modal. The original formulation of the GAN objective function (Goodfellow et al., 2014) has been shown to lead to unimodal pdfs, even if the training set pdf itself is multi-modal (Goodfellow, 2016). The behavior of a generator to represent multi-modal pdfs by a pdf with fewer modes is called mode-collapse (Goodfellow, 2016). This behavior may occur due to the fact that there is no incentive for diversity in GAN training.

Visually the images generated by the presented GAN model are nearly indistinguishable from their real counterparts (Fig. 4.5). Minkowski functionals and statistical parameters allow us to perform an evaluation of the reconstruction quality. Nevertheless, this does not rule out the fact that the generator may be memorizing the training set, show mode-collapse behavior or output a low diversity of synthetic samples. A generator showing one or more of these behaviors will falsely indicate low errors in the Minkowski functionals, statistical and effective properties.

By visual inspection of the validation set generated by the GAN model, no evidence of identical or repeated features in the generated images could be found.

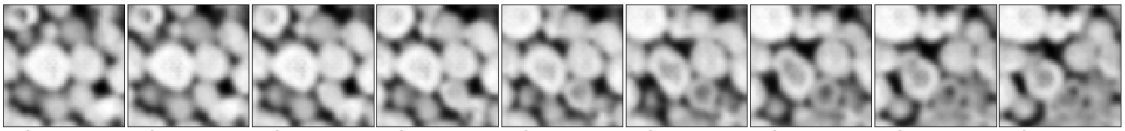

Figure 4.12: Interpolation in the latent space $\mathbf{z}$ performed for the evaluated generator G_θ shows a smooth interpolation between the start latent random vector $\mathbf{z}_{start}$ ($\beta = 1$) and the end point $\mathbf{z}_{end}$ ($\beta = 0$). An example feature of this can be seen by a bright calcite grain present in the left most image slowly being transformed into a spherical grain with significant micro-porosity.

Following the approach by Radford et al., (2015) we perform an interpolation between two points in the latent space $\mathbf{z}$:

$$\mathbf{z}_{start}, \mathbf{z}_{end} \in \mathcal{N}(0,1)^{512\times1\times1\times1}, \ \beta \in [0,1] \tag{4.18a}$$

$$\mathbf{z}_{inter} = \beta \ \mathbf{z}_{start} + (1-\beta) \ \mathbf{z}_{end} \tag{4.18b}$$

where β is a range of numbers from zero to one. This provides evidence of the generator's ability to learn meaningful representations and shows the absence of memorization.

The smooth transition between the starting image $G_\theta(\mathbf{z}_{start})$ and the endpoint $G_\theta(\mathbf{z}_{end})$ shown in Fig. 4.12 indicates that the generator has not memorized the training set and has instead learned a lower-dimensional representation $\mathbf{z}$ that results in meaningful features of the pore-grain micro-structure. Definition of GAN training objectives compatible with high-diversity samples showing no mode-collapse and stable training remains an open problem. Che et al., (2016) have presented a summary of recent advances to counteract mode-collapse,and have proposed a regularization method to improve GAN output variety. Reformulations of the GAN training criterion (Eq. 4.4) based on the Wasserstein distance (WGAN-GP) (Gulrajani et al., 2017) and other training approaches to GANs such as EBGAN (Zhao et al., 2016) or DRAGAN (Kodali et al., 2017) show the ability to model multi-modal densities and allow stable training.

It is important to note that the output of the generator is parameterized by the stochastic latent random vector and can be optimized due to the differentiable nature of the generative neural network. This is a powerful concept that has been leveraged in a number of applications in computer vision. Inpainting is the task of creating semantically meaningful content where missing data exists. Commonly this is a task performed where objects are occluded or only partially visible. In micro-structural applications and often at larger geological scales lower dimensional information may be more readily available than acquiring a full three-dimensional image e.g. thin-sections of porous media. Constraining images to this data is referred to as conditioning and can be reformulated as an inpainting problem. Yeh et al., (2016) introduced a framework for inpainting using GANs where the latent random vector can be optimized with regards to a perceptual objective function determined by the discriminator and a mismatch between the observed data and the output of the generator. In other work we have shown that the method of Yeh et al., (2016) can be applied and produces stochastic three-dimensional samples that honor the given two- and one-dimensional conditioning data (Mosser et al., 2018a).

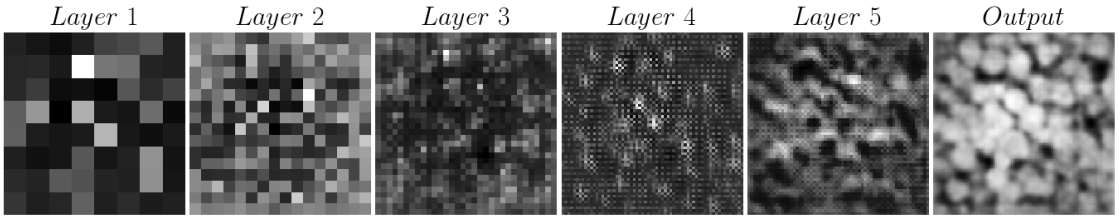

Figure 4.13: Representations of the noise prior $\mathbf{z}$ as it is propagated through the generator G_θ. Each layer adds to a multi-scale reconstruction of the final image $G_\theta(\mathbf{z})$. The shallow Layers 1, 2 and 3 introduce global features of the final image whereas deeper layers add high fidelity details to the output image. Significant noise is still present in Layer 3 due to the use of transposed convolution operations, but reduced by the convolution in Layer 5.

While the input and output to the GAN generator and discriminator is well defined, the interior mechanics of the neural network that result in high-quality reconstructions are not well understood. Rather than treating GANs as a black-box mechanism, it is of interest to evaluate the behavior of the generator and discriminator in more detail. In Fig. 4.13 we have extracted the generator's output after each layer's activation function (following the convolution operation and batch normalization).

Based on the consecutive upsampling of the noise prior $\mathbf{z}$ by each transposed convolution in the generator, we observe a multi-scale feature representation of the final image. Early layers, where the spatial dimensions of the images are small, can be related to global features in the generator output. The final layers create highly detailed representations of the structural features of the reconstructed images. This view of the generator's behavior also helps identify deficiencies in the network's architecture. In Layers 3 and 4 we see repeated noise that appears to be following a grid like structure. This is due to the transposed convolutional operation and in parts is diminished by the additional convolution operation prior to the last upsampling operation. This could be alleviated by the use of other convolution-based upsampling layers such as the sub-pixel convolution operation (Shi et al., 2016) or interpolation upsampling (nearest neighbor, bilinear, trilinear).

The discriminator's role is simply to label images as real or "fake", but it also is a critical component in the ability of the generator to learn features in the original image space. The discriminator, in order to distinguish GAN-generated from real training images, needs to learn a unique set of features that distinguish real samples from fake ones. As such, for future work, it may be of interest to use a GAN trained discriminator for classification or feature extraction (Arora and Zhang, 2017).

Nevertheless, we can perform a similar operation as for the generator and inspect some of the features learned by the discriminator. Figure 4.14 shows a set of 5 learned filters applied to an image of the Ketton training set. At shallow layers we find that the discriminator has learned to identify the pore-space (Layer 1, second row) as well as a number of edges. Deeper layers in the network represent more abstract features and after Layer 2 no original feature of the pore-space is distinguishable.

Considering that the samples used to evaluate the statistical and effective properties were not chosen by hand but represent a random group of generated images based on the GAN model, further improvement can be obtained in the reconstruction results. The discriminator may be used as an evaluation criterion for samples where

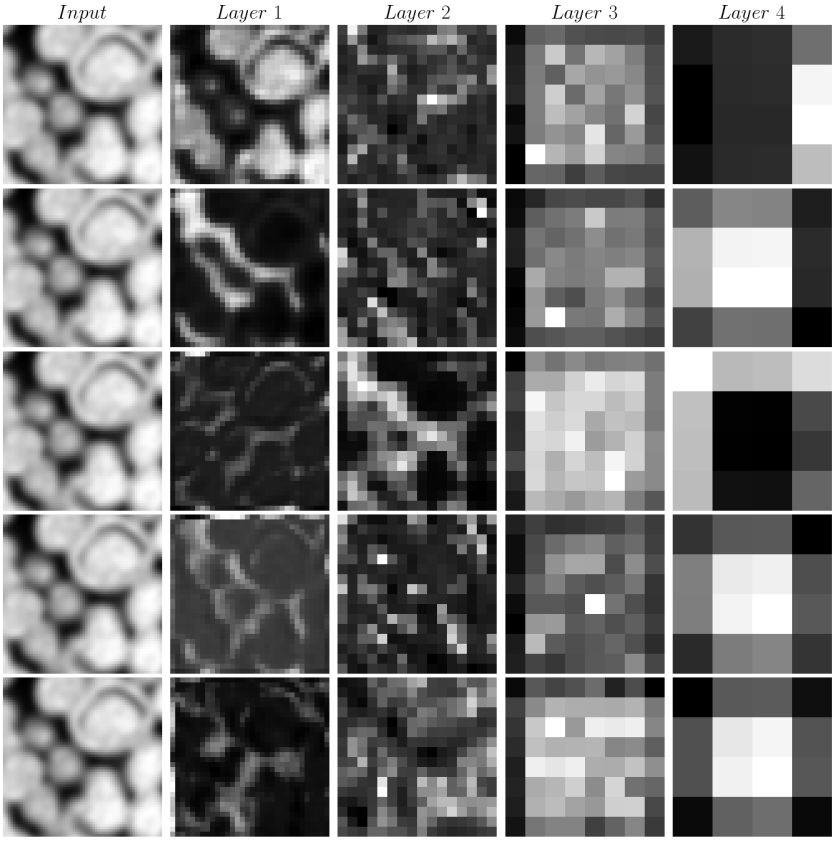

Figure 4.14: An inspection of the behavior of the discriminator's learned feature representations for a training sample of the original Ketton training image. Each column represents one layer of the discriminator network. Each row represents one learned filter kernel in each layer applied to the input (leftmost column).

higher values obtained from the discriminator $D(G_\theta(\mathbf{z}))$ indicate that the samples are closer to the real training image dataset. In this way, high-quality reconstructions may be "cherry-picked" by choosing representations that score values $D(\mathbf{x})$ close to one (real label) from a much larger set of reconstructions.

The computational effort to perform image reconstruction using GANs can be split into two parts: training time and generation time. The training time is the total time required to find a set of parameters of the generator that allows generation at sufficient image quality. We define generation time as the total time required to initialize a neural network and the associated parameters obtained during the training phase and the generation of the images by passing a latent random vector $\mathbf{z}$ through the generator to obtain an image $\mathbf{x} \sim G_\theta(\mathbf{z})$. To create one realization from a GAN it is necessary to train the generator-discriminator pairing only once, therefore training time is a fixed computational cost. Once trained, the generator can simply be reused for each new realization.

We have performed benchmarking of our GAN model in terms of the computational time. Training was performed on eight Nvidia K40 GPUs and the total training time was 8 hours. We evaluate the generation time of 100 realizations based on this set of pre-trained parameters. Each benchmark consists of the following steps: initialization of the generator parameters, sampling and initializing a latent random vector in GPU memory and finally applying the generator to the latent random

Table 4.3: Comparison of reported computational run-times of recent reconstruction methods. The computational cost for running 100 realizations is estimated based on the smallest reported run-time for each algorithm. An unbiased comparison is difficult as each method has been evaluated on a different dataset and image sizes.

Computational Runtime Comparison

Authors	Method	Size $[voxels^3]$	Runtime ($\times 1$)	Runtime($\times 100$)
Pant (Pant, 2016)	Simulated Annealing	300^3	22-47 [hours]	2200 [hrs]
Čapek et. al (Čapek et al., 2009)	Simulated Annealing	320^3	160-400 [hrs]	16000 [hrs]
Tahmasebi et. al (Tahmasebi et al., 2017)	Patch-Based	$1000^2 \times 300$	0.1 [hrs]	11 [hrs]
Okabe and Blunt (Okabe and Blunt, 2004)	MPS	150^3	12 [hrs]	1200 [hrs]
Current Work	GAN	450^3	**8** [hrs]	**8.1** [hrs]

vector $\mathbf{x} \sim G_\theta(\mathbf{z})$ to create a realization with 450^3 voxels. When sampling 100 realizations the first step, the initialization of the pre-trained generator parameters is only required once and is not repeated for subsequent sampling operations. We have repeated this benchmarking exercise ten times on an NVIDIA V100 GPU and have quoted the average total run times. Our benchmark shows that the average run time to perform sampling of 100 realizations with 450^3 voxels is 100 seconds.

The main limitations in computational effort come from two factors: the training time and available GPU memory. In the future we expect the training time to decrease, due to greater performance of GPUs and development of novel GAN training methods that allow faster convergence. Furthermore GAN-based image synthesis for large spatial domains requires large amounts of GPU memory, for example reconstruction with 450^3 voxel requires more than 10 gigabytes of GPU memory.

Recently, a number of algorithms have been developed to perform high-quality reconstruction of porous media based on training images (Jiao et al., 2009; Zachary and Torquato, 2011; Tahmasebi et al., 2017). While considering the resulting image quality to be equal, one possible differentiation of these methods is computational run time. Reported run times are heavily dependent on a number of criteria such as the simulated image size, software implementation or hardware used. Table 4.3 presents a summary of measured computational time reported for a number of recent reconstruction methods as well as their respective simulated image sizes.

Most methods reported in Table 4.3 incur a high computational cost per generated realization, with the exception of the method of Tahmasebi et al., (2017). We refer to these methods as proportional-cost methods as the computational cost scales strongly with the number of created realizations. Training-based methods, such as the presented GAN-based approach have a high initial computational cost due to the required training phase. Our method, once training is completed, has a very small generation time per realization. It is possible to determine an amortization time, when the use of one approach, considering all other factors equal, becomes beneficial.

Figure 4.15 presents a comparison of the computational cost induced by different methods as a function of the number of realizations at a fixed image size. The amortization time, where the two curves intersect, correspond to the number of realizations at which training-based methods, such as GANs, become faster.

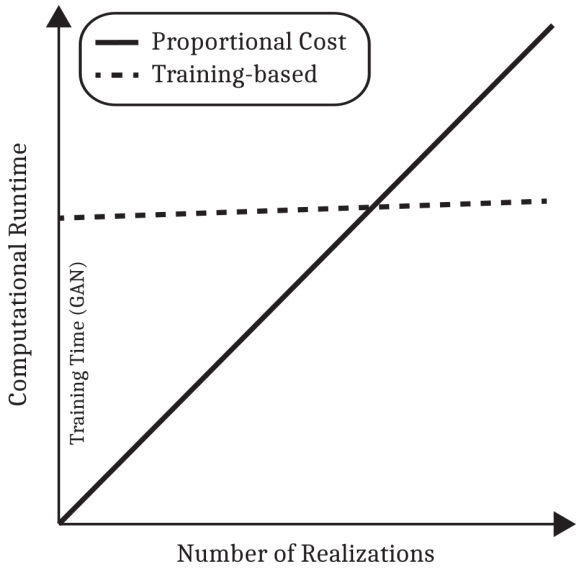

Figure 4.15: Comparison of the computational cost for two stochastic reconstruction methods at fixed image size. Proportional cost-based methods are associated with a high run time per realization. Training-based methods, such as the presented GAN method, have a high initial computational cost due to their training phase and a small cost per generated realization afterward.

4.6 Conclusions

We have presented a method to reconstruct micro-structures of porous media based on grayscale image representations of volumetric porous media. By creating a GAN-based model of an oolitic Ketton limestone, we have shown that GANs can learn to represent the statistical and effective properties of segmented representations of the pore-space as well as their Minkowski functionals as a function of the image gray-level.

In addition to the effective permeability which is associated with a global average of the velocity field, we show that the pore-scale velocity statistical distributions have been recovered by the synthetic GAN-based models. We highlight the roles of the discriminator and generator function of the GAN and show that the GAN learns a multi-scale representation of the pore-space based on inference from a latent random prior. Large hyper-parameter searches involved in the deep neural network architectures and learning instabilities make the training of GANs difficult. The high computational cost involved in training GANs is made good use of for applications when very large or many stochastic reconstructions are required. The differentiable nature of the generative network parameterized by the latent random vector provides a powerful framework in the context of gradient-based optimization and inversion techniques.

Future work will focus on creating GAN-based methodologies that ensure a valid representation of the underlying data distribution allowing application of GANs for uncertainty quantification and inversion of effective material properties.

Chapter 5

Conditioning of Three-Dimensional Generative Adversarial Networks for Pore and Reservoir-scale Models

5.1 Abstract

Geostatistical modeling of petrophysical properties is a key step in modern integrated oil and gas reservoir studies. This contribution leverages the differentiable nature of neural networks to extend GANs to the conditional simulation of three-dimensional pore- and reservoir-scale models. Based on the previous work of Yeh et al., (2016), we use a content loss to constrain to the conditioning data and a perceptual loss obtained from the evaluation of the GAN discriminator network. The technique is tested on the generation of three-dimensional micro-CT images of a Ketton limestone, as presented in the previous chapter, constrained by two-dimensional cross-sections, and on the simulation of the Maules Creek alluvial aquifer constrained by one-dimensional sections. Our results show that GANs represent a powerful method for sampling conditioned pore and reservoir samples for stochastic reservoir evaluation workflows.[1]

5.2 Introduction

The spatial distribution of rock properties within a reservoir can have a critical impact on hydrocarbon recovery. In recent years a number of geostatistical methods have been developed to generate earth models given sparse information.

Different approaches exist to model the distribution of geological facies and petrophysical properties. Variograms quantify geological and petrophysical variations using so-called two-point statistics (Matheron, 1975; Pyrcz and Deutsch, 2014). Anisotropic behavior is incorporated by introducing orientation-dependent variograms. Truncated Gaussian simulation obtains facies models by truncation of variogram-based models of Gaussian fields (Armstrong et al., 2011). In contrast, multiple-point statistical (MPS) methods (Strebelle, 2002) evaluate the dependency of the facies occurrence at a given location based on statistics available on multiple-point templates. These statistics are provided by training images, that represent the conceptual

[1] We have released our code for conditioning at `http://github.com/LukasMosser/geogan`

geological knowledge and act as a discrete prior on the geological understanding of the subsurface reservoir.

Recent methods such as direct sampling (Mariethoz et al., 2010) have led to significant reduction in the computational overhead of MPS methods, allowing faster sampling of three-dimensional reservoir models. Object-based methods populate model domains with predefined parameterized geometric representations of geobodies. This allows realistic representation of geological features, but the conditioning to well data is challenging when the size of geological objects is large compared to the well spacing.

At the reservoir scale (meters to 10s of kilometers) we have no knowledge of the true subsurface distribution of reservoir properties, except at discrete well locations; on the other hand, at the scale of individual pores of the reservoir rock, direct imaging methods such as micro-computed tomography allow images of the pore-grain structure to be made (Blunt et al., 2013; Berg et al., 2017). These images are often limited in size. Where large spatial domains are required for e.g. upscaling tasks, statistical models enable statistical and physical representations of the pore-grain structure. Due to the abundance of two-dimensional thin sections compared to three-dimensional CT measurements, models are often conditioned to match existing two-dimensional images.

In Chapters 3 and 4 we have shown that GANs are able to generate very realistic stochastic representations of pore-scale structures (Mosser et al., 2017; Mosser et al., 2018c). Chan and Elsheikh, (2017) have shown that GANs are able to produce parametric geological representations. Laloy et al., (2017) incorporated GANs in a Markov Chain Monte-Carlo approach to create representations conditional to dynamic hydraulic data. Laloy et al., (2017a) used conditional MPS simulations as a training set for variational autoencoders to sample conditioned geological representations. Yeh et al., (2016) proposed a combined "content + perceptual" loss approach and leverage the differentiable and parametric nature of the deep neural networks used to condition GAN simulations to pre-existing data.

5.3 Theory

Generative adversarial networks (GANs) (Goodfellow et al., 2014; Goodfellow, 2016) are a recent methodology developed in deep learning that allows modeling and sampling from a data distribution represented by a set of training examples. GANs consist of two differentiable functions; a generator $G_\theta(\mathbf{z})$ that maps samples obtained from a multivariate standardized normal distribution to an image $\mathbf{x}$ and a discriminator $D_\omega(\mathbf{x})$ that takes on the role of a classifier to distinguish between simulations $\mathbf{x} \sim G_\theta(\mathbf{z})$ created by the generator and the training images. Both networks are trained in an alternating two-step procedure that optimizes a min-max objective function

$$\min_\theta \max_\omega \{\mathbb{E}_{\mathbf{x}\sim p_{data}(\mathbf{x})}[log(D_\omega(\mathbf{x}))] + \mathbb{E}_{\mathbf{z}\sim p_{\mathbf{z}}(\mathbf{z})}[log(1 - D_\omega(G_\theta(\mathbf{z})))]\} \tag{5.1}$$

In this setting, the two functions have distinct and opposing objectives: the discriminator's goal is to distinguish between real training images and samples obtained from the generator, whereas the generator tries to create samples that the discriminator falsely classifies as being a sample from the data distribution. Due

to this two-player game between the generator and discriminator the training of GANs is inherently unstable. Using the Wasserstein distance as a surrogate objective function (Arjovsky et al., 2017) has proven to be a successful way of stabilizing the GAN training process. This contribution uses Wasserstein-GANs combined with a single-sided gradient penalty to train the generator-discriminator pairing (Petzka et al., 2017; Gulrajani et al., 2017). Generative adversarial networks were trained on the three-dimensional Maules Creek reservoir dataset (Mariethoz and Caers, 2014) and on a grayscale micro-CT image of a Ketton limestone. The Ketton dataset serves as a pore-scale example for model conditioning (Menke et al., 2017). An overview of the image processing performed on the Ketton dataset can be found in the previous chapter (Mosser et al., 2018c).

The generator and discriminator of each trained GAN are represented by a deep convolutional neural network (DCGAN) (Radford et al., 2015; Jetchev et al., 2016). Due to the differentiable nature of the deep neural network we can perform gradient-based optimization of the latent variables $\mathbf{z}$ with respect to an objective function $\mathcal{L}(\mathbf{z})$ acting on the output of the GAN generator $G_\theta(\mathbf{z})$

$$\mathbf{z}^* = \underset{\mathbf{z}}{\operatorname{argmin}}\ \mathcal{L}(\mathbf{z}) \tag{5.2}$$

while keeping the parameters θ of the GAN generator constant.

We solve Eq. 5.2 by gradient-descent, where we first obtain an initial sample $\mathbf{z}^0 \sim p(\mathbf{z})$ from the prior distribution of the latent-variables, in our case a multivariate Gaussian and perform gradient-descent on the latent variables

$$\mathbf{z}^{t+1} = \mathbf{z}^t - \eta \frac{\partial \mathcal{L}(\mathbf{z})}{\partial \mathbf{z}^t} \tag{5.3}$$

where the gradients of the objective function with respect to the latent variables $\frac{\partial \mathcal{L}(\mathbf{z})}{\partial \mathbf{z}^t}$ are obtained using backpropagation (Rumelhart et al., 1988; LeCun et al., 1998) and η represents the step-size.

To constrain a single realization $\mathbf{x} = G_\theta(\mathbf{z})$ of the generator to a set of continuous conditioning data, we first define the objective function $\mathcal{L}(\mathbf{z})$ referred to as the content loss in the work of Yeh et al., (2016), that is the masked mean-squared error between the generator output $G_\theta(\mathbf{z})$ and the conditioning data $\mathbf{M} \odot \mathbf{y}$

$$\mathcal{L}_{content}(\mathbf{z}) = \|\mathbf{M} \odot G(\mathbf{z}) - \mathbf{M} \odot \mathbf{y}\|_2^2 \tag{5.4}$$

where the elements of the mask $\mathbf{M} = M_i$ indicate the location of the conditioning data ($M_i = 1$) and the regions of missing data ($M_i = 0$) which we aim to reconstruct.

By minimizing the content loss according to Eq. 5.2 we obtain a realization $\mathbf{x}^* = G_\theta(\mathbf{z}^*)$ that matches the conditioning data ($M_i = 1$) and provides a synthetic output by the generator where data is missing ($M_i = 0$).

In the case where a GAN has been trained to represent binary indicator variables, the objective function $\mathcal{L}_{content}(\mathbf{z})$ is the binary cross-entropy between the values of the binary indicator variable at the conditioning locations and the GAN output. The output of the GAN, conditioned on a set of binary indicator data, represents the probability of each value of the GANs output to belong to one of two classes. After conditioning, a binary representation of the conditioned GAN output is obtained by thresholding the probability values at the 0.5 level. Yeh et al., (2016) showed that only minimizing the content loss does not lead to visual realistic results and

therefore introduced a so-called perceptual loss that is given by the discriminator's evaluation of the generator's output $D_\omega(G_\theta(\mathbf{z}))$.

$$\mathcal{L}_{perceptual}(\mathbf{z}) = \log(1 - D_\omega(G_\theta(\mathbf{z}))) \tag{5.5}$$

This perceptual loss evaluates the similarity of patterns observed on the generated samples to the samples of the training set.

The perceptual loss is therefore added to the content loss in our optimization procedure and weighted by a user-defined factor λ.

$$\mathcal{L}_{total}(\mathbf{z}) = \mathcal{L}_{content}(\mathbf{z}) + \lambda \mathcal{L}_{perceptual}(\mathbf{z}) \tag{5.6}$$

The conditioned realizations presented in this chapter therefore minimize the total loss

$$\mathbf{z}^* = \underset{\mathbf{z}}{\operatorname{argmin}}\ \mathcal{L}_{total}(\mathbf{z}) \tag{5.7}$$

by optimizing the set of latent variables $\mathbf{z}$ and keeping the parameters of the generator and discriminator fixed.

For continuous grayscale images we stop optimizing when the content loss has reached values less than 10^{-3}, whereas for binary indicator models we use a unit accuracy i.e. exact matching of the indicator variables at the conditioning locations after thresholding as the convergence criterion.

$$Accuracy = \frac{True\ Positives + True\ Negatives}{Positives + Negatives} \tag{5.8}$$

5.4 Results

We evaluate the ability of GANs to generate conditional samples by conditioning the trained GAN networks to lower-dimensional data. Conditioning the three-dimensional output of the Ketton generator network to two-dimensional micro-CT data is performed using orthogonal intersections centered at the origin (Okabe and Blunt, 2004). The location of the orthogonal conditioning planes can be seen (black) in Fig. 5.1. Although the conditioning data is lower-dimensional, this has a spatial influence on the resulting realizations in the third dimension (Fig. 5.1).

For the Maules Creek dataset we condition to a single well in the center of the domains shown in Fig. 5.2. Conditioning of 1024 Maules Creek three-dimensional simulations was performed on a single GPU in 8 hours. We present the mean and standard deviation of the ensemble of realizations in Fig. 5.2b-c. An elliptical influence of the conditioning data is observed showing that the anisotropy of the training images is properly captured by the GAN generator. Good variation in the samples is shown by the high variance of the conditioned model ensembles. Each conditioned sample matches the indicator data at the well exactly.

5.5 Conclusions

Generative adversarial networks are a new powerful machine learning approach for generating three-dimensional simulations of porous media at the reservoir and pore scale. One pore-scale and one reservoir-scale case study have shown that conditional simulation can also be performed, constrained by lower-dimensional well data at the reservoir scale or cross-sections at the pore scale.

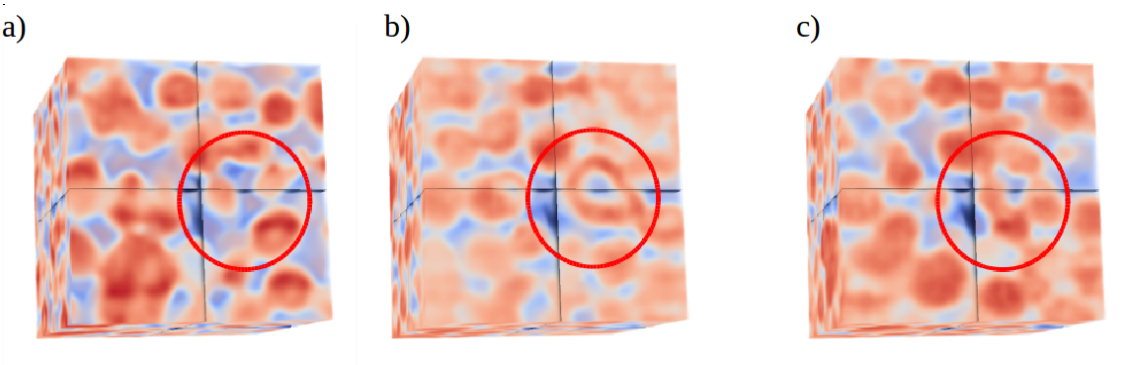

Figure 5.1: (a) A subset of the Ketton limestone training image has been used to extract orthogonal two-dimensional cross-sections used as conditioning data (black planes). (b-c) Two conditioned simulations of a GAN trained on the Ketton dataset are shown. The same lower-dimensional conditioning data has a different volumetric expression away from the conditioning planes.

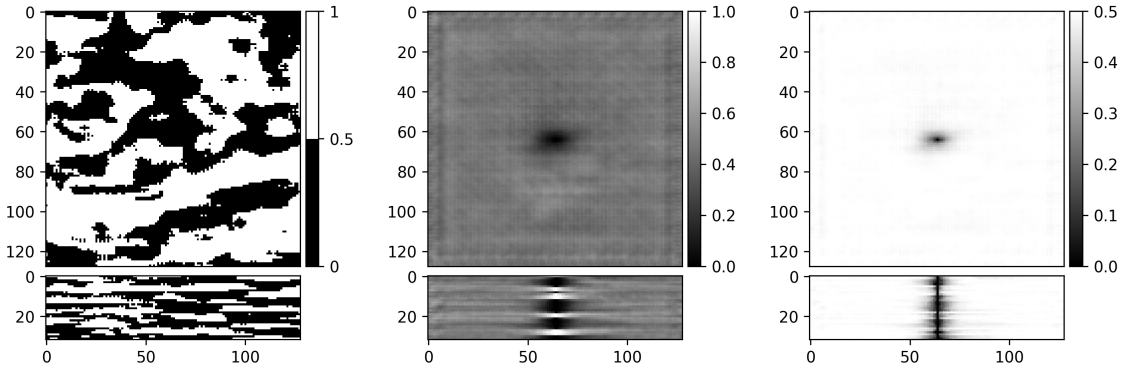

Figure 5.2: (a) A GAN was used to create simulations of the Maules Creek training image conditional to a single centered well. (b-c) Mean and standard deviation maps and cross-sections of 1024 conditioned simulations were created by optimizing the latent vectors of GAN samples trained on the Maules Creek training image. An elliptical region of influence can be observed around the conditional well data. Boundary artifacts observed in the mean of all samples occur due to transposed convolutional layers used in the generator architecture.

Chapter 6

Stochastic Seismic Waveform Inversion using Generative Adversarial Networks as a Geological Prior

6.1 Abstract

We present an application of deep generative models in the context of partial-differential equation constrained inverse problems. We combine a generative adversarial network representing an a priori model that generates geological heterogeneities and their petrophysical properties, with the numerical solution of the partial-differential equation governing the propagation of acoustic waves within the earth's interior. We perform Bayesian inversion using an approximate Metropolis-adjusted Langevin algorithm to sample from the posterior distribution of earth models given seismic observations. Gradients with respect to the model parameters governing the forward problem are obtained by solving the adjoint of the acoustic wave-equation. Gradients of the mismatch with respect to the latent variables are obtained by leveraging the differentiable nature of the deep neural network used to represent the generative model. We show that approximate Metropolis-adjusted Langevin sampling allows efficient Bayesian inversion of model parameters obtained from a prior represented by a deep generative model, obtaining a diverse set of realizations that reflect the observed seismic response.

6.2 Introduction

Solving an inverse problem means finding a set of model parameters that best fit observed data (Tarantola, 2005). The observed data or measurements are often noisy and/or sparse and therefore lead to an ill-posed inverse problem where numerous realizations of the underlying model parameters may lead to a model response that matches observed data (Kabanikhin, 2008). Additionally, the model used to describe how the observed data are generated, the so-called forward model, may be uncertain (Hansen and Cordua, 2017).

Based on natural observations or an understanding of the underlying data gen-

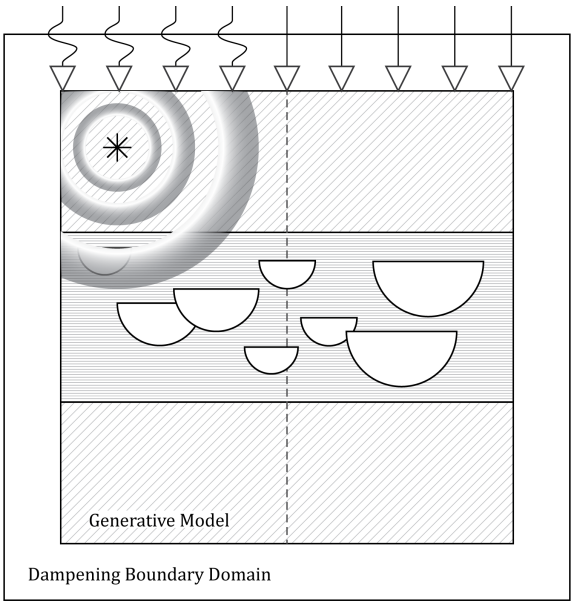

Figure 6.1: Computational domain for the acoustic inversion problem. Acoustic recording devices are placed on the surface (∇) and record incoming acoustic waves reflected from geological structures and emanating from an artificial source ($*$). The computational domain is embedded within a dampened boundary domain to emulate lateral and vertical dissipation of the wave-source. The generative model $G_\theta(\mathbf{z})$ creates the underlying spatially distributed P-wave velocity. Additional lower-dimensional constraints (dashed vertical line representing a well) can be placed on the generative model, by incorporating loss terms. The vertical axis of the computational domain has been rescaled by a factor of 10 for visualization purposes.

erating process we may have a preconception about possible or impossible states of the model parameters. We may formulate this knowledge as a prior probability distribution function (pdf) of our model parameters and use Bayesian inference to obtain a posterior pdf of the model parameters given the observations (Tarantola, 2005).

Seismic inversion involves modeling the physical process of waves radiating through the earth's interior (Fig. 6.1). By comparing the simulated synthetic measurements to actual acoustic recordings of reflected waves we can modify model parameters and minimize the misfit between synthetic data and measurements. The adjoint of the partial differential equation (PDE) represents the gradient of the data mismatch with respect to the parameters, leading to a gradient-based optimization of the model parameters (Plessix, 2006). In the most general case, which has been used in this study, these gradients are obtained by backpropagating the full wavefield in time an approach commonly referred to as full-waveform inversion (FWI). The set of parameters represented by the spatial distribution of the acoustic velocity of the rocks within the earth can easily exceed 10^6 values depending on the resolution of the simulation grid and the observed data. Large three-dimensional seismic observations may require millions of parameters to be inverted for, demanding enormous computational resources (Akcelik et al., 2003).

Direct observations of the earth's interior, boreholes, may have been drilled for hydrocarbon exploration/development or hydrological measurements. These represent a quasi one-dimensional source of information of spatially sparse nature. Typical borehole sizes are on the order of 10s of centimeters in diameter whereas the vertical resolution of seismic observations is usually on the order of 10s of meters.

We can deduce prior knowledge of the earth's interior from observations of analog outcrops or subsurface reservoirs. This geological knowledge can be incorporated into prior distributions of physical properties of rocks, such as the acoustic P-wave velocity, or into the distribution of geological features such as geological facies and fault distributions within the earth.

Efficient parameterizations (Akcelik et al., 2002; Kadu et al., 2016) that allow a dimensionality-reduced representation of the high-dimensional parameter space of possible models have been shown to reduce computational cost and increase spatial resolution. Due to the high computational cost incurred by full-waveform inversion (Modrak and Tromp, 2015; Akcelik et al., 2003), probabilistic ensembles of models that match observed data are rarely generated and often only a single model that satisfies pre-defined quality criteria is created and used for interpretation and decision making processes. We parameterize the earth model by a deep generative model that creates stochastic realizations of possible model parameters. The probabilistic distribution of model parameters is parameterized by a lower-dimensional set of multi-Gaussian distributed latent variables. Combined with a generative deep neural network this represents a differentiable prior on the possible model parameters. We combine this differentiable generative model with the numerical solution of the acoustic wave equation to produce synthetic acoustic observations of the earth's interior (Louboutin et al., 2017). Using the adjoint method (Plessix, 2006), we compute a gradient of the mismatch between real and synthetic data with respect to model parameters not only in the high-dimensional model space, but also in the much smaller set of latent variables. These gradients are required to perform a Metropolis-adjusted Langevin (MALA) sampling of the posterior of the model

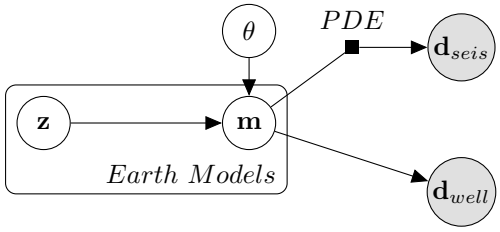

Figure 6.2: Graphical model of the geological inversion problem. The set of possible earth models is represented by a generative model with parameters θ (the parameters of the generator $\mathbf{m} \sim G_\theta(\mathbf{z})$). We obtain model observations of the acoustic waves $\mathbf{d}_{seis}$ via the deterministic PDE as well as partial observation of the model parameters $\mathbf{m}$ from local information at e.g. boreholes $\mathbf{d}_{well}$.

parameters given the observed seismic data. Performing MALA sampling allows us to obtain a diverse ensemble of model parameters that match the observed seismic data. Additional constraints on the generative model, such as information located at existing boreholes, are readily incorporated and included in the MALA sampling procedure.

We summarize our contributions as follows:

(i) We combine a differentiable generative model controlled by a set of latent variables with the solution of a PDE-constrained numerical solution of a physical forward problem.

(ii) We use gradients obtained from the adjoint method and from neural network backpropagation to perform approximate MALA sampling of the posterior in the lower-dimensional set of latent variables.

(iii) We illustrate the proposed inversion framework using a simple synthetic seismic inversion problem and evaluate the resulting ensemble of model parameters.

(iv) The framework allows integration of additional information, such as the knowledge of geological facies along one-dimensional vertical boreholes.

(v) The proposed approach may readily be extended to a number of inverse problems where gradients of the objective function with respect to input parameters can be calculated.

The code, data and trained weights of the neural networks have been made available under an open-source license[1].

6.3 Related Work

Tarantola, (2005) cast the geophysical seismic inversion problem in a Bayesian framework. Mosegaard, (1995) presented a general methodology to perform probabilistic inversion using Monte Carlo sampling. They used a Metropolis rule combined with a sampling of the prior to obtain the posterior distribution. In a similar manner, Sen

[1]Code Repository: https://github.com/LukasMosser/Stochastic_Seismic_Waveform_Inversion

and Stoffa, (1996) evaluated the use of Gibbs sampling to obtain *a posteriori* model parameters and evaluate parameter uncertainties. Mosegaard, (1998) showed that the general Bayesian inversion approach of Mosegaard, (1995) also gives information on the ability to resolve geological features. Geostatistical models allow spatial relationships and dependencies of the petrophysical parameters to be modeled and incorporated into a stochastic inversion framework (Bortoli et al., 1993; Haas and Dubrule, 1994). Bayesian linear inversion has successfully been applied to infer petrophysical property distributions (Grana and Della Rossa, 2010). Buland and Omre, (2003) developed an approach to perform Bayesian inversion for elastic petrophysical properties in a linearized setting. Grana et al., (2017) used a Gaussian mixture model for Bayesian linear inversion from seismic and well data. Stochastic sampling of petrophysical properties conditioned to well-log data allows petrophysical property distributions to be inferred using an appropriate sampling strategy such as Markov-Chain-Monte-Carlo (MCMC) (Bosch et al., 2009). A fully integrated stochastic inversion method that allows direct inversion from seismic Amplitude versus Angle (AVA) data creates a direct link between observed seismic data and underlying rock physics models (Azevedo et al., 2019). Geological modeling using multiple-point statistics (Guardiano and Srivastava, 1993) can be used for inversion from seismic data (González et al., 2007) where geological features are represented by a set of representative training images. For a more extensive review of statistical inversion approaches we refer to Bosch et al., (2010) and the comprehensive overviews of Dubrule, (2003), Doyen, (2007), and Azevedo and Soares, (2017).

In the case of non-linear physics-based inversion schemes such as FWI, computation of the solution to the forward problem is very expensive. Therefore, computationally efficient approximations to the full solution of the wave equation may allow efficient solutions to complex geophysical inversion problems. Neural networks have been shown to be universal function approximators (Hornik et al., 1989) and as such lend themselves as possible proxy-models for solutions to the geophysical forward and inverse problem (Hansen and Cordua, 2017).

The early work by Röth and Tarantola, (1994) presents an application of neural networks to invert from acoustic time-domain seismic amplitude responses to a depth profile of acoustic velocity in a supervised setting. They used pairs of synthetic data and velocity models to train a multi-layer feed-forward neural network with the goal of predicting acoustic velocities from recorded data only. They showed that neural networks can produce high resolution approximations to the solution of the inverse problem based on representations of the input model parameters and resulting synthetic waveforms alone. In addition, they showed that neural networks can invert for geophysical parameters in the presence of significant levels of acoustic noise.

Representing the geophysical model parameters at each point in space quickly leads to a large number of model parameters especially in the case of three-dimensional problems. Berg and Nyström, (2017) represented the spatially varying coefficients that govern the solution of a PDE by a neural network. The neural network acts as an approximation to the spatially varying coefficients characterized by the weights of the neural network. The weights of the individual neurons are modified by leveraging the adjoint-state equation in the reduced-dimensional space of network-parameters rather than at each spatial location of the computational grid.

Hansen and Cordua, (2017) replaced the solution of the partial differential

equation by a neural network allowing fast computation of forward models and facilitating a solution to the inversion problem by Monte-Carlo sampling. Araya-Polo et al., (2018) used deep neural networks to perform a mapping between seismic features and the underlying P-wave velocity domain; they validated their approach based on synthetic examples. Recently, a number of applications of deep generative priors have been presented in the context of computer vision for image reconstruction, linear (Chang et al., 2017) and bilinear (Asim et al., 2018) inverse problems, as well as compressed sensing (Bora et al., 2017). Mosser et al., (2017) proposed GANs to generate three-dimensional stochastic realizations of porous media from binary and gray-scale computer-tomography images (Mosser et al., 2018d). These deep generative models can further be conditioned to honor lower-dimensional features such as cross-sections or borehole data (Dupont et al., 2018; Mosser et al., 2018a; Chan and Elsheikh, 2018b). For more general subsurface inverse problems, Laloy et al., (2017) used a GAN to create geological models for hydrological inversion. Inversion was performed using an adapted Markov-Chain Monte-Carlo (MCMC) (Laloy and Vrugt, 2012) algorithm where the generative model was used as an unconditional prior to sample hydrological model parameters. Chan and Elsheikh, (2017) evaluated the applicability of Wasserstein-GANs to parameterize geological models for uncertainty propagation.

Mosser et al., (2018b) used a GAN with cycle-constraints (cycle-GAN) (Zhu et al., 2017) to perform seismic inversion formulating the inversion task as a domain-transfer problem. Their work used a cycle-GAN to map between the seismic amplitude domain and P-wave velocity models. The cycle-constraint ensures that models obtained by transforming from the amplitude to P-wave velocity representation and back to the amplitude domain are consistent. Due to the P-wave velocity models and seismic amplitudes being represented as a function of depth, rather than depth and time respectively, this approach lends itself to stratigraphic inversion, where a pre-existing velocity model is used to perform time-depth conversion of the seismic amplitudes. Richardson, (2018) showed that a quasi-Newton method can optimize model parameters in the latent-space of a pre-trained GAN for a synthetic salt-body benchmark dataset.

6.4 Problem Definition

6.4.1 Bayesian Inversion

In the Bayesian framework of inverse problems we aim to find the posterior of latent variables $\mathbf{z}$ given the observed data $\mathbf{d}_{obs}$ (Fig. 6.2). The joint probability of the latent variables $\mathbf{z}$ and observed data $\mathbf{d_{obs}}$ is

$$p(\mathbf{z}, \mathbf{d}_{obs}) = p(\mathbf{d}_{obs}|\mathbf{z})p(\mathbf{z}) \tag{6.1}$$

Furthermore, by applying Bayes rule, we define the posterior over the latent variables $\mathbf{z}$ given the observed seismic data $\mathbf{d}_{obs}$

$$p(\mathbf{z}|\mathbf{d}_{obs}) = \frac{p(\mathbf{d}_{obs}|\mathbf{z})p(\mathbf{z})}{p(\mathbf{d}_{obs})} \propto p(\mathbf{d_{obs}}|\mathbf{z})p(\mathbf{z}) \tag{6.2}$$

We express the observed data by assuming conditional independence between the observed seismic data $\mathbf{d}_{seis}$ and data observed at the wells $\mathbf{d}_{well}$

$$p(\mathbf{d}_{obs}|\mathbf{z}) = p(\mathbf{d}_{seis}|\mathbf{z})p(\mathbf{d}_{well}|\mathbf{z}) \tag{6.3}$$

We represent the observed seismic data by

$$\mathbf{d}_{seis} = S(\mathbf{m}) + \varepsilon, \varepsilon \sim \mathcal{N}(0, \sigma^2_{seis}\mathbf{I}) \tag{6.4}$$

where $S(\mathbf{m}) = S(m(\mathbf{x})) = S(G_\theta(\mathbf{z}))$, denoting the spatial model coordinates by $\mathbf{x}$, the seismic forward modeling operator by S, and the generative model by $G_\theta(\mathbf{z})$ with parameters θ. We assume a normally distributed noise term ε with zero mean and standard deviation σ_{seis} equal to 25% of the standard deviation of the reference model seismic amplitude data. The geological facies $\mathbf{m}^{facies}$, the P-wave velocity $\mathbf{m}^{V_p}$, and the rock density $\mathbf{m}^{\rho}$ represent the set of model parameters $\mathbf{m}$. The model parameter $\mathbf{m}^{facies}$ represents the probability of a geological facies to occur at a spatial location $\mathbf{x}$.

The aim is to generate samples of the posterior $\mathbf{z} \sim p(\mathbf{z}|\mathbf{d}_{obs})$. We reformulate the approach using an iterative approximate Metropolis-adjusted Langevin sampling rule (MALA-approx) with iteration number t as follows (Roberts and Tweedie, 1996; Roberts and Rosenthal, 1998; Nguyen et al., 2016)

$$\mathbf{z}_{t+1} = \mathbf{z}_t + \gamma_t \nabla \log\ p(\mathbf{z}_t|\mathbf{d}_{obs}) + \eta_t \tag{6.5a}$$

$$\mathbf{z}_{t+1} = \mathbf{z}_t + \gamma_t \nabla \log[p(\mathbf{d}_{seis}|\mathbf{z}_t)p(\mathbf{d}_{well}|\mathbf{z}_t)p(\mathbf{z}_t)] + \eta_t \tag{6.5b}$$

$$\mathbf{z}_{t+1} = \mathbf{z}_t + \gamma_t \{\nabla \log p(\mathbf{d}_{seis}|\mathbf{z}_t) + \nabla \log p(\mathbf{d}_{well}|\mathbf{z}_t) + \nabla \log p(\mathbf{z}_t)\} + \eta_t \tag{6.5c}$$

where $\eta_t \sim \mathcal{N}(0, 2\gamma_t\mathbf{I})$ is a sample from a Gaussian distribution with variance proportional to the step-size γ_t at MALA iteration t. Assuming a Gaussian log-likelihood of the seismic data given the latent variables $\log p(\mathbf{d}_{seis}|\mathbf{z}_t) \propto -\|S(G_\theta(\mathbf{z}_t)) - \mathbf{d}_{seis}\|_2^2$ leads to the proposal rule of the MALA approximation (Nguyen et al., 2016) for the case when only seismic observations $\mathbf{d}_{seis}$ are considered

$$\mathbf{z}_{t+1} = \mathbf{z}_t - \gamma_t \frac{\partial \|S(G_\theta(\mathbf{z}_t)) - d_{seis}\|_2^2}{\partial G_\theta(\mathbf{z}_t)} \frac{\partial G_\theta(\mathbf{z}_t)}{\partial \mathbf{z}_t} + \gamma_t \nabla \log p(\mathbf{z}_t) + \eta_t \tag{6.6}$$

Using this sampling approach requires gradients of the data mismatch with respect to model parameters, which are obtained by the adjoint-state method which will be presented in the following section. The gradients of the model parameters $\frac{\partial G_\theta(\mathbf{z}_t)}{\partial \mathbf{z}_t}$ with respect to the latent variables are obtained by traditional neural network backpropagation. The gradient of the log-probability of the Gaussian prior distribution of latent variables $\nabla \log p(\mathbf{z}_t)$ can be interpreted as a regularization of the latent variables against deviation from the Gaussian prior assumption (Creswell and Bharath, 2018).

We follow the MALA step-proposal algorithm using an initial step size $\gamma_{t=0} = 10^{-2}$ for every model inference (Xifara et al., 2013). To obtain valid samples of the posterior we furthermore anneal the step size from the initial value of $\gamma_{t=0} = 10^{-2}$ to $\gamma_{t=200} = 10^{-5}$ over 200 iterations.

Where lower-dimensional information is available, such as at boreholes, the geological models should honor not only the seismic response but also this additional

lower-dimensional information. In this study we additionally find samples of the posterior that reflect observed geological facies indicators $\mathbf{d}_{well} = \mathbf{m}_{well}^{facies}$ at a one-dimensional borehole. When including borehole information the step-proposal corresponds to

$$\begin{aligned} \mathbf{z}_{t+1} = \mathbf{z}_t - \gamma_t \frac{\partial \|S(G_\theta(\mathbf{z}_t)) - d_{seis}\|_2^2}{\partial G_\theta(\mathbf{z}_t)} \frac{\partial G_\theta(\mathbf{z}_t)}{\partial \mathbf{z}_t} \\ + \gamma_t \frac{\partial \log\ p(\mathbf{d}_{well} = \mathbf{m}_{well}^{facies} | \mathbf{z}_t)}{\partial \mathbf{z}_t} + \gamma_t \nabla \log p(\mathbf{z}_t) + \eta_t \end{aligned} \tag{6.7}$$

where we obtain samples of the posterior given the observed seismic data $\mathbf{d}_{seis}$ and geological facies at the wells $\mathbf{d}_{well} = \mathbf{m}_{well}^{facies}$.

The additional term $\log\ p(\mathbf{d}_{well} = \mathbf{m}_{well}^{facies} | \mathbf{z}_t)$ in Eq. 6.7 represents the assumption of a Bernoulli distribution for the facies as derived from the generator and observed at the borehole.

6.4.2 Adjoint-State Method

We perform numerical solutions of the time-dependent acoustic wave equation given a set of model parameters

$$F(\mathbf{u}, \mathbf{m}^{V_p}) = \frac{1}{m^{V_p}(\mathbf{x})^2} \frac{d^2 u(\mathbf{x}, t)}{dt^2} - \Delta u(\mathbf{x}, t) + \eta \frac{du(\mathbf{x}, t)}{dt} - q(\mathbf{x}, \mathbf{x_s}, t) = 0 \tag{6.8}$$

where $u(\mathbf{x}, t)$ is the unknown wave-field and $m^{V_p}(\mathbf{x})$ is the acoustic P-wave velocity. The dampening term $\eta \frac{du(\mathbf{x},t)}{dt}$ prevents reflections from domain boundaries and ensures that waves dissipate laterally. We refer to the evaluation of $F(\mathbf{u}, \mathbf{m}^{V_p}) = 0$ (Eq. 6.8) as the forward problem.

Time-dependent source wavelets $q(\mathbf{x}, \mathbf{x_s}, t)$ are introduced at locations $\mathbf{x}_s$. We emulate the seismic acquisition process by placing regularly spaced acoustic receivers that record the incoming wave-field at the top edge of the simulation domain (Fig. 6.1). To show the impact of adding additional information from the acoustic forward problem to the posterior pdf of models, we perform Bayesian inversion using the proposed approach in a number of scenarios where we increase the number of acoustic shot data from 2 to 27 acoustic sources.

To perform sampling according to the MALA algorithm presented in Eq. 6.6, we seek to obtain a gradient of the following functional

$$J(m^{V_p}(\mathbf{x})) = \sum_{i=1}^{n_{sources}} \|\mathbf{d}_{seis}^{pred}(m^{V_p}(\mathbf{x})), q_i) - \mathbf{d}_{seis}\|_2^2 \tag{6.9}$$

where $\mathbf{d}_{seis}^{pred}$ and $\mathbf{d}_{seis}$ are the predicted and observed seismic observations respectively.

We augment the functional $J(m^{V_p}(\mathbf{x}))$ by forming the Lagrangian

$$\mathcal{L}(\mathbf{m}^{V_p},\ \mathbf{u},\ \lambda) = J(\mathbf{m}^{V_p}) - \langle \lambda,\ F(\mathbf{u}, \mathbf{m}^{V_p}) \rangle \tag{6.10}$$

Differentiating $\mathcal{L}(\mathbf{m}^{V_p},\ \mathbf{u},\ \lambda)$ with respect to λ leads to the state equation (Eq. 6.8), but differentiation with respect to the acoustic wave-field $\mathbf{u}$ leads to the adjoint state equations (Plessix, 2006):

$$(\frac{\partial F(\mathbf{u}, \mathbf{m}^{V_p})}{\partial \mathbf{u}})^T\ \lambda = (\mathbf{d}_{seis}^{pred} - \mathbf{d}_{seis}) \tag{6.11}$$

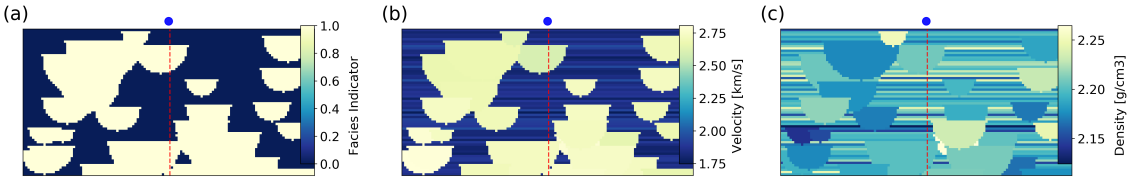

Figure 6.3: Overview of the object-based model realization used as a reference model for evaluating the inversion procedure. Geological facies (a) distinguish between river channel bodies (light) and shale (dark). (b) Acoustic P-wave velocity V_p and (c) rock density ρ is constant within river-channels and varies by layer within shale.

showing that we obtain a similar back-propagation equation as that used to derive gradients in neural networks (LeCun et al., 1988): the data mismatch is backpropagated thanks to a linear equation in the adjoint state vector λ. By differentiating the Lagrangian in Eq. 6.10 with respect to $m(\mathbf{x})$ we obtain

$$\frac{\partial J}{\partial \mathbf{m}^{V_p}} = \frac{\partial \|S(G_\theta(\mathbf{z}_t)) - \mathbf{d}_{seis}\|_2^2}{\partial G_\theta(\mathbf{z}_t)} = \langle \lambda, \ \frac{\partial F(\mathbf{u}, \ \mathbf{m}^{V_p})}{\partial \mathbf{m}^{V_p}} \rangle \tag{6.12}$$

which is the gradient required to perform MALA sampling of the posterior distribution of latent variables, Eq. 6.6.

We perform numerical solution of the acoustic wave equation and the respective adjoint computation using the domain-specific symbolic language Devito (Kukreja et al., 2016; Louboutin et al., 2017). The numerical solution is performed using a fourth-order finite-difference scheme in space and 2nd-order in time.

6.5 Generative Model

We use a generative model to sample realizations of spatially varying model parameters $m(\mathbf{x}) \sim G_\theta(\mathbf{z})$. These realizations are obtained by sampling a number of latent variable vectors $\mathbf{z}$. The associated model representations represent the *a priori* knowledge about the spatially varying properties of the geological structures in the subsurface.

We model the prior distribution of the spatially varying model parameters $m(\mathbf{x})$ (Section 6.4.1) by a generative adversarial network (GAN) (Goodfellow et al., 2014). GANs represent a generative model where the underlying probability density function is implicitly defined by a set of training examples. To train GANs two functions are required: a generator $G_\theta(\mathbf{z})$ and a discriminator $D_\omega(\mathbf{m})$. The role of the generator is to create random samples of an implicitly defined probability distribution that are statistically indistinguishable from a set of training examples. The discriminator's role is to distinguish real samples from those created by the generator. Both functions are trained in a competitive two-player min-max game where the overall loss is defined by

$$\begin{aligned} \min_\theta \max_\omega \{ & \mathbb{E}_{\mathbf{m} \sim p_\mathbf{m}}[\log \ D_\omega(\mathbf{m})] \\ & + \mathbb{E}_{\mathbf{z} \sim p_\mathbf{z}}[\log \ (1 - D_\omega(G_\theta(\mathbf{z})))] \} \end{aligned} \tag{6.13}$$

Due to the opposing nature of the objective functions, training GANs is inherently unstable and finding stable training methods remains an open research problem.

Nevertheless, a number of training methods have been proposed that allow more stable training of GANs. In this work we use a so-called Wasserstein-GAN (Arjovsky et al., 2017; Gulrajani et al., 2017; Chan and Elsheikh, 2017), that seeks to minimize the Wasserstein distance between the generated and real probability distribution. We use a Lipschitz penalty term proposed by Petzka et al., (2017) to stabilize training of the Wasserstein-GAN. For the discriminator, keeping the parameters θ of the generator fixed, we minimize

$$\begin{aligned}\min_{\omega}\{&\mathbb{E}_{\mathbf{z}\sim p_{\mathbf{z}}}[D_{\omega}(G_{\theta}(\mathbf{z}))] - \mathbb{E}_{\mathbf{m}\sim p_{\mathbf{m}}}[D_{\omega}(\mathbf{m})] \\ &+\lambda_{LP}\mathbb{E}_{\hat{\mathbf{m}}\sim p_{\hat{\mathbf{m}}}}[(\max\{0, \|\nabla D_{\omega}(\hat{\mathbf{m}})\| - 1\})^2]\}\end{aligned} \tag{6.14}$$

where $\hat{\mathbf{m}}$ is linear combination between a real and generated sample controlled by a random variable τ (Petzka et al., 2017). For the generator, keeping the parameters of the discriminator ω constant, we minimize

$$-\mathbb{E}_{\mathbf{z}\sim p_{\mathbf{z}}}[D_{\omega}(G_{\theta}(\mathbf{z}))] \tag{6.15}$$

In our work we set $\lambda_{LP} = 200$ to train the generative model. We represent both the generator and discriminator[2] function by deep convolutional neural networks (Appendix Table 6.1). The generator uses a number of convolutional layers followed by so-called pixel-shuffle transformations to create output models (Shi et al., 2016). The latent vector is parameterized as a multivariate standardized normal distribution:

$$\mathbf{z} \sim \mathcal{N}(0, \mathbf{I})^{50\times 1\times 2} \tag{6.16a}$$

$$G_{\theta} : \mathbf{z} \rightarrow \mathbb{R}^{3\times 64\times 128} \tag{6.16b}$$

Due to the geological properties represented in our dataset, namely, geological facies indicators $\mathbf{m}^{facies}$, acoustic P-wave velocity $\mathbf{m}^{V_p}$ and density $\mathbf{m}^{\rho}$, the generator must output three data channels. We represent the geological facies as a probability of a spatial location belonging to a sandstone facies. To facilitate numerical stability of the GAN training process we apply a hyperbolic tangent activation function and convert to a probability $\mathbf{m}^{facies}$ for subsequent computation (Eq. 6.7). We apply a hyperbolic tangent activation function to model the output distribution of the P-wave model parameters $\mathbf{m}^{V_p}$. For rock density $\mathbf{m}^{\rho}$ a soft-plus activation function is used to ensure positive values (Appendix 6.10.1). In this study, only the facies indicator $\mathbf{m}^{facies}$ and acoustic P-wave velocity $\mathbf{m}^{V_p}$ are used in the inversion process.

The generator-discriminator pairing is trained on the set of training images described in Section 6.6. GAN training required approximately eight hours on eight NVIDIA K80 graphics processing units. A set of samples obtained from the GAN prior are presented in Appendix Fig. 6.9. After training, the generator $G_{\theta}(\mathbf{z})$ and the forward modeling operator $S(\mathbf{m})$ are arranged in a fully differentiable computational graph. To accommodate the sources and receivers of the acoustic forward modeling process described in Section 6.4, we pad the output of the generator by a domain of constant P-wave velocity.

[2]In the Wasserstein-GAN literature the discriminator is also termed a "critic".

6.6 Dataset

To demonstrate the proposed inversion method we will use a model of a fluvial-dominated system consisting of highly porous sandstones embedded in a fine grained shaly material. Object-based models are commonly used to model such geological systems (Deutsch and Wang, 1996). They represent the fluvial environment as a set of randomly located geometric objects following various size, shape and property distributions. We train a set of GANs on a dataset of ten thousand realizations of two-dimensional cross-sections of fluvial object-based models.

The individual cross-sections are created with an object-based model, where half-circle sand-bodies follow a uniform width distribution. P-wave velocity and density are constant within each channel-body and their values are sampled independently from a Gaussian distribution for each individual channel-body. The locations of the channel-bodies are determined by a uniform distribution in spatial location. The fine-grained material surrounding the river systems is made of layers of single-pixel thickness where each layer has a constant value of acoustic P-wave velocity and density which varies randomly from one layer to another and is sampled from a Gaussian distribution. We use a binary indicator variable to distinguish the two facies regions, river channel vs shale matrix. The ratio of how much of a given cross-section is filled with river channels compared to the overall area of the geological domain is a key property in understanding the geological nature of these structures. This ratio follows a uniform distribution from 30 to 60% in our dataset and river channels are placed at random until a cross-section meets the randomly sampled ratio.

A total of ten thousand training images were created as a training set for the GAN. A further four thousand images were retained as a test set to evaluate the inversion technique. While training the generative model outlined in Section 6.5 we monitor image quality and output distributions for each of the modeled properties. The reference model (Fig. 6.3) used to evaluate the Bayesian inversion approach was chosen randomly from the test-set of object-based models. Figure 6.3 shows a comparison of the distribution of the three-modeled properties, geological facies indicator, acoustic P-wave velocity and rock density for the reference model.

6.7 Results

We evaluate the proposed method of inversion by sampling a set of latent variables $\mathbf{z}$ determining the output of the generative model $G_\theta(\mathbf{z})$ (Section 6.6, Fig. 6.3). First, we evaluate the generative model as a prior for representing possible earth models and generating $N = 100$ unconditional samples (Fig. 6.4-1, Appendix Fig. 6.9).

Two cases of inversion are considered: inversion for the acoustic P-wave velocity V_p and combined inversion of acoustic velocity and of geological facies along a borehole. In all the cases presented we assume that density is a constant. For all tests we perform inversion using the approximate MALA scheme. For the additional borehole constraint we require an accuracy above 95% of geological facies to be accepted as a valid inverted sample. While lower errors in seismic mismatch and borehole accuracy can be achieved, evaluating the forward problem and adjoint of the partial differential equation comes at a high computational cost and therefore a cost-effectiveness trade-off was necessary.

For the first case of seismic inversion without borehole constraints we perform

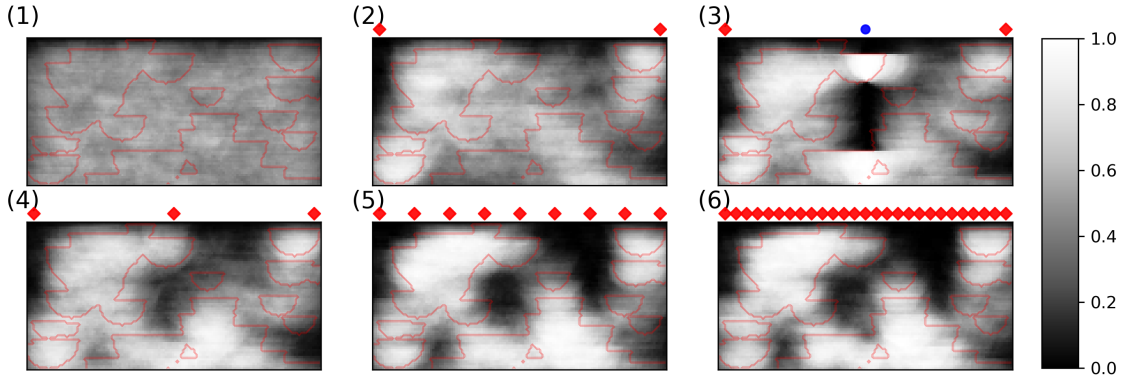

(a) Pixel-wise mean for 100 samples obtained from the prior and inferred models.

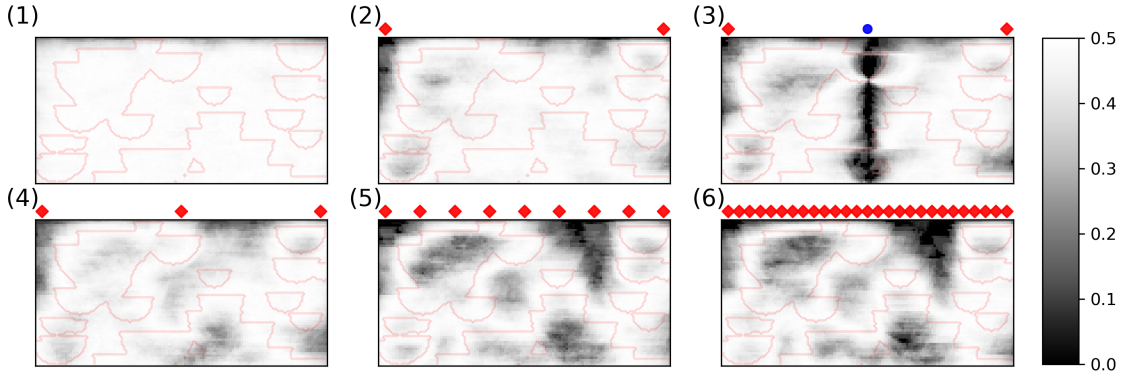

(b) Pixel-wise standard deviation for 100 samples obtained from the prior and inferred models.

Figure 6.4: (a) Pixel-wise mean and (b) standard deviation of the ensemble of 100 models sampled unconditionally from the prior (1) represented by the generator function $G_\theta(\mathbf{z})$. Posterior ensemble of geological indicator variables matched to the seismic representation of the reference model shown in Fig. 6.3 for (2) 2 sources, (3) 2 sources and a single borehole, (4) 3 sources, (5) 9 sources, (6) 27 sources. Source locations are indicated by red diamonds and the borehole location by a blue circle. The reference model is indicated by red contours.

simulations where the number of acoustic sources are increased. Fewer acoustic sources means that less of the domain is properly imaged, leading to high uncertainty in areas where no incoming waves have been reflected and recorded by the receivers on the surface. The acoustic sources and 128 receivers are equally spaced across the top edge of the domain.

In Fig. 6.4 we show the pixel-wise mean (Fig. 6.4a) and standard deviation (Fig. 6.4b) of 100 inferred models for an increasing number of acoustic sources (2 sources up to 27 sources). As the total number of acoustic sources increases we obtain a lower standard deviation for the resulting model ensembles. In the case of two acoustic sources (Fig. 6.4b-2) we find that close to the sources there is small variation amongst the inferred models (dark shades) whereas the central area where no acoustic source has been placed shows a very high degree of variation. This is confirmed by the three source case where a central acoustic source has been placed in addition to the sources on the borders of the domain. Lower variability in the inverted ensemble can be observed. This correlates well with the Bayesian interpretation of

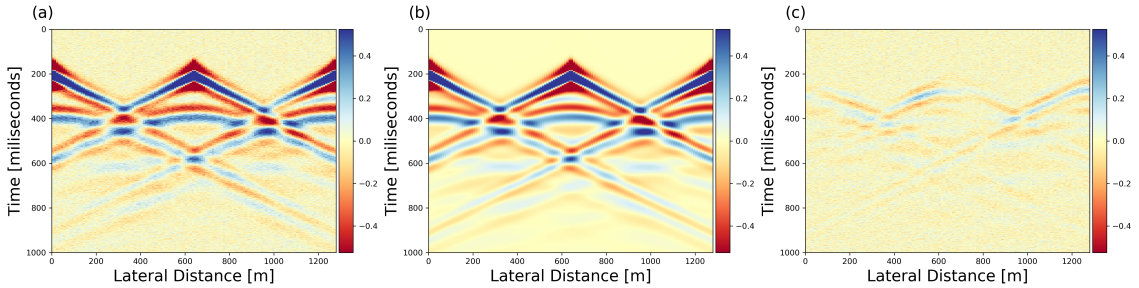

Figure 6.5: Comparison of (a) the seismic waveform based on the reference model acoustic velocity with (b) the waveform of an inferred model with three seismic sources. The difference (c) in amplitude of the two waveforms. Colormaps are scaled based on one standard deviation in amplitudes of the reference model wave-form (a).

the inverse problem; where acoustic sources allow the subsurface to be imaged we arrive at a low standard deviation in the posterior ensemble of geological models, whereas within regions that are only sparsely sampled by the acoustic sources we expect the prior, the unconditional generative model, to be more prevalent, leading to a higher variability of geological features. As expected, when we increase the number of sources, we find overall smaller variability in the resulting ensemble of inverted earth models. We observe only marginal reduction in variability between the cases with nine and twenty-seven sources (Figs. 6.4b, 5-6). For all inversion scenarios considered we present samples from the posterior in Appendix Figs. 6.10-6.13.

In the case where lower-dimensional information such as a borehole was included as an additional objective function constraining the generative model (Fig. 6.4b-3), we find a lower standard deviation around this borehole. The standard deviation along the well is close to zero due to the per-realization 95% accuracy constraint. Furthermore there is a region of influence where the borehole constrains lateral features such as channel bodies. This is shown by channel shaped features of low standard deviation at the top and bottom of the domain. Comparison with the reference model (Fig. 6.3a) shows that two channel bodies can be found along the one-dimensional feature.

For each generated realization we have recorded the ratio of the squared error norm (Eq. 6.6) and the squared norm of the noise in the seismic data (Fig. 6.6) at each MALA sampling iteration. The global minimum of the data mismatch in the presence of Gaussian noise is reached when the objective function value is equal to the squared norm of the noise in the data i.e. at a ratio equal to one (Fig. 6.6). In practice we find that performing 200 MALA iterations leads to a sufficient reduction in the mismatch of the seismic data and as required by the approximate MALA algorithm the error stabilizes as the step-size is reduced.

Due to the fact that modern FWI methods come at very high computational cost for two and possibly three-dimensional inversion, a small number of required iterations is a necessity. In further tests, reducing the number of iterations of the MALA approximation or simply optimizing by gradient descent, as performed by Richardson, (2018), allows convergence to small errors, but this approach has been shown to lead to reduced sample diversity (Nguyen et al., 2016).

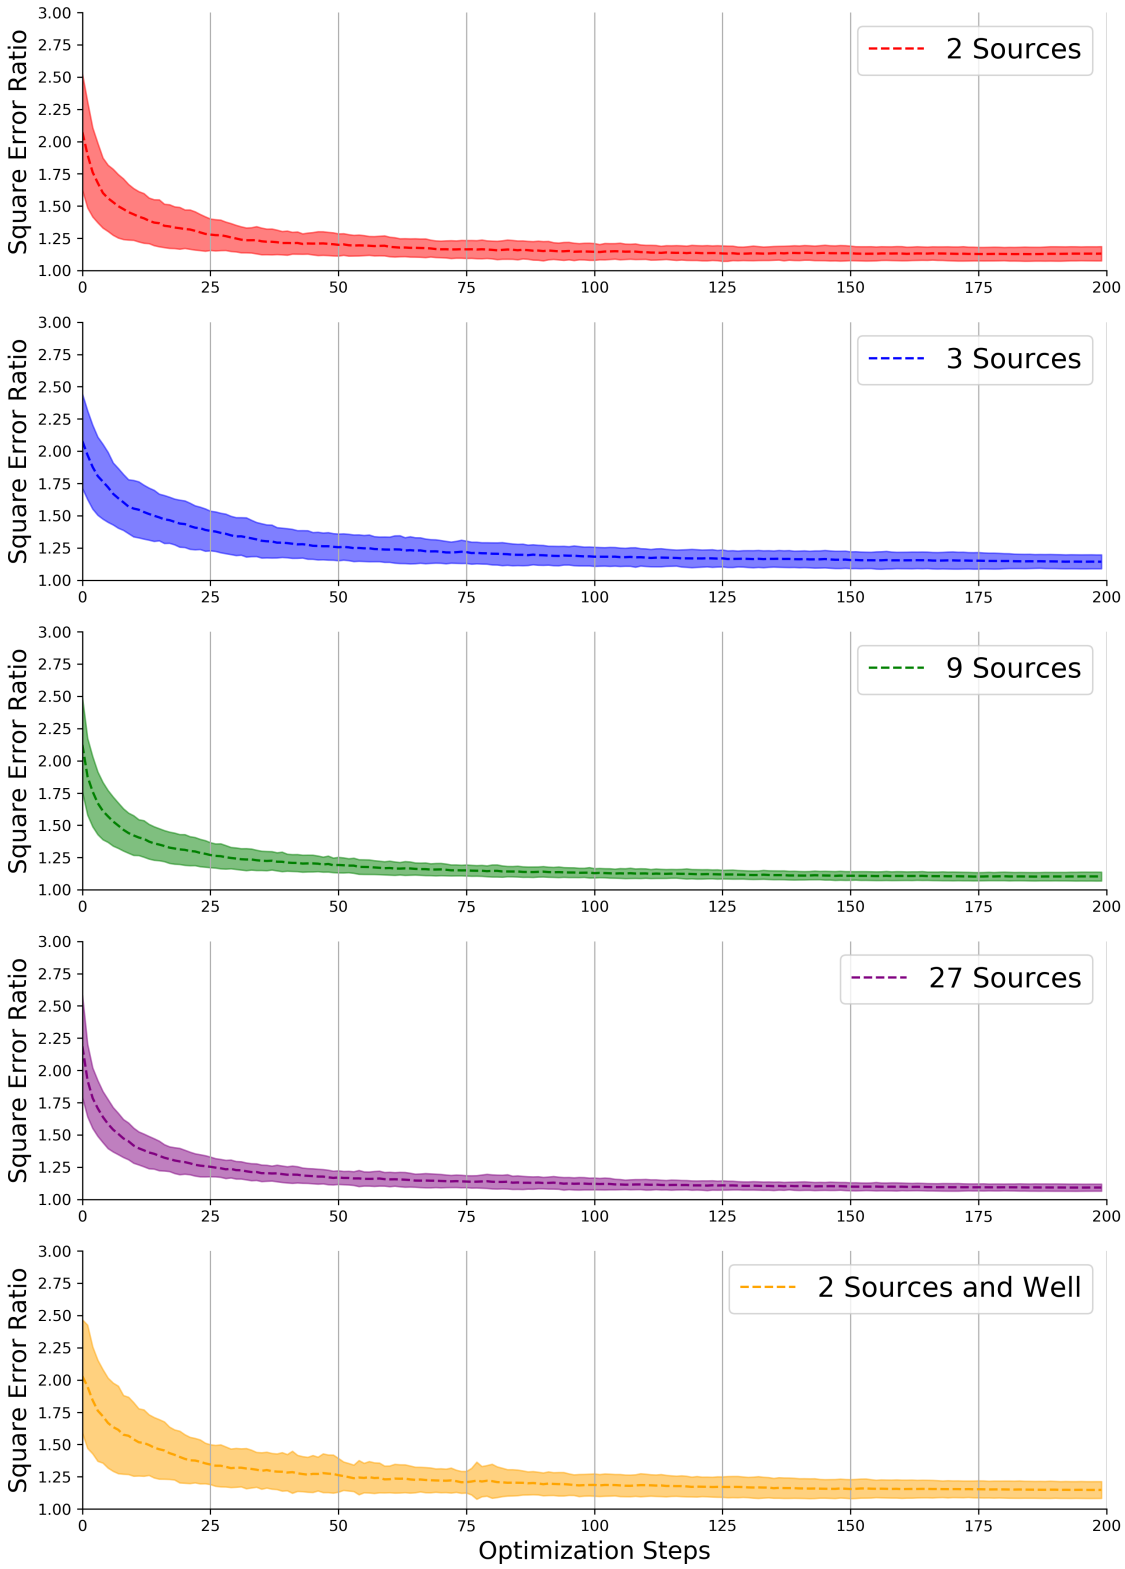

Figure 6.6: Comparison of the ratio of the squared error norm and the squared norm of the Gaussian noise. The global minimum is reached at values of one. Shaded regions indicate $\pm\,\sigma$ of the squared error ratio. We perform 200 approximate MALA iterations to obtain samples of the posterior given seismic observations only, as well as where borehole information and seismic observations are included. The step-size was annealed to very small values leading to a stabilization of the squared error norm at the end of the sampling procedure.

6.8 Discussion

We have shown that it is possible to obtain posterior realizations inferred from the latent space of a GAN generator that honor seismic and well-bore data by using an approximate Bayesian sampling method. A number of open questions remain concerning the generative model and the posterior distribution of models that are obtained.

A common challenge with GANs specifically is their so-called mode-collapse behavior where the distribution represented by the generative model has collapsed to one or a few modes of the distribution represented implicitly by the set of training images. GANs do not represent the density explicitly and therefore it is not possible to evaluate the ability of a GAN to represent the distribution by e.g. evaluating the likelihood of a set of test images given the model parameters $p(\mathbf{m}|\theta)$. Theis et al., (2015) have shown that evaluating sample quality and diversity of generative adversarial networks is difficult. Nevertheless a number of heuristic approaches such as the Inception Score (IS) (Salimans et al., 2016) or the Frechet Inception distance (FID) (Heusel et al., 2017) have been proposed, and while these methods are popular for evaluating GANs trained on natural images, they may not be representative measures to compare GANs as shown by Barratt and Sharma, (2018). Arora and Zhang, (2017) propose a method to empirically evaluate the support of the distribution represented by a GAN.

Another common failure case of GANs occurs when the generator only memorizes the images of the training set and does not learn a representation of the entire distribution. In this case, it should only be possible to infer models which are part of the training set and which match the well and seismic data associated with the reference model. In the following we investigate whether the ensemble of models obtained by solving the inverse problem represent new stochastic realizations of the underlying distribution represented implicitly by the training images.

We have evaluated the mean-squared-error (MSE) and the structural similarity index (SSIM) (Wang et al., 2004) between pairs of binary facies models. A perfect agreement between two models is reached for a MSE of zero and an SSIM of one. The MSE, while being a common measure to compare pairs of data, is very sensitive to small translations of the models that are compared. The structural similarity index attempts to capture perceptual similarity and is less sensitive to pixel-wise differences in the two compared models (Wang and Bovik, 2009).

In Fig. 6.7 we show kernel-density-estimates for the distributions of the two image similarity measures. First, we compare the distribution of the MSE and SSIM between the reference model and the 10^4 models in the training set (Ref.-TI), with that between the reference model and 10^5 models sampled from the GAN prior (Ref.-Prior). We find that the two distributions match closely. This confirms that images drawn from the GAN prior and from the training set are statistically similar and that none of the images from the training set and prior are likely to be identical with the reference model. This finding is a good indication that the GAN does not seem to have collapsed to a few modes but it does not exclude the possibility of our generative model having memorized the training set as in this case we would expect the distributions between Ref.-Prior and Ref.-TI to match.

In a second step we now compare the reference model to the models inferred by our Bayesian inversion approach using the GAN as a prior. We find that the

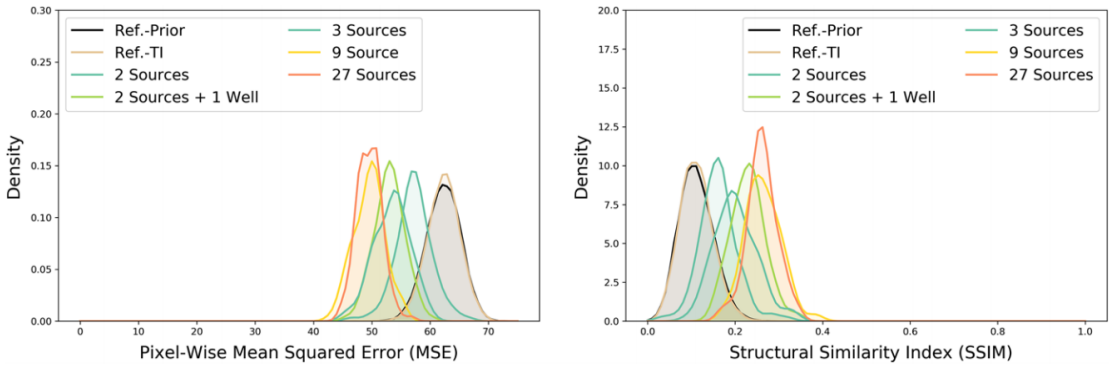

Figure 6.7: Kernel density estimates of the distributions of the (left) mean-squared-error and (right) structural similarity index (SSIM) with respect to the reference model for models sampled from the GAN prior, and inferred models obtained by Bayesian inversion.

distributions are all consistently shifted to regions of higher similarity to the reference model i.e. lower MSE and higher SSIM for models inferred when considering the seismic data as well as seismic and well data. This shows that our inversion, when the number of data is increased, tends towards models that are more and more similar to the reference model. When 9 and 27 acoustic sources are used we find that inversion leads to models that on average have a SSIM that has very low probability under the Ref.-TI. and Ref.-Prior distributions showing that our GAN is able to create images outside the set of training images. If the generator had only memorized the training set, we should not be able to infer models with higher similarity as the number of data increased.

In Fig. 6.8a we show models from the training set, samples from the GAN prior and models inferred with the highest SSIM when compared to the reference case (Fig. 6.3). In Fig. 6.8b we show models that have an SSIM close to the mode of the SSIM distributions and find that the model from the posterior inferred by inversion using 27 acoustic sources is visually more similar to the reference case (Fig. 6.3) than the samples obtained from the prior and from the training set.

It is important to note that the evaluation of the inferred models with respect to a reference model is only possible in the case of synthetic data. In subsurface applications it is not possible to obtain the entire reference model. Furthermore, models that are structurally very different can be valid solutions of the ill-posed inverse problem. These models, which represent possible solutions of the inverse problem, may be associated with different modes of the prior distribution. In the case of GANs, the generator may be able to represent all of these modes or only a subset (mode-collapse). If mode-collapse has occurred the posterior ensemble only represents solutions obtained from the modes represented by the generator. Therefore, checking for the occurrence of mode-collapse is key for practical applications as mode-collapse may significantly affect the ensemble of obtained solutions and possibly lead to underestimated uncertainty.

For future work, evaluating other deep generative models based on explicit density representations (Kingma and Welling, 2013; Dinh et al., 2016; Oord et al., 2016a; Kingma and Dhariwal, 2018), which can calculate the likelihood of a set of test images, may help to improve the representation of the prior distribution and mitigate

(a) Samples with the highest SSIM with the reference model from the (1) training set, (2) models from the GAN prior, and (3) models obtained by Bayesian inversion using 27 sources.

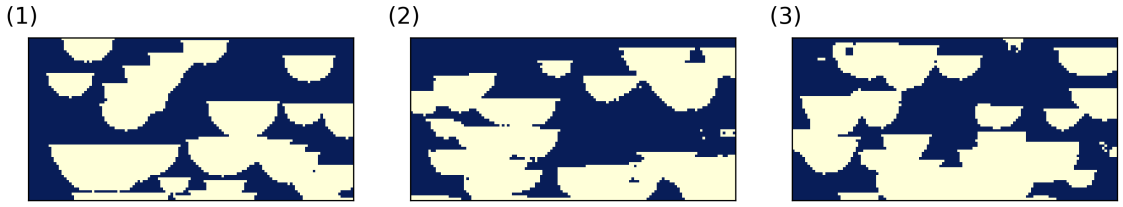

(b) P50 realizations of the distributions of the SSIM for the (1) training set, (2) models from the GAN prior, and (3) models obtained by Bayesian inversion using 27 sources.

Figure 6.8: Overview of models from the training set, GAN prior, and inferred models using the MALA approach that (a) show the highest similarity to the reference model (Fig. 6.3) measured by the SSIM metric and (b) represent realizations close to the mode of the distribution of the SSIM.

the effect of mode-collapse on inversion.

6.9 Conclusions

Inversion of subsurface geological heterogeneities from acoustic reflection seismic data is a classical method to aid the understanding of the earth's interior. The inference of model parameters from measured acoustic properties is often performed in the very high-dimensional space of model properties leading to very CPU-intensive optimization (Akcelik et al., 2003).

We apply a method that combines a generative model of geological heterogeneities efficiently parameterized by a lower-dimensional set of latent variables, with a numerical solution of the acoustic inverse problem for seismic inversion using the adjoint method. Leveraging the adjoint of the studied partial differential equation we deduce gradients that are consequently used to sample from the posterior over the latent variables given the mismatch of the observed seismic data by following an approximate MALA scheme (Nguyen et al., 2016).

While the proposed application was illustrated on a simple geophysical inversion this method may find use in other domains where spatial property models control the evolution of physical systems such as in fluid flow in porous media or materials science. The combination of a deep generative model parameterized by a lower-dimensional set of latent variables and gradients obtained by the adjoint method may lead to new efficient techniques for solving high-dimensional inverse problems. This is explored in the next chapter where we study the problem of reservoir history matching.

6.10 Chapter Appendix

6.10.1 Generative Model Network Architectures

Table 6.1: Generator and discriminator network architectures used to create synthetic geological structures. Binary indicators of geological facies and corresponding P-wave velocities are represented by a bi-variate Gaussian distribution and a hyperbolic tangent activation function is used to represent the two families of properties. Rock density shows a Gaussian distribution. A soft-plus activation function ($f(x) = \frac{1}{\beta}\log(1 + exp(\beta\ x)),\ \beta = 1$) is used to ensure positive values of density. Notation for convolutional layers: LayerType(Number of filters), k=kernel size, s=stride, p=padding. BN=BatchNorm, PS=PixelShuffle

Latent Variables $z \in \mathbf{R}^{50\times1\times2}$
Conv2D(512)k3s1p1, BN, ReLU, PSx2
Conv2D(256)k3s1p1, BN, ReLU, PSx2
Conv2D(128)k3s1p1, BN, ReLU, PSx2
Conv2D(64)k3s1p1, BN, ReLU, PSx2
Conv2D(64)k3s1p1, BN, ReLU, PSx2
Conv2D(64)k3s1p1, BN, ReLU, PSx2
Conv2D(3)k3s1p1
Tanh (0,1) | Softplus (2)

(a) Multi-Channel Generator

Geological Properties $m \in \mathbf{R}^{3\times64\times128}$
Conv2D(64)k5s2p2, ReLU
Conv2D(64)k5s2p1, ReLU
Conv2D(128)k3s2p1, ReLU
Conv2D(256)k3s2p1, ReLU
Conv2D(512)k3s2p1, ReLU
Conv2D(512)k3s2p1, ReLU
Conv2D(1)k3s1p1, ReLU

(b) Discriminator for multi-channel GAN

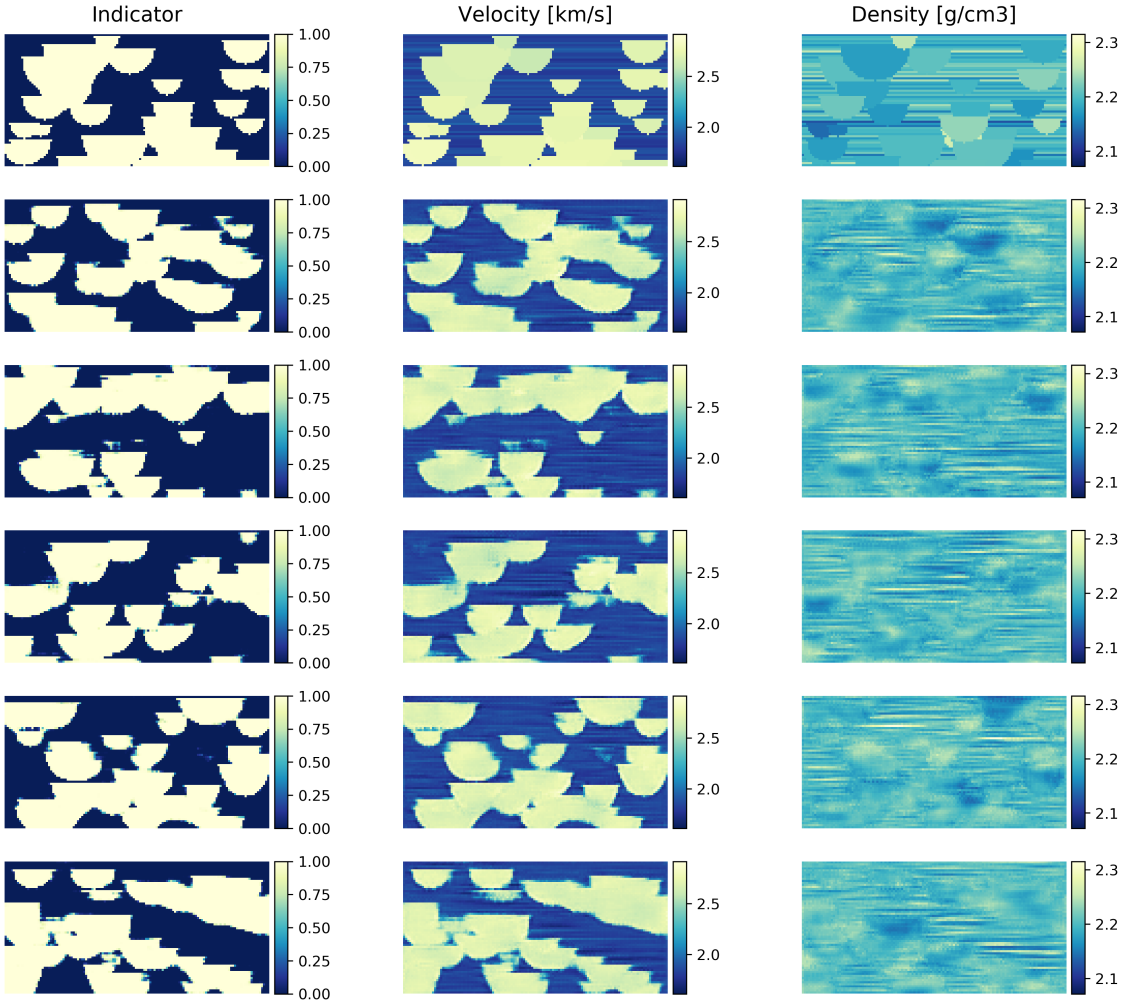

Figure 6.9: Samples from the prior distribution of models obtained from the GAN with the reference model (Fig. 6.3) shown in the first row.

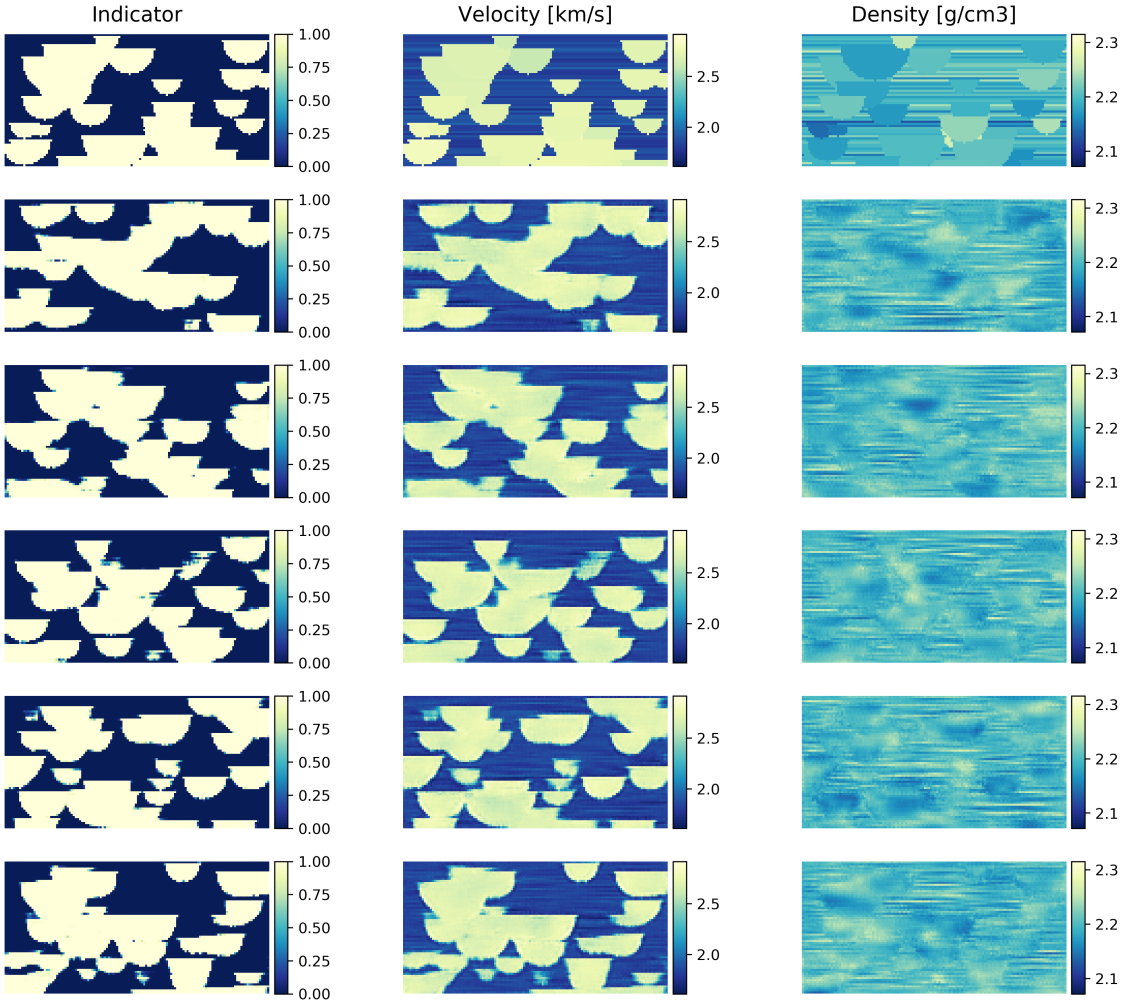

Figure 6.10: Samples obtained from latent space optimization with 2 acoustic sources with the reference model (Fig. 6.3) shown in the first row.

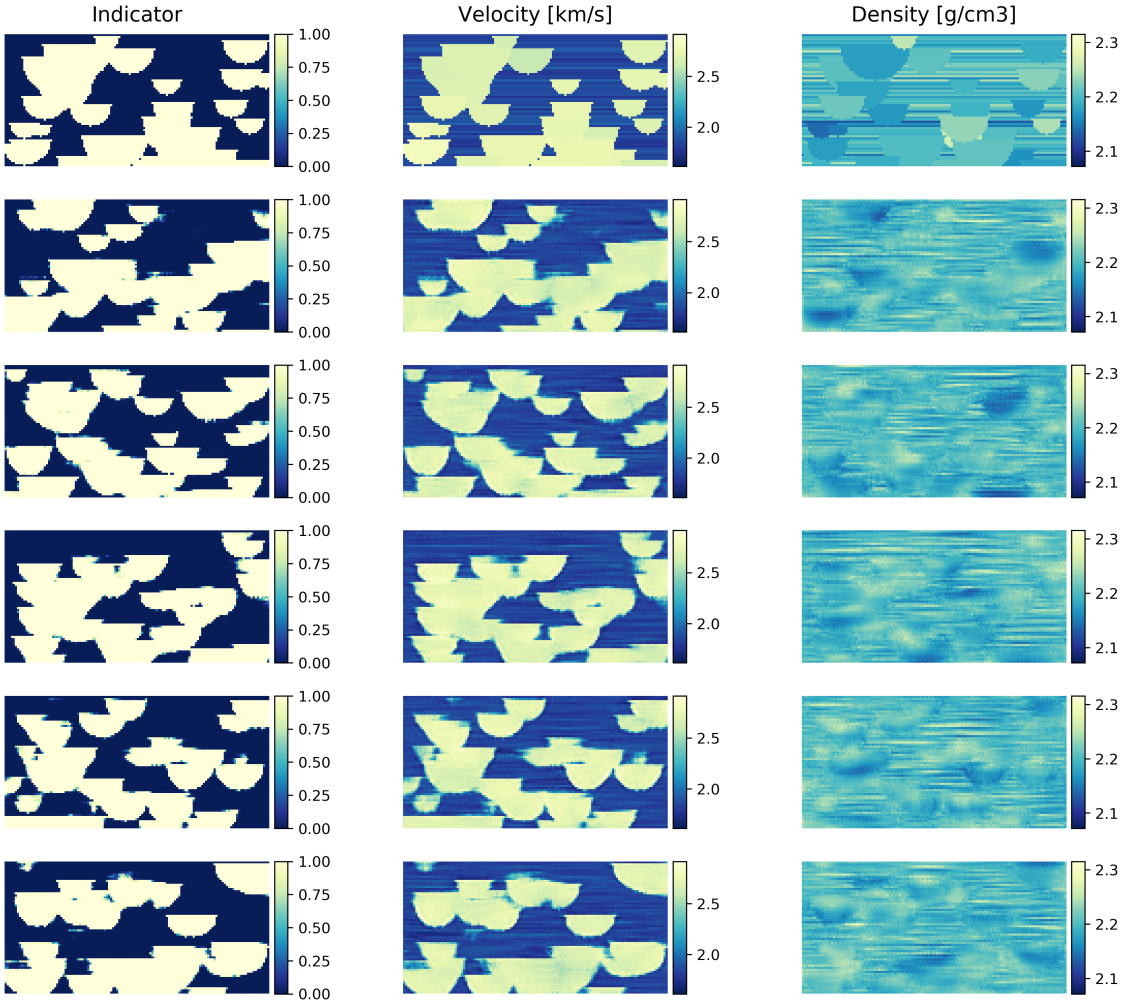

Figure 6.11: Samples obtained from latent space optimization with 3 acoustic sources with the reference model (Fig. 6.3) shown in the first row.

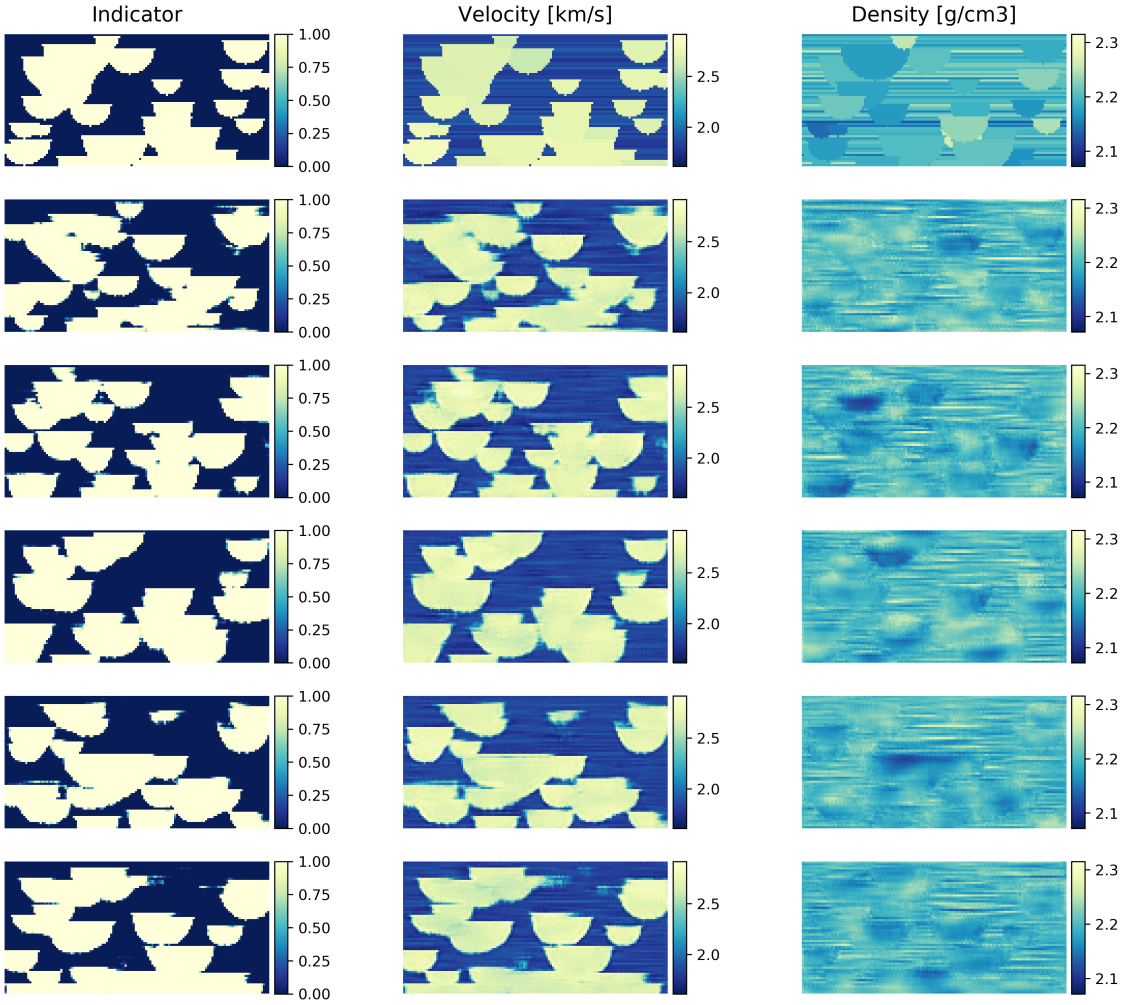

Figure 6.12: Samples obtained from latent space optimization with 9 acoustic sources with the reference model (Fig. 6.3) shown in the first row.

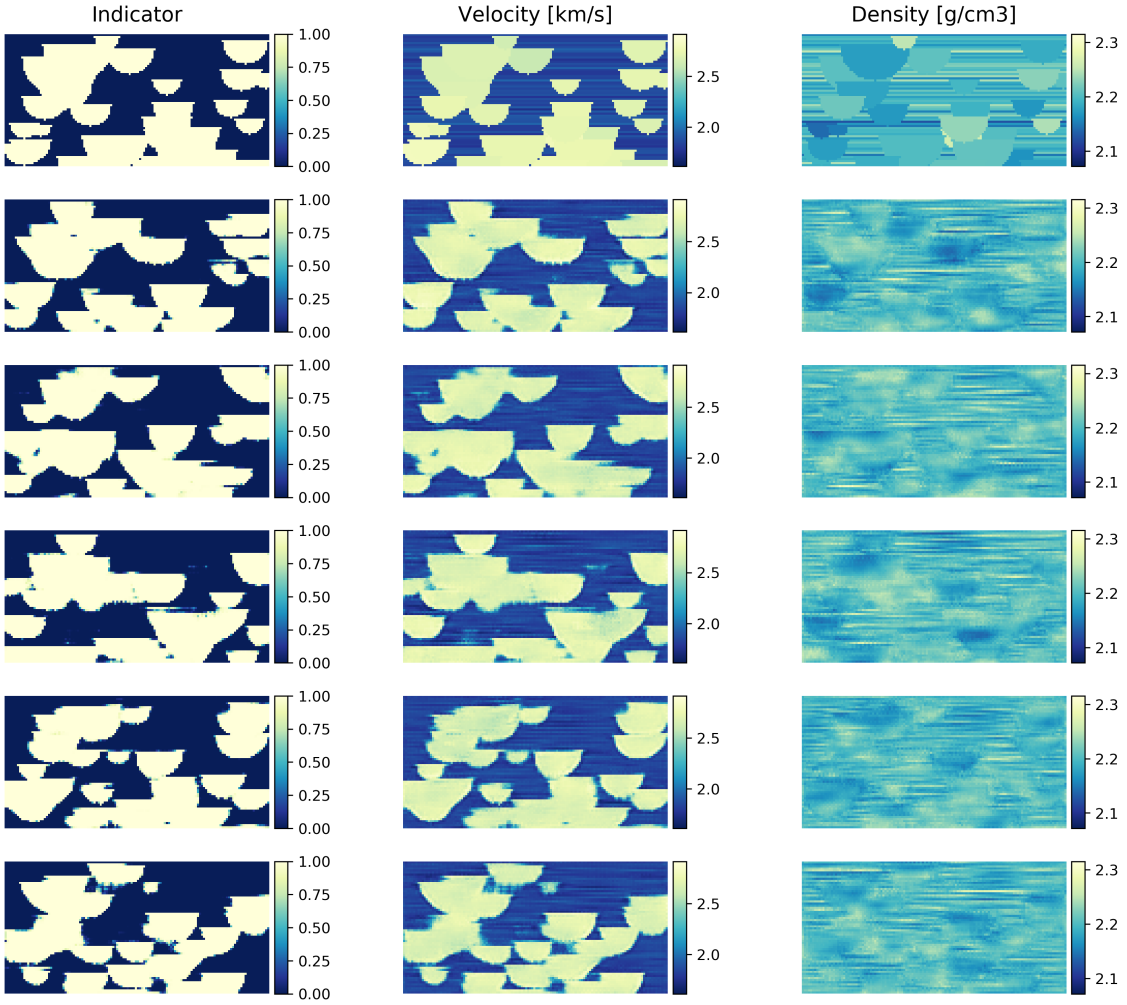

Figure 6.13: Samples obtained from latent space optimization with 27 acoustic sources with the reference model (Fig. 6.3) shown in the first row

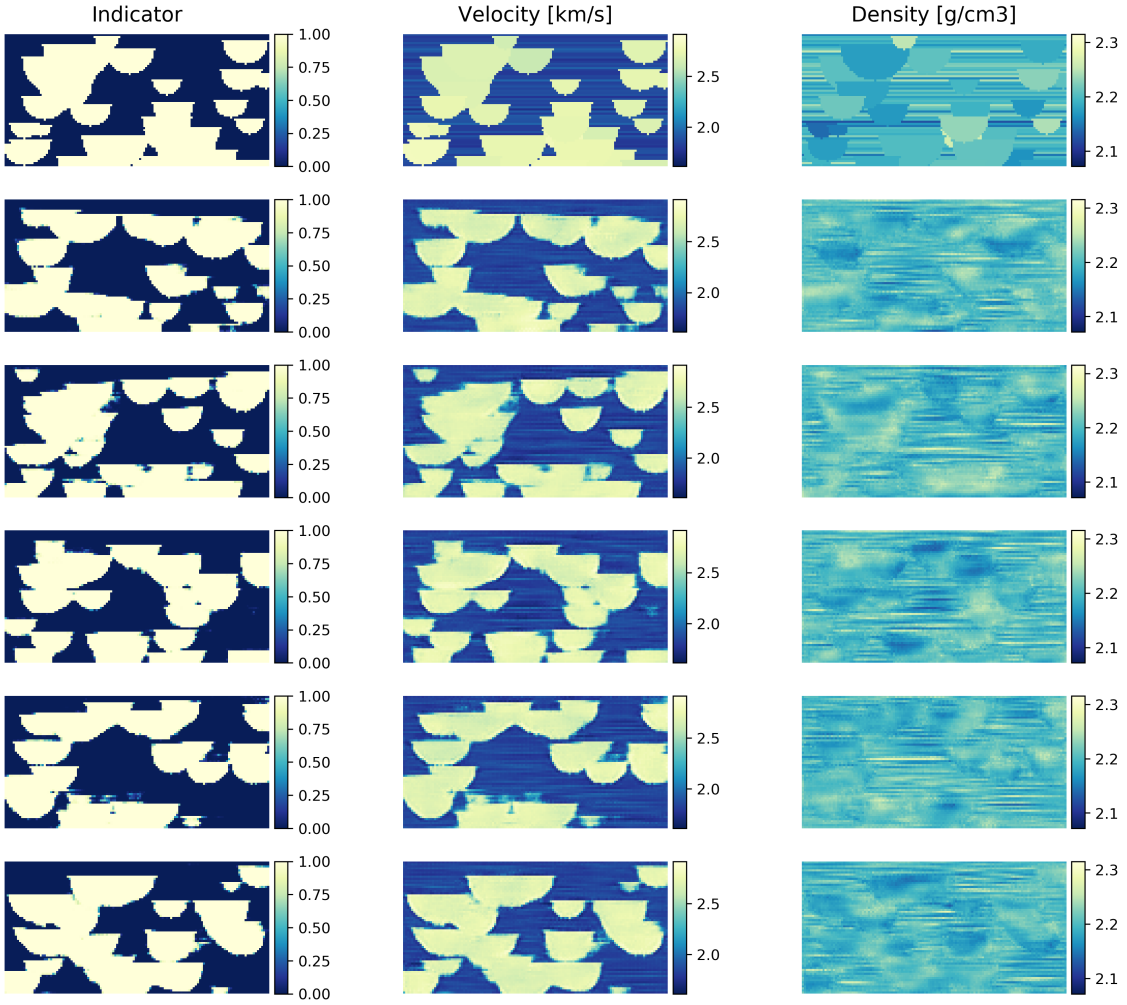

Figure 6.14: Samples obtained from latent space optimization with 2 acoustic sources and one borehole with the reference model (Fig. 6.3) shown in the first row

Chapter 7

DeepFlow: History Matching in the Space of Deep Generative Models

7.1 Abstract

The calibration of a reservoir model with observed transient data of fluid pressures and rates is a key task in obtaining a predictive model of the flow and transport behavior of the earth's subsurface. The model calibration task, commonly referred to as "history matching", can be formalized as an ill-posed inverse problem where we aim to find the underlying spatial distribution of petrophysical properties that explain the observed dynamic data. We use a generative adversarial network pre-trained on geostatistical object-based models to represent the distribution of rock properties for a synthetic model of a hydrocarbon reservoir. The dynamic behavior of the reservoir fluids is modeled using a transient two-phase slightly compressible Darcy formulation. We invert for the underlying reservoir properties by first modeling property distributions using the pre-trained generative model then using the adjoint equations of the forward problem to perform gradient descent on the latent variables that control the output of the generative model. In addition to the dynamic observation data, we include well rock-type constraints by introducing an additional loss function. Our contribution shows that for a synthetic test case, we are able to obtain solutions to the inverse problem by optimizing in the latent variable space of a deep generative model, given a set of transient observations of a non-linear forward problem.

7.2 Introduction

Understanding the flow of fluids within the earth's subsurface is a key issue in many practical applications such as the understanding of how pollutants in aqueous phases are transported and affect the environment (Yong et al., 1992; Bear and Cheng, 2010), hydrogeological models (Domenico et al., 1998; Fetter, 2018), the formation of mineral deposits (Garven, 1985; Sibson et al., 1975), carbon-capture and sequestration (CCS) (Lackner, 2003; Holloway, 2005), and the migration and production of hydrocarbons from subsurface reservoirs (England et al., 1987). When modeling subsurface architectures a spatial gridded representation is usually built

for modeling physical flow processes (Aziz and Others, 1993). The flow of fluids, including aqueous and non-aqueous fluids such as supercritical CO_2 is simulated by solving a set of partial differential equations that describe the transient behavior of the flowing liquid phases (Gerritsen and Durlofsky, 2005).

In many cases the ability for fluids to flow within a porous material is dominated by two key properties: the effective porosity ϕ defines how much pore-space is available to flow as a proportion of the bulk volume of the material (Bear, 2013). The permeability $\mathbf{k}$ is a property that is proportional to the flow conductance (Darcy, 1856; Muskat, 1938). We refer to Blunt et al., (2013) for a detailed review of the principles of multiphase fluid flow within porous media. Obtaining direct measurements of these two properties, permeability and porosity, is difficult and in many cases these measurements are sparse and obtained at various spatial scales, making their integration within the modeling workflow difficult (Wu et al., 2002; Farmer, 2002).

Nevertheless, representing the spatial distribution of the properties of porous media within the earth's subsurface is necessary to build reliable predictive models. In applications such as the injection of supercritical CO_2 into hosting rock formations for long-term storage (Chadwick and Noy, 2010; Martens et al., 2012) or the production of hydrocarbons and water from subsurface reservoirs we obtain, in addition to well and seismic data, measurements of the variables dependent on these flow properties ($\mathbf{k}$ and ϕ) such as fluid phase pressures, as well as produced and injected volumes over time (Muskat, 1949). Calibration of a spatial model of the petrophysical parameters with available physical measurements is commonly referred to as "history matching" (Oliver and Chen, 2011), and more generally as data assimilation (Reichle et al., 2002). Determining the underlying spatially distributed parameters $\mathbf{x}$ from observed parameters $\mathbf{y}$ such as fluid pressures and rates, can be formalized as an an ill-posed inverse problem.

For an inverse problem to be well-posed it must satisfy three criteria (Tikhonov and Arsenin, 1977):

1. Existence: for every set of observed parameters $\mathbf{y}$ there must exist a solution.
2. Uniqueness: the obtained solution has to be unique.
3. Stability: the solution is stable with regards to small changes in the observed parameters i.e. a small change in the observed parameters $\mathbf{y}_{obs}$ induces only a small change in the solution parameters $\mathbf{x}$

If these three criteria are not fulfilled the problem is considered ill-posed. It has been shown that inverse problems with spatially distributed parameters are often ill-posed because they don't fulfill the uniqueness or stability criteria (Kravaris and Seinfeld, 1985). A detailed analysis of the existence criterion for distributed parameter systems of the heat, transport, and wave equations is provided by Chavent, (1979).

Reducing the number of parameters used to represent the solution space $\mathbf{m}$ helps the inverse problem to be well-defined. Coats et al., (1970) assumed a constant value for large spatial regions of the underlying parameters (Cooley and Sinclair, 1976; Yeh and Tauxe, 1971). Leo et al., (1986) used B-splines to represent spatial distributions of porosity and permeability. Recently, a number of approaches for model parameterization have been developed based on principal component analysis (PCA) (Vo and Durlofsky, 2014) and improvements thereof using kernel and optimization-based

PCA. Demyanov et al., (2011) use kernel learning to find a parameterization of the petrophysical parameter space.

In a Bayesian setting of the inverse problem (Tarantola, 2005) the gradual deformation method applies a continuous perturbation of a set of model parameters $\mathbf{m}$ obtained from a prior distribution to honor the observed data of the forward problem (Roggero et al., 1998; Caers, 2007). Bayesian approaches to solve the reservoir history matching problems, such as the Ensemble Kalman Filter (EnKF) (Evensen, 1994; Evensen, 2003) have been applied in numerous history-matching case studies (Lorentzen et al., 2005). Emerick and Reynolds, (2012) present an ensemble method based on EnKF, that shows improved performance in the case of highly non-linear problems.

Recent developments in deep learning (LeCun et al., 2015; Goodfellow et al., 2016) have motivated the use of neural-network-based parameter representations. Canchumuni et al., (2018) parameterized geological training images using a Deep Belief Network (DBN) (Hinton et al., 2006; Hinton and Salakhutdinov, 2006) and used an ensemble smoother with multiple data assimilation (ES-MDA) (Emerick and Reynolds, 2013) to perform history matching. Variational autoencoders (Laloy et al., 2017a) and spatial generative adversarial networks (SGAN) (Jetchev et al., 2016; Laloy et al., 2017b) have been used to generate a set of channelized training images to perform probabilistic inversion using adaptive Markov-Chain Monte-Carlo methods (Laloy and Vrugt, 2012). Chan and Elsheikh, (2018a) used generative adversarial networks trained using patch-based kernel discrepancy measures to synthesize geological models from training images. Realizations conditional to spatial observations of the underlying model parameters can be obtained by gradient-based image inpainting methods (Yeh et al., 2016; Mosser et al., 2018a; Dupont et al., 2018) or direct learning of conditional image distributions using autoregressive generative models such as PixelCNNs (Oord et al., 2016a; Dupont and Suresha, 2018).

Adjoint-state methods allow gradients of the forward problem with respect to the solution parameters $\mathbf{m}$ to be obtained. These gradients can then be further backpropagated through a deep generative models such as generative adversarial networks (GANs) (Goodfellow et al., 2014; Kingma and Welling, 2013; Richardson, 2018; Mosser et al., 2018d) to perform inversion in the space of latent variable models and have been used for probabilistic inversion in the case of acoustic waveform inversion using a Metropolis-adjusted Langevin algorithm (MALA) (Roberts and Stramer, 2002; Mosser et al., 2018d). Numerous studies have evaluated the use of GANs for linear inverse problems (Shah and Hegde, 2018) in the context of compressed sensing (Bora et al., 2018; Mardani et al., 2019) and have analyzed the convergence of inversion algorithms that use GANs (Bora et al., 2017; Shah and Hegde, 2018).

A number of recent approaches have combined methods developed in the context of deep neural networks with traditional approaches to solve (ill-posed) inverse problems such as learning the regularization itself (Li et al., 2018; Lunz et al., 2018), using iterative deep neural networks (Adler and Öktem, 2017), or introducing deep neural networks as proxys for non-linear forward operators in reconstruction algorithms (Adler and Öktem, 2018).

Following the work of Chapter 6 we parameterize a family of geological models and associated permeability and porosity values using a GAN. We solve the transient two-phase immiscible Darcy flow equations without gravity and under isothermal

conditions using a finite-difference approach and obtain gradients of the mismatch between observed and simulated data with respect to the grid block permeability and porosity using the adjoint-state method.

Starting from random initial locations within the underlying Gaussian latent space controlling the GAN's output, we use gradient-descent to minimize the misfit between the flow data and the forward model output by modifying the set of latent variables. A regularization scheme of the optimization problem is derived from a Bayesian inversion setting. Based on a two-dimensional synthetic case of a channelized reservoir system we show that gradient-based inversion using GANs as a model parameterization can be achieved, honoring the observed physical quantities of the non-linear forward problem as well as borehole geological measurements. The trained models and code are available as open-source software[1].

7.3 Methodology

We consider the problem of determining the flow behavior of a subsurface hydrocarbon reservoir from which oil and water are being produced from a well. To displace the oil phase we inject water at a constant injection rate in another well and produce oil by maintaining a constant bottom-hole pressure. The forward problem $\mathcal{F}(\cdot)$ is therefore defined as solving the two-phase (oil and water) Darcy flow equations with a slightly compressible oil phase with no gravity effects and under isothermal conditions. The observed variables of the system y_{obs} are the injection pressure $p_{inj}(\mathbf{x}_i, t)$ at the location of the injection well $\mathbf{x}_i$ as a function of time t, as well as the produced volumes of water $q_w(\mathbf{x}_p, t)$ and oil $q_o(\mathbf{x}_p, t)$ at the production well $\mathbf{x}_p$.

$$\begin{aligned} \mathbf{y}_{obs,flow} &= \{p_{inj}(\mathbf{x}_i, t), q_w(\mathbf{x}_p, t), q_o(\mathbf{x}_p, t)\} \\ \mathbf{y}_{obs,well} &= \{\mathbb{1}_r(\mathbf{x}_{wells})\} \end{aligned} \tag{7.1}$$

We consider all fluid parameters of the forward problem, the oil and water viscosity $\mu_{o,w}$, density $\rho_{o,w}$, compressibility $c_{o,w}$, as well as the relative-permeability behavior, represented by a Brooks-Corey model, as constants (Blunt, 2017). A table of fluid properties is provided in Appendix Table 7.2. Therefore, the set of model parameters $\mathbf{m}(\mathbf{x})$ are the spatial distribution of the reservoir rock type indicator function $\mathbb{1}_r(\mathbf{x})$, the single-phase permeability $\mathbf{k}(\mathbf{x})$, and porosity $\phi(\mathbf{x})$.

$$\mathbf{m}(\mathbf{x}) = \{\mathbb{1}_r(\mathbf{x}), \mathbf{k}(\mathbf{x}), \phi(\mathbf{x})\} \tag{7.2}$$

We assume a binary set of reservoir rock-types where $\mathbb{1}_r(\mathbf{x}) = 0$ corresponds to a low permeability shale and $\mathbb{1}_r(\mathbf{x}) = 1$ indicates the high permeability sandstone of a river-channel.

Given the non-linear forward operator $\mathcal{F}(\cdot)$ this allows us to define the forward problem as

$$\mathbf{y}_{\mathbf{m}} = \mathcal{F}(\mathbf{m}) \tag{7.3a}$$

$$\mathbf{y}_{obs,flow} = \mathbf{y}_{\mathbf{m}} + \varepsilon,\ \varepsilon \sim \mathcal{N}(\mu_\epsilon, \sigma_\epsilon^2 \mathbf{I}) \tag{7.3b}$$

[1] https://github.com/LukasMosser/DeepFlow

We represent the distribution of possible model parameters $\mathbf{m}$ by a (deep) generative latent-variable model

$$\begin{aligned} \mathbf{m} &= G_\theta(\mathbf{z}) \\ \mathbf{z} &\sim p(\mathbf{z}) = \mathcal{N}(0, \mathbf{I}) \end{aligned} \tag{7.4}$$

with latent variables $\mathbf{z}$ sampled from a multi-Gaussian distribution (Eq. 7.4). Here we assume that the generative model has been trained prior to the inversion process, and is therefore constant.
We can then write

$$\mathbf{y}_{obs,flow} = \mathbf{y}_\mathbf{m} + \varepsilon = \mathcal{F}(G_\theta(\mathbf{z})) + \varepsilon \tag{7.5}$$

and from Bayes theorem

$$p(\mathbf{z}|\mathbf{y}_{obs,flow}) \propto p(\mathbf{y}_{obs,flow}|\mathbf{z})p(\mathbf{z}) \tag{7.6}$$

where for the observed flow data $\mathbf{y}_{obs,flow}$ the likelihood function $p(\mathbf{y}_{obs,flow}|\mathbf{z})$ is a Gaussian with mean $\mathcal{F}(G_\theta(\mathbf{z}))$ and variance σ_ϵ^2. We have assumed a Bernoulli likelihood for the observed facies at each well sample (Eq. 7.2).

Furthermore, we assume conditional independence of the likelihood of the observed flow and facies data at the wells

$$p(\mathbf{y}_{obs}|\mathbf{z}) = p(\mathbf{y}_{obs,flow}|\mathbf{z})p(\mathbf{y}_{obs,well}|\mathbf{z}) \tag{7.7}$$

In the case of GANs only individual samples can be obtained and therefore estimation of the posterior is restricted to point-estimates. We derive a regularization scheme for the history matching problem from finding point estimates of the maxima of the posterior distribution (MAP)

$$\begin{aligned} &\underset{\mathbf{z}}{\operatorname{argmax}}\{p(\mathbf{y}_{obs,flow}|\mathbf{z})p(\mathbf{y}_{obs,well}|\mathbf{z})p(\mathbf{z})\} && \text{(7.8a)} \\ &= \underset{\mathbf{z}}{\operatorname{argmax}}\{\log p(\mathbf{y}_{obs,flow}|\mathbf{z}) + \log p(\mathbf{y}_{obs,well}|\mathbf{z}) + \log p(\mathbf{z})\} && \text{(7.8b)} \\ &= \underset{\mathbf{z}}{\operatorname{argmin}}\{-\log p(\mathbf{y}_{obs,flow}|\mathbf{z}) - \log p(\mathbf{y}_{obs,well}|\mathbf{z}) - \log p(\mathbf{z})\} && \text{(7.8c)} \\ &= \underset{\mathbf{z}}{\operatorname{argmin}}\{\mathcal{L}_{flow}(\mathbf{z}) - \mathcal{L}_{wells}(\mathbf{z}) - \mathcal{L}_{prior}(\mathbf{z})\} && \text{(7.8d)} \\ &= \underset{\mathbf{z}}{\operatorname{argmin}}\{\frac{1}{2\sigma_\epsilon^2}\|\mathbf{y}_{obs,flow} - \mathcal{F}(G_\theta(\mathbf{z}))\|_2^2 && \text{(7.8e)} \\ &- \sum_{Cells}^{Well} [\mathbb{1}_r(\mathbf{x}_{wells}) \log(p(\mathbb{1}_r(\mathbf{x}_{wells}))) + (1 - \mathbb{1}_r(\mathbf{x}_{wells})) \log(1 - p(\mathbb{1}_r(\mathbf{x}_{wells})))] \\ &- \log p(\mathbf{z})\} \end{aligned}$$

where the squared-error norm of the difference between the observed data $\mathbf{y}_{obs,flow}$ and the modeled data $\mathbf{y}_\mathbf{m} = \mathcal{F}(G_\theta(\mathbf{z}))$ is a result of the assumption of a Gaussian likelihood in Eq. 7.3 and the well loss $\mathcal{L}_{wells}(\mathbf{z})$ corresponds to the binary cross-entropy between the observed facies indicator $\mathbb{1}_r(\mathbf{x}_{wells})$ and the facies probability $p(\mathbb{1}_r)$ of each grid-block at the wells.

The derived regularization scheme can be framed as a minimization problem over the latent variables i.e. we minimize Eq. 7.8 by modifying the latent variables $\mathbf{z}$. Specifically we use a squared-error functional that measures the difference between

the observed pressures and rates at the well locations averaged over the total duration T of the observed data

$$\mathcal{L}_{flow}(\mathbf{z}) = \frac{1}{T}\sum_{t=0}^{T}\Bigg[\left(\frac{\|q_w^{obs}(\mathbf{x}_p,\ t) - q_w^{\mathbf{m}}(\mathbf{x}_p,\ t)\|}{\sigma_q}\right)^2 + \left(\frac{\|q_o^{obs}(\mathbf{x}_p,\ t) - q_o^{\mathbf{m}}(\mathbf{x}_p,\ t)\|}{\sigma_q}\right)^2 + \left(\frac{\|p^{obs}(\mathbf{x}_i,\ t) - p^{\mathbf{m}}(\mathbf{x}_i,\ t)\|}{\sigma_p}\right)^2\Bigg] \tag{7.9}$$

where $\mathbf{x}_p$ and $\mathbf{x}_i$ are the locations of the production and injection wells respectively and measurement uncertainties have been quantified by the standard deviation of the oil and water rate

$$\sigma_q = 0.03\ q_{inj} \tag{7.10}$$

and the standard deviation of the pressure measurements

$$\sigma_p = 0.05\ p_{ref} \tag{7.11}$$

used to generate the set of observed pressure and rate data.

Using the adjoint-state method (Plessix, 2006; Suwartadi et al., 2012; Krogstad et al., 2015) and traditional neural network backpropagation (Rumelhart et al., 1988) we can obtain gradients $\frac{\partial \mathcal{L}_{flow}(\mathbf{z})}{\partial \mathbf{z}}$ of the error functional $\mathcal{L}_{flow}(\mathbf{z})$ (Eq. 7.9) with respect to the latent variables.

We therefore seek to minimize the total loss

$$\mathcal{L}(\mathbf{z}) = \mathcal{L}_{flow}(\mathbf{z}) - \mathcal{L}_{well}(\mathbf{z}) - \mathcal{L}_{prior}(\mathbf{z}) \tag{7.12}$$

consisting of the flow loss $\mathcal{L}_{flow}(\mathbf{z})$ (Eq. 7.9), the well-data loss $\mathcal{L}_{wells}(\mathbf{z})$, and the prior loss $\mathcal{L}_{prior}(\mathbf{z}) = \log p(\mathbf{z})$ (Eq. 7.8). We note that this is similar to the approach by Creswell and Bharath, (2018) which includes an additional loss term corresponding to the prior distribution of the latent variables when inverting the generator function of a GAN. The system of partial differential equations of the non-linear forward problem as well as the likelihood function (Eq. 7.9) have been previously implemented in the open-source MATLAB/GNU-Octave reservoir simulation framework MRST (Lie, 2019). An overview of the end-to-end coupling between the latent variables, the deep generative model G_θ, and the numerical forward model $\mathcal{F}(\cdot)$ is shown in Fig. 7.1.

To evaluate the proposed approach we create synthetic two-dimensional vertical cross-sections of a stacked river-channel system using an object-based approach. Channel-bodies are represented by half-circles with centres following uniform distributions in space. The full object-based modeling workflow is detailed in Mosser et al., (2018d). For the case of reservoir history matching (Eq. 7.3) each object-based model is associated with the three model parameters $\mathbf{m}$ following Eq. 7.2 (Fig. 7.2).

We train a deep generative adversarial network (GAN) (Goodfellow et al., 2014) on a set of 10^4 object-based realizations to learn the generative model $G_\theta(\mathbf{z})$ (Eq. 7.4) prior to obtaining MAP samples (Eq. 7.8) by optimizing Eq. 7.12. Specifically, we train a so-called Wasserstein-GAN (Arjovsky et al., 2017; Gulrajani et al., 2017), that minimizes the Wasserstein distance between generated and real probability

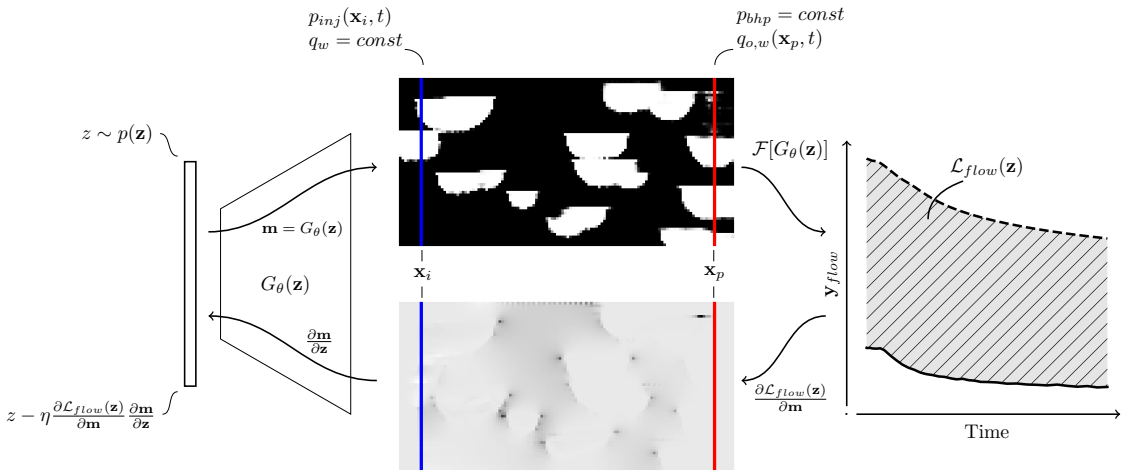

Figure 7.1: Overview of a single forward and backward-pass through the combined generative network G_θ and forward problem of two-phase Darcy flow $\mathcal{F}(\cdot)$. A set of model parameters $\mathbf{m} = G_\theta(\mathbf{z})$ is generated from a forward-pass through the generative model G_θ. The generative model creates spatial properties such as the permeability $\mathbf{k}$ shown here (top). The forward problem is solved where water is injected in the injection well (blue) and oil and water are produced from the production well (red). The error functional between observed parameters $\mathbf{y}_{obs}$ and modeled observations $\mathbf{y}_\mathbf{m} = \mathcal{F}(\mathbf{m}) = \mathcal{F}(G_\theta(\mathbf{z}))$ is evaluated. A spatially distributed gradient (bottom) with respect to the underlying properties is computed using the adjoint state method. The obtained gradients are then backpropagated by traditional neural-network backpropagation to compute a gradient with respect to the latent variables. The error functional is minimized by performing gradient descent using the obtained gradient.

distributions. To stabilize training of the Wasserstein-GAN we use the Lipschitz penalty introduced by Petzka et al., (2017). A detailed overview of the training procedure for the deep generative model is presented in Mosser et al., (2018d). Network architectures for the generator $G_\theta(\mathbf{z})$ and discriminator D_ω are detailed in Appendix Table 7.1.

The generator network $G_\theta(\mathbf{z})$ is a deep convolutional neural network that maps a 100-dimensional noise-sample from a standardized multi-Gaussian distribution $z \sim \mathcal{N}(0, \mathbf{I})$ to a spatial distribution of the model parameters $\mathbf{m}$ (Eq. 7.2) with 128 pixels in the x- and 64 pixels in the z-direction

$$\mathbb{R}^{50\times2\times1} \rightarrow \mathbb{R}^{3\times128\times64} \tag{7.13}$$

where the three output features correspond to the grid-block rock-type probability, permeability and porosity (Eq. 7.2). Due to the continuous nature of the generative networks sigmoid activation function, the output of the generative network corresponding to a representation of the reservoir rock-type $\mathbb{1}_r$ can be interpreted as a rock-type probability. The grid-block permeability and porosity values are obtained by transforming the generated rock-type probability from the interval $(0-1)$ to a permeability by a linear transformation (Appendix Table 7.1a) to rescale the output of the generative model to a range of values that correspond to admissible petrophysical values for sandstone and shale. The finite-difference approximation of the forward model is computed on the regular grid given by the pixel-based output

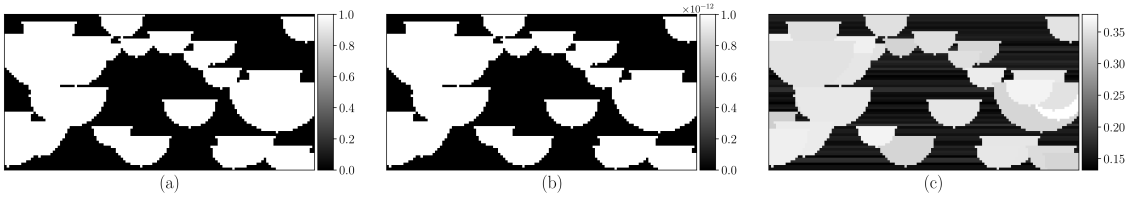

Figure 7.2: Reference object-based model used to generate observed data. Each object-based model has three associated properties: (a) the rock-type (shale 0 - sandstone 1) indicator $\mathbb{1}_r$, (b) the grid block permeability $\mathbf{k}$, (c) and porosity ϕ. The model was obtained from a test-set of object-based model realizations that the generator network $G_\theta(\mathbf{z})$ was not trained on.

of the generative model $G_\theta(\mathbf{z})$ and hence each pixel-based property corresponds to the grid-block property used in the numerical evaluation of the forward problem.

We have selected a reference realization from a test-set of object-based realizations that were generated independently of the models used to train the generative network G_θ. The model shown in Fig. 7.2 was used to generate a set of synthetic observed data $\mathbf{y}_{obs}$ that served as a reference case used to validate the proposed inversion approach. Gaussian noise was added to the production rates and pressures based on assumed measurement uncertainties (Eqs. 7.10-7.11). The transient pressure and rates of oil and water were recorded for a duration of 600 days and the full history was used in the gradient-based optimization outlined in Eq. 7.8. To minimize the total loss (Eq. 7.12) we used the ADAM optimizer (Kingma and Ba, 2014) using a fixed step size and β-parameters ($\eta = 3 \times 10^{-2}$, $\beta_1 = 0.9$, $\beta_2 = 0.999$). The simulation parameters used for the numerical evaluation of the forward problem are provided in Appendix Table 7.2.

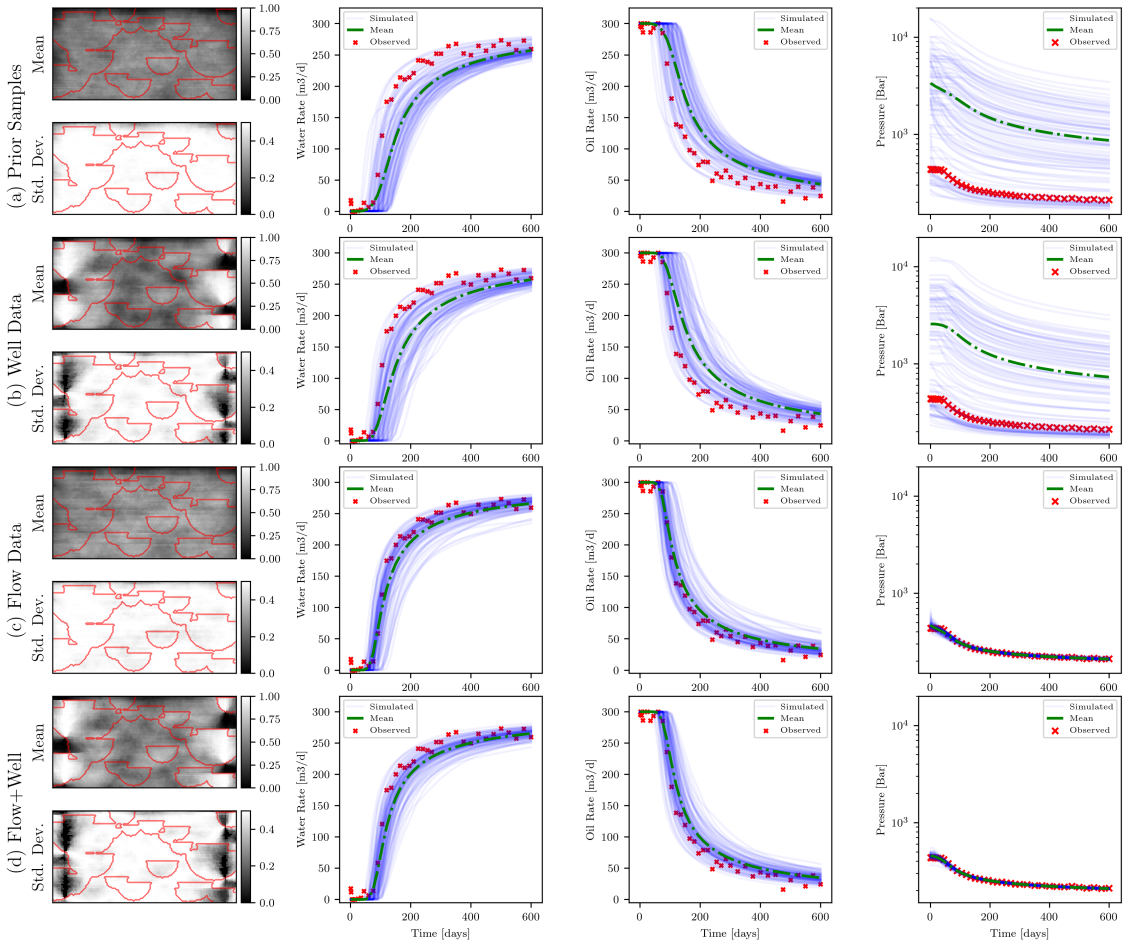

Figure 7.3: (a) Overview of the dynamic flow behavior, as well as average and standard deviation of the grid block rock-types for $N = 100$ samples obtained from the unconditional prior distribution i.e. samples drawn at random from the generative model. (b) Model realizations after optimizing only the grid block rock-type at wells, (c) optimizing only the dynamic flow behavior, and (d) optimizing both well rock-type and flow behavior. Red contours on the mean and standard deviation of rock-type distributions reflect the reference model of object-based channel-bodies shown in Fig. 7.2.

7.4 Results

Four history matching scenarios were considered:

- Scenario 1: Samples of the unconditional prior were obtained ($N = 1000$) and a single evaluation of the forward problem was performed. This allowed us to evaluate the prior distribution of the dynamic flow behavior of the underlying generative model $\mathbf{y}_{prior} = \mathcal{F}(\mathbf{m}) = \mathcal{F}[G_\theta(\mathbf{z})]$ (Fig. 7.3-a).

- Scenario 2: Inversion was performed by considering only the well data loss $\mathcal{L}_{well}(\mathbf{z})$ at the injection and production wells and the prior loss $\mathcal{L}_{prior}(\mathbf{z})$ to honor the prior distribution of the latent variables $\mathbf{z}$ (Fig. 7.3-b).

- Scenario 3: We performed inversion by considering only the flow-based loss $\mathcal{L}_{flow}(\mathbf{z})$ due to the mismatch between simulated and observed dynamic two-phase flow data, as well as the prior loss on the latent variables $\mathcal{L}_{prior}(\mathbf{z})$ (Fig. 7.3-c).

- Scenario 4: All three losses were combined: the flow loss $\mathcal{L}_{flow}(\mathbf{z})$, the well-data loss $\mathcal{L}_{well}(\mathbf{z})$, and the prior loss $\mathcal{L}_{prior}(\mathbf{z})$, corresponding to the total loss $\mathcal{L}(\mathbf{z})$, as outlined in Eq. 7.12 (Fig. 7.3-d).

For Scenario 2 to 4, we performed $N = 100$ inversions, starting from the same set of 100 prior random samples of latent variables $\mathbf{z}$. Each optimization procedure was run for a total of 500 ADAM optimization steps. For Scenarios 3 and 4, we choose the iteration with the lowest total loss as the output of each run. Early stopping was performed in Scenario 2 where only the well data loss was considered when the facies accuracy at the well reached 100%. Figure 7.3 shows a comparison of these four scenarios in terms of their dynamic flow behavior i.e. pressures and rates at production and injection wells as a function of the simulation time, as well as the mean and standard deviation of the thresholded rock-type probability maps. The threshold for rock-type probability maps was set at 0.5 to distinguish between sand (1) and shale (0).

The prior distribution of the observed variables $\mathbf{y}_{prior}$ (Scenario 1, Fig. 7.3-a, right) shows a large spread in the distribution of the observed oil-water production rates and injection pressure. The observed liquid rates of the reference case (red) lie at the edge of the unconditional data distribution. The pressure behavior of the prior models varies across 2 orders of magnitude in relation with the pressure behavior of the reference case at the low end of the pressure versus time distribution.

The unconditional prior samples of the model parameters, shown here as the rock-type indicator function $\mathbb{1}_r$, are shown as cross-sections of the mean and standard deviation of the $N = 100$ models (Scenario 1, Fig. 7.3-a, left). As shown by the standard deviation map, nearly all values are close to a value of 0.5 as expected from an indicator function of mean 0.5.

Considering the likelihood of the grid-block rock-type indicator function at the wells only (Scenario 2) where the binary cross-entropy loss at the injection and production well $\mathcal{L}_{well}(\mathbf{z})$, as well as the prior loss $\mathcal{L}_{prior}(\mathbf{z})$ were minimized, we find a large spread in the production behavior of the generated samples closely resembling the unconditional prior distribution (Scenario 2, Fig. 7.3-b, right). Inspecting the mean and standard deviation maps of the rock-type indicator function of the inverted

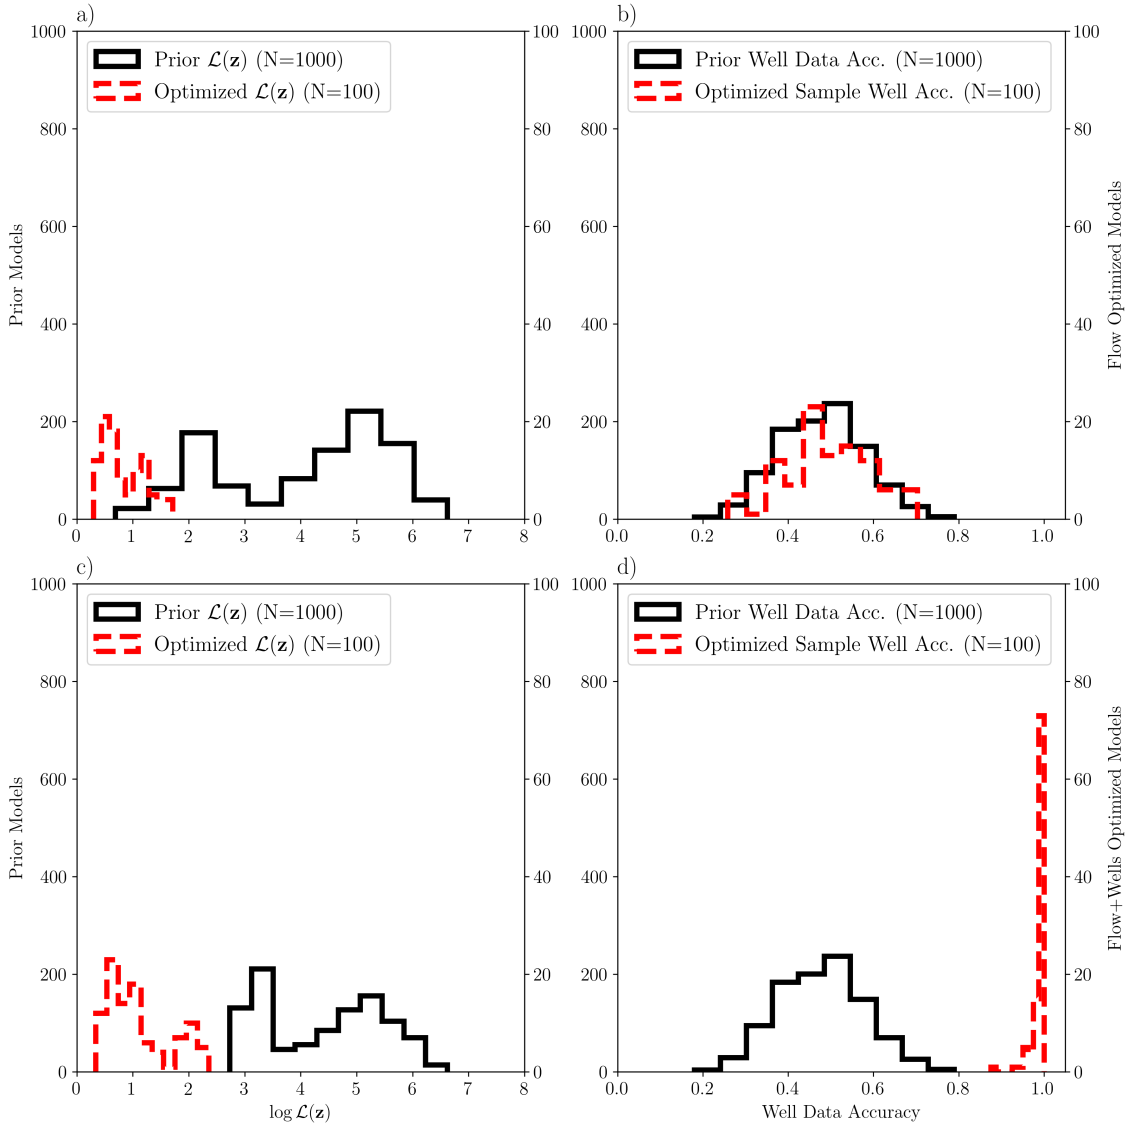

Figure 7.4: Histograms of the total loss functional $\log \mathcal{L}(\mathbf{z})$ (a, c) and the grid-block rock-type accuracy at the wells (b, d). Models ($N = 100$) optimized on the observed flow data only (a-b, red) and optimized on both flow and well rock type (c-d, red) are compared to models ($N = 1000$) sampled unconditionally from the prior distribution (a-d, black). In both cases, we observe that optimizing the flow-based objective leads to realizations that match the observed data closer than sampling unconditionally from the prior. The accuracy of the well rock type for models optimized only on the dynamic flow data follow the prior distribution of models, whereas the models obtained by optimizing flow and well-data loss functions have high accuracy at the well locations.

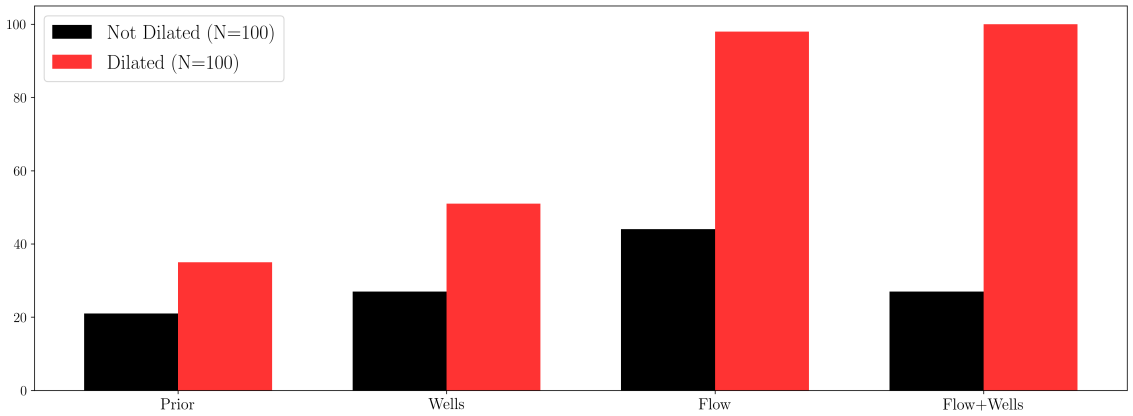

Figure 7.5: Comparison of the number of models that have a connected cluster of channel-bodies that connects the injection and production wells in each individual sample obtained in the inversion process for the four scenarios outlined in Section 7.3. After Boolean dilation, nearly all models that were optimized with respect to the observed dynamic flow data form a connected cluster between the two wells, whereas half of the samples obtained from honoring only the well data remained disconnected.

samples matching the well data only, we find that these samples match the well data while their flow behavior is consistent with the unconditional prior. This means that constraining only by well data has very little effect on the facies model connectivity between the wells. A lateral effect of the grid-block well data can be observed comparing the outline of the mean rock-type maps and the red contours of the reference model (Fig. 7.2) used to generate the reference observed data $\mathbf{y}_{obs,flow}$ and $\mathbf{y}_{obs,well}$.

When only the flow loss $\mathcal{L}_{flow}(\mathbf{z})$ was considered in the optimization, the production data of the generated realizations were tightly constrained to the observed liquid rates and pressures (Scenario 3, Fig. 7.3-c, right). The mean and standard deviation maps for the obtained distribution of the rock-type indicator functions $\mathbb{1}_r$ do not indicate any clear structural features and resemble the prior maps shown in the first row of Fig. 7.3. We have obtained samples of model parameters $\mathbf{m}$ that all honor the observed dynamic flow data $\mathbf{y}_{obs,flow}$, but which are not anchored at the wells by the observed well rock types. Combining the flow loss $\mathcal{L}_{flow}(\mathbf{z})$, the well data loss $\mathcal{L}_{well}(\mathbf{z})$ and prior loss $\mathcal{L}_{prior}(\mathbf{z})$ as the total loss (Eq. 7.12) (Scenario 4, Fig. 7.3-d) we obtain samples that match the observed dynamic behavior closer than Scenario 3, which only considered the flow loss. In addition to honoring the dynamic observed data $\mathbf{y}_{obs,flow}$, the facies data $\mathbb{1}_r$ of the obtained samples at the wells match the observed grid-block data at the production and injection well $\mathbf{y}_{obs,well}$. This is indicated by the low standard deviation map of the $N = 100$ inverted models (Scenario 4, Fig. 7.3-d, left). Figure 7.4 compares the distribution of the total loss ($\mathcal{L}(\mathbf{z})$, Eq. 7.12) for the $N = 1000$ models obtained from the unconditional prior distribution i.e. sampling a set of model parameters $\mathbf{m}$ (Eq. 7.2) from the generative model $G_\theta(\mathbf{z})$ and solving the forward problem once, (Eq. 7.3) and the models of Scenario 3 (flow and prior loss only) and 4 (flow, well and prior loss). For both scenarios considered (Fig. 7.4, a and c respectively) the distribution of the total loss is shifted to lower values when compared to the total loss distribution of the models obtained by sampling from the prior distribution. The distribution of well rock type accuracy shows that history

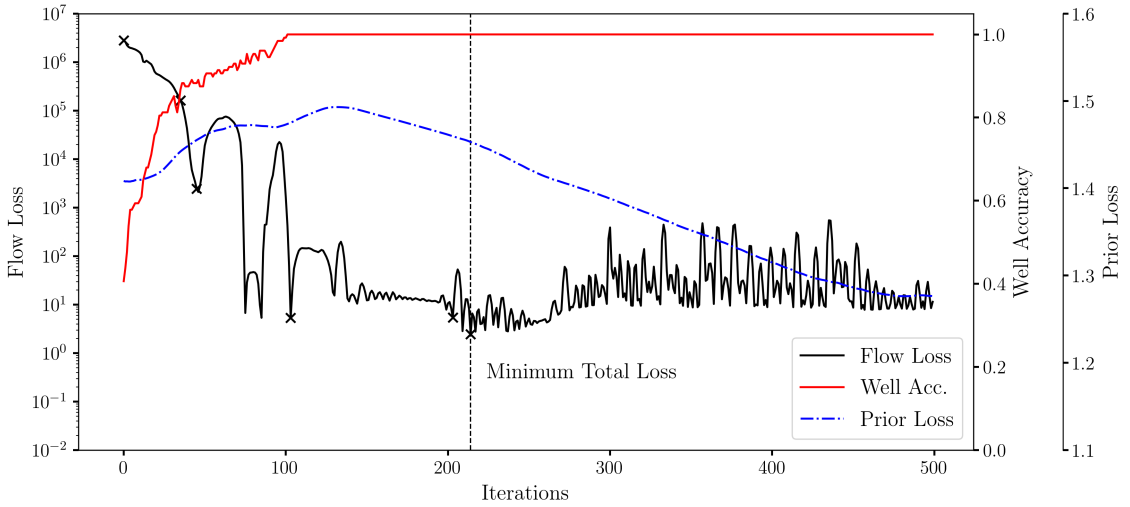

Figure 7.6: Optimization trajectory for a selected inverted sample based on flow, well-data and prior losses. The flow loss $\mathcal{L}(\mathbf{z})$ shows a highly non-linear optimization problem with many local minima leading up to the smallest error being achieved after 214 iterations. The well data accuracy at the production and injection wells quickly reaches values greater than 90% accuracy. Cross-marks indicate the loss values of the intermediate realizations presented in Fig. 7.8.

matching only flow data (Fig. 7.4 b) does not improve the rock type accuracy. When the well-data loss $\mathcal{L}_{well}(\mathbf{z})$ is optimized in addition to the flow-loss (Fig. 7.4 d), we observe that the obtained models closely match the rock type at the wells in addition to the dynamic flow data.

Due to the high computational cost of the forward model, reaching a good match with the observed data, within a reasonable number of forward-model evaluations, is necessary. Appendix Fig. 7.10 shows the distribution of the number of optimization iterations of Eq. 7.8 necessary to reach a given threshold value of the total loss $\mathcal{L}(\mathbf{z})$ for the case where all three loss functions are used in the history matching process (Scenario 4). All $N = 100$ inverted models achieve a total-loss less than 1×10^3 (Fig. 7.10 c) which corresponds to the low end of the loss values obtained by sampling from the prior distribution only (Fig. 7.4 c). In Fig. 7.3 we have shown that including the flow loss $\mathcal{L}_{flow}(\mathbf{z})$ leads to high variance in the generated distribution of the model properties $\mathbf{m}$ while honoring the observed flow data (Scenario 3, Fig. 7.3-c, right), while only constraining samples to honor the well data reduces the variance near the wells, but does not constrain to observed dynamic data (Scenario 2, Fig. 7.3-b, right). A smaller injection pressure observed for the samples obtained from Scenario 3 compared to those that match only the wells in Scenario 2 would indicate a better connected system of channel bodies. After closer inspection of the inverted samples it was found that nearly all models obtained from Scenario 3 and 4 are connected (Fig. 7.5 and Appendix Fig. 7.14-7.15). We evaluate for all four scenarios whether any of the clusters of river-channel bodies connects the injection and production wells. Less than half of the models for all scenarios show a connection between injector and producer wells (Fig. 7.5 - black). However after Boolean dilation of the largest cluster and performing the connectivity analysis again, half of the samples that honor the well data only (Scenario 2) show connectivity, whereas all models obtained that were optimized to honor flow (Scenario 3), as well as flow and well data (Scenario 4) form a

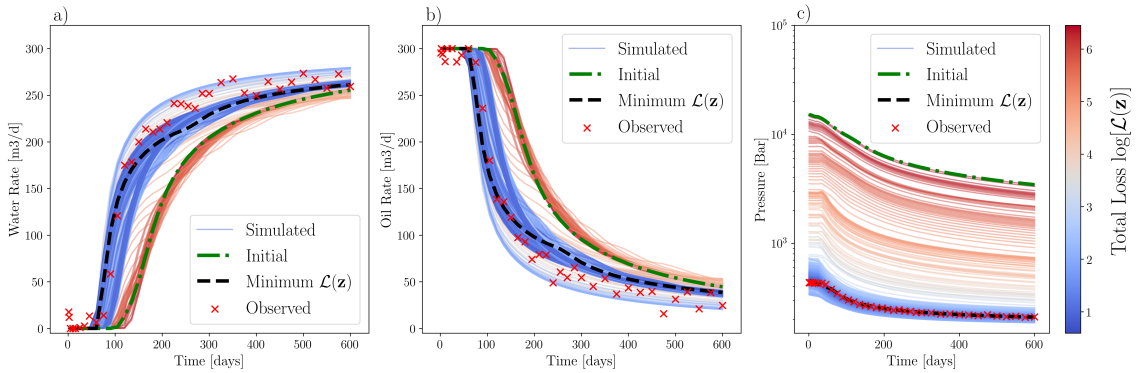

Figure 7.7: Comparison of the observed dynamic pressure and rate data (red) obtained for a single inversion case. The observed dynamic data from the initial starting model (green) is optimized by honoring flow and well data for 500 iterations. The intermediate pressure and rate data of the 500 iteration steps are colored according to their total loss $\mathcal{L}(\mathbf{z})$.

connected cluster between the injector and producer well (Fig. 7.5 - red). A possible explanation of this result is that due to the finite-difference approximation (two-point flux approximation) used to solve the forward-problem the transmissibilities between two neighbouring grid blocks are computed as a weighted average of the permeability values i.e. a single grid block spacing between two high-permeable channel-bodies does not provide an effective flow barrier.

A single model obtained from optimization using flow and well losses, while honoring the prior distribution of the latent-variables (Scenario 4), was chosen as an example to highlight the optimization process for a single realization. This specific example was chosen as it showed the largest reduction in the total loss $\mathcal{L}(\mathbf{z})$ from the initial starting model to the iteration with the lowest total loss out of 500 optimization steps. Figure 7.6 shows the evolution of the total loss $\mathcal{L}(\mathbf{z})$ (black), the accuracy in matching the rock-type indicator function $\mathbb{1}_r$ at the well locations (red), and the prior loss evolution. A highly non-linear optimization with many local minima can be observed. The accuracy of matching the well data can be observed to quickly reach values above 90% and staying at 100% for the entirety of the optimization process. The smallest total loss for this model is achieved after 214 iterations.

For the example model, we have recorded the observed dynamic flow behavior as well as the distribution of the flow parameters at intermediate steps in the optimization procedure, where the total loss $\mathcal{L}(\mathbf{z})$ achieved a new local minimum. Figure 7.7 shows the evolution of the dynamic flow behavior in the course of performing the inversion. The model initially shows a very high injection pressure (Fig. 7.7 c -blue), which is significantly and continuously reduced to values close to the observed pressure data (red). Similar behavior is observed for the liquid water and oil rates at the production wells (Fig. 7.7 a-b). The solution found by the optimization strategy (Eq. 7.8) with the lowest error value (Fig. 7.7 a-c, black) closely matches the observed dynamic production data and the rock-type indicator function $\mathbb{1}_r$ at the wells.

The intermediate spatial distributions of the rock-type indicator $\mathbb{1}_r$ for each of the obtained local minima in the optimization process are shown in Fig. 7.8. Initially the channel-bodies were completely disconnected (Fig. 7.8 a). As the optimization process progresses more channel-bodies were introduced gradually (Fig. 7.8 b-e) and

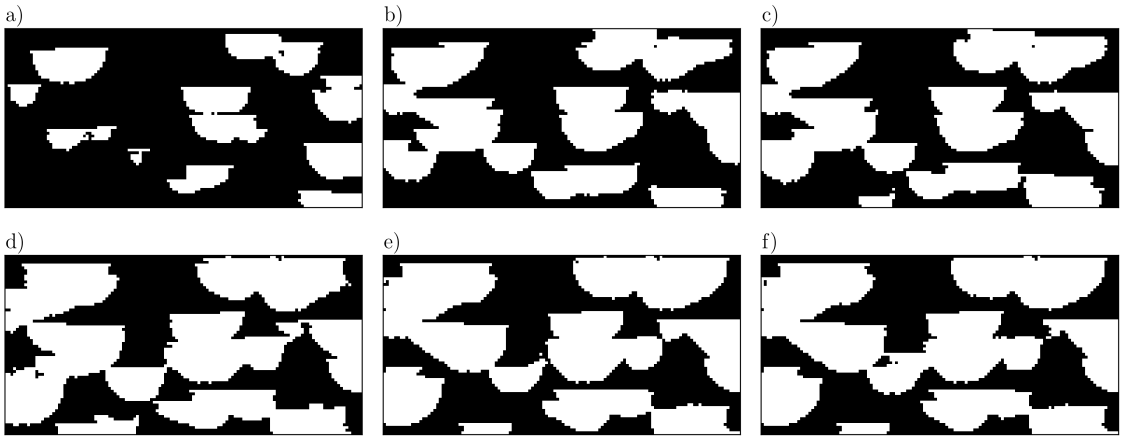

Figure 7.8: Evolution of the spatial distribution of the rock-type indicator $\mathbb{1}_r$ for the selected model.

form a spanning cluster that matches the grid-block data at the wells (Fig. 7.8 f).

7.5 Discussion

We have shown that solutions of the ill-posed inverse problem of reservoir history matching can be obtained by inversion in the space of a deep generative adversarial network for a synthetic case. The model parameter distributions $\mathbf{m}$ match the obtained dynamic observed flow data $\mathbf{y}_{obs}$ as well as honor the grid-block properties at the wells with adequate solutions found on average in less than 100 optimization iterations. To evaluate whether the obtained samples of the posterior distribution are local maxima we followed the approach by Zhang et al., (2019) and performed a cyclical spherical linear interpolation (White, 2016) in the latent space between three obtained realizations. For each interpolated latent vector we have solved the forward model and evaluated the loss functions. Figures 7.9a and 7.9b indicate that the posterior distribution $p(\mathbf{z}|\mathbf{y}_{obs})$ is highly multi-modal, where each of the obtained realizations lie at a minimum of the total loss $\mathcal{L}(\mathbf{z})$ i.e. near a local maximum of the posterior. While the samples obtained match the flow and well data it is important to note that in a high-dimensional setting, finding the maxima of the posterior may not be a sufficient description of the full posterior distribution and this needs to be considered when using the obtained history-matched ensemble of reservoir models for forecasts of future reservoir performance. To estimate posterior parameters such as the cumulative production, a harmonic mean averaging approach (Green, 1995; Raftery et al., 2006) could be used where individual samples are weighted by their corresponding likelihood (Zhang et al., 2019). Furthermore, a number of open challenges are still to be addressed in the context of solving inverse problems with deep generative models such as GANs.

The synthetic case presented is a simplified representation of a system of river-channel bodies. The features encountered in real reservoir systems are more diverse and heterogeneous, and it is possible that many other reservoir features should be included in the training set of generative models. The types of models that may be needed include three-dimensional multi-rock-type models i.e. more than two rock-types in a generated reservoir model. Furthermore, it may be of interest to

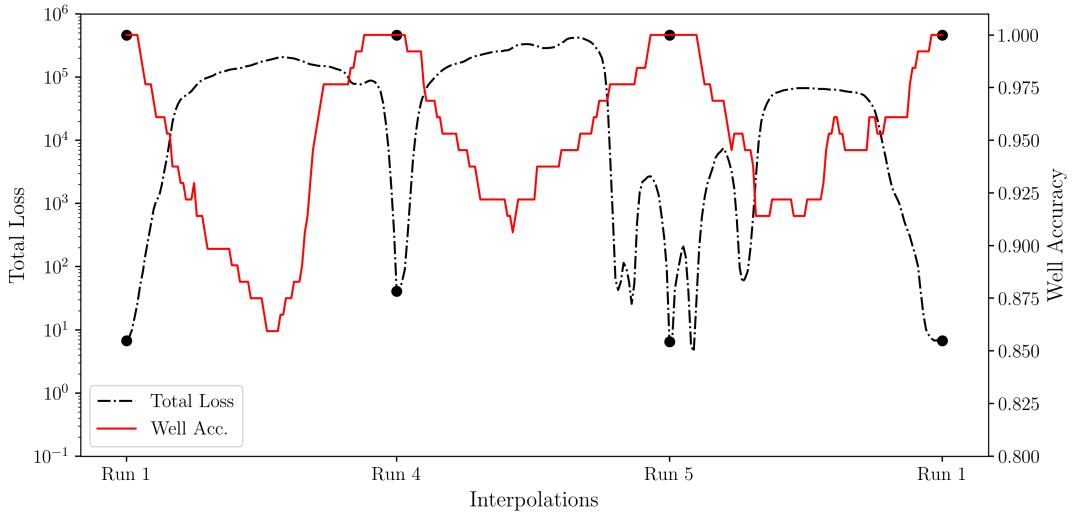

(a) Total Loss $\mathcal{L}(\mathbf{z})$ and well facies accuracy for 100 realizations interpolated between three obtained maxima of the posterior $p(\mathbf{z}|\mathbf{y}_{obs})$.

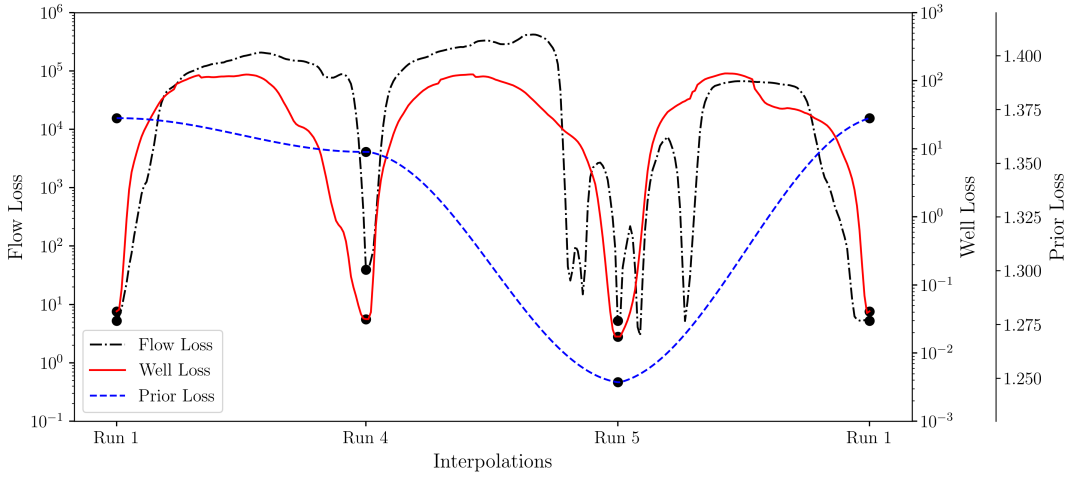

(b) Flow Loss $\mathcal{L}_{flow}(\mathbf{z})$, well loss $-\mathcal{L}_{well}(\mathbf{z})$, and prior loss $-\mathcal{L}_{prior}(\mathbf{z})$ of the interpolated latent variables between three maxima obtained by optimizing the total loss (Scenario 4).

Figure 7.9: Values of (a) the total loss and well accuracy, and (b) flow, well and prior losses for interpolations of the latent variables between realizations obtained from MAP sampling (bold circles) for Scenario 4 (Section 7.4).

combine a number of datasets into a single generative model that is conditioned on the type of reservoir architecture $G_\theta(\mathbf{z}, \mathbf{c})$, where $\mathbf{c}$ is a vector of conditional variables that influence the output of the deep generative model using discrete conditionals such as depositional environment e.g. fluvial versus lacustrine environments or using continuous conditioning variables such as reservoir net-to-gross. This type of conditional generative model could be facilitated by making use of architectures such as that proposed by Miyato and Koyama, (2018). Empirically, it has been shown that introducing class-conditional variables to GANs, so-called cGANs, result in improved training and distributional representation behavior compared to their unconditional counter-parts (Gauthier, 2014). This may speak for training a single general reservoir model generator compared to creating a library of standalone pre-trained generator networks for each depositional environment.

An important aspect to be considered here is the evaluation of the generated models in terms of model quality, mode-collapse, and representation of spatial statistics as highlighted in Mosser et al., (2017) and Laloy et al., (2017b), where two-point correlation functions are used to quality control the individual generated reservoir models. The use of established GAN quality control measures such as Inception Score (IS) (Salimans et al., 2016) or Frechet Inception Distance (FID) (Heusel et al., 2017) may be challenging due to the induced distribution shift between ImageNet (Deng et al., 2009) pre-trained networks and the features observed in typical reservoir architectures. Kernel-based methods such as the maximum mean discrepancy (MMD) have already been successfully applied to train GANs on reservoir models (Chan and Elsheikh, 2018b) which could be repurposed to quality control generated reservoir models.

While the synthetic case presented herein has shown that a gradient-based approach to solving the ill-posed inverse problem is feasible, other strategies such as using Markov-Chain Monte-Carlo methods as demonstrated by Laloy et al., (2017a), may lead to successful inversion results, but at high computational cost. Combinations with other established inversion techniques, such as Ensemble Kalman-Filters (Gu et al., 2005; Aanonsen et al., 2009) or Ensemble Smoothers (ES-MDA) (Emerick and Reynolds, 2013) should also be evaluated. The presented approach, together with a set of suitable generative models should be applied to benchmark studies such as the PUNQ or SPE comparative solutions datasets to investigate the applicability of inversion with deep generative models in the presence of measurement errors and noise. While the present study only considers inverting for one model at a time i.e. using a batch-size of one, it is possible to perform inversion using a batch-size greater than one, as in Creswell and Bharath, (2018) who show that the gradient contributions to the latent variables can be computed independently for each element of a batch of latent-variables and hence for a set of model parameters $\mathbf{m}$. By computing the numerical solution forward problem and the associated adjoint for each set of model-parameters within a batch, the proposed method could be integrated into existing workflows for ensemble reservoir modeling.

Using deep latent variable models, such as GANs, to represent the space of solutions for ill-posed inverse problems opens up interesting avenues with regards to the theoretical basis of the ill-posed nature of these problems. Investigation of the existence and stability criteria (Section 7.2) (Tikhonov and Arsenin, 1977) of deep latent variable inversion schemes could lead to new promising theoretical insights into the convergence behavior of inversion methods using deep generative models. Furthermore, the geometry of the loss landscape and representation manifolds within the space of latent variables may allow for faster convergence and higher quality inversion results. Chen et al., (2017) and Arvanitidis et al., (2017) have investigated the geometric structure of the manifold represented by deep generative models. Shao et al., (2018) have shown that the manifold in the data space obtained from VAEs are non-linear with near-zero curvature and linear paths in latent space correspond to geodesics on the data manifold. This indicates that the geometry of the learned manifolds is accounted for in operations performed in the latent space of deep generative models.

We have shown that GANs may be used as generative models for parameterization of geological models, and can provide a solution space for ill-posed inverse problems. Nevertheless GANs, due to their challenging training, evaluation and complex latent-

to-parameter-space relationship may not be the optimal reparameterization choice in the case of ill-posed inverse problems. Deep generative models such as variational (Kingma and Welling, 2013) and disentangling autoencoders (Higgins et al., 2017; Burgess et al., 2018), as well as flow-based generative models, such as RealNVP (Dinh et al., 2016) or GLOW (Kingma and Dhariwal, 2018), that provide a bijective and invertible mapping between latent space and model space may be better suited for the inversion task. This has been demonstrated by Ardizzone et al., (2018) for ill-posed inverse problems.

7.6 Conclusions

We have presented an application of deep generative models in the context of the ill-posed history matching inverse problem. Based on a two-dimensional synthetic cross-section of a river-channel system we have trained a GAN to represent the prior distribution of the subsurface properties $\mathbf{m}$, permeability $\mathbf{k}$, porosity ϕ and the rock-type indicator $\mathbb{1}_r$. We find solutions to the two-phase slightly-compressible Darcy flow problem of oil and water systems commonly used to describe the transient flow behavior in hydrocarbon reservoirs. By incorporating a finite-difference-based numerical simulator with adjoint-state capabilities (Krogstad et al., 2015) in an end-to-end differentiable framework, we perform inversion using gradient-based optimization of the mismatch between observed dynamic data (Eq. 7.9) and grid-block-scale well data (Fig. 7.3), while honoring the prior distribution of the latent variables (Eq. 7.12). By using a momentum-accelerated first-order gradient descent scheme (Kingma and Ba, 2014), the method converges to a solution despite the highly non-linear and non-convex loss landscape (Fig. 7.6). Future work will focus on applications to real reservoirs and history matching benchmark studies such as the PUNQ or SPE comparative solutions models, as well as the evaluation of other deep latent-variable generative models and their theoretical benefits for finding solutions to ill-posed inverse problems.

7.7 Chapter Appendix

7.7.1 Generative Model Architecture and Quality Control

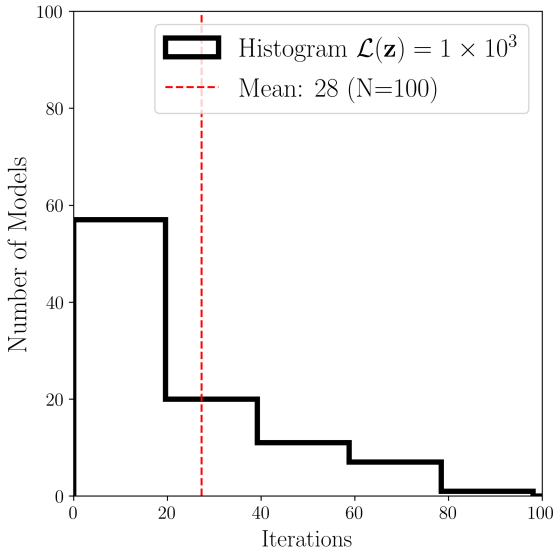

Figure 7.10: Distribution of the number of optimization iterations required to reach a total loss $\mathcal{L}(\mathbf{z})$ (Eq. 7.9) of less than or equal to 1×10^3 in the case of matching flow and well rock type data (Scenario 4). This threshold corresponds to the lower end of the total-loss distribution of the prior distribution (Fig. 7.4 c). Optimization converges on average within the first 30 iterations.

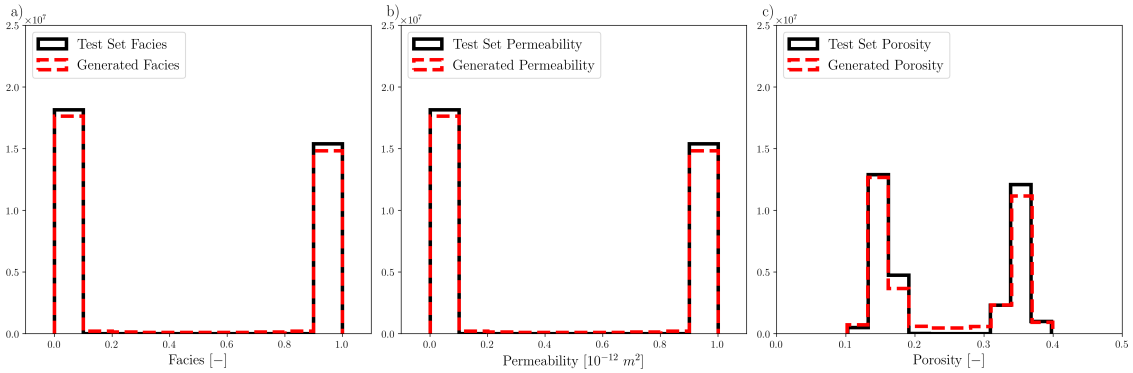

Figure 7.11: Comparison of the pdfs of $N = 1 \times 10^4$ samples of the test set and $N = 1 \times 10^4$ samples obtained from the generative network $G_\theta(\mathbf{z})$ trained to represent the river-channel body system. The reservoir rock-type (a) and permeability (b) show a near binary distribution whilst porosity follows a bi-modal pdf. The pdfs of the generated samples match the test set closely.

Table 7.1: Generator and discriminator network architectures used to represent the distribution of model parameters $\mathbf{m}$ used in the inversion process (Eq. 7.2). Binary indicators of geological facies and derived permeability are represented by a bi-variate Gaussian distribution with a hyperbolic tangent activation function used to create the binary output distribution. We use a linear transformation layer to renormalize the output of the GAN from the $(-1, 1)$ interval of the tanh activation to a known range of permeability and porosity values. Permeability transform parameters: $a = 1 \times 10^{-3}$, $b = 1 \times 10^{-12}$
Porosity transform parameters: $c = 0.3$, $d = 0.1$
Notation for convolutional layers: LayerType(Number of filters), k=kernel size, s=stride, p=padding. BN=BatchNorm, PS=PixelShuffle

Latent Variables $z \in \mathbf{R}^{50\times2\times1}$
Conv2D(512)k3s1p1, BN, ReLU, PSx2
Conv2D(256)k3s1p1, BN, ReLU, PSx2
Conv2D(128)k3s1p1, BN, ReLU, PSx2
Conv2D(64)k3s1p1, BN, ReLU, PSx2
Conv2D(64)k3s1p1, BN, ReLU, PSx2
Conv2D(64)k3s1p1, BN, ReLU, PSx2
Conv2D(2)k3s1p1
$m_0 = Tanh(Channel\ 0) \cdot 0.5 + 0.5 = p(\mathbb{1}_r)$ $x_1 = Tanh(Channel\ 1) \cdot 0.5 + 0.5$
$m_1 = (a + m_0) \cdot b =$ permeability $\mathbf{k}$ $m_2 = c \cdot x_1 + d =$ porosity ϕ

(a) Multi-Property Generator

Model Parameters $\mathbf{m} \in \mathbf{R}^{2\times128\times64}$
Conv2D(64)k5s2p2, ReLU
Conv2D(64)k5s2p1, ReLU
Conv2D(128)k3s2p1, ReLU
Conv2D(256)k3s2p1, ReLU
Conv2D(512)k3s2p1, ReLU
Conv2D(512)k3s2p1, ReLU
Conv2D(1)k3s1p1, ReLU

(b) Discriminator

Table 7.2: Fluid and simulation parameters for oil and water phases used in the numerical solution of the two-phase flow forward problem.

Parameter		Value	Unit
Normalized Water Saturation	$\bar{S}_w$	$\frac{S_w - S_{w,cr}}{1 - S_{or} - S_{w,cr}}$	–
Normalized Oil Saturation	$\bar{S}_o$	$\frac{S_o - S_{o,cr}}{1 - S_{wc} - S_{o,cr}}$	–
Water, Oil Viscosity	$(\mu_w,\ \mu_o)$	$(3 \times 10^{-4},\ 5 \times 10^{-3})$	$Pa \cdot s$
Water, Oil Density	$(\rho_w,\ \rho_o)$	$(1000,\ 700)$	$kg \cdot m^{-3}$
Brooks Corey Exponents	$(n_w,\ n_o)$	$(2,\ 2)$	–
Water, Oil Compressibility	$(c_w,\ c_o)$	$(0.0,\ 1 \times 10^{-5})$	bar^{-1}
Connate, Init., Crit. Water Saturation	$(S_{wc},\ S_{w,i},\ S_{w,cr})$	$(0.10,\ 0.15,\ 0.15)$	–
Residual, Init., Crit. Oil Saturation	$(S_{or},\ S_{o,i},\ S_{o,cr})$	$(0.10,\ 0.15,\ 0.12)$	–
Water Injection Rate	$q_{w,inj}$	300	$m^3 \cdot day^{-1}$
Bottomhole Pressure	p_{bhp}	150	bar
Reference Pressure	p_{ref}	200	bar

7.7.2 Inverted Samples

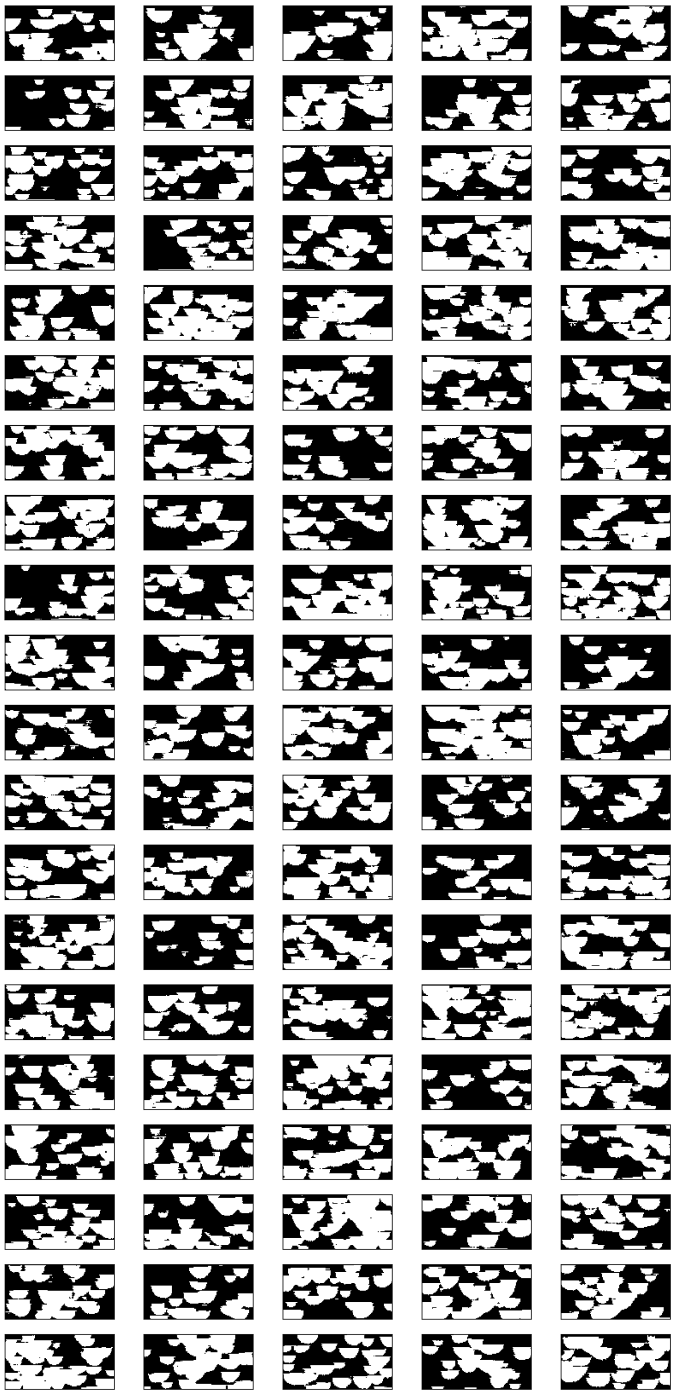

Figure 7.12: Unconditional samples ($N = 100$) obtained from sampling the prior distribution of the generative model (Section 7.4, Scenario 1)

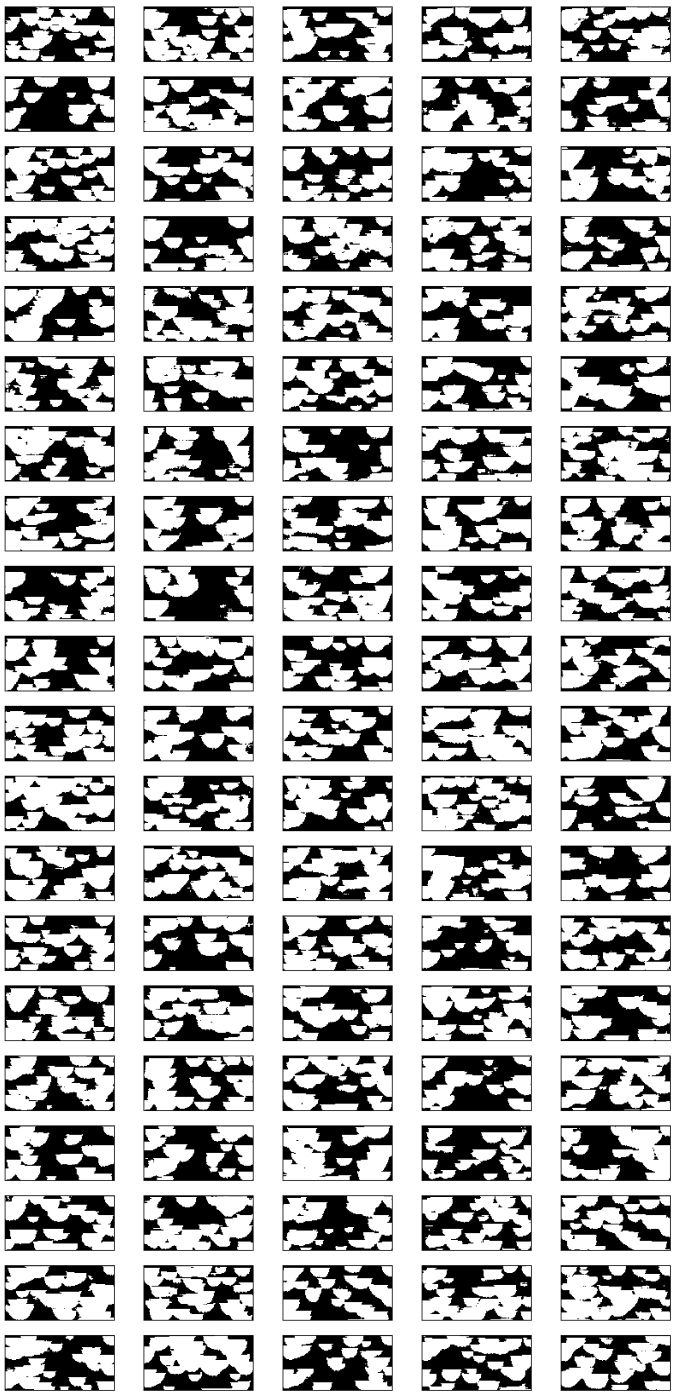

Figure 7.13: Samples obtained by minimizing the sum of the well and prior losses (Section 7.4, Scenario 2)

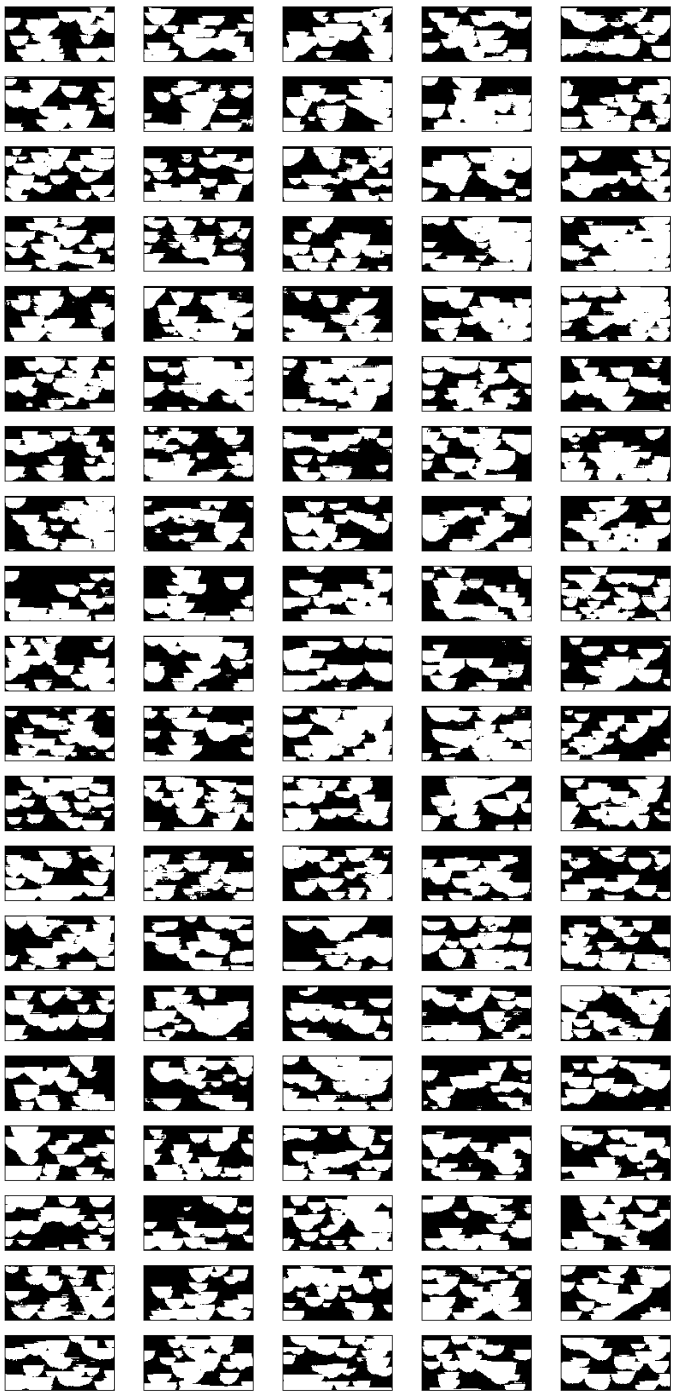

Figure 7.14: Samples obtained by minimizing the sum of the flow and prior losses (Section 7.4, Scenario 3)

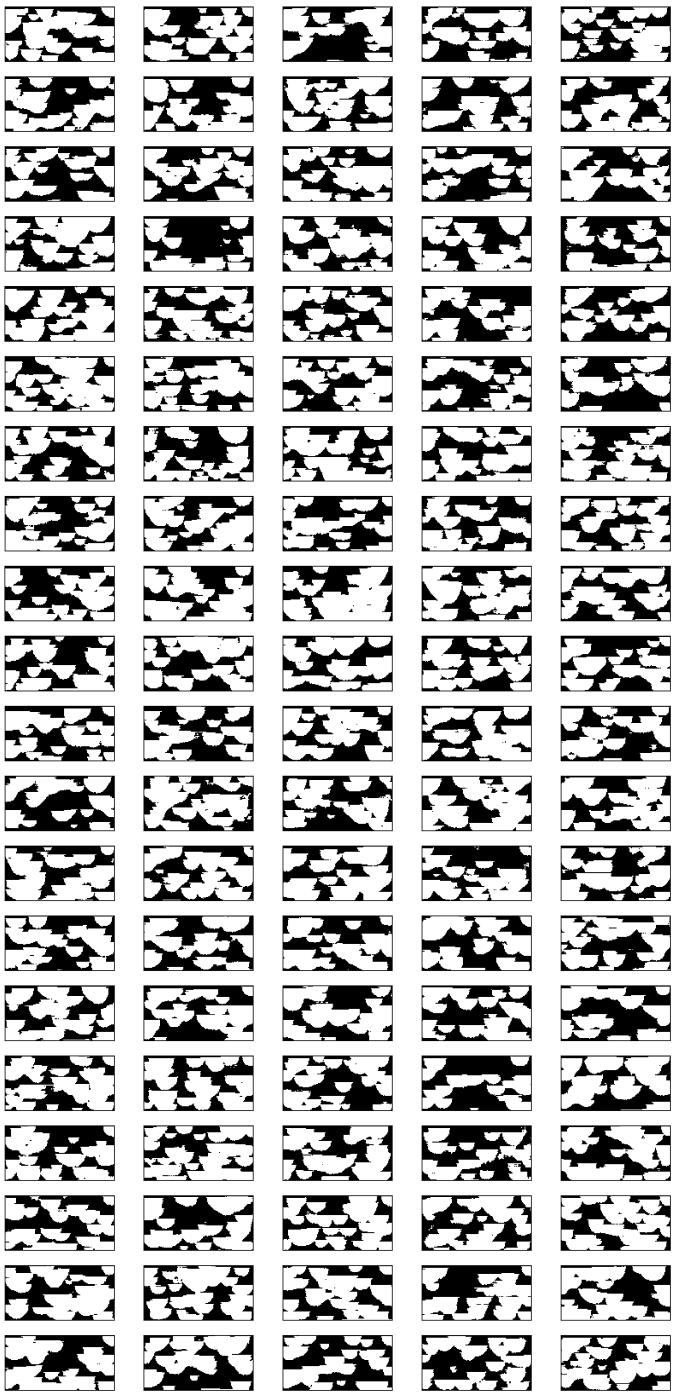

Figure 7.15: One hundred samples obtained by minimizing the sum of the flow, well, and prior losses (Section 7.4, Scenario 4)

Chapter 8

Discussion and Recommendations

Prior to the start of the research presented in this thesis GANs had been applied to two-dimensional datasets of small natural images such as human faces or handwritten-digits (LeCun et al., 1995) as for example in the seminal papers on GANs by Goodfellow et al., (2014) and on DCGAN by Radford et al., (2015). The most prominent work on three-dimensional datasets using GANs was presented by Wu et al., (2016) where a three-dimensional convolutional GAN was separately trained on each class of the IKEA dataset (Lim et al., 2013).

Since then, there have been enormous improvements in the field of generative modeling and GANs specifically, not only in terms of the ability to model higher resolution images (Karras et al., 2017) and diverse distributions (Miyato and Koyama, 2018; Brock et al., 2018), but also in the theoretical understanding of GANs (Mescheder et al., 2017; Mescheder et al., 2018; Arora et al., 2017; Arora and Zhang, 2017).

Considering these advances, I will now present a number of related approaches to deep generative modeling, how we can use these to build realistic prior distributions at pore and reservoir-scale, as well as create a link between the Bayesian framework presented in Chapters 6-7 and a much wider field of statistical learning approaches that consider deep generative priors for inverse problems.

8.1 Recent similar Work

Concurrently to my work on generative models at pore and reservoir-scale GANs have been trained on datasets of two and three-dimensional object-based and geostatistical models. Chan and Elsheikh, (2017) have used a GAN to parameterize geological models based on training sets of two-dimensional patches extracted from a single larger geostatistical training image and evaluate the use of the trained generator for uncertainty propagation for a two-dimensional flow problem. Conditioning to well-data was performed by using a separate neural network to predict a latent-vector which once passed through the generator results in a realization conditioned to the observed data (Chan and Elsheikh, 2018b). Laloy et al., (2017a) trained a VAE (Kingma and Welling, 2013) on two-dimensional geostatistical training images. A specialized GAN architecture by Jetchev et al., (2016) was also extended to 3-D and after training used for hydrological inversion using MCMC (Laloy et al., 2017b; Laloy and Vrugt, 2012).

Concurrently to my work on conditioning pore and reservoir-scale models to available data (Chapter 5 and Mosser et al., (2018a)) Dupont et al., (2018) presented

a gradient-based conditioning approach based on image-inpainting using GANs (Yeh et al., 2016) and showed that generative models can be trained on two-dimensional training images of fluvial systems and conditioned to multiple independent sets of well-data. Zhang et al., (2019) trained GANs on two-dimensional object-based training images for fluvial and non-stationary deltaic systems, three-dimensional object-based models of river-channels with three facies types, and on non-stationary three-dimensional carbonate systems.

One of the key benefits of training a deep generative model compared to other modeling approaches is the amortization of computational cost (Fig. 4.15). Once a large set of training images has been acquired generative models only need to be trained and evaluated once and can be used for fast sampling thereafter.

8.2 Availability of Training Images

The availability of training images and the scale at which training is performed is an equally important factor at the reservoir-scale even more so as obtaining a set of valid training images is a common topic of discussion in training image-based methods such as MPS. While it is impossible to obtain three-dimensional direct observations of the earth's subsurface that could serve as training images, there may still be significant benefit to training deep generative models for reservoir-scale applications.

In many reservoir modeling applications object-based models are used to represent subsurface geo-bodies and their properties. In the case where a large set of object-based models or other synthetic training images for example from process-based models can be created, deep generative models can act as a proxy for the modeling workflow. While such a training set can not enumerate the space of all possible object-based models, generative models aim to represent the underlying probability distribution and hence are able to synthesize new realizations that are not part of the training set. This workflow has been used in Chapters 6 and 7 based on two-dimensional object-based models for river channels.

8.3 Non-stationarity Issues of the Prior

To be able to use deep generative models such as GANs for stochastic modeling, or solving inverse problems as presented in Chapters 6 and 7, a basic premise is that it is possible to train a model such as a GAN on a given dataset. Deep generative models are data-driven and have been shown to scale well in terms of image quality and diversity with the number of available training images (Brock et al., 2018).

At the pore-scale, acquisition of large three-dimensional images is possible with modern micro-CT technology and therefore acquiring a large training image is feasible. Nevertheless, I have made the assumption that there exists a representative volume element that can be defined at a smaller scale than the original CT-image to be able to extract a large number of training patches.

In Chapter 3 I have introduced a method to upsample images by using a fully-convolutional generator network. The generator was trained on patches of 64^3 voxels (Mosser et al., 2017). After training, latent vectors with larger spatial dimensions are used to sample larger spatial domains and the network performs an implicit

upsampling due to the fully convolutional nature of the GAN generator. This implicitly makes the assumption that there exists a representative volume from which we can statistically extrapolate to create larger images. For complex porous media, such as carbonates, making such an assumption may not be possible and applications to complex non-stationary porous media using the approach presented in Chapters 3 and 4 may be difficult.

Deep generative models such as GANs do not require stationarity of the training images e.g. it is possible to train on datasets of highly diverse natural images and objects. Therefore if it is possible to acquire a sufficient amount of training data for complex porous media at the scale of interest, then GANs should be able to represent the distribution of these non-stationary training images without the need for latent variable upscaling as presented in Chapters 3-4.

8.4 Challenges with GANs and how to address them

In my introduction to deep generative models in Chapter 2 I have introduced a number of methods that have been developed to model high-dimensional probability distributions. While GANs specifically have been used throughout this thesis as the choice of deep generative model to represent geological features, GANs suffer from a number of challenges in their application. In Chapters 3 and 4 (Mosser et al., 2017; Mosser et al., 2018c) I highlight that the property distributions of the realizations sampled from the trained GAN show lower diversity than the properties of the test-set images. This behavior is symptomatic of GANs and is referred to as mode-collapse i.e. when the generator has not learned to represent all of the modes of the distribution defined implicitly by the set of training images.

Since the first publication on using GANs for three-dimensional pore-scale stochastic modeling (Mosser et al., 2017) the field of deep generative modeling and understanding of GANs has improved significantly. Arora et al., (2017) have shown that the ability of a GAN generator to represent a distribution from a set of training images is inherently bounded by the size of the discriminator. This motivates the use of very large discriminative networks (Arora and Zhang, 2017; Brock et al., 2018). The use of highly overparameterized discriminators in turn can lead to overfitting which can be alleviated by limiting the number of weight updates of the discriminator per update of the generator parameters and by using specialized gradient penalties that incentivizes generalization in GANs (Thanh-Tung et al., 2019).

Numerous types of convolutional neural networks have been proposed to represent the generator and discriminator networks (Radford et al., 2015) leading to improved representations at higher resolutions (Karras et al., 2017) and for very large datasets (Miyato et al., 2018; Miyato and Koyama, 2018; Brock et al., 2018). While the original GAN algorithm by Goodfellow et al., (2014) has been shown to be locally stable (Nagarajan and Kolter, 2017) a number of alternative training frameworks have been developed around so-called f-divergences (Nowozin et al., 2016) and Wasserstein-GANs (Arjovsky et al., 2017; Gulrajani et al., 2017) which are related to the field of optimal transport (Genevay et al., 2017; Salimans et al., 2018).

In the course of this thesis I have been able to benefit from some of these advances to improve training of generative networks in practice. GANs used for the initial publications on pore-scale stochastic modeling (Chapters 3-4) were trained using the original GAN formulation which requires a number of heuristic modifications

to enable stable training such as adding noise to the generator (Zhao et al., 2016; Salimans et al., 2016) or label flipping (Sønderby et al., 2016).

For the later chapters where GANs were applied to solve ill-posed inverse problems (Chapters 6-7) I used Wasserstein-GANs which at the time were shown to suffer less from mode-collapse and require fewer hyperparameters and heuristics than the original GAN formulation. Many of the more recent advances such as the use of spectral normalization (Miyato et al., 2018), self-attention (Zhang et al., 2018) or the use of the two-timescale update rule presented by Heusel et al., (2017) are promising new methods to further reduce the challenges involved in training GANs. A comparative overview of these learning approaches for GANs was provided by Mescheder et al., (2018).

8.5 Evaluating Deep Generative Models

While choosing a different learning approach for GANs may help to alleviate some of the challenges in their training and representation capabilities, diagnosing mode-collapse and evaluating the distributions represented by GANs is a challenging task. Since GANs do not explicitly define a density that can be evaluated such as in the case of explicit tractable density methods (Section 2.3.3) or provide a lower-bound on the log-likelihood as in the case of VAEs (Section 2.3.3), we cannot use these measures to evaluate their representative capacity.

In Chapter 3 I have proposed the use of two-point correlation functions and numerical evaluations of flow properties to evaluate the quality of the generated stochastic realizations. This was further expanded in Chapter 4 to gray-scale images and the use of Minkowski functionals as a means to characterize random materials. Minkowski functionals have also been used to quantify the quality of GAN generated images in cosmological simulations by Mustafa et al., (2019).

The main reason for me to train on gray-scale images was to avoid introducing errors due to image segmentation into the generative model representation. While learning a representation of the full gray-scale features may be more difficult, this allows a practitioner to reuse the same set of parameters irrespective of the segmentation method used to separate the pore and grain-space.

The motivation for the use of these statistical and image morphological properties was to be comparative with existing stochastic reconstruction methods. A number of approaches have been developed to evaluate the quality and diversity of images created by deep generative models. The Inception score (IS) (Salimans et al., 2016) and Frechet inception distance (FID) (Heusel et al., 2017) are measures derived using correlations obtained from the feature vectors of a pretrained deep neural network. The FID has been developed as a replacement for the IS as it has been shown to be sensitive to the weights used to compute the IS (Barratt and Sharma, 2018). Nevertheless, both the IS and FID measures are commonly found concurrently in GAN literature as a means to quantify the image quality and output diversity of GANs. Ledig et al., (2016) trained GANs for image super-resolution and compared the use of MSE, SSIM, and a mean opinion score based on human evaluations to quantify the quality of the generated high-resolution images. Arora and Zhang, (2017) propose a method based on the birthday paradox to find the support size of the distribution represented by a GAN.

A number of these approaches should be combined to give an estimate of the

image quality and diversity of a trained GAN. It is unclear whether the FID and IS will generalize to pore-scale images and reservoir models. The main reason for this is that the neural networks used to obtain the feature representation for the FID and IS scores are typically trained on the Imagenet dataset (Deng et al., 2009). To apply the FID and IS score it may therefore be necessary to train on a large dataset of porous media at the pore-scale or various types of reservoir facies models specifically to obtain an appropriate set of neural network weights. Where possible the FID and IS scores, should be combined with domain specific quality measures as those proposed in Chapters 3 and 4 (Mosser et al., 2017; Mosser et al., 2018d).

8.6 Interpolating in Latent Space

Radford et al., (2015) perform a linear interpolation $\mathbf{z}(\beta) = \beta\mathbf{z}_1 + (1-\beta)\mathbf{z}_2$ between two samples $\{\mathbf{z}_1,\ \mathbf{z}_2\}$ of the latent space and map the interpolated latent variables to the image domain $\mathbf{m}_\beta = G_\theta(\mathbf{z}(\beta))$. They obtain smooth transitions between the images created by the GAN generator.

In Chapter 4 I present such an interpolation between a set of latent variables and their associated representations of three-dimensional images of the Ketton limestone (Fig. 4.12) and show that smooth continuous transformations are obtained by linear interpolation in the latent-space of the trained generator network. Chapters 5-7 are based on the assumption that it is possible to interpolate or move in latent-space while obtaining physically meaningful representations in the generated image (Chapter 5) or model parameter space (Chapters 6-7).

White, (2016) showed that linear interpolation between two samples obtained from a multivariate Gaussian distribution leads to interpolated latent variables $\mathbf{z}(\beta)$ with smaller variance than the start and end-point samples $\{\mathbf{z}_1, \mathbf{z}_2\}$ obtained from the Gaussian prior. For the case of a Gaussian prior distribution White, (2016) propose to perform spherical linear interpolation which corresponds to a great-circle path between the two points in the high-dimensional space of latent-variables.

In Chapters 6-7 I have presented a Bayesian approach to finding a set of model parameters created by a pre-trained GAN generator. In both Chapters 6 and 7 the prior distribution of latent-variables is honored by incorporating the prior distribution in the Bayesian inference approach. In Chapter 5 I have followed the approach of Yeh et al., (2016) where the discriminator was used as an additional measure for a sample obtained from the generator to be close to the data distribution (Eq. 5.5). Honoring the prior distribution of latent variables or using the discriminator to evaluate a set of obtained images can both be interpreted as a regularization technique to obtain samples from the GAN generator that are close to the prior distribution of images. Nevertheless it remains an open question whether this additional regularization is sufficient to guarantee that samples obtained from the GAN generator and by extension from the Bayesian inference approach presented in Chapters 6-7 are statistically valid.

Makhzani et al., (2015) have shown that in the case of VAEs the space of latent-variables is represented by a continuous multivariate distribution which can have a number of gaps where no valid image representation can be found. This indicates that obtaining valid image representations is not only dependent on the sampling strategy used but also on the deep generative model itself.

Furthermore, if the distribution of latent variables associated with a GAN contains

regions for which interpolation should not be possible this needs to be accounted for explicitly by considering GAN architectures that allow for a discontinuous space of latent variables or by incorporating categorical variables using so-called conditional GANs (Mirza and Osindero, 2014).

While for many applications where only visual plausibility is considered e.g. generating realistic faces of humans (Karras et al., 2019) statistically valid representations obtained from the GAN generator may not be of concern, nevertheless when considering solving inverse-problems with deep generative models such as GANs future work will need to address how we can ensure that the solutions obtained represent physically meaningful solutions from the prior distribution of model parameters. In the case of ill-posed inverse problems the non-physical solutions obtained from a deep generative model may still represent valid solutions of the forward problem that match the data, but which have zero probability under the the prior distribution of model parameters.

8.7 Conditional GANs

In this thesis GANs have been used to create probabilistic generative models that given a random sample from a latent-noise distribution $\mathbf{z}$ create a new stochastic realization $\mathbf{m} = G_\theta(\mathbf{z}),\ \mathbf{z} \sim p(\mathbf{z})$. The family of these latent-variable models can be expanded to conditional-GANs (Mirza and Osindero, 2014) which define a joint distribution $p(\mathbf{m}, c)$ where $\mathbf{m}$ are the data we wish to model and c is an additional observed random variable. The conditional variable is often a categorical random variable that describes class-labels for a set of training images e.g. $c = \{\text{house}, \text{cat}, \text{dog}\}$ or a continuous property. Incorporating such additional information when training a GAN has been shown to help stabilize training and lead to higher image quality and diversity. Additional information of the porous media or reservoir models that we wish to model is often known a priori for a set of training images e.g. lithology as a categorical variable or porosity and permeability as continuous properties at the pore-scale - environment of deposition as a categorical variable and facies proportions as a continous property at the reservoir-scale.

Conditional GANs have found use in particle physics where numerical computation of the scatter of atomic particles on the detector of the Large Hadron Collider (LHC) is very computational demanding (Oliveira et al., 2017). A conditional GAN has been trained on a limited set of numerical simulations with known parameters (Paganini et al., 2017). The particle energy and type of particle then serve as conditional variables to the GAN generator which is able to create random realizations of energy maps of the particle scatter for a specified energy and particle type. The conservation laws of energy can be encoded as a differentiable criterion that has to be satisfied for a given generated image i.e. energy of the incoming particle is equal to the total energy of the particle scatter, in addition to the GAN losses (Oliveira et al., 2017). Where possible such differentiable constraints should be added as additional penalties to ensure that the generative model honors a-priori known relationships within the modeled images.

Many properties at the pore-scale such as the porosity or the two-point correlation function require a binary image representation which means that an image has to be thresholded prior to computation of these quantities. In unpublished work I have attempted to develop differentiable Minkowski functional representations and

two-point correlation functions which could be directly included as a constraint on the GAN learning objective, but ultimately I did not pursue this due to the non-differentiable nature of the segmentation process. Nevertheless, Feng et al., (2019) have shown that including pore-morphology derived constraints in the GAN training objective is possible for pore-scale images. These types of conditioning approaches are not specific to the use of GANs, as they represent a differentiable objective function that is included as a regularization on the training objective. In addition to using a discriminator in the GAN framework, Singh et al., (2017) incorporate a loss that checks for known invariances and was applied to a training set of Cahn-Hillard simulations. Cang et al., (2018) have used a more general style-based approach where features are extracted from the activations of a VGG-type neural network to regularize the visual appearance of the generated micro-structures created by a VAE trained on two-dimensional pore-scale cross-sections.

8.8 Other Deep Generative Models

In Chapters 3-5 I have shown that GANs can be a very general and flexible framework to model data-distributions such as images at the pore-scale, but it is not clear whether GANs themselves are the ideal deep generative model for stochastic pore and reservoir-scale modeling. The challenges in quality control and training stability as I have outlined here may outweigh their ability to perform fast sampling at high-resolution. As we look beyond the family of implicit density models such as GANs there exist a number of recent advances in explicit deep generative models that provide comparable image quality while allowing for stable training and quantitative metrics of model quality and diversity.

Explicit or approximate density models such as VAEs (Kingma and Welling, 2013) or PixelRNN (Oord et al., 2016a) allow us to measure or estimate a lower bound of the log-likelihood (ELBO, Chapter 2) of the data which helps compare different models after training or even families of generative models trained on the same dataset. Furthermore these networks are typically trained using maximum likelihood which is quite stable using stochastic gradient descent compared to the unstable adversarial learning framework of GANs. VAEs (Kingma and Welling, 2013) as compared to GANs do not suffer from mode-collapse as they have a mode-covering behavior due to the form of the KL-divergence in their objective function.

Some of the early criticisms of explicit density models compared to implicit models such as GANs have been the lack of high-resolution details in generated realizations as samples often appear blurred, or long sampling and training times are required as is the case of sequential factorized models such as PixelRNN (Oord et al., 2016a). Both of these criticisms have been refuted recently.

Specialized neural network architectures inspired by PixelRNN have been developed to encode sequential features to represent large-scale correlations in acoustic data for speech synthesis (Wavenet (Oord et al., 2016c)). Fast training and sampling for tractable explicit density models such as PixelRNN (Oord et al., 2016a) has been made possible due to the use of masked convolution operations in PixelCNN (Oord et al., 2016b) and PixelCNN++ (Salimans et al., 2017). State of the art explicit density models combine sequential generation of a latent representation with convolutional decoder networks (VQ-VAE, Oord et al., (2017) and Razavi et al., (2019)) and results are comparable with state-of-the-art GAN methods such as BigGAN

(Brock et al., 2018) in terms of image quality and diversity.

Self-normalizing flows (Rezende and Mohamed, 2015; Dinh et al., 2016) represent a new family of generative models based on change-of-variables (Eq. 2.21) for probability distributions and have been recently extended by Kingma and Dhariwal, (2018) to model high-resolution images using invertible neural networks. Flow-based generative models are based on reversible neural networks which allow exact inference of latent variables for a new data-point and because they are based on an explicit representation of the density function they can be evaluated using the exact log-likelihood. Furthermore, Kingma and Dhariwal, (2018) have shown that the latent variables correlate with semantically meaningful attributes of images, meaning that the network learns to group images with different features in different parts of the latent space. Additional computational benefits come from fast sampling and inference in reversible neural networks as well as reduced computational cost during training.

Currently their main drawback is the need for specialized reversible neural network architectures, which limits the set of possible network architectures, and a lower log-likelihood compared to other methods such as PixelCNNs. Nevertheless, the lower log-likelihood of flow-based generative models should not exclude them from being used for high-fidelity image generation. Theis et al., (2015) have shown that a higher likelihood of deep generative models is not indicative of higher image fidelity. The observations by Nalisnick et al., (2018) and Choi and Jang, (2018) indicate that the Watanabe-Akaike Information Criterion (WAIC) (Watanabe, 2010; Gelman et al., 2013) - an extension to the Akaike Information Criterion (AIC) (Akaike, 1998) - may provide a more robust approach to evaluate deep generative models.

8.9 Engineering Considerations

The task of training generative models for pore and reservoir modeling should also be approached from an engineering point of view with a set of well-defined goals and metrics in mind. To practically train GANs or any other type of deep generative model I have implemented a number of computational libraries that allow automatic computation of the key properties of interest during the training process.

In preparation of Chapter 3 I have monitored the two-point correlation function as a function of the training iterations to allow for an additional quantitative measure of convergence in addition to visual appearance during training. This was expanded to monitoring the Minkowski functionals as a convergence metric in Chapter 4. For different values of the hyperparameters e.g. learning rate or number of convolutional filters in each network layer, a model was trained and the metrics compared to obtain an optimal selection of the hyperparameters. While the metrics used in this thesis are inspired by pore and reservoir-scale features, additional criteria such as the time to sample a new realization may be of importance in practical applications. While in this thesis only GANs were evaluated as a generative model, the choice of model or training formulation itself should be regarded as a hyperparameter and large-scale studies performed to find an optimal model given the metrics of interest. These studies could leverage Bayesian-optimization for hyperparameter tuning of deep neural networks (Snoek et al., 2012; Bergstra et al., 2011) or automated-machine-learning methods (Auto-ML) (Hutter et al., 2019) such as neural architecture search (NAS) (Elsken et al., 2018).

8.10 Combining Geostatistics and Deep Learning

Finally, a promising avenue is to find and investigate relationships between methods from traditional geostatistics or random material theory (Torquato, 2013) and deep generative modeling. Recently Alanov et al., (2019) have proposed a combination of a texture synthesis approach with GANs that has been directly compared to Mosser et al., (2017) and shows a significant improvement in the ability to synthesize textural features. In the future it may be worth investigating the similarities between deep generative models and sequential simulation (Daly, 2005; Parra and Ortiz, 2011) or MPS methods to develop improved geostatistical simulation methods.

A first attempt at developing such approaches has been made recently by Avalos and Ortiz, (2019) where a recurrent convolutional network is used to predict lithologies from a training image and data at wells. It may equally be beneficial for the machine learning community to explore methods developed in geostatistical simulation as there exist many parallel developments in both fields as similarities between sequential simulation techniques (Daly, 2005; Parra and Ortiz, 2011) and tractable density models such as PixelRNN (Oord et al., 2016a) show. Notably Feng et al., (2018) have used convolutional neural networks to accelerate MPS simulation.

Beyond the field of pore-scale reconstruction and geostatistical simulation GANs have gained attention for the design of new micro-structures (Yang et al., 2018). These approaches first train a deep generative model such as a GAN (Yang et al., 2018) or a VAE (Gómez-Bombarelli et al., 2016) on a set of training examples and then optimize the latent variables to maximize a derived material property. Optimizing semantically meaningful latent variables of a deep generative model to optimize a specific physical or chemical quantity may lead to the discovery of efficient batteries and pharmaceuticals (Liu et al., 2018; Gómez-Bombarelli et al., 2016).

8.11 Using Deep Generative Models for Inverse Problems

Optimizing and performing Bayesian inference of the latent variables of a GAN forms the basis of Chapters 5-7 (Mosser et al., 2018a; Mosser et al., 2018d; Mosser et al., 2019) and the latter two chapters consider applications where the function that is being evaluated includes a physics-derived forward-modeling operator. This is a specialization of using deep generative models as an approach to represent the prior distribution in inverse problems that was first developed in the field of compressed sensing (CS) and specifically using GANs in the linear case by Bora et al., (2017).

In CS the general problem definition consists of trying to recover an underlying image $\mathbf{m}^*$ from a set of linear measurements $\mathbf{d}$ which is smaller than the number of parameters n in $\mathbf{m}^*$

$$\mathbf{d} = \mathbf{A}\mathbf{m}^* + \varepsilon \tag{8.1}$$

where $\mathbf{A}$ is the linear measurement operator and ε corresponds to additive noise. This represents an underdetermined linear system and therefore additional assumptions on the behavior of $\mathbf{m}^*$ have to be made. This is analogous to solving an ill-posed inverse problem with a linear forward operator.

The most widely used technique to solve Eq. 8.1 is the *Lasso*-approach, which is equivalent to an L_1 regularization on $\mathbf{m}^*$. Direct reconstruction of a fully observed

image corresponds to a special case of CS where the measurement matrix is a diagonal unit matrix. In a similar way image inpainting which forms the basis of Chapter 5 can be defined in the context of CS where measurements correspond to observed pixels of a masked image (Bora et al., 2017). Bora et al., (2017) have used a GAN to represent the space of images $\mathbf{m} \sim p_{data}(\mathbf{m})$ and find reconstructions by minimizing

$$\mathcal{L}(\mathbf{z}) = \|\mathbf{A}G_\theta(\mathbf{z}) - \mathbf{d}\|_2^2 \tag{8.2}$$

which without any additional constraints on the latent variables $\mathbf{z}$ is equivalent to maximum-likelihood estimation of $\mathbf{m}^*$ in the Gaussian case. More importantly the authors provide a proof of convergence for specific neural network representations of GAN generators and measurement matrices $\mathbf{A}$ which considerably outperforms the Lasso method on the same dataset (Bora et al., 2017). Deep generative priors have recently also been successful at solving the phase retrieval problem in compressed sensing (Hand et al., 2018; Işíl et al., 2019) and progress has been made in finding theoretical convergence guarantees for these types of generative priors (Hand and Voroninski, 2017; Shah and Hegde, 2018; Heckel and Hand, 2018; Ma et al., 2018).

8.12 Bayesian Inference for Deep Latent Variable Models

In Chapters 6-7 I set the inversion problem in a Bayesian context using GANs as a prior distribution of the model parameters. In Bayesian inversion the goal is to find the posterior distribution of the model parameters given the observations $p(\mathbf{m}|\mathbf{d}_{obs})$. When we introduce a deep latent variable model such as a GAN as a prior distribution over the model parameters $\mathbf{m}$ we aim to infer the posterior distribution of the latent variables $p(\mathbf{z}|\mathbf{d}_{obs})$ given the observations.

In Chapter 6 I have used an approximate MCMC-based sampling method, based on previous work on so-called Plug-and-Play Generative Networks (PPGN), to obtain samples from the posterior of the latent variables (Nguyen et al., 2016). Plug-and-play generative networks use an approximate Metropolis-adjusted Langevin algorithm (MALA) (Roberts and Tweedie, 1996) derived from stochastic gradient Langevin dynamics (SGLD) (Welling and Teh, 2011) to obtain posterior samples. Compared to MALA, SGLD has no Metropolis rejection step and assumes that the step-sizes approach zero, which in practice is implemented by assuming small step sizes that are annealed to very small values compared to the initial step size (Welling and Teh, 2011). In contrast to the PPGN approach where the likelihood is represented as a neural network trained to classify natural images, the likelihood in Chapters 6-7 requires the numerical solution of a set of partial differential equations that describe the acoustic wave equation and two-phase flow problem respectively. To be able to apply the MALA formalism, gradients of the log-density need to be computed. In both cases, seismic waveform inversion and reservoir history matching, I have obtained gradients using the adjoint-state method. The approximate MALA algorithm as defined by Nguyen et al., (2016) includes a number of hyperparameters, the learning rate ε_1, the noise variance ε_2, and the weight-decay factor λ. The influence of the choice of these hyper-parameters was only evaluated qualitatively by Nguyen et al., (2016).

In Chapter 6 (Mosser et al., 2018d) I consider the case where inversion of acoustic properties is performed, assuming a Gaussian likelihood for the observed seismic data, and a Bernoulli likelihood for data at the wells. The prior distribution of the latent variables is given by a multi-Gaussian distribution. Combining the likelihood functions and the prior for the latent variables, as well as coupling the learning rate to the noise variance according to Roberts and Tweedie, (1996) reduces the set of hyperparameters for approximate-MALA to the step-size and the total number of iterations. Approximate MALA sampling obtains samples from the posterior in the limiting case where the step size approaches zero when in practise I have annealed the step size to a finite and very small amount (Welling and Teh, 2011). Because the prior represented by the GAN generator is multi-modal and the forward-problem non-linear it is likely that the posterior distribution is also multi-modal. By starting from a number of random locations in latent-space and performing approximate MALA sampling I aim to obtain an ensemble of samples from the multi-modal posterior distribution.

In Chapter 7 (Mosser et al., 2019) I chose to simplify the inference scheme further by finding samples corresponding to local maxima of the posterior distribution. MAP sampling from the posterior starting from the GAN prior amounts to gradient descent with a contribution from the Gaussian prior on the latent variables. By starting from random initial starting locations in latent-space I aim to obtain samples close to the modes of the posterior distribution. An open question is whether an approximate MCMC algorithm like MALA actually converges to samples of the posterior as well as how sufficient exploration of the posterior distribution can be ensured. Zhang et al., (2019) have shown that samples obtained using SGLD represent only a subset of the modes of a multi-modal density and propose an extension using a cyclical step-size to sample from all the modes of a density. Future work should therefore investigate the impact of various posterior sampling strategies.

For both the seismic inversion and reservoir history matching scenarios the computational graph that represents the prior and likelihood is fully differentiable and therefore sophisticated Monte-Carlo methods such as Hamiltonian-Monte-Carlo (Duane et al., 1987; Neal, 2011; Betancourt, 2017) and their stochastic variants (Stochastic Gradient-Hamiltonian-Monte-Carlo (Chen et al., 2014)) should be considered as possible sampling strategies.

It is important to note at this point that using a deep generative model as a prior representation for inverse problems is a very recent development and theoretical and practical applications of these models are at a very early stage. The main driving factors currently come from the field of compressed sensing where a number of proofs have been developed for certain fully connected and convolutional network architectures. The aim of the work that I have performed and presented in Chapters 6 and 7 was not to attempt to solve the open and challenging problem of finding general solutions or theoretical guarantees of convergence as was done by Bora et al., (2017) for the limited framework of a linear inverse problem, fully connected networks and under global optimization, but rather to frame the problem of finding solutions to ill-posed inverse problems using deep generative priors in a Bayesian context and to apply these methods to important inverse problems in geoscience. Nevertheless, I want to emphasize that future work will need to investigate convergence behavior and find new algorithms that give theoretical guarantees to solving inverse problems with deep generative priors (Heckel and Hand, 2018; Ma et al., 2018).

The notion of creating a data-driven prior also represents a paradigm shift in the context of inverse problems and compressed sensing which is driven by efforts to collect large high-quality datasets of high-quality data in the field of machine learning. In areas where such efforts are possible such as in medicine (Wang et al., 2017), data-driven priors for generative models have large potential (Tezcan et al., 2018) and have already been shown to outperform traditional methods such as Lasso (Mardani et al., 2017).

8.13 Challenges in representing the Prior for large-scale Inverse Problems

Bayesian inversion using GAN priors has been applied to synthetic acoustic (Chapter 6) and to reservoir history matching (Chapter 7) cases where the prior was represented by a GAN trained on a set of object-based models of fluvial river systems. For practical applications of the outlined inversion scheme to real datasets we need to consider the representation of the prior, the implementation of the forward model and the sampling strategy.

For the case of seismic inversion, the domain where properties were inferred is a subdomain of the overall computational domain which emulates an inversion being performed only in the reservoir interval. Creating a high-fidelity prior represented by a deep generative model that is able to capture high-resolution details of a realistic basin-scale inversion problem is most likely not feasible and not fit-for-purpose. Richardson, (2018) has shown that large-scale inversion of salt-domes with a prior represented by a GAN is possible which is a problem that lies between the reservoir- and basin-scale.

When considering large-scale full-waveform inversion problems one of the challenges is to find a good initial model that allows convergence of the FWI inversion algorithm. It may therefore be more beneficial to train deep generative models on much simpler relationships of P-wave velocity and density with respect to depth and using such a differentiable generative model to find good starting models for a traditional FWI approach. Where geological knowledge and well data justify reservoir-scale representations more detailed prior models based on deep generative models as is presented in Chapter 6 can be applied.

Where enumerating the space of possible geological scenarios is not possible or collecting a sufficiently large dataset to train a deep generative prior is not possible, other generative priors that do not rely at all on a training set may be more successful. The so-called deep-image-prior (DIP) (Ulyanov et al., 2018) considers a CNN with random weights, which are fixed during the inversion process, and a random latent variable code that serves as the input to the CNN and which is optimized to find a solution to the inverse problem. One of the drawbacks of DIP is that the networks considered are highly over-paramterized compared to the space of images they aim to represent. This means that these networks tend to overfit to the input data and most importantly to the noise present within natural images. Therefore early stopping has to be applied to prevent the networks from overfitting (Van Veen et al., 2018). In contrast under-paramerterized neural networks have been shown to perform equally well compared to DIP and without the need for early stopping because the networks do not have enough degrees of freedom to overfit to high-frequency noise (Heckel and

Hand, 2018). Due to their simplicity they also allow theoretical investigation into the ability for these approaches to be used for inverse problems (Heckel and Hand, 2018).

In the case of reservoir history matching, the space of conceptual geological models may be much smaller and can be reduced to a well-defined prior by an integrated reservoir development team. Therefore creating deep generative priors based on existing object-based or geostatistical approaches may be feasible.

8.14 Representing the Forward Problem

Practical application of the presented approaches relies on efficient methods to compute the forward model. Both applications of deep generative priors for seismic inversion and reservoir history matching use a numerical solution of the system of partial differential equations that governs the forward problem as well as computation of the adjoint to calculate gradients for the considered gradient-based inference methods, MALA and MAP sampling.

Numerical solution of the governing equations on real datasets is computationally demanding and may therefore be prohibitive to applying generative priors for inversion in a geophysical context. The requirement of fast forward and adjoint-solvers motivates the use of proxy-models to replace the numerical solution and where required by the inference method, also the adjoint.

In the case of linear acoustic waveform tomography such approximate models can be learned using neural networks (Holm-Jensen and Hansen, 2018; Holm-Jensen and Hansen, 2019). Specialized CNNs developed in the field of speech synthesis (Oord et al., 2016c) have been adopted as proxy models for acoustic wave tomography (Moseley et al., 2018). Applying direct mappings to pre-stack seismic data (Röth and Tarantola, 1994; Araya-Polo et al., 2018; Yang and Ma, 2019; Li et al., 2019) or post-stack seismic data (Mosser et al., 2018b) to perform acoustic inversion have been proven successful on synthetic datasets.

Fluid flow modeling at pore and reservoir-scale can be accelerated using a similar approach such as learning temporally coherent representations of fluid flow (Tompson et al., 2017; Xie et al., 2018; Kim et al., 2019) and accelerating lattice-Boltzmann simulations (Hennigh, 2017). Directly incorporating operators of partial differential equations into so-called physics-informed neural networks (PINN) (Raissi et al., 2019; Wu et al., 2018) can further help to regularize the solution space of neural networks to be in agreement with the governing physical laws.

In conclusion, it is clear that using data-driven deep generative models as prior distributions shows promise for solving inverse problems. The work I have presented in this thesis represents a contribution to this new and emerging field where I first considered using deep generative models as an unconditional prior distribution (Chapters 3-4) and consequently used deep generative prior distributions in the context of Bayesian inversion (Chapters 6-7).

Nevertheless, future work will need to address some of the theoretical limitations of our understanding of these methods before applying deep generative priors in a practical context. For a limited subset of problems in the field of compressed sensing there have been early successes at showing that this new family of methods can outperform traditional approaches (Bora et al., 2017; Ulyanov et al., 2018; Heckel and Hand, 2018). Due to the sparse nature of the information of the subsurface, the

challenges in computational complexity, as well as the ability to obtain representative data to train deep generative models as prior distributions motivates the future development of specific techniques and algorithms at the intersection of inverse problems, geostatistics, and deep generative modeling.

Chapter 9

Conclusions

Recent advances in deep generative modeling (Goodfellow et al., 2014; Kingma and Welling, 2013; Oord et al., 2016a), have enabled the representation of complex multi-modal distributions of natural images which are parameterized by deep neural networks. In this thesis, I have introduced generative adversarial networks (Goodfellow et al., 2014) as a new approach to create stochastic realizations of geological features at pore and reservoir-scale. GANs consist of two competing neural networks that learn a generative function to sample from high-dimensional probability distributions that are implicitly defined by a set of training images.

At the pore-scale, a number of training images were obtained from micro-CT imaging and segmented into pore and grain domains. These binary training images were then subsequently divided into smaller subvolumes or patches to create a training-set of images to train a GAN generator-discriminator network pair. In Chapter 3 I have used a fully convolutional neural network architecture for the GAN generator that is first trained on patches of images and after training can be used to generate significantly larger samples of porous media. I have evaluated the generated stochastic realizations in terms of spatial statistical and image morphological properties as well as their fluid flow properties and good agreement was found between the statistical and physical properties of a held-out test-set of micro-CT images and GAN generated realizations.

While the GAN generated samples honor the average statistical and physical behavior the sampled realizations show a lower diversity and range of variance than properties measured on test-set images. This most likely stems from the mode-seeking behavior of GANs where the generator only learns to represent a subset of the modes of the modeled multi-modal distributions implicitly defined by the set of training images. Since the publication of the work presented in Chapters 3-4 a number of improved GAN formulations (Arjovsky et al., 2017; Miyato and Koyama, 2018; Zhang et al., 2018) have been proposed to counteract mode-collapse behavior. One of these improved methods, so-called Wasserstein GANs (Gulrajani et al., 2017; Petzka et al., 2017), have been used in the later chapters of this thesis. Future work on GANs for stochastic modeling at pore- and reservoir-scale should evaluate the use of new measures of image quality and diversity (Arora and Zhang, 2017) such as the Inception Score (Salimans et al., 2016) or Frechet Inception Distance (Heusel et al., 2017).

Chapter 5 shows that GANs can be readily conditioned to observed data by using a gradient-based image inpainting approach (Yeh et al., 2016; Dupont et al.,

2018). Optimizing the latent variables $\mathbf{z}$ to create an output from the generator that matches the observed data, and including a perceptual feature loss based on the discriminator's evaluation of the generated images, allows high-fidelity realizations to be obtained that honor the observed data. By using a second neural network that is trained to infer the latent variables, conditional realizations can be obtained directly from a GAN (Chan and Elsheikh, 2018a) and other deep generative models such as PixelCNNs (Oord et al., 2016a) that use a per-pixel sequential generation of new random samples can be used to generate samples without the need for gradient-based optimization of latent variables.

The probabilistic nature of deep generative models such as GANs allows their use as a prior distribution in a Bayesian inversion framework. In Chapters 6-7 I first created synthetic object-based models that resemble a fluvial channel body system embedded in a shale matrix. Each object-based model represents continuous petrophysical property distribution such as the P-wave velocity, permeability or porosity as well as the spatial distribution of the rock-type facies, sand or shale. A GAN parameterized by deep convolutional neural networks was then trained to sample new stochastic realizations of these object-based models.

In Chapter 6 I considered the case of Bayesian inversion of acoustic rock properties of a synthetic fluvial reservoir system. The GAN prior was coupled with a forward model of the acoustic wave equation in a full-waveform inversion scheme. Gradients of the data-mismatch with respect to the grid-block scale property distributions generated by the GAN were obtained using the adjoint-state method. Facies distributions at the wells were incorporated in the inversion approach by assuming conditional independence between the likelihood of the observed seismic data and a Bernoulli likelihood assumed for the facies indicator variable. Bayesian inference of the latent variables that control the output of the GAN prior was performed using an approximate MCMC method (Roberts and Rosenthal, 1998). Inference under a varying number of acoustic sources was performed to determine the influence of increasing the available seismic information on the posterior distribution. Incorporating the Bernoulli likelihood for facies data observed at wells, allowed realizations to be conditioned to well and seismic data.

At the reservoir-scale, building a single conditional GAN to output numerous reservoir architectures (Mirza and Osindero, 2014; Miyato and Koyama, 2018) or training a number of GANs on e.g. training images corresponding to different environments of deposition (Chapter 8) may be feasible. I therefore considered the ill-posed inverse problem of reservoir history matching as a second application of Bayesian inversion with a GAN prior. A two-dimensional flow-problem of water injection into an oil-saturated reservoir with an injector-producer pair was considered. Similar to Chapter 6 and Mosser et al., (2018d) a fully differentiable graph was created to allow MAP sampling from a number of random starting locations in latent space using gradient-descent. The proposed inference method, based on the Bayesian inversion framework outlined in Chapter 6 and Mosser et al., (2018d) converges readily to a local maximum of the posterior that matches observed production rates and pressures as well as facies data at the wells.

Nevertheless, the non-linear mapping between the latent-space and the model domain makes optimization using gradient-based techniques challenging. Rigorous theoretical proofs of convergence have only recently been obtained for a limited set of neural network models and optimization strategies (Bora et al., 2017; Shah and

Hegde, 2018; Heckel and Hand, 2018).

The computational challenge of calculating the forward problem in inverse problems as well as obtaining gradients using the adjoint-state method motivates the development of new proxy models. By first generating a number of numerical simulations and training deep neural networks (Holm-Jensen and Hansen, 2019) as a proxy to map from model parameters to a solution of the forward problem allows a differentiable representation of the forward problem to be obtained.

Physics informed neural networks show the potential to constrain overparameterized neural networks to a solution space that honors the governing physical relationships (Raissi et al., 2019; Wu et al., 2018). In future work, neural network-based parameterizations of the forward-problem could accelerate the proposed Bayesian inversion approach presented in Chapters 6-7.

Furthermore, generative models based on invertible neural networks and self-normalizing flows (Rezende and Mohamed, 2015; Dinh et al., 2016) have been used to solve ill-posed inverse problems in a fully data-driven framework (Ardizzone et al., 2018) and could in future work be used to map between the geophysical parameter space and observed data without an explicit representation of the likelihood.

The recent advances in the theoretical understanding of GANs (Mescheder et al., 2017; Mescheder et al., 2018; Arora and Zhang, 2017), the development of new families of deep generative models (Ardizzone et al., 2018; Dinh et al., 2016; Kingma and Dhariwal, 2018), and the successful applications of GANs in compressed sensing (Hand et al., 2018), represent numerous opportunities to develop new subsurface inversion techniques using deep generative models to represent prior distributions.

Bibliography

Aanonsen, S. I., G. Nævdal, D. S. Oliver, A. C. Reynolds, B. Vallès, and Others (2009). "The ensemble Kalman filter in reservoir engineering–a review". In: *SPE Journal* 14.03, pp. 393–412 (cit. on p. 117).

Adler, J. and O. Öktem (2017). "Solving ill-posed inverse problems using iterative deep neural networks". In: *Inverse Problems* 33.12, p. 124007 (cit. on p. 103).

— (2018). "Learned primal-dual reconstruction". In: *IEEE transactions on medical imaging* 37.6, pp. 1322–1332 (cit. on p. 103).

Akaike, H. (1998). "Information theory and an extension of the maximum likelihood principle". In: *Selected papers of Hirotugu Akaike*. Springer, pp. 199–213 (cit. on p. 133).

Akcelik, V, G Biros, and O Ghattas (2002). "Parallel multiscale Gauss-Newton-Krylov methods for inverse wave propagation". In: *Supercomputing, ACM/IEEE 2002 Conference*. IEEE, p. 41 (cit. on p. 79).

Akcelik, V, J Bielak, G Biros, I Epanomeritakis, A Fernandez, O Ghattas, E. J. Kim, J Lopez, D O'Hallaron, and T Tu (2003). "High resolution forward and inverse earthquake modeling on terascale computers". In: *Supercomputing, 2003 ACM/IEEE Conference*. IEEE, p. 52 (cit. on pp. 79, 93).

Alabert, F. (1987). "The practice of fast conditional simulations through the LU decomposition of the covariance matrix". In: *Mathematical Geology* 19.5, pp. 369–386 (cit. on p. 12).

Alanov, A., M. Kochurov, D. Volkhonskiy, D. Yashkov, E. Burnaev, and D. Vetrov (2019). "User-controllable multi-texture synthesis with generative adversarial networks". In: *ArXiv e-prints*. arXiv: 1904.04751 (cit. on p. 134).

Alhashmi, Z, M. J. Blunt, and B Bijeljic (2016). "The impact of pore structure heterogeneity, transport, and reaction conditions on fluid–fluid reaction rate studied on images of pore space". In: *Transport in Porous Media* 115.2, pp. 215–237 (cit. on p. 64).

Araya-Polo, M, J Jennings, A Adler, and T Dahlke (2018). "Deep-learning tomography". In: *The Leading Edge* 37.1, pp. 58–66 (cit. on pp. 82, 138).

Ardizzone, L., J. Kruse, S. Wirkert, D. Rahner, E. W. Pellegrini, R. S. Klessen, L. Maier-Hein, C. Rother, and U. Köthe (2018). "Analyzing inverse problems with invertible neural networks". In: *ArXiv e-prints*. arXiv: 1808.04730 (cit. on pp. 118, 142).

Arjovsky, M., S. Chintala, and L. Bottou (2017). "Wasserstein GAN". In: *ArXiv e-prints*. arXiv: 1701.07875 (cit. on pp. 49, 53, 74, 86, 106, 128, 140).

Armstrong, M., A. Galli, H. Beucher, G. Loc'h, D. Renard, B. Doligez, R. Eschard, and F. Geffroy (2011). *Plurigaussian simulations in geosciences*. Springer Science & Business Media (cit. on pp. 2, 12, 72).

Arns, C. H., M. A. Knackstedt, and K. R. Mecke (2009). "Boolean reconstructions of complex materials: Integral geometric approach". In: *Physical Review E* 80.5, p. 051303 (cit. on pp. 13, 22).

Arora, S. and Y. Zhang (2017). "Do GANs actually learn the distribution? An empirical study". In: *ArXiv e-prints*. arXiv: 1706.08224 (cit. on pp. 68, 91, 126, 128, 129, 140, 142).

Arora, S., R. Ge, Y. Liang, T. Ma, and Y. Zhang (2017). "Generalization and Equilibrium in Generative Adversarial Nets (GANs)". In: *ArXiv e-prints*. arXiv: 1703.00573 (cit. on pp. 126, 128).

Arvanitidis, G., L. K. Hansen, and S. Hauberg (2017). "Latent space oddity: On the curvature of deep generative models". In: *ArXiv e-prints*. arXiv: 1710.11379 (cit. on p. 117).

Asim, M., F. Shamshad, and A. Ahmed (2018). "Blind image deconvolution using deep generative priors". In: *ArXiv e-prints*. arXiv: 1802.04073 (cit. on p. 82).

Avalos, S. and J. M. Ortiz (2019). "Geological modeling using a recursive convolutional neural networks approach". In: *ArXiv e-prints*. arXiv: 1904.12190 (cit. on p. 134).

Azevedo, L. and A. Soares (2017). *Geostatistical methods for reservoir geophysics*. Springer (cit. on p. 81).

Azevedo, L., D. Grana, and C. Amaro (2019). "Geostatistical rock physics AVA inversion". In: *Geophysical Journal International* 216.3, pp. 1728–1739 (cit. on p. 81).

Aziz, K. and Others (1993). "Reservoir simulation grids: opportunities and problems". In: *Journal of Petroleum Technology* 45.07, pp. 658–663 (cit. on p. 102).

Barratt, S. and R. Sharma (2018). "A note on the inception score". In: *ArXiv e-prints*. arXiv: 1801.01973 (cit. on pp. 91, 129).

Bear, J. (2013). *Dynamics of fluids in porous media*. Courier Corporation (cit. on pp. 47, 102).

Bear, J. and A. H.-D. Cheng (2010). *Modeling groundwater flow and contaminant transport*. Vol. 23. Springer Science & Business Media (cit. on p. 101).

Beaumont, M. A., W. Zhang, and D. J. Balding (2002). "Approximate Bayesian computation in population genetics". In: *Genetics* 162.4, pp. 2025–2035 (cit. on p. 9).

Bengio, Y., E. Laufer, G. Alain, and J. Yosinski (2014). "Deep generative stochastic networks trainable by backprop". In: *International Conference on Machine Learning*, pp. 226–234 (cit. on p. 3).

Berg, C. F., O. Lopez, and H. Berland (2017). "Industrial applications of digital rock technology". In: *Journal of Petroleum Science and Engineering* 157.Supplement C, pp. 131–147 (cit. on p. 73).

Berg, J. and K. Nyström (2017). "Neural network augmented inverse problems for PDEs". In: *ArXiv e-prints*, pp. 1–22. arXiv: 1712.09685 (cit. on pp. 51, 81).

Bergstra, J. S., R. Bardenet, Y. Bengio, and B. Kégl (2011). "Algorithms for hyper-parameter optimization". In: *Advances in Neural Information Processing Systems*, pp. 2546–2554 (cit. on p. 133).

Berryman, J. G. (1987). "Relationship between specific surface area and spatial correlation functions for anisotropic porous media". In: *Journal of Mathematical Physics* 28.1, pp. 244–245 (cit. on p. 60).

Betancourt, M. (2017). "A conceptual introduction to Hamiltonian Monte Carlo". In: *ArXiv e-prints*. arXiv: 1701.02434 (cit. on pp. 10, 136).

Bierkens, M. F. P. and P. A. Burrough (1993). "The indicator approach to categorical soil data". In: *Journal of Soil Science* 44.2, pp. 361–368 (cit. on pp. 2, 12).

Bishop, C. M. (2006). *Pattern recognition and machine learning*. Springer (cit. on pp. 3, 8, 10, 14).

Bissell, R., O Dubrule, P Lamy, P Swaby, O Lepine, et al. (1997). "Combining geostatistical modelling with gradient information for history matching: The pilot point method". In: *SPE Annual Technical Conference and Exhibition*. Society of Petroleum Engineers (cit. on p. 2).

Blei, D. M., A. Kucukelbir, and J. D. McAuliffe (2017). "Variational inference: A review for statisticians". In: *Journal of the American Statistical Association* 112.518, pp. 859–877 (cit. on pp. 10, 11, 16).

Blunt, M. J. (2017). *Multiphase flow in permeable media: A pore-scale perspective*. Cambridge University Press (cit. on pp. 1, 19, 22, 28, 104).

Blunt, M. J., B. Bijeljic, H. Dong, O. Gharbi, S. Iglauer, P. Mostaghimi, A. Paluszny, and C. Pentland (2013). "Pore-scale imaging and modelling". In: *Advances in Water Resources* 51, pp. 197–216 (cit. on pp. 50, 73, 102).

Bora, A., A. Jalal, E. Price, and A. G. Dimakis (2017). "Compressed sensing using generative models". In: *ArXiv e-prints*. arXiv: 1703.03208 (cit. on pp. 82, 103, 134–136, 138, 141).

Bora, A., E. Price, and A. G. Dimakis (2018). "Ambient GAN: Generative models from lossy measurements". In: *International Conference on Learning Representations (ICLR)* (cit. on p. 103).

Bortoli, L.-J., F Alabert, A Haas, and A Journel (1993). "Constraining stochastic images to seismic data". In: *Geostatistics Tróia'92*. Springer, pp. 325–337 (cit. on p. 81).

Bosch, M., C. Carvajal, J. Rodrigues, A. Torres, M. Aldana, and J. Sierra (2009). "Petrophysical seismic inversion conditioned to well-log data: Methods and application to a gas reservoir". In: *Geophysics* 74.2, pp. 1–15 (cit. on p. 81).

Bosch, M., T. Mukerji, and E. F. Gonzalez (2010). "Seismic inversion for reservoir properties combining statistical rock physics and geostatistics: A review". In: *Geophysics* 75.5, 75A165–75A176 (cit. on p. 81).

Bottou, L. (2010). "Large-scale machine learning with stochastic gradient descent". In: *Proceedings of COMPSTAT'2010*. Springer, pp. 177–186 (cit. on pp. 3, 14).

Bretheau, T. and D. Jeulin (1989). "Caractéristiques morphologiques des constituants et comportement à la limite élastique d'un matériau biphasé Fe/Ag". In: *Revue de Physique Appliquée* 24.9, pp. 861–869 (cit. on pp. 13, 22).

Brock, A., J. Donahue, and K. Simonyan (2018). "Large scale GAN training for high fidelity natural image synthesis". In: *ArXiv e-prints*. arXiv: 1809.11096 (cit. on pp. 126–128, 133).

Buland, A and H Omre (2003). "Bayesian linearized AVO inversion". In: *Geophysics* 68.1, pp. 185–198 (cit. on pp. 2, 81).

Burgess, C. P., I. Higgins, A. Pal, L. Matthey, N. Watters, G. Desjardins, and A. Lerchner (2018). "Understanding disentangling in beta-VAE". In: *ArXiv e-prints*. arXiv: 1804.03599 (cit. on p. 118).

Caers, J. (2001). "Geostatistical reservoir modelling using statistical pattern recognition". In: *Journal of Petroleum Science and Engineering* 29.3-4, pp. 177–188 (cit. on pp. 13, 51).

Caers, J. (2007). "Comparing the gradual deformation with the probability perturbation method for solving inverse problems". In: *Mathematical Geology* 39.1, pp. 27–52 (cit. on p. 103).

Caers, J. and T. Zhang (2004). "Multiple-point geostatistics: a quantitative vehicle for integrating geologic analogs into multiple reservoir models". In: *AAPG Memoir* 80, pp. 383–394 (cit. on pp. 13, 22).

Canchumuni, S. W. A., A. A. Emerick, and M. A. C. Pacheco (2018). "History matching channelized facies models using ensemble smoother with a deep learning parameterization". In: *ECMOR XVI-16th European Conference on the Mathematics of Oil Recovery* (cit. on p. 103).

Cang, R., H. Li, H. Yao, Y. Jiao, and Y. Ren (2018). "Improving direct physical properties prediction of heterogeneous materials from imaging data via convolutional neural network and a morphology-aware generative model". In: *Computational Materials Science* 150, pp. 212–221 (cit. on p. 132).

Čapek, P., V. Hejtmánek, L. Brabec, A. Zikánová, and M. Kočiřík (2009). "Stochastic reconstruction of particulate media Using simulated annealing: Improving pore connectivity". In: *Transport in Porous Media* 76.2, pp. 179–198 (cit. on p. 70).

Chadwick, R. A. and D. J. Noy (2010). "History-matching flow simulations and time-lapse seismic data from the Sleipner CO2 plume". In: *Geological Society, London, Petroleum Geology Conference series*. Vol. 7. Geological Society of London, pp. 1171–1182 (cit. on p. 102).

Chan, S. and A. H. Elsheikh (2017). "Parametrization and generation of geological models with generative adversarial networks". In: *ArXiv e-prints*. arXiv: 1708.01810 (cit. on pp. 73, 82, 86, 126).

— (2018a). "Exemplar-based synthesis of geology using kernel discrepancies and generative neural networks". In: *ArXiv e-prints*. arXiv: 1809.07748 (cit. on pp. 103, 141).

— (2018b). "Parametric generation of conditional geological realizations using generative neural networks". In: *ArXiv e-prints*. arXiv: 1807.05207 (cit. on pp. 82, 117, 126).

Chang, J. H. R., C.-L. Li, B. Poczos, B. V. K. V. Kumar, and A. C. Sankaranarayanan (2017). "One network to solve them all — Solving linear inverse problems using deep projection models". In: *ArXiv e-prints*. arXiv: 1703.09912 (cit. on p. 82).

Chavent, G. (1979). "Identification of distributed parameter systems: about the output least square method, its implementation, and identifiability". In: *IFAC Proceedings Volumes* 12.8, pp. 85–97 (cit. on p. 102).

Che, T., Y. Li, A. P. Jacob, Y. Bengio, and W. Li (2016). "Mode regularized generative adversarial networks". In: *ArXiv e-prints*. arXiv: 1612.02136 (cit. on p. 67).

Chen, N., A. Klushyn, R. Kurle, X. Jiang, J. Bayer, and P. van der Smagt (2017). "Metrics for deep generative models". In: *ArXiv e-prints*. arXiv: 1711.01204 (cit. on p. 117).

Chen, T., E. Fox, and C. Guestrin (2014). "Stochastic gradient hamiltonian monte carlo". In: *International Conference on Machine Learning*, pp. 1683–1691 (cit. on p. 136).

Chiles, J.-P. and P. Delfiner (2009). *Geostatistics: modeling spatial uncertainty*. Vol. 497. John Wiley & Sons (cit. on pp. 2, 12).

Chiu, S. N., D. Stoyan, W. S. Kendall, and J. Mecke (2013). *Stochastic geometry and its applications*. John Wiley & Sons (cit. on p. 51).

Choi, H. and E. Jang (2018). "Generative ensembles for robust anomaly detection". In: *ArXiv e-prints*. arXiv: 1810.01392 (cit. on p. 133).

Coats, K. H., J. R. Dempsey, J. H. Henderson, and Others (1970). "A new technique for determining reservoir description from field performance data". In: *Society of Petroleum Engineers Journal* 10.01, pp. 66–74 (cit. on p. 102).

Comunian, A., P Renard, J Straubhaar, and P Bayer (2011). "Three-dimensional high resolution fluvio-glacial aquifer analog - Part 2: Geostatistical modeling". In: *Journal of Hydrology* 405.1-2, pp. 10–23 (cit. on p. 23).

Cooley, R. L. and P. J. Sinclair (1976). "Uniqueness of a model of steady-state groundwater flow". In: *Journal of Hydrology* 31.3-4, pp. 245–269 (cit. on p. 102).

Creswell, A. and A. A. Bharath (2018). "Inverting the generator of a generative adversarial network". In: *IEEE transactions on neural networks and learning systems* (cit. on pp. 83, 106, 117).

Csilléry, K., M. G. B. Blum, O. E. Gaggiotti, and O. François (2010). "Approximate Bayesian computation (ABC) in practice". In: *Trends in Ecology & Evolution* 25.7, pp. 410–418 (cit. on p. 9).

Daly, C. (2005). "Higher order models using entropy, Markov random fields and sequential simulation". In: *Geostatistics Banff 2004*. Springer, pp. 215–224 (cit. on p. 134).

Darcy, H. P. G. (1856). *Les Fontaines publiques de la ville de Dijon. Exposition et application des principes à suivre et des formules à employer dans les questions de distribution d'eau, etc.* V. Dalamont (cit. on p. 102).

Davis, M. W. (1987). "Production of conditional simulations via the LU triangular decomposition of the covariance matrix". In: *Mathematical geology* 19.2, pp. 91–98 (cit. on p. 12).

de Witt Jr, W. (1970). "Age of the Bedford Shale, Berea Sandstone, and Simbury Shale in the Appalachian and Michigan Basins, Pennsylvania, Ohio, and Michigan". In: *U.S. Geological Survey Bulletin* 1294 (cit. on p. 31).

Debye, P., H. R. Anderson, and H. Brumberger (1957). "Scattering by an inhomogeneous solid. II. The correlation function and its application". In: *Journal of Applied Physics* 28.6, pp. 679–683 (cit. on p. 27).

Demyanov, V., L. Foresti, M. A. Christie, M. Kanevski, and Others (2011). "Reservoir modelling with feature selection: kernel learning approach". In: *SPE Reservoir Simulation Symposium*. Society of Petroleum Engineers (cit. on p. 103).

Deng, J., W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei (2009). "Imagenet: A large-scale hierarchical image database". In: (cit. on pp. 3, 117, 130).

Denton, E., S. Chintala, A. Szlam, and R. Fergus (2015). "Deep generative image models using a Laplacian pyramid of aversarial networks". In: *ArXiv e-prints*. arXiv: 1506.05751 (cit. on p. 23).

Deutsch, C. V. and L Wang (1996). "Hierarchical object-based stochastic modeling of fluvial reservoirs". In: *Mathematical Geology* 28.7, pp. 857–880 (cit. on p. 87).

Deutsch, C. V. and A. G. Journel (1992). "Geostatistical software library and user's guide". In: *New York* 119, p. 147 (cit. on p. 12).

Dinh, L., J. Sohl-Dickstein, and S. Bengio (2016). "Density estimation using Real NVP". In: *ArXiv e-prints*. arXiv: 1605.08803 (cit. on pp. 16, 92, 118, 133, 142).

Domenico, P. A., F. W. Schwartz, and Others (1998). *Physical and chemical hydrogeology*. Vol. 506. Wiley New York (cit. on p. 101).

Doyen, P. (2007). *Seismic reservoir characterization: An earth modelling perspective*. Vol. 2. EAGE Publications Houten (cit. on p. 81).

Doyen, P. M. (1988). "Porosity from seismic data: A geostatistical approach". In: *Geophysics* 53.10, pp. 1263–1275 (cit. on p. 12).

Duane, S., A. D. Kennedy, B. J. Pendleton, and D. Roweth (1987). "Hybrid Monte Carlo". In: *Physics Letters B* 195.2, pp. 216–222 (cit. on pp. 10, 136).

Dubrule, O. (2003). *Geostatistics for seismic data integration in earth models*. Society of Exploration Geophysicists and European Association of Geoscientists (cit. on pp. 2, 13, 81).

Dumoulin, V. and F. Visin (2016). "A guide to convolution arithmetic for deep learning". In: *ArXiv e-prints*, pp. 1–28. arXiv: 1603.07285 (cit. on p. 56).

Dupont, E. and S. Suresha (2018). "Probabilistic semantic inpainting with pixel constrained CNNs". In: *ArXiv e-prints*. arXiv: 1810.03728 (cit. on p. 103).

Dupont, E., T. Zhang, P. Tilke, L. Liang, and W. Bailey (2018). "Generating realistic geology conditioned on physical measurements with generative adversarial networks". In: *ArXiv e-prints*. arXiv: 1802.03065 (cit. on pp. 82, 103, 126, 140).

Efros, A. A. and W. T. Freeman (2001). "Image quilting for texture synthesis and transfer". In: *Proceedings of the 28th Annual Conference on Computer graphics and Interactive Techniques*. ACM, pp. 341–346 (cit. on pp. 13, 51).

Elsken, T., J. H. Metzen, and F. Hutter (2018). "Neural architecture search: A survey". In: *ArXiv e-prints*. arXiv: 1808.05377 (cit. on p. 133).

Emerick, A. A. and A. C. Reynolds (2012). "History matching time-lapse seismic data using the ensemble Kalman filter with multiple data assimilations". In: *Computational Geosciences* 16.3, pp. 639–659 (cit. on p. 103).

— (2013). "Ensemble smoother with multiple data assimilation". In: *Computers & Geosciences* 55, pp. 3–15 (cit. on pp. 103, 117).

Emery, D (1988). *The origin of late spar cements in the Lincolnshire limestone, Jurassic of central England: 145(4)*. Vol. 145. 4. Journal of the Geological Society, pp. 621–633 (cit. on p. 32).

Emery, X. (2007). "Simulation of geological domains using the plurigaussian model: new developments and computer programs". In: *Computers & Geosciences* 33.9, pp. 1189–1201 (cit. on p. 12).

England, W. A., A. S. Mackenzie, D. M. Mann, and T. M. Quigley (1987). "The movement and entrapment of petroleum fluids in the subsurface". In: *Journal of the Geological Society* 144.2, pp. 327–347 (cit. on p. 101).

Evensen, G. (1994). "Sequential data assimilation with a nonlinear quasi-geostrophic model using Monte Carlo methods to forecast error statistics". In: *Journal of Geophysical Research: Oceans* 99.C5, pp. 10143–10162 (cit. on p. 103).

— (2003). "The ensemble Kalman filter: Theoretical formulation and practical implementation". In: *Ocean Dynamics* 53.4, pp. 343–367 (cit. on pp. 2, 103).

Farmer, C. L. (2002). "Upscaling: a review". In: *International journal for numerical methods in fluids* 40.1-2, pp. 63–78 (cit. on p. 102).

Feng, J., Q. Teng, X. He, and X. Wu (2018). "Accelerating multi-point statistics reconstruction method for porous media via deep learning". In: *Acta Materialia* 159, pp. 296–308 (cit. on p. 134).

Feng, J., X. He, Q. Teng, C. Ren, H. Chen, and Y. Li (2019). "Accurate and fast reconstruction of porous media from extremely limited information using

conditional generative adversarial network". In: *ArXiv e-prints*. arXiv: 1905.02135 (cit. on p. 132).

Fetter, C. W. (2018). *Applied hydrogeology*. Waveland Press (cit. on p. 101).

Finney, J. L. (1970). "Random packings and the structure of simple liquids. I. The geometry of random close packing". In: *Proceedings of the Royal Society of London A: Mathematical, Physical and Engineering Sciences*. Vol. 319. 1539. The Royal Society, pp. 479–493 (cit. on p. 31).

Flannery, B. P., H. W. Deckman, W. G. Roberge, and K. L. D'Amico (1987). "Three-dimensional X-ray microtomography". In: *Science* 237.4821, pp. 1439–1444 (cit. on pp. 13, 51).

Garven, G. (1985). "The role of regional fluid flow in the genesis of the Pine Point deposit, Western Canada sedimentary basin". In: *Economic Geology* 80.2, pp. 307–324 (cit. on p. 101).

Gauthier, J. (2014). "Conditional generative adversarial nets for convolutional face generation". In: *Class Project for Stanford CS231N: Convolutional Neural Networks for Visual Recognition* (cit. on pp. 47, 116).

Gelfand, A. E. and A. F. M. Smith (1990). "Sampling-based approaches to calculating marginal densities". In: *Journal of the American Statistical Association* 85.410, pp. 398–409 (cit. on p. 10).

Gelman, A., H. S. Stern, J. B. Carlin, D. B. Dunson, A. Vehtari, and D. B. Rubin (2013). *Bayesian data analysis*. Chapman and Hall/CRC (cit. on pp. 9, 10, 133).

Genevay, A., G. Peyré, and M. Cuturi (2017). "GAN and VAE from an optimal transport point of view". In: *ArXiv e-prints*. arXiv: 1706.01807 (cit. on p. 128).

Gerritsen, M. G. and L. J. Durlofsky (2005). "Modeling fluid flow in oil reservoirs". In: *Annual Reviews Fluid Mechchanics* 37, pp. 211–238 (cit. on p. 102).

Gilks, W. R., S. Richardson, and D. Spiegelhalter (1995). *Markov chain Monte Carlo in practice*. Chapman and Hall/CRC (cit. on p. 9).

Gómez-Bombarelli, R., J. N. Wei, D. Duvenaud, J. M. Hernández-Lobato, B. Sánchez-Lengeling, D. Sheberla, J. Aguilera-Iparraguirre, T. D. Hirzel, R. P. Adams, and A. Aspuru-Guzik (2016). "Automatic chemical design using a data-driven continuous representation of molecules". In: *ArXiv e-prints*, pp. 1–26. arXiv: 1610.02415 (cit. on pp. 3, 134).

González, E. F., T. Mukerji, and G. Mavko (2007). "Seismic inversion combining rock physics and multiple-point geostatistics". In: *Geophysics* 73.1, R11–R21 (cit. on p. 81).

Goodfellow, I. (2016). "NIPS 2016 tutorial: Generative adversarial networks". In: *ArXiv e-prints*. arXiv: 1701.00160 (cit. on pp. 16, 18, 23–26, 52, 66, 73).

Goodfellow, I., Y. Bengio, and A. Courville (2016). *Deep learning*. MIT Press (cit. on pp. 3, 10, 15, 103).

Goodfellow, I. J., J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio (2014). "Generative Adversarial Networks". In: *ArXiv e-prints*. arXiv: 1406.2661 (cit. on pp. 3, 18, 23, 24, 52, 66, 73, 85, 103, 106, 126, 128, 140).

Grana, D. and E. Della Rossa (2010). "Probabilistic petrophysical-properties estimation integrating statistical rock physics with seismic inversion". In: *Geophysics* 75.3, O21–O37 (cit. on p. 81).

Grana, D., T. Fjeldstad, and H. Omre (2017). "Bayesian Gaussian mixture linear inversion for geophysical inverse problems". In: *Mathematical Geosciences* 49.4, pp. 493–515 (cit. on p. 81).

Green, P. J. (1995). "Reversible jump Markov chain Monte Carlo computation and Bayesian model determination". In: *Biometrika* 82.4, pp. 711–732 (cit. on p. 115).

Gu, Y., D. S. Oliver, and Others (2005). "History matching of the PUNQ-S3 reservoir model using the ensemble Kalman filter". In: *SPE Journal* 10.02, pp. 217–224 (cit. on p. 117).

Guardiano, F. B. and R. M. Srivastava (1993). "Multivariate geostatistics: Beyond bivariate moments". In: *Geostatistics Tróia '92: Volume 1*. Ed. by A. Soares. Dordrecht: Springer Netherlands, pp. 133–144 (cit. on pp. 2, 13, 51, 81).

Gulrajani, I., F. Ahmed, M. Arjovsky, V. Dumoulin, and A. Courville (2017). "Improved Training of Wasserstein GANs". In: *ArXiv e-prints*. arXiv: 1704.00028 (cit. on pp. 67, 74, 86, 106, 128, 140).

Haas, A and O Dubrule (1994). "Geostatistical inversion-a sequential method of stochastic reservoir modelling constrained by seismic data". In: *First Break* 12.11, pp. 561–569 (cit. on p. 81).

Hand, P. and V. Voroninski (2017). "Global guarantees for enforcing deep generative priors by empirical risk". In: *ArXiv e-prints*. arXiv: 1705.07576 (cit. on p. 135).

Hand, P., O. Leong, and V. Voroninski (2018). "Phase retrieval under a generative prior". In: *Advances in Neural Information Processing Systems*. Ed. by S Bengio, H Wallach, H Larochelle, K Grauman, N Cesa-Bianchi, and R Garnett. Curran Associates, Inc., pp. 9136–9146 (cit. on pp. 135, 142).

Hansen, T. M. and K. S. Cordua (2017). "Efficient Monte-Carlo sampling of inverse problems using a neural network-based forward-applied to GPR crosshole traveltime inversion". In: *Geophysical Journal International* 211.3, pp. 1524–1533 (cit. on pp. 77, 81).

Hansen, T. M., A. G. Journel, A. Tarantola, and K. Mosegaard (2006). "Linear inverse Gaussian theory and geostatistics". In: *Geophysics* 71.6, R101–R111 (cit. on p. 2).

Hastings, W. K. (1970). "Monte Carlo sampling methods using Markov chains and their applications". In: *Biometrika* 57.1, pp. 97–109 (cit. on p. 9).

Heckel, R. and P. Hand (2018). "Deep decoder: Concise image representations from untrained non-convolutional networks". In: *ArXiv e-prints*. arXiv: 1810.03982 (cit. on pp. 135–138, 142).

Hennigh, O. (2017). "Lat-Net: Compressing lattice Boltzmann flow simulations using deep neural networks". In: *ArXiv e-prints*. arXiv: 1705.09036 (cit. on p. 138).

Heusel, M., H. Ramsauer, T. Unterthiner, B. Nessler, and S. Hochreiter (2017). "GANs trained by a two time-scale update rule converge to a local nash equilibrium". In: *Advances in Neural Information Processing Systems*, pp. 6626–6637 (cit. on pp. 91, 117, 129, 140).

Higgins, I., L. Matthey, A. Pal, C. Burgess, X. Glorot, M. Botvinick, S. Mohamed, and A. Lerchner (2017). "beta-VAE: Learning basic visual concepts with a constrained variational framework". In: *International Conference on Learning Representations* (cit. on p. 118).

Hinton, G. E. and R. R. Salakhutdinov (2006). "Reducing the dimensionality of data with neural networks". In: *Science* 313.5786, pp. 504–507 (cit. on p. 103).

Hinton, G. E., S. Osindero, and Y.-W. Teh (2006). "A fast learning algorithm for deep belief nets". In: *Neural computation* 18.7, pp. 1527–1554 (cit. on p. 103).

Hinton, G. E., V. Nair, and G. E. Hinton (2010). "Rectified linear units improve restricted Boltzmann machines". In: *Proceedings of the 27th International Conference on Machine Learning (ICML-10)*. 3, pp. 807–814 (cit. on p. 30).

Hoffman, B. T., X.-H. Wen, S. B. Strebelle, J. K. Caers, and Others (2005). "Geologically consistent history matching of a deepwater turbidite reservoir". In: *SPE Annual Technical Conference and Exhibition*. Society of Petroleum Engineers (cit. on p. 2).

Hoffman, M. D. and A. Gelman (2014). "The No-U-Turn sampler: adaptively setting path lengths in Hamiltonian Monte Carlo." In: *Journal of Machine Learning Research* 15.1, pp. 1593–1623 (cit. on p. 10).

Holden, L., R. Hauge, Ø. Skare, and A. Skorstad (1998). "Modeling of fluvial reservoirs with object models". In: *Mathematical Geology* 30.5, pp. 473–496 (cit. on p. 13).

Holloway, S (2005). "Underground sequestration of carbon dioxide-a viable greenhouse gas mitigation option". In: *Energy* 30.11-12, pp. 2318–2333 (cit. on p. 101).

Holm-Jensen, T. and T. M. Hansen (2018). "Fast full wavefield inversion of cross-hole tomographic data using machine learning methods". In: *80th EAGE Conference and Exhibition 2018* (cit. on p. 138).

Holm-Jensen, T. and T. M. Hansen (2019). "Linear waveform tomography inversion using machine learning algorithms". In: *Mathematical Geosciences* (cit. on pp. 138, 142).

Hornik, K, M Stinchcombe, and H White (1989). "Multilayer feedforward networks are universal approximators". In: *Neural Networks* 2.5, pp. 359–366 (cit. on p. 81).

Huang, T., X. Li, T. Zhang, and D.-T. Lu (2013). "GPU-accelerated direct sampling method for multiple-point statistical simulation". In: *Computers & Geosciences* 57, pp. 13–23 (cit. on pp. 13, 23).

Hutter, F., L. Kotthoff, and J. Vanschoren (2019). *Automated machine learning-Mmthods, systems, challenges*. Springer (cit. on p. 133).

Işíl, Ç., F. S. Oktem, and A. Koç (2019). "Deep iterative reconstruction for phase retrieval". In: *ArXiv e-prints*. arXiv: 1904.11301 (cit. on p. 135).

Ioffe, S. and C. Szegedy (2015). "Batch normalization: Accelerating deep network training by reducing internal covariate shift". In: *ArXiv e-prints*. arXiv: 1502.03167 (cit. on pp. 30, 56).

Jetchev, N., U. Bergmann, and R. Vollgraf (2016). "Texture synthesis with spatial generative adversarial networks". In: *ArXiv e-prints*. arXiv: 1611.08207 (cit. on pp. 54, 74, 103, 126).

Jeulin, D. (2000). "Random texture models for material structures". In: *Statistics and Computing* 10.2, pp. 121–132 (cit. on pp. 13, 22).

Jiao, Y, F. H. Stillinger, and S Torquato (2008). "Modeling heterogeneous materials via two-point correlation functions. II. Algorithmic details and applications". In: *Physical Review E* 77.3, p. 031135 (cit. on pp. 12, 51, 60).

— (2009). "A superior descriptor of random textures and its predictive capacity". In: *Proceedings of the National Academy of Sciences* 106.42, pp. 17634–17639 (cit. on p. 70).

Jordan, M. I., Z. Ghahramani, T. S. Jaakkola, and L. K. Saul (1999). "An introduction to variational methods for graphical models". In: *Machine Learning* 37.2, pp. 183–233 (cit. on p. 10).

Joshi, M. Y. (1974). *A class of stochastic models for porous media*. University of Kansas, Chemical and Petroleum Engineering (cit. on p. 51).

Juanes, R, E. J. Spiteri, F. M. Orr, and M. J. Blunt (2006). "Impact of relative permeability hysteresis on geological CO2 storage". In: *Water Resources Research* 42.12 (cit. on p. 22).

Kabanikhin, S. I. (2008). "Definitions and examples of inverse and ill-posed problems". In: *Journal of Inverse and Ill-Posed Problems* 16.4, pp. 317–357 (cit. on p. 77).

Kadu, A., T. Van Leeuwen, and W. Mulder (2016). "A parametric level-set approach for seismic full-waveform inversion". In: *SEG Technical Program Expanded Abstracts 2016*. Society of Exploration Geophysicists, pp. 1146–1150 (cit. on p. 79).

Kang, Q., P. C. Lichtner, H. S. Viswanathan, and A. I. Abdel-Fattah (2010). "Pore scale modeling of reactive transport involved in geologic CO 2 sequestration". In: *Transport in Porous Media* 82.1, pp. 197–213 (cit. on p. 22).

Karras, T., T. Aila, S. Laine, and J. Lehtinen (2017). "Progressive growing of GANs for improved quality, stability, and variation". In: *ArXiv e-prints*. arXiv: 1710.10196 (cit. on pp. 126, 128).

Karras, T., S. Laine, and T. Aila (2019). "A style-based generator architecture for generative adversarial networks". In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 4401–4410 (cit. on p. 131).

Kiefer, J and J Wolfowitz (1952). "Stochastic estimation of the maximum of a regression function". In: *The Annals of Mathematical Statistics* 23.3, pp. 462–466 (cit. on p. 3).

Kim, B., V. C. Azevedo, N. Thuerey, T. Kim, M. Gross, and B. Solenthaler (2019). "Deep fluids: A generative network for parameterized fluid simulations". In: *Computer Graphics Forum*. Vol. 38. 2. Wiley Online Library, pp. 59–70 (cit. on p. 138).

Kingma, D. P. and J. Ba (2014). "Adam: A Method for Stochastic Optimization". In: *ArXiv e-prints*. arXiv: 1412.6980 (cit. on pp. 30, 108, 118).

Kingma, D. P. and M. Welling (2013). "Auto-Encoding Variational Bayes". In: *ArXiv e-prints*. arXiv: 1312.6114 (cit. on pp. 3, 17, 92, 103, 118, 126, 132, 140).

Kingma, D. P. and P. Dhariwal (2018). "Glow: Generative flow with invertible 1x1 convolutions". In: *Advances in Neural Information Processing Systems*, pp. 10236–10245 (cit. on pp. 17, 92, 118, 133, 142).

Kodali, N., J. Abernethy, J. Hays, and Z. Kira (2017). "On Convergence and Stability of GANs". In: *ArXiv e-prints*. arXiv: 1705.07215 (cit. on p. 67).

Koller, D. and N. Friedman (2009). *Probabilistic graphical models: principles and techniques*. The MIT Press (cit. on p. 10).

Kravaris, C. and J. H. Seinfeld (1985). "Identification of parameters in distributed parameter systems by regularization". In: *SIAM Journal on Control and Optimization* 23.2, pp. 217–241 (cit. on p. 102).

Krizhevsky, A., I. Sutskever, and G. E. Hinton (2012). "Imagenet classification with deep convolutional neural networks". In: *Advances in Neural Information Processing Systems*, pp. 1097–1105 (cit. on pp. 3, 14, 15).

Krogstad, S., K.-A. Lie, O. Møyner, H. M. Nilsen, X. Raynaud, B. Skaflestad, and Others (2015). "MRST-AD–an open-source framework for rapid prototyping and evaluation of reservoir simulation problems". In: *SPE Reservoir Simulation Symposium*. Society of Petroleum Engineers (cit. on pp. 106, 118).

Kukreja, N., M. Louboutin, F. Vieira, F. Luporini, M. Lange, and G. Gorman (2016). "Devito: automated fast finite difference computation". In: *ArXiv e-prints*. arXiv: 1608.08658 (cit. on p. 85).

Lackner, K. S. (2003). "A guide to CO2 sequestration". In: *Science* 300.5626, pp. 1677–1678 (cit. on p. 101).

Laloy, E and J. A. Vrugt (2012). "High-dimensional posterior exploration of hydrologic models using multiple-try DREAM (ZS) and high-performance computing". In: *Water Resources Research* 48.1 (cit. on pp. 82, 103, 126).

Laloy, E., R. Hérault, D. Jacques, and N. Linde (2017). "Efficient training-image based geostatistical simulation and inversion using a spatial generative adversarial neural network". In: *ArXiv e-prints*. arXiv: 1708.04975 [stat.ML] (cit. on pp. 73, 82).

Laloy, E., R. Hérault, J. Lee, D. Jacques, and N. Linde (2017a). "Inversion using a new low-dimensional representation of complex binary geological media based on a deep neural network". In: *Advances in Water Resources* 110.June, pp. 387–405 (cit. on pp. 73, 103, 117, 126).

Laloy, E., R. Hérault, D. Jacques, and N. Linde (2017b). "Training-image based geostatistical inversion using a spatial generative adversarial neural network". In: *ArXiv e-prints*. arXiv: 1708.04975 (cit. on pp. 103, 117, 126).

Lang, C, J Ohser, and R Hilfer (2001). "On the analysis of spatial binary images." In: *Journal of Microscopy* 203, pp. 303–313 (cit. on p. 28).

LeCun, Y, D Touresky, G Hinton, and T Sejnowski (1988). "A theoretical framework for back-propagation". In: *Proceedings of the 1988 Connectionist Models Summer School*. CMU, Pittsburgh, Pa: Morgan Kaufmann, pp. 21–28 (cit. on p. 85).

LeCun, Y, L Bottou, G. B. Orr, and K.-R. Müller (1998). "Efficient backprop". In: *Neural Networks: Tricks of the Trade*. Springer, pp. 9–50 (cit. on pp. 3, 14, 56, 74).

LeCun, Y., Y. Bengio, and Others (1995). "Convolutional networks for images, speech, and time series". In: *The handbook of brain theory and neural networks* 3361.10, p. 1995 (cit. on pp. 15, 126).

LeCun, Y., Y. Bengio, and G. Hinton (2015). "Deep learning". In: *Nature* 521.7553, p. 436 (cit. on p. 103).

Ledig, C., L. Theis, F. Huszar, J. Caballero, A. Cunningham, A. Acosta, A. Aitken, A. Tejani, J. Totz, Z. Wang, and W. Shi (2016). "Photo-realistic single image super-resolution using a generative adversarial network". In: *ArXiv e-prints*. arXiv: 1609.04802 (cit. on p. 129).

Legland, D., I. Arganda-Carreras, and P. Andrey (2016). "MorphoLibJ: integrated library and plugins for mathematical morphology with ImageJ". In: *Bioinformatics* 32.22, btw413 (cit. on p. 29).

Leo, T.-y., C. Kravaria, J. H. Seinfeld, and Others (1986). "History matching by spline approximation and regularization in single-phase areal reservoirs". In: *SPE Reservoir Engineering* 1.05, pp. 521–534 (cit. on p. 102).

Li, H., J. Schwab, S. Antholzer, and M. Haltmeier (2018). "NETT: Solving inverse problems with deep neural networks". In: *ArXiv e-prints*. arXiv: 1803.00092 (cit. on p. 103).

Li, S., B. Liu, Y. Ren, Y. Chen, S. Yang, Y. Wang, and P. Jiang (2019). "Deep learning inversion of seismic data". In: *ArXiv e-prints*. arXiv: 1901.07733 (cit. on p. 138).

Lie, K.-A. (2019). *An introduction to reservoir simulation using MATLAB/GNU Octave: User guide for the MATLAB reservoir simulation toolbox (MRST)*. Cambridge University Press (cit. on p. 106).

Lim, J. J., H. Pirsiavash, and A. Torralba (2013). "Parsing Ikea objects: Fine pose estimation". In: *Proceedings of the IEEE International Conference on Computer Vision*, pp. 2992–2999 (cit. on p. 126).

Lin, C and M. H. Cohen (1982). "Quantitative methods for microgeometric modeling". In: *Journal of Applied Physics* 53.6, pp. 4152–4165 (cit. on pp. 13, 22).

Liu, Q., M. Allamanis, M. Brockschmidt, and A. Gaunt (2018). "Constrained graph variational autoencoders for molecule design". In: *Advances in Neural Information Processing Systems*, pp. 7795–7804 (cit. on p. 134).

Lorentzen, R. J., G. Naevdal, B. Valles, A. Berg, and A.-A. Grimstad (2005). "Analysis of the ensemble Kalman filter for estimation of permeability and porosity in reservoir models". In: *SPE Annual Technical Conference and Exhibition*. Vol. SPE-96375-. Society of Petroleum Engineers (cit. on pp. 2, 103).

Louboutin, M, P Witte, M Lange, N Kukreja, F Luporini, G Gorman, and F. J. Herrmann (2017). "Full-waveform inversion, Part 1: Forward modeling". In: *The Leading Edge* 36.12, pp. 1033–1036 (cit. on pp. 79, 85).

Lunz, S., O. Öktem, and C.-B. Schönlieb (2018). "Adversarial Regularizers in Inverse Problems". In: *ArXiv e-prints* NeurIPS. arXiv: 1805.11572 (cit. on p. 103).

Ma, F., U. Ayaz, and S. Karaman (2018). "Invertibility of convolutional generative networks from partial measurements". In: *Advances in Neural Information Processing Systems*. Ed. by S Bengio, H Wallach, H Larochelle, K Grauman, N Cesa-Bianchi, and R Garnett. Curran Associates, Inc., pp. 9628–9637 (cit. on pp. 135, 136).

Maas, A. L. and A. Y. Ng (2013). "Rectifier nonlinearities improve neural network acoustic models". In: 28 (cit. on p. 30).

Makhzani, A., J. Shlens, N. Jaitly, I. Goodfellow, and B. Frey (2015). "Adversarial Autoencoders". In: *ArXiv e-prints*. arXiv: 1511.05644 [cs.LG] (cit. on p. 130).

Mao, X., Q. Li, H. Xie, R. Y. K. Lau, Z. Wang, and S. P. Smolley (2016). "Least squares generative adversarial networks". In: *ArXiv e-prints*. arXiv: 1611.04076 (cit. on p. 49).

Mardani, M., E. Gong, J. Y. Cheng, S. Vasanawala, G. Zaharchuk, M. Alley, N. Thakur, S. Han, W. Dally, J. M. Pauly, et al. (2017). "Deep generative adversarial networks for compressed sensing automates MRI". In: *ArXiv e-prints*. arXiv: 1706.00051 (cit. on p. 137).

Mardani, M., E. Gong, J. Y. Cheng, S. S. Vasanawala, G. Zaharchuk, L. Xing, and J. M. Pauly (2019). "Deep generative adversarial neural networks for compressive sensing MRI". In: *IEEE transactions on medical imaging* 38.1, pp. 167–179 (cit. on p. 103).

Mariethoz, G and J Caers (2014). *Multiple-point geostatistics: stochastic modeling with training images*. John Wiley & Sons (cit. on pp. 1, 13, 23, 51, 74).

Mariethoz, G., P. Renard, and J. Straubhaar (2010). "The Direct Sampling method to perform multiple-point geostatistical simulations". In: *Water Resources Research* 46.11, p. 11 (cit. on pp. 13, 22, 73).

Marsily, G de (1978). "On the Calibration of Hydrologic Systems". In: *Université Paris 6: Paris, France* (cit. on p. 2).

Martens, S., T. Kempka, A. Liebscher, S. Lüth, F. Möller, A. Myrttinen, B. Norden, C. Schmidt-Hattenberger, M. Zimmer, M. Kühn, and Others (2012). "Europe's longest-operating on-shore CO2 storage site at Ketzin, Germany: a progress report after three years of injection". In: *Environmental Earth Sciences* 67.2, pp. 323–334 (cit. on p. 102).

Matheron, G (1971). *Elements pour une Theorie des Milieux Poreux*. Paris: Masson (cit. on pp. 13, 22, 26, 27).

Matheron, G. (1975). *Random sets and integral geometry* (cit. on pp. 13, 22, 51, 72).

Mecke, K. R. (2000). "Additivity, convexity, and beyond: applications of Minkowski Functionals in statistical physics". In: *Statistical Physics and Spatial Statistics*. Springer, pp. 111–184 (cit. on p. 60).

Meerschman, E., G. Pirot, G. Mariethoz, J. Straubhaar, M. Van Meirvenne, and P. Renard (2013). "A practical guide to performing multiple-point statistical simulations with the Direct Sampling algorithm". In: *Computers & Geosciences* 52, pp. 307–324 (cit. on pp. 13, 22).

Menke, H. P., B Bijeljic, and M. J. Blunt (2017). "Dynamic reservoir-condition microtomography of reactive transport in complex carbonates: Effect of initial pore structure and initial brine pH". In: *Geochimica et Cosmochimica Acta* 204, pp. 267–285 (cit. on pp. 54, 62, 74).

Mescheder, L., S. Nowozin, and A. Geiger (2017). "The numerics of GANs". In: *Advances in Neural Information Processing Systems*, pp. 1825–1835 (cit. on pp. 126, 142).

Mescheder, L., A. Geiger, and S. Nowozin (2018). "Which training methods for GANs do actually converge?" In: *ArXiv e-prints*. arXiv: 1801.04406 (cit. on pp. 126, 129, 142).

Metropolis, N., A. W. Rosenbluth, M. N. Rosenbluth, A. H. Teller, and E. Teller (1953). "Equation of state calculations by fast computing machines". In: *The Journal of Chemical Physics* 21.6, pp. 1087–1092 (cit. on p. 9).

Mirza, M. and S. Osindero (2014). "Conditional generative adversarial nets". In: *ArXiv e-prints*. arXiv: 1411.1784 (cit. on pp. 131, 141).

Miyato, T. and M. Koyama (2018). "cGANs with projection discriminator". In: *ArXiv e-prints*. arXiv: 1802.05637 (cit. on pp. 116, 126, 128, 140, 141).

Miyato, T., T. Kataoka, M. Koyama, and Y. Yoshida (2018). "Spectral normalization for generative adversarial networks". In: *ArXiv e-prints*. arXiv: 1802.05957 (cit. on pp. 128, 129).

Modrak, R. and J. Tromp (2015). "Computational efficiency of full waveform inversion algorithms". In: *SEG Technical Program Expanded Abstracts 2015*. Society of Exploration Geophysicists, pp. 4838–4842 (cit. on p. 79).

Mosegaard, K (1998). "Resolution analysis of general inverse problems through inverse Monte Carlo sampling". In: *Inverse problems* 14.3, p. 405 (cit. on p. 81).

Mosegaard, K. (1995). "Monte Carlo sampling of solutions to inverse problems". In: *Journal of Geophysical Research* 100447.12, pp. 431–12 (cit. on pp. 80, 81).

Moseley, B., A. Markham, and T. Nissen-Meyer (2018). "Fast approximate simulation of seismic waves with deep learning". In: *ArXiv e-prints*. arXiv: 1807.06873 (cit. on p. 138).

Mosser, L., O. Dubrule, and M. J. Blunt (2017). "Reconstruction of three-dimensional porous media using generative adversarial neural networks". In: *Physical Review E* 96.4, p. 043309 (cit. on pp. 4, 52, 60, 73, 82, 117, 127, 128, 130, 134).

Mosser, L., O. Dubrule, and M. J. Blunt (2018a). "Conditioning of generative adversarial networks for pore and reservoir scale models". In: *80th EAGE Conference and Exhibition 2018* (cit. on pp. 4, 67, 82, 103, 126, 134).

Mosser, L., W. Kimman, J. Dramsch, S. Purves, A De la Fuente Briceño, and G. Ganssle (2018b). "Rapid seismic domain transfer: Seismic velocity inversion and modeling using deep generative neural networks". In: *80th EAGE Conference and Exhibition 2018* (cit. on pp. 82, 138).

Mosser, L., O. Dubrule, and M. J. Blunt (2018c). "Stochastic reconstruction of an oolitic limestone by generative adversarial networks". In: *Transport in Porous Media* 125.1, pp. 81–103 (cit. on pp. 4, 73, 74, 128).

— (2018d). "Stochastic seismic waveform inversion using generative adversarial networks as a geological prior". In: *ArXiv e-prints*. arXiv: 1806.03720 (cit. on pp. 4, 82, 103, 106, 107, 130, 134, 136, 141).

— (2019). "DeepFlow: History matching in the space of deep generative models". In: *ArXiv e-prints*. arXiv: 1905.05749 (cit. on pp. 4, 134, 136).

Mostaghimi, P., M. J. Blunt, and B. Bijeljic (2013). "Computations of absolute permeability on micro-CT images". In: *Mathematical Geosciences* 45.1, pp. 103–125 (cit. on pp. 22, 29, 62).

Murphy, K. P. (2012). *Machine learning: a probabilistic perspective*. MIT press (cit. on p. 7).

Muskat, M. (1938). "The flow of homogeneous fluids through porous media". In: *Soil Science* 46.2, p. 169 (cit. on p. 102).

— (1949). *Physical principles of oil production*. McGraw-Hill Book Co., New York (cit. on p. 102).

Mustafa, M., D. Bard, W. Bhimji, Z. Lukić, R. Al-Rfou, and J. M. Kratochvil (2019). "CosmoGAN: Creating high-fidelity weak lensing convergence maps using generative adversarial networks". In: *Computational Astrophysics and Cosmology* 6.1, p. 1 (cit. on p. 129).

Nagarajan, V. and J. Z. Kolter (2017). "Gradient descent GAN optimization is locally stable". In: *Advances in Neural Information Processing Systems*, pp. 5585–5595 (cit. on p. 128).

Nalisnick, E., A. Matsukawa, Y. W. Teh, D. Gorur, and B. Lakshminarayanan (2018). "Do deep generative models know what they don't know?" In: *ArXiv e-prints*. arXiv: 1810.09136 (cit. on p. 133).

Neal, R. M. (2011). "MCMC using Hamiltonian dynamics". In: *Handbook of Markov Chain Monte Carlo* 2.11, p. 2 (cit. on pp. 10, 136).

Nguyen, A., J. Clune, Y. Bengio, A. Dosovitskiy, and J. Yosinski (2016). "Plug & Play generative networks: Conditional iterative generation of images in latent space". In: *ArXiv e-prints*. arXiv: 1612.00005 (cit. on pp. 10, 83, 89, 93, 135).

Nowozin, S., B. Cseke, and R. Tomioka (2016). "F-GAN: Training generative neural samplers using variational divergence minimization". In: *Advances in Neural Information Processing Systems*, pp. 271–279 (cit. on p. 128).

Okabe, H. and M. J. Blunt (2004). "Prediction of permeability for porous media reconstructed using multiple-point statistics". In: *Physical Review E* 70.6, p. 066135 (cit. on pp. 1, 13, 23, 51, 70, 75).

— (2005). "Pore space reconstruction using multiple-point statistics". In: *Journal of Petroleum Science and Engineering* 46.1-2, pp. 121–137 (cit. on pp. 13, 23, 51).

Okabe, H. and M. J. Blunt (2007). "Pore space reconstruction of vuggy carbonates using microtomography and multiple-point statistics". In: *Water Resources Research* 43.12, pp. 3–7 (cit. on pp. 1, 13, 23, 51).

Oliveira, L. de, M. Paganini, and B. Nachman (2017). "Learning particle physics by example: Location-aware generative adversarial networks for physics synthesis". In: *ArXiv e-prints*. arXiv: 1701.05927 (cit. on p. 131).

Oliver, D. S. (1995). "Moving averages for Gaussian simulation in two and three dimensions". In: *Mathematical Geology* 27.8, pp. 939–960 (cit. on p. 12).

Oliver, D. S. and Y. Chen (2011). "Recent progress on reservoir history matching: a review". In: *Computational Geosciences* 15.1, pp. 185–221 (cit. on pp. 2, 102).

Oord, A. van den, N. Kalchbrenner, L. Espeholt, O. Vinyals, A. Graves, and Others (2016a). "Conditional image generation with PixelCNN decoders". In: *Advances in Neural Information Processing Systems*, pp. 4790–4798 (cit. on pp. 16, 92, 103, 132, 134, 140, 141).

Oord, A. van den, N. Kalchbrenner, and K. Kavukcuoglu (2016b). "Pixel recurrent neural networks". In: *ArXiv e-prints*. arXiv: 1601.06759 (cit. on pp. 16, 132).

Oord, A. van den, S. Dieleman, H. Zen, K. Simonyan, O. Vinyals, A. Graves, N. Kalchbrenner, A. Senior, and K. Kavukcuoglu (2016c). "WaveNet: A generative model for raw audio". In: *ArXiv e-prints*. arXiv: 1609.03499 (cit. on pp. 16, 132, 138).

Oord, A. van den, O. Vinyals, and Others (2017). "Neural discrete representation learning". In: *Advances in Neural Information Processing Systems*, pp. 6306–6315 (cit. on p. 132).

Øren, P.-E. and S. Bakke (2002). "Process based reconstruction of sandstones and prediction of transport properties". In: *Transport in Porous Media* 46.2-3, pp. 311–343 (cit. on p. 1).

— (2003). "Reconstruction of Berea sandstone and pore-scale modelling of wettability effects". In: *Journal of Petroleum Science and Engineering* 39.3-4, pp. 177–199 (cit. on pp. 1, 22, 52).

Otsu, N. (1975). "A threshold selection method from gray-level histograms". In: *Automatica* 11.285-296, pp. 23–27 (cit. on pp. 31, 54).

Paganini, M., L. de Oliveira, and B. Nachman (2017). "Accelerating science with generative adversarial networks: An application to 3D particle showers in multilayer calorimeters". In: *ArXiv e-prints*. arXiv: 1705.02355 (cit. on p. 131).

Pant, L. (2016). "Stochastic characterization and reconstruction of porous media". PhD thesis (cit. on pp. 1, 12, 23, 51, 70).

Parra, A. and J. M. Ortiz (2011). "Adapting a texture synthesis algorithm for conditional multiple point geostatistical simulation". In: *Stochastic environmental research and risk assessment* 25.8, pp. 1101–1111 (cit. on p. 134).

Pepper, J. F. (1954). *Geology of the Bedford Shale and Berea Sandstone in the Appalachian Basin: A Study of the Stratigraphy, Sedimentation and Paleogeography of Rocks of Bedford and Berea Age in Ohio and Adjacent States*. Vol. 259. US Government Printing Office (cit. on p. 31).

Petzka, H., A. Fischer, and D. Lukovnicov (2017). "On the regularization of Wasserstein GANs". In: *ArXiv e-prints*. arXiv: 1709.08894 (cit. on pp. 74, 86, 107, 140).

Plessix, R.-E. (2006). "A review of the adjoint-state method for computing the gradient of a functional with geophysical applications". In: *Geophysical Journal*

International 167.2, pp. 495–503. arXiv: 0402594v3 [arXiv:cond-mat] (cit. on pp. 79, 84, 106).

Pyrcz, M. J. and C. V. Deutsch (2014). *Geostatistical reservoir modeling*. Oxford University Press (cit. on pp. 22, 72).

Quiblier, J. A. (1984). "A new three-dimensional modeling technique for studying porous media". In: *Journal of Colloid and Interface Science* 98.1, pp. 84–102 (cit. on p. 51).

Radford, A., L. Metz, and S. Chintala (2015). "Unsupervised representation learning with deep vonvolutional generative adversarial networks". In: *ArXiv e-prints*. arXiv: 1511.06434 (cit. on pp. 29, 30, 54, 58, 67, 74, 126, 128, 130).

Raftery, A. E., M. A. Newton, J. M. Satagopan, and P. N. Krivitsky (2006). "Estimating the integrated likelihood via posterior simulation using the harmonic mean identity". In: (cit. on p. 115).

Raissi, M., P. Perdikaris, and G. E. Karniadakis (2019). "Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations". In: *Journal of Computational Physics* 378, pp. 686–707 (cit. on pp. 138, 142).

Ravanbakhsh, S., F. Lanusse, R. Mandelbaum, J. Schneider, and B. Poczos (2016). "Enabling dark energy science with deep generative models of galaxy images". In: *ArXiv e-prints*. arXiv: 1609.05796 (cit. on p. 47).

Razavi, A., A. van den Oord, and O. Vinyals (2019). "Generating diverse high-fidelity images with VQ-VAE-2". In: *ArXiv e-prints*. arXiv: 1906.00446 (cit. on p. 132).

Reichle, R. H., D. B. McLaughlin, and D. Entekhabi (2002). "Hydrologic data assimilation with the ensemble Kalman filter". In: *Monthly Weather Review* 130.1, pp. 103–114 (cit. on p. 102).

Rezende, D. J. and S. Mohamed (2015). "Variational inference with normalizing flows". In: *ArXiv e-prints*. arXiv: 1505.05770 (cit. on pp. 16, 17, 133, 142).

Richardson, A. (2018). "Generative adversarial networks for model order reduction in seismic full-waveform inversion". In: *ArXiv e-prints*. arXiv: 1806.00828 (cit. on pp. 82, 89, 103, 137).

Rikvold, P. A. and G. Stell (1985). "Porosity and specific surface for interpenetrable-sphere models of two-phase random media". In: *The Journal of Chemical Physics* 82.2, pp. 1014–1020 (cit. on pp. 13, 22).

Robbins, H. and S. Monro (1951). "A stochastic approximation method". In: *The Annals of Mathematical Statistics*, pp. 400–407 (cit. on p. 3).

Roberts, G. O. and J. S. Rosenthal (1998). "Optimal scaling of discrete approximations to Langevin diffusions". In: *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* 60.1, pp. 255–268 (cit. on pp. 10, 83, 141).

Roberts, G. O. and R. L. Tweedie (1996). "Exponential convergence of Langevin distributions and their discrete approximations". In: *Bernoulli* 2.4, pp. 341–363 (cit. on pp. 10, 83, 135, 136).

Roberts, G. O. and O. Stramer (2002). "Langevin diffusions and Metropolis-Hastings algorithms". In: *Methodology and Computing in Applied Probability* 4.4, pp. 337–357 (cit. on p. 103).

Roggero, F, L. Y. Hu, and Others (1998). "Gradual deformation of continuous geostatistical models for history matching". In: *SPE Annual Technical Conference and Exhibition*. Society of Petroleum Engineers (cit. on p. 103).

Röth, G and A Tarantola (1994). "Neural networks and inversion of seismic data". In: *Journal of Geophysical Research: Solid Earth* 99.B4, pp. 6753–6768 (cit. on pp. 81, 138).

Rudkiewicz, J. L., D Guérillot, A Galli, and Others (1990). "An integrated software for stochastic modelling of reservoir lithology and property with an example from the Yorkshire Middle Jurassic". In: *North Sea oil and gas reservoirs-II*. Springer, pp. 399–406 (cit. on p. 12).

Rumelhart, D. E., G. E. Hinton, R. J. Williams, and Others (1988). "Learning representations by back-propagating errors". In: *Cognitive modeling* 5.3, p. 1 (cit. on pp. 3, 14, 74, 106).

Salimans, T., I. Goodfellow, W. Zaremba, V. Cheung, A. Radford, and X. Chen (2016). "Improved techniques for training GANs". In: *Advances in Neural Information Processing Systems*, pp. 2226–2234 (cit. on pp. 26, 30, 91, 117, 129, 140).

Salimans, T., A. Karpathy, X. Chen, and D. P. Kingma (2017). "PixelCNN++: Improving the PixelCNN with discretized logistic mixture likelihood and other modifications". In: *ArXiv e-prints*. arXiv: 1701.05517 (cit. on p. 132).

Salimans, T., H. Zhang, A. Radford, and D. Metaxas (2018). "Improving GANs using optimal transport". In: *ArXiv e-prints*. arXiv: 1803.05573 (cit. on p. 128).

Schawinski, K., C. Zhang, H. Zhang, L. Fowler, and G. K. Santhanam (2017). "Generative adversarial networks recover features in astrophysical images of galaxies beyond the deconvolution limit". In: *ArXiv e-prints*. arXiv: 1702.00403 (cit. on p. 47).

Schmähling, J. (2006). "Statistical characterization of technical surface microstructure". PhD thesis (cit. on p. 62).

Scholz, C., F. Wirner, M. A. Klatt, D. Hirneise, G. E. Schröder-Turk, K. Mecke, and C. Bechinger (2015). "Direct relations between morphology and transport in Boolean models". In: *Physical Review E* 92.4, p. 043023 (cit. on p. 28).

Sen, M. K. and P. L. Stoffa (1996). "Bayesian inference, Gibbs' sampler and uncertainty estimation in geophysical inversion". In: *Geophysical Prospecting* 44.2, pp. 313–350 (cit. on p. 80).

Serra, J (1980). "The Boolean model and random sets". In: *Computer Graphics and Image Processing* 12.2, pp. 99–126 (cit. on pp. 13, 22).

Shah, V. and C. Hegde (2018). "Solving linear inverse problems using GAN priors: An algorithm with provable guarantees". In: *2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, pp. 4609–4613 (cit. on pp. 103, 135, 141).

Shao, H., A. Kumar, and P Thomas Fletcher (2018). "The Riemannian geometry of deep generative models". In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*, pp. 315–323 (cit. on p. 117).

Shi, W., J. Caballero, F. Huszár, J. Totz, A. P. Aitken, R. Bishop, D. Rueckert, and Z. Wang (2016). "Real-time single image and video super-resolution using an efficient sub-pixel convolutional neural network". In: *ArXiv e-prints*. arXiv: 1609.05158 (cit. on pp. 68, 86).

Sibson, R. H., J. M. M. Moore, and A. H. Rankin (1975). "Seismic pumping-a hydrothermal fluid transport mechanism". In: *Journal of the Geological Society* 131.6, pp. 653–659 (cit. on p. 101).

Siddique, N, A. Salehi, and F. Liu (2012). "Stochastic reconstruction and electrical transport studies of porous cathode of Li-ion batteries". In: *Journal of Power Sources* 217.Supplement C, pp. 437–443 (cit. on p. 50).

Singh, K., H. Menke, M. Andrew, Q. Lin, C. Rau, M. J. Blunt, and B. Bijeljic (2017). "Dynamics of snap-off and pore-filling events during two-phase fluid flow in permeable media". In: *Scientific Reports* 7 (cit. on pp. 50, 132).

Smith, W. E., H. H. Barrett, and R. G. Paxman (1983). "Reconstruction of objects from coded images by simulated annealing". In: *Optics Letters* 8.4, p. 199 (cit. on p. 23).

Snoek, J., H. Larochelle, and R. P. Adams (2012). "Practical Bayesian optimization of machine learning algorithms". In: *Advances in Neural Information Processing Systems*, pp. 2951–2959 (cit. on p. 133).

Sønderby, C. K., J. Caballero, L. Theis, W. Shi, and F. Huszár (2016). "Amortised MAP inference for image super-resolution". In: *ArXiv e-prints*. arXiv: 1610.04490 (cit. on pp. 26, 30, 57, 129).

Straubhaar, J., P. Renard, G. Mariethoz, R. Froidevaux, and O. Besson (2011). "An improved parallel multiple-point algorithm using a list approach". In: *Mathematical Geosciences* 43.3, pp. 305–328 (cit. on pp. 13, 23).

Strebelle, S. (2002). "Conditional simulation of complex geological structures Using multiple-point statistics". In: *Mathematical Geology* 34.1, pp. 1–21 (cit. on p. 72).

Suwartadi, E., S. Krogstad, and B. Foss (2012). "Nonlinear output constraints handling for production optimization of oil reservoirs". In: *Computational Geosciences* 16.2, pp. 499–517 (cit. on p. 106).

Svergun, D. (1999). "Restoring low resolution structure of biological macromolecules from solution scattering using simulated annealing". In: *Biophysical Journal* 76.6, pp. 2879–2886 (cit. on p. 23).

Tahmasebi, P. and M. Sahimi (2012). "Reconstruction of three-dimensional porous media using a single thin section". In: *Physical Review E* 85.6, p. 66709 (cit. on pp. 13, 51).

— (2013). "Cross-correlation function for accurate reconstruction of heterogeneous media". In: *Physical Review Letters* 110.7, p. 78002 (cit. on pp. 13, 51).

— (2016). "Enhancing multiple-point geostatistical modeling: 1. Graph theory and pattern adjustment". In: *Water Resources Research* 52.3, pp. 2074–2098 (cit. on p. 1).

Tahmasebi, P., A. Hezarkhani, and M. Sahimi (2012). "Multiple-point geostatistical modeling based on the cross-correlation functions". In: *Computational Geosciences* 16.3, pp. 779–797 (cit. on pp. 13, 51).

Tahmasebi, P., M. Sahimi, and J. Caers (2014). "MS-CCSIM: Accelerating pattern-based geostatistical simulation of categorical variables using a multi-scale search in Fourier space". In: *Computers & Geosciences* 67, pp. 75–88 (cit. on pp. 13, 23).

Tahmasebi, P., M. Sahimi, and J. E. Andrade (2017). "Image-based modeling of granular porous media". In: *Geophysical Research Letters* 44.10, pp. 4738–4746 (cit. on pp. 51, 70).

Tanner, M. A. and W. H. Wong (1987). "The calculation of posterior distributions by data augmentation". In: *Journal of the American statistical Association* 82.398, pp. 528–540 (cit. on p. 10).

Tarantola, A (1987). *Inverse problem theory: Methods for data fitting and parameter estimation* (cit. on p. 7).

Tarantola, A (2005). *Inverse problem theory and methods for model parameter estimation*. Vol. 89. SIAM (cit. on pp. 2, 6, 77, 79, 80, 103).

Tezcan, K. C., C. F. Baumgartner, R. Luechinger, K. P. Pruessmann, and E. Konukoglu (2018). "MR image reconstruction using deep density priors". In: *IEEE transactions on medical imaging* (cit. on p. 137).

Thanh-Tung, H., T. Tran, and S. Venkatesh (2019). "Improving generalization and stability of generative adversarial networks". In: *ArXiv e-prints*. arXiv: 1902.03984 (cit. on p. 128).

Theis, L., A. van den Oord, and M. Bethge (2015). "A note on the evaluation of generative models". In: *ArXiv e-prints*. arXiv: 1511.01844 (cit. on pp. 91, 133).

Tikhonov, A. N. and V. I. A. Arsenin (1977). *Solutions of ill-posed problems*. Scripta series in mathematics. Winston (cit. on pp. 102, 117).

Tompson, J., K. Schlachter, P. Sprechmann, and K. Perlin (2017). "Accelerating Eulerian fluid simulation with convolutional networks". In: *Proceedings of the 34th International Conference on Machine Learning-Volume 70*. JMLR. org, pp. 3424–3433 (cit. on p. 138).

Torquato, S and F Lado (1985). "Characterisation of the microstructure of distributions of rigid rods and discs in a matrix". In: *Journal of Physics A: Mathematical and General* 18.1, p. 141 (cit. on pp. 27, 60).

Torquato, S. (2013). *Random heterogeneous materials: microstructure and macroscopic properties*. Vol. 16. Springer Science & Business Media (cit. on pp. 13, 22, 27, 51, 134).

Ulyanov, D., A. Vedaldi, and V. Lempitsky (2018). "Deep image prior". In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 9446–9454 (cit. on pp. 137, 138).

Van Veen, D., A. Jalal, M. Soltanolkotabi, E. Price, S. Vishwanath, and A. G. Dimakis (2018). "Compressed sensing with deep image prior and learned regularization". In: *ArXiv e-prints*. arXiv: 1806.06438 (cit. on p. 137).

Vo, H. X. and L. J. Durlofsky (2014). "A new differentiable parameterization based on principal component analysis for the low-dimensional representation of complex geological models". In: *Mathematical Geosciences* 46.7, pp. 775–813 (cit. on p. 102).

Vogel, H.-J., U Weller, and S Schlüter (2010). "Quantification of soil structure based on Minkowski functions". In: *Computers & Geosciences* 36.10, pp. 1236–1245 (cit. on p. 62).

Wainwright, M. J., M. I. Jordan, and Others (2008). "Graphical models, exponential families, and variational inference". In: *Foundations and Trends in Machine Learning* 1.1–2, pp. 1–305 (cit. on p. 10).

Wang, X., Y. Peng, L. Lu, Z. Lu, M. Bagheri, and R. M. Summers (2017). "Chestx-ray8: Hospital-scale chest x-ray database and benchmarks on weakly-supervised classification and localization of common thorax diseases". In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 2097–2106 (cit. on pp. 3, 137).

Wang, Z. and A. C. Bovik (2009). "Mean squared error: Love it or leave it? A new look at signal fidelity measures". In: *IEEE signal processing magazine* 26.1, pp. 98–117 (cit. on p. 91).

Wang, Z., A. C. Bovik, H. R. Sheikh, E. P. Simoncelli, and Others (2004). "Image quality assessment: from error visibility to structural similarity". In: *IEEE transactions on image processing* 13.4, pp. 600–612 (cit. on p. 91).

Watanabe, S. (2010). "Asymptotic equivalence of Bayes cross validation and widely applicable information criterion in singular learning theory". In: *Journal of Machine Learning Research* 11.Dec, pp. 3571–3594 (cit. on p. 133).

Weller, H. G., G Tabor, H Jasak, and C Fureby (1998). "A tensorial approach to computational continuum mechanics using object-oriented techniques". In: *Computers in Physics* 12.6, p. 620 (cit. on p. 29).

Welling, M. and Y. W. Teh (2011). "Bayesian learning via stochastic gradient Langevin dynamics". In: *Proceedings of the 28th International Conference on Machine Learning (ICML-11)*, pp. 681–688 (cit. on pp. 135, 136).

White, T. (2016). "Sampling generative networks". In: *ArXiv e-prints*. arXiv: 1609.04468 (cit. on pp. 115, 130).

Wu, J., C. Zhang, T. Xue, W. T. Freeman, and J. B. Tenenbaum (2016). "Learning a probabilistic latent space of object shapes via 3D generative-adversarial modeling". In: *ArXiv e-prints*. arXiv: 1610.07584 (cit. on pp. 23, 126).

Wu, J.-L., H. Xiao, and E. Paterson (2018). "Physics-informed machine learning approach for augmenting turbulence models: A comprehensive framework". In: *Physical Review Fluids* 3.7, p. 74602 (cit. on pp. 138, 142).

Wu, X.-H., Y Efendiev, and T. Y. Hou (2002). "Analysis of upscaling absolute permeability". In: *Discrete and Continuous Dynamical Systems Series B* 2, pp. 185–204 (cit. on p. 102).

Xie, Y., E. Franz, M. Chu, and N. Thuerey (2018). "tempoGAN: A temporally coherent, volumetric GAN for super-resolution fluid flow". In: *ACM Transactions on Graphics (TOG)* 37.4, p. 95 (cit. on p. 138).

Xifara, T., C. Sherlock, S. Livingstone, S. Byrne, and M. Girolami (2013). "Langevin diffusions and the Metropolis-adjusted Langevin algorithm". In: *ArXiv e-prints*. arXiv: 1309.2983 (cit. on p. 83).

Yang, F. and J. Ma (2019). "Deep-learning inversion: a next generation seismic velocity-model building method". In: *arXiv e-prints*, arXiv:1902.06267. arXiv: 1902.06267 [physics.geo-ph] (cit. on p. 138).

Yang, Z., X. Li, L. C. Brinson, A. N. Choudhary, W. Chen, and A. Agrawal (2018). "Microstructural materials design via deep adversarial learning methodology". In: *Journal of Mechanical Design* 140.11, p. 111416 (cit. on p. 134).

Yao, T. (1999). "Automatic covariance modeling and conditional spectral simulation with fast Fourier transform." In: (cit. on p. 12).

Yeh, R. A., C. Chen, T. Y. Lim, A. G. Schwing, M. Hasegawa-Johnson, and M. N. Do (2016). "Semantic image inpainting with deep generative models". In: *ArXiv e-prints*. arXiv: 1607.07539 (cit. on pp. 20, 53, 67, 72–74, 103, 127, 130, 140).

Yeh, W. W.-G. and G. W. Tauxe (1971). "Quasilinearization and the identification of aquifer parameters". In: *Water Resources Research* 7.2, pp. 375–381 (cit. on p. 102).

Yeong, C. L. Y. and S Torquato (1998). "Reconstructing random media". In: *Physical Review E* 57.1, pp. 495–506 (cit. on pp. 1, 12, 13, 22, 51).

Yong, R. N., A.-M. O. Mohamed, B. P. Warkentin, and Others (1992). *Principles of contaminant transport in soils*. Elsevier Science Publishers (cit. on p. 101).

Zachary, C. E. and S Torquato (2011). "Improved reconstructions of random media using dilation and erosion processes". In: *Physical Review E* 84.5, p. 56102 (cit. on p. 70).

Zhang, H., T. Xu, H. Li, S. Zhang, X. Wang, X. Huang, and D. Metaxas (2016). "StackGAN: Text to photo-realistic image synthesis with stacked generative adversarial networks". In: *ArXiv e-prints*. arXiv: 1612.03242 (cit. on p. 48).

Zhang, H., I. Goodfellow, D. Metaxas, and A. Odena (2018). "Self-attention generative adversarial networks". In: *ArXiv e-prints*. arXiv: 1805.08318 (cit. on pp. 129, 140).

Zhang, R., C. Li, J. Zhang, C. Chen, and A. G. Wilson (2019). "Cyclical stochastic gradient MCMC for Bayesian deep learning". In: *ArXiv e-prints*. arXiv: 1902.03932 (cit. on pp. 115, 127, 136).

Zhao, J., M. Mathieu, and Y. LeCun (2016). "Energy-based generative adversarial network". In: *ArXiv e-prints*. arXiv: 1609.03126 (cit. on pp. 67, 129).

Zhu, J.-Y., T. Park, P. Isola, and A. A. Efros (2017). "Unpaired image-to-image translation using cycle-consistent adversarial networks". In: *ArXiv e-prints*. arXiv: 1703.10593 (cit. on p. 82).