Bayesian approach to time-resolved tomography

Glenn R. Myers,[1,*] **Matthew Geleta,**[1,2] **Andrew M. Kingston,**[1] **Benoit Recur**[1] **and Adrian P. Sheppard**[1]

[1]*The Australian National University, Canberra ACT 0200, Australia*
[2]*The University of Melbourne, 1-100 Grattan Street, Parkville VIC 3010, Australia*

[*]*glenn.myers@anu.edu.au*

Abstract: Conventional X-ray micro-computed tomography (μCT) is unable to meet the need for real-time, high-resolution, time-resolved imaging of multi-phase fluid flow. High signal-to-noise-ratio (SNR) data acquisition is too slow and results in motion artefacts in the images, while fast acquisition is too noisy and results in poor image contrast. We present a Bayesian framework for time-resolved tomography that uses priors to drastically reduce the required amount of experiment data. This enables high-quality time-resolved imaging through a data acquisition protocol that is both rapid and high SNR. Here we show that the framework: (i) encompasses our previous, algorithms for imaging two-phase flow as limiting cases; (ii) produces more accurate results from imperfect (i.e. real) data, where it can be compared to our previous work; and (iii) is generalisable to previously intractable systems, such as three-phase flow.

OCIS codes: (100.6950) Tomographic image processing; (110.6955) Tomographic imaging; (170.6920) Time-resolved imaging.

References and links

1. C. Caubit, G. Hamon, A. P. Sheppard, and P. E. Øren, "Evaluation of the reliability of prediction of petrophysical data through imagery and pore network modelling," in "22nd International Symposium of the Society of Core Analysts," (Society of Core Analysts, Abu Dhabi, UAE, 2008). SCA2008-33.
2. J. E. Aravena, M. Berli, T. A. Ghezzehei, and S. W. Tyler, "Effects of root-induced compaction on rhizosphere hydraulic properties–X-ray microtomography imaging and numerical simulations," Environ. Sci. Technol. **45**, 425–431 (2011).
3. H. Derluyn, J. Dewanckele, M. N. Boone, V. Cnudde, D. Derome, and J. Carmeliet, "Crystallization of hydrated and anhydrous salts in porous limestone resolved by synchrotron X-ray microtomography," Nucl. Instrum. Meth. B **324**, 102–112 (2014).
4. G. Myers, A. Kingston, T. Varslot, M. Turner, and A. Sheppard, "Dynamic tomography with a priori information," Appl. Opt. **50**, 3685–3690 (2011).
5. R. L. OHalloran, Z. Wen, J. H. Holmes, and S. B. Fain, "Iterative projection reconstruction of time-resolved images using highly-constrained back-projection (hypr)," Magn. Reson. Med. **59**, 132–139 (2008).
6. C. A. Mistretta, "Sub-nyquist acquisition and constrained reconstruction in time resolved angiography," Med. Phys. **38**, 2975–2985 (2011).
7. Z. Tian, X. Jia, B. Dong, Y. Lou, and S. B. Jiang, "Low-dose 4DCT reconstruction via temporal nonlocal means," Med. Phys. **38**, 1359 (2011).
8. Y. Shi and W. C. Karl, "Level set methods for dynamic tomography," in *IEEE International Symposium on Biomedical Imaging* (2004), pp. 620–623.
9. S. Dubsky, S. B. Hooper, K. K. W. Siu, and A. Fouras, "Synchrotron-based dynamic computed tomography of tissue motion for regional lung function measurement," J. R. Soc. Interface **9**, 2213–2224 (2012).

10. G. Chen, J. Tang, and S. Leng, "Prior image constrained compressed sensing (piccs): a method to accurately reconstruct dynamic ct images from highly undersampled projection data sets," Med. Phys. Lett. **35**, 660–663 (2008).
11. R. T. Armstrong, A. Georgiadis, H. Ott, D. Klemin, and S. Berg, "Critical capillary number: desaturation studied with fast x-ray computed microtomography," Geophys. Res. Lett. **41**, 55–60 (2014).
12. F. Natterer, *The Mathematics of Computerized Tomography* (Society for Industrial and Applied Mathematics, 2001).
13. A. C. Kak and M. Slaney, *Principles of Computerized Tomographic Imaging* (SIAM, 2001).
14. K. Lange and J. A. Fessler, "Globally convergent algorithms for maximum a posteriori transmission tomography," IEEE Trans. Image Process. **4**, 1430–1438 (1995).
15. C. A. Bouman and K. Sauer, "A unified approach to statistical tomography using coordinate descent optimization," IEEE Trans. Image Process. **5**, 480–492 (1996).
16. S. V. Venkatakrishnan, C. A. Bouman, and B. Wohlberg, "Plug-and-Play priors for model based reconstruction," in "IEEE Global Conference on Signal and Information Processing (GlobalSIP)" (2013), pp. 945–948.
17. G. Myers, M. Geleta, A. Kingston, B. Recur, and A. Sheppard, "Improving dynamic tomography, through maximum a posteriori estimation," Proc. SPIE, **9212**, 921211 (2014).
18. G. Myers, A. Kingston, T. Varslot, M. Turner, and A. Sheppard, "Dynamic x-ray micro-tomography for real time imaging of drainage and imbibition processes at the pore scale," in Procedings of the International Symposium of the Society of Core Analysts 2011 **2011**, SCA2011–27 (2011).
19. G. Myers, T. Varslot, A. Kingston, A. Herring, and A. Sheppard, "Ground-truth verification of dynamic x-ray micro-tomography images of fluid displacement," Proc. SPIE, **8506**, 85060P (2012).
20. A. Kingston, A. Sakellariou, T. Varslot, G. R. Myers, and A. Sheppard, "Reliable automatic alignment of tomographic projection data by passive auto-focus," Med. Phys. **38**, 4934 (2011).
21. G. Myers, A. Kingston, T. Varslot, and A. Sheppard, "Extending reference scan drift correction to high-magnification high-cone-angle tomography," Opt. Lett. **36**, 4809–4811 (2011).
22. A. Sheppard, S. Latham, J. Middleton, A. Kingston, G. Myers, T. Varslot, A. Fogden, T. Sawkins, R. Cruikshank, M. Saadatfar, N. Francois, C. Arns, and T. Senden, "Techniques in helical scanning, dynamic imaging and image segmentation for improved quantitative analysis with x-ray micro-ct," Nucl. Instrum. Meth. B **324**, 49–56 (2014).
23. K. J. Batenburg and J. Sijbers, "Adaptive thresholding of tomograms by projection distance minimization," Pattern Recognition **42**, 2297–2305 (2009).
24. J.-B. Thibault, K. D. Sauer, C. A. Bouman, and J. Hsieh, "A three-dimensional statistical approach to improved image quality for multislice helical CT," Med. Phys. **34**, 4526 (2007).
25. W. B. Haines, "Studies in the physical properties of soil. v. the hysteresis effect in capillary properties, and the modes of moisture distribution associated therewith," J. Agr. Sci. **20**, 97–116 (1930).
26. S. Berg, H. Ott, S. A. Klapp, A. Schwing, R. Neiteler, N. Brussee, A. Makurat, L. Leu, F. Enzmann, J. O. Schwarz, M. Kersten, S. Irvine, and M. Stampanoni, "Real-time 3d imaging of haines jumps in porous media flow," Proc. Nat. Acad. Sci. **110**, 3755–3759 (2013).
27. D. A. DiCarlo, J. I. G. Cidoncha, and C. Hickey, "Acoustic measurements of pore-scale displacements," Geophys. Res. Lett. **30**, 1901 (2003).
28. R. G. Larson, H. T. Davis, and L. E. Scriven, "Displacement of residual nonwetting fluid from porous media," Chem. Eng. Sci. **36**, 75–85 (1981).
29. K. K. Mohanty, H. T. Davis, and L. E. Scriven, "Physics of oil entrapment in water-wet rock," SPE Reservoir Eng. **2**, 113–128 (1987).

1. Introduction

Multi-phase fluid flow in complex media is an important problem in fluid mechanics, as well as oil recovery (oil/air displacement and trapping by brine) [1], agriculture (fluid transport in the soil surrounding plant roots) [2], carbon sequestration, and materials science (salt deposition and crystallisation in limestone) [3]. Unfortunately, testing numerical models of multi-phase flow is extremely difficult due to the lack of easily obtainable experimental data. Although X-ray micro-computed tomography (μCT) has increasingly been used to conduct time-resolved (i.e. 3D + time) studies of dynamic, time-evolving samples [4–11], standard μCT techniques for non-destructive 3D study of static samples do not generalise to the dynamic case, and "native" dynamic techniques are restricted to simple, or very specific physical systems.

At a typical μCT facility, a source of penetrating X-rays illuminates a sample of interest. The sample attenuates the X-rays, and the intensity of the transmitted X-rays is recorded by a downstream 2D position-sensitive detector. The sample is rotated through a variety of viewing

positions, and a sequence of radiographs is collected. These radiographs are then reconstructed into a tomogram: a 3D map of the sample's linear X-ray attenuation coefficient. Standard reconstruction methods include filtered backprojection (FBP), simultaneous iterative reconstruction technique (SIRT), and maximum likelihood (MLTR) methods [12–14]. At a lab-based μCT facility, acquiring the radiographs will take several hours. Synchrotron facilities can acquire data much more rapidly, but time at these facilities is correspondingly more valuable. The time required to collect a full set of radiographs places a lower bound on the time-resolution of static CT imaging techniques: if the sample changes during acquisition, then inconsistencies between the radiographs will cause motion blurring in the reconstructed tomogram.

Methods for dynamic μCT reduce the time required to collect a set of radiographs by up to an order of magnitude, using *a priori* information about the sample to replace the need for measured data [4–8, 10]. In the medical context, dynamic processes are often assumed to be cyclical (e.g. breathing, heartbeat, etc.) [10], so radiographs can be collected across multiple identical cycles using sophisticated gating algorithms. Methods based on temporal non-local means filtering assume a degree of self-similarity is present throughout the 4D process [7]. Similarly, 4D Bayesian Model-Based Iterative Reconstruction (MBIR) methods typically assume a degree of local self-similarity [15]. Neither of these assumptions are applicable to the non-cyclic multiphase flow systems we wish to study. In particular, we avoid explicit assumptions concerning local self-similarity, and structure in the reconstructed volume, as they may bias calculation of inter-facial curvature in multi-phase flow. 4D angiographic techniques such as HYPR and its successors decompose the reconstructed volume into spatially distinct static and dynamic regions, and assume that changes will only display limited overlap in the radiographs data [5, 6]. Although we cannot assume limited overlap in the radiographs, the framework presented in this paper uses the same core idea: we conceptually separate the sample into a static equilibrium component, and a dynamic component that is restricted by the physics underlying the time-evolution of the sample.

We first present a Bayesian framework for finding the most probable solution to the dynamic tomography reconstruction problem, subject to the measured data and *a priori* physical constraints. This *maximum a posteriori* estimate is expressed such that each conditional probability distribution captures a separate piece of the underlying physics. Our goal is to develop a modular framework, into which we can explicitly build different *a priori* physical constraints on a case-by-case basis. This "plug-and-play" nature of Bayesian tomographic reconstruction has been demonstrated in the static case [16]. We use this Bayesian framework to derive a dynamic tomography reconstruction algorithm, and: (i) demonstrate that it encompasses our previous algorithms for imaging two-phase flow (see Myers et. al. [4, 17]) as limiting cases, placing that work into a broader mathematical context; (ii) quantitatively and qualitatively compare its performance to our previous algorithms for imaging two-phase flow; and (iii) show that unlike our previous efforts, this algorithm has potential for imaging three-phase flow.

2. Framework for dynamic tomography

In this section we present a general framework for the time-resolved μCT imaging of complex, non-periodic, constrained processes, from highly under-sampled data. This framework forms the basis of the reconstruction algorithms derived and tested in later sections.

We characterise a time-evolving sample using its linear X-ray attenuation coefficient, which is the sum of a static (equilibrium) component $\mu_s(\mathbf{x})$ and a dynamic (time-evolving) component $\mu(\mathbf{x},t)$, parameterised by 3D Cartesian coordinates $\mathbf{x}$ and time t [4, 17]. Solving the time-resolved μCT reconstruction problem thus involves estimating $\mu(\mathbf{x},t)$ at a set of discrete times $t \in T_\mu$, producing data that will resemble a 3D movie. We build our framework around finding the *maximum a posteriori* (MAP) estimate $\hat{\mu}(\mathbf{x}, t \in T_\mu)$, defined as the most probable solution

to the reconstruction problem, subject to measured data and any *a priori* constraints. The remainder of this section will discuss how the MAP estimate may be formulated in terms of the measured data and several latent variables. The usefulness of the latent variables flows from practical constraints on acquiring and processing time-resolved μCT data.

2.1. Measured data

Following the experimental procedure laid out by Myers et. al. [4, 17–19], the sample is placed in a standard μCT imaging system. To isolate the static/equilibrium component $\mu_s(\mathbf{x})$, a conventional 3D μCT image is collected either: before the process of interest begins (e.g. before switching on a syringe pump), or after the process is complete and the sample has reached equilibrium. To study the dynamic evolution of the sample we collect 2D radiographs at a sequence of discrete times $t \in T_i$ and viewing positions θ. These radiographs are pre-processed to remove the signal due to the equilibrium component of the sample, as well as mis-alignments, source drift, clear and dark field signals, etc. [20–22].

Our measured data thus consists of: (i) the static linear attenuation coefficient $\mu_s(\mathbf{x})$, estimated from a standard CT scan of the object in its equilibrium state; and (ii) the time-sequence of intensity images $I(\mathbf{r}, \theta, t \in T_i)$ containing only signal due to the dynamic component of the sample, parameterised by the 2D Cartesian detector coordinates $\mathbf{r}$. The physics underlying the time-evolution of the sample may constrain $\mu(\mathbf{x}, t)$ directly (e.g. imaging an incompressible fluid), or via its interactions with the static component (e.g. fluid flow through an impermeable scaffold) [19]. Thus, the *maximum a posteriori* estimate may be written as [17]:

$$\hat{\mu}(\mathbf{x}, t \in T_\mu) = \underset{\mu(\mathbf{x}, t \in T_\mu)}{\arg\max} P[\mu(\mathbf{x}, t \in T_\mu) | I(\mathbf{r}, \theta, t \in T_i), \mu_s(\mathbf{x})]. \tag{1}$$

In later equations, functional dependencies are sometimes omitted for the sake of clarity.

2.2. Latent variables

As part of taking an expectation-maximisation (EM) approach to finding the MAP estimate [Eq. (1)], we introduce two latent variables. We select these variables so that we can be more explicit about our *a priori* assumptions, and because they correspond to quantities of interest.

Eq. (1) requires us to relate a radiograph measured at time $t_0 \in T_i$, to the reconstructed volume at some different time $t_1 \in T_\mu$. The first latent variable quantifies our confidence that we can predict a given radiograph from a 3D frame in our reconstruction $\mu(\mathbf{x}, t_1)$. The rate at which radiographs can be collected (and thus T_i) is largely dictated by the capabilities of our camera (frame rate, etc.), X-ray source (available flux), rotation stage (rotation speed), and sample (maximum allowable flux/dose). In contrast, computational constraints (e.g. RAM, processing time) limit the number of frames in our 4D reconstruction (and thus T_μ). This suggests a disparity between radiograph collection times (T_i) and reconstruction frame times (T_μ), so we require a confidence measure $\tau(t_0 \in T_i, t_1 \in T_\mu)$, such that $P[\tau(t_0 \in T_i, t_1 \in T_\mu) | I, \mu]$ quantifies our confidence that a radiograph at time t_0 can be estimated from a volume at time t_1 [17]. This measure is independent of μ_s.

Depending on how it is constructed, τ may encode implicit assumptions regarding temporal self-similarity of the volume. This can be illustrated by considering two extreme cases. For example, if we construct $\tau(t_0 \in T_i, t_1 \in T_\mu)$ to be independent of t_1, then we have assumed the time-evolution of the sample is not encoded in the measured data. This will lead to a high degree of self-similarity in the reconstructed volume. On the opposite extreme, there are many ways to construct τ such that each projection relates to only one 3D frame in the reconstruction (i.e. for a given radiograph time in T_i, $P(\tau)$ is non-zero for at most one element of T_μ). In this case, τ leads to no implicit assumptions regarding temporal self-similarity, as each volume is

reconstructed from an entirely distinct set of radiographs. In constructing τ we are thus able to codify any *a priori* information about how we believe the sample changes between adjacent reconstruction frames (e.g. Haines-jump dominated flow vs. smooth flow) [17]. We see below [Eq. (2)] that if $P(\tau)$ is chosen to be independent of μ, then these probabilities play a role analogous to temporal interpolation weights in our framework.

The second latent variable encodes information of interest concerning sample composition. Because X-ray attenuation coefficient is rarely a quantity of interest in-and-of itself, the first step in analysis of attenuation-contrast CT images is often segmentation: the automated or manual classification of each voxel in the image, according to the materials it contains. Recent work by (for example) Batenburg et al. indicates that segmentation is best done as part of the reconstruction process [23]. This motivates the inclusion of a classification map $m(\mathbf{x}, t \in T_\mu) \in D$, where materials are chosen from some dictionary D. Proper construction of the dictionary D allows us to incorporate *a priori* information regarding the composition of the sample (e.g. imaging water-air flow in a quartz scaffold). Our choice of $P(m|\mu_s, \mu)$ corresponds to a choice of segmentation method.

2.3. MAP-EM

We incorporate the latent variables m and τ (discussed above) into Eq. (1) using an expectation-maximisation algorithm

$$
\begin{aligned}
f^{(n)}(\mu|\mu^{(n)}, I, \mu_s) &= \sum_\tau \sum_m P(\tau, m|\mu^{(n)}, I, \mu_s) \ln[P(\tau, m, I, \mu_s|\mu)], \\
\mu^{(n+1)} &= \arg\max_\mu \left[f^{(n)}(\mu|\mu^{(n)}, I, \mu_s) + \ln[P(\mu)] \right],
\end{aligned}
$$

and employ Bayes theorem to find the following iterative framework [17]:

$$
\begin{aligned}
\mu^{(n+1)} = \arg\max_\mu \quad & \sum_\tau P(\tau|I, \mu^{(n)}) \ln[P(I|\tau, \mu)] \\
+ & \sum_\tau P(\tau|I, \mu^{(n)}) \ln[P(\tau|\mu)] \\
+ & \sum_m P(m|\mu_s, \mu^{(n)}) \ln[P(m|\mu_s, \mu)] \\
+ & \ln[P(\mu_s|\mu)] + \ln[P(\mu)],
\end{aligned}
\tag{2}
$$

which will converge as $n \to \infty$ to the maximum a posteriori estimate $\hat{\mu}(\mathbf{x}, t \in T_\mu)$. Eq. (2) serves to define $\mu^{(n)}$ in terms of μ; if one takes an iterative approach to solving Eq. (2) then μ will be varied in the inner loop, and converge to $\mu^{(n)}$ which will be varied in the outer EM loop [see, e.g. the psudocode in Eq. (5)]. This formulation has been chosen because each term encodes a specific element of the physics underlying the sample and imaging system.

The first term on the right hand side of Eq. (2) is a data-consistency term, ensuring that the solution is consistent with: (i) the time-sequence of measured radiographs; (ii) our *a priori* information $P(\tau|I, \mu^{(n)})$ about how the sample relates to radiographs taken between adjacent reconstruction frames (see above); and (iii) a physical model $P(I|\tau, \mu)$ of the CT imaging system that includes noise, and potentially more complex behaviours such as partial coherence and refraction. As discussed above, the second term vanishes if $P(\tau)$ is chosen to be independent of μ. The third term in Eq. (2) incorporates a segmentation model, and the fourth incorporates any *a priori* compositional information that allows the static component to constrain the dynamic component of the sample. Finally, the fifth term represents *a priori* information un-related to the latent variables (e.g. a smoothness or sparsity constraint), which can be used to regularise the reconstruction problem.

3. Reconstruction algorithm for multi-phase flow

The MAP-EM framework allows us to build more physically accurate models of the sample and imaging system, and avoid physically-unrealistic "absolute" constraints such as the strict binary thresholding used in Myers et. al. [4]. In order to establish the value of the MAP-EM framework we examine the case of micro-scale multi-phase flow in porous materials, where results can be compared against both our previous algorithms, and ground-truth reconstructions. We consider a typical fine-focus, lab-based X-ray micro-CT setup. Assuming that the effects of partial coherence and refraction are negligible, and that noise profile is dominated by quantum noise at the detector, the system is characterised by the conditional probability distribution [24]:

$$P[I(\mathbf{r},\theta,t_0 \in T_i)|\tau(t_0,t_1 \in T_\mu),\mu(\mathbf{x},t \in T_\mu)] \quad \propto \quad \frac{\{I_s(\mathbf{r},\theta)\exp[-(\mathscr{P}_\theta\mu)(\mathbf{r},t_1)]\}^{I_s(\mathbf{r},\theta)I(\mathbf{r},\theta,t_0)}}{[I_s(\mathbf{r},\theta)I(\mathbf{r},\theta,t_0)]!}$$
$$\times \quad \exp\{-I_s(\mathbf{r},\theta)\exp[-(\mathscr{P}_\theta\mu)(\mathbf{r},t_1)]\},$$

where $\mathscr{P}_\theta$ is the X-ray projection operator at viewing position θ, and I_s is the signal due to the static component of the sample

$$I_s(\mathbf{r},\theta) \propto \exp[-(\mathscr{P}_\theta\mu_s)(\mathbf{r})].$$

This noise model (and associated probability distribution) will give rise to an MLTR-based algorithm.

Fluid flow in micro-porous materials is known to mostly consist of Haines jumps [25]. These are rapid, taking only microseconds, and typically involve a cascade through multiple connected pores [26–29]. As a Haines jump can significantly alter the fluid distribution in the sample, we choose a confidence measure that decays rapidly when a large Haines jump occurs [17]:

$$P[\tau(t_0 \in T_i, t_1 \in T_\mu)|I,\mu^{(n)}] \quad = \quad \Gamma(t_0,t_1),$$
$$1 - \Gamma(t_0,t_1) \quad \propto \quad |t_0 - t_1|, \text{ if } t_0 \text{ and } t_1 \text{ are not separated by a large Haines jump,}$$
$$\Gamma(t_0,t_1) \quad = \quad 0, \text{ otherwise.}$$

In this case, we have defined a "large" Haines jump as one that causes $\partial_t||I(\mathbf{r},\theta,t \in T_i)||_{L_1}$ to deviate more than two standard deviations outside the mean [17].

Our materials dictionary consists of several mobile incompressible fluids, void, and a static impermeable rock scaffold. We describe μ and μ_s using a Gaussian mixture model. Using σ_s and σ to denote the standard deviation of Gaussian distributions in μ_s and μ respectively, we find:

$$P[m|\mu_s,\mu^{(n)}] \quad \propto \quad \exp\left[-\frac{(\mu_s - \mu_{\text{rock}})^2}{2\sigma_s^2} - \frac{\mu^{(n)2}}{2\sigma^2}\right], \text{ if } m = \text{rock,}$$
$$\propto \quad \exp\left[-\frac{\mu_s^2}{2\sigma_s^2} - \frac{(\mu^{(n)} - \mu_m)^2}{2\sigma^2}\right], \text{ otherwise.} \quad (3)$$

This conditional probability is the materials classification (i.e. segmentation) model required by the third term in Eq. 2. Although we assume the attenuation of the fluids (μ_{m}) and scaffold (μ_{rock}) can be calculated *a priori*, an additional expectation-maximisation loop can be used to estimate quantities.

The fourth term in Eq. 2 allows the static component of the sample to restrict the dynamic component. As the rock scaffold is the only non-mobile, non-permeable phase in the sample,

the static component of the sample only restricts the dynamic component where the static attenuation value is consistent with rock:

$$\ln P(\mu_s | \mu) \propto -\frac{\mu^2}{2\sigma^2} \exp\left[-\frac{(\mu_s - \mu_{\text{rock}})^2}{2\sigma_s{}^2}\right].$$

Finally, the sparsity constraint may be enforced by setting the regularisation prior equal to the Laplace density function with parameter a:

$$P[\mu(\mathbf{x}, t \in T_\mu)] \propto \exp\left[-\frac{|\mu(\mathbf{x}, t \in T_\mu)|}{a}\right]. \tag{4}$$

Employing a nested gradient ascent loop (iteration superscript l) to calculate $\mu^{(n+1)}$, we arrive at the following reconstruction algorithm:

1. From the dynamic radiographs $I(\mathbf{r}, \theta, t_0 \in T_i)$, locate large Haines jumps to determine the interpolation coefficients $\Gamma(t_0, t_1)$.

2. Iteratively solve Eq. (2) using a gradient ascent loop (iteration l) nested within an expectation maximisation loop (iteration n):

   (a) A gradient descent step (with tunable step sizes α_j, and backprojection operator $\mathscr{B}$) to fit the measured data and compositional priors

$$\mu^{(n:l+0.5)}(\mathbf{x}, t_1 \in T_\mu) = \mu^{(n:l)}(\mathbf{x}, t_1) \tag{5}$$

$$- \alpha_1 \sum_{t_0 \in T_i} \Gamma(t_0, t_1) \mathscr{B}\left\{ I_s(\mathbf{r}, \theta)[I(\mathbf{r}, \theta, t_0) - \sum_{t_1' \in T_\mu} \Gamma(t_0, t_1') e^{(-\mathscr{P}_\theta \mu^{(n:l)})(\mathbf{r}, t_1')} \right\}$$

$$- \alpha_2 \frac{\mu}{\sigma^2} \exp\left[-\frac{(\mu_s - \mu_{\text{rock}})^2}{2\sigma_s{}^2} - \frac{\mu^{(n)^2}}{2\sigma^2}\right] - \alpha_3 \frac{\mu}{\sigma^2} \exp\left[-\frac{(\mu_s - \mu_{\text{rock}})^2}{2\sigma_s{}^2}\right].$$

$$- \alpha_2 \sum_{m \neq \text{rock}} \frac{\mu - \mu_m}{\sigma^2} \exp\left[-\frac{\mu_s{}^2}{2\sigma_s{}^2} - \frac{(\mu^{(n)} - \mu_m)^2}{2\sigma^2}\right].$$

   (b) A compressed sensing step to satisfy our *a priori* sparsity constraint, using the soft-thresholding operator $\mathscr{T}$

$$\mu^{(n:l+1)}(\mathbf{x}, t_1 \in T_\mu) = [\mathscr{T}\mu^{(n:l+0.5)}](\mathbf{x}, t_1).$$

   (c) If the gradient descent loop has converged, update the outer loop

$$\mu^{(n+1:l=0)}(\mathbf{x}, t_1 \in T_\mu) = \mu^{(n:l \to \infty)}(\mathbf{x}, t_1 \in T_\mu)$$

As discussed in the introduction, we have three key objectives for this algorithm: (i) it should place our previous algorithms within a larger context; (ii) it should improve reconstruction performance with imperfect (i.e. real) data; and (iii) it should generalise to more complex, previously-intractable physical systems. In the remaining sections of this paper we will address each of these in turn.

4. Placing previous theoretical results in context

In this section we will discuss how our previous algorithms for time-resolved tomography of two-phase flow in porous materials [4, 17], arise as limiting cases of this algorithm. The empirically-derived SIRT-based algorithm of Myers et. al. [4] can be obtained as a limiting case of Eq. (2) if one assumes: (i) an approximately monochromatic, attenuation-contrast cone-beam CT imaging system where the linearised signal is dominated by Gaussian noise; (ii) a sample characterised by smooth change between evenly spaced reconstruction frames; (iii) a materials dictionary consisting of a mobile incompressible fluid, void, and a known static impermeable rock scaffold; (iv) a materials classification/segmentation scheme that leads to strict binary thresholding; and (v) a regularising sparsity constraint on the reconstruction. Assumptions (iii), (iv) and (v) are explicit, whilst (ii) is implicit in the use of linear interpolation along the time axis [4].

Assumption (i) describes the following linearised CT imaging system

$$P[I(\mathbf{r},\theta,t_0 \in T_i)|\tau(t_0,t_1 \in T_\mu),\mu(\mathbf{x},t \in T_\mu)] \propto \exp\left\{-\frac{[\ln I(\mathbf{r},\theta,t_0)-(\mathscr{P}_\theta\mu)(\mathbf{r},t_1)]^2}{2\sigma^2}\right\}, \quad (6)$$

where σ is the standard distribution of the noise. This leads to a SIRT-based algorithm, in the same sense that the assumption of Poisson noise at the detector led to an MLTR-based algorithm in the previous section. If the frames in the reconstruction are evenly spaced along the time axis with separation Δ, then assumption (ii) is consistent with the confidence measure

$$P[\tau(t_0 \in T_i, t_1 \in T_\mu)|I,\mu^{(n)}] \propto \max(0,\Delta-|t_1-t_0|). \quad (7)$$

Under assumptions (iii) and (iv), denoting the known support of the rock scaffold as Ω, we have

$$\begin{aligned}
\ln[P(\mu_s|\mu)] \quad &+ \quad \sum_m P[m|\mu_s,\mu^{(n)}]\ln[P(m|\mu_s,\mu)] \\
&\propto \quad -\mu(\mathbf{x},t \in T_\mu)^2, \text{ if } \mathbf{x} \in \Omega \text{ or } \mu(\mathbf{x},t \in T_\mu) < \varepsilon, \\
&\propto \quad -\left\{\frac{\int_\Omega \mu(\mathbf{x},t \in T_\mu)d\mathbf{x}dt}{\int_\Omega d\mathbf{x}dt} - \mu(\mathbf{x},t \in T_\mu)\right\}^2, \text{ otherwise.} \quad (8)
\end{aligned}$$

This will draw the solution μ to zero within the scaffold support Ω, and where the present guess at the solution is below some threshold ε chosen to separate fluid and void. This prior is similar to the "Discrete Reconstruction" prior of Venkatakrishnan et. al. [16], with the contribution of neighbouring voxels tuned to zero. Again using the Laplace density function [Eq. (4)] to enforce a sparsity constraint, substituting Eq. (6) through (8) into our MAP-EM framework in Eq. (2) and taking a gradient ascent approach to evaluating $\mu^{(n+1)}$, we can derive Eqs. 10 and 11 of Myers et. al. [4], which constitute a SIRT-based two-phase-flow reconstruction algorithm.

In order to derive the 2-phase flow reconstruction algorithm from Myers et. al. [17], we make assumptions (i), (iii) and (v), and define the interpolation weights $\Gamma(t_0,t_1)$ as per the previous section. Assumption (iii) restricts us to a single fluid phase, and requires a known rock scaffold. If a binary segmentation of the rock scaffold is used as μ_s, then this corresponds to taking the limit $\sigma_s \ll \mu_{\text{rock}}$ in equation 3.

5. Reconstruction of 2-phase flow

To allow verification and comparison of the MAP-EM algorithm [Eq. (5)] against prior work, we return to the ground-truth two-phase-flow data set described in Myers et. al. [19]. In this experiment the objective was to obtain data in the "slow-flow" limit, allowing enough data to

Fig. 1. 3D renderings of five sequential volume frames (proceeding top-to-bottom). Each frame is calculated as a difference image between the initial and current reconstructed states of the sample [i.e. $\mu(\mathbf{x},t)$]: in a drainage experiment this difference is primarily due to air infiltrating the pore space. The middle column shows the results of the reference FBP reconstruction (Visualization 1). The left column shows the reconstruction produced by our empirical SIRT-based algorithm [4]. The right column shows reconstructions performed using Eq. (5) (Visualization 2).

be collected so that conventional reconstruction algorithms (i.e. FBP) could be used to calculate a "ground-truth" reference reconstruction. The measured data could then be sub-sampled (e.g. keeping only every 10th radiograph) to simulate a more rapid experiment. A 5mm diameter Bentheimer quartz plug was saturated with 1M potassium iodide. and then drained as slowly as possible: 0.025μL/min for 15 hours. The imaging setup was designed to maximise signal-to-noise ratio at the cost of field-of-view and spatial resolution, with a cone-angle of 52 degrees, detector resolution of 512^2 pixels, and a minimally filtered polychromatic beam. After geometric magnification, the voxel size was 12.8μm. We acquired 133,920 radiographs during drainage, with angular separation 0.5 degrees and exposure time of 0.4s.

Firstly, we calculated a "ground truth" reference reconstruction using conventional FBP. We then simulated a more rapid flow experiment by sub-sampling the data along the time axis, keeping only every tenth radiograph. This under-sampled data set was used to reconstruct the infiltration of air into the sample during drainage with: our prior empirical SIRT-based algorithm [4]; our prior statistical SIRT-based algorithm [17]; and the MAP-EM, MLTR-like algorithm shown in Eq. (5).

The majority of drainage occurs over 30 revolutions, corresponding to 30 frames in the reconstruction. Volume renderings of the air distribution in the sample are presented and compared in Fig. 1. The exterior ring present in these renderings is due to fluid surrounding the sample that has drained almost immediately. The additional external ring in the FBP reconstruction is due to a sub-pixel shift in the sample container; this is removed in the dynamic tomography reconstructions by the regularisation term. Horizontal slices through these volume renderings are presented in Fig. 2. Qualitative inspection shows that the MAP-EM algorithm displays improved performance compared to the empirical SIRT-based algorithm: (i) in small regions where our assumptions concerning the object are violated, such as partially-filled pores (see, e.g. upper right of 2D slices in Fig. 2); (ii) for small features where signal is very low, and a more physically realistic reconstruction model is important (see, e.g. features towards the bottom of the 2D slices in Fig. 2); and (iii) where a Haines jump has occurred whilst the object is revolving (see, e.g. the features towards the bottom of the 3D renderings in Fig. 2, which are reconstructed differently be each method).

Qualitative inspection of individual slices does not reveal substantial differences in performance between the algorithm presented in this paper, and our prior statistical SIRT-based algorithm [17]. Accordingly, we perform quantitative error analysis on our dynamic tomography reconstructions from under-sampled data, taking the FBP reference reconstruction as a ground truth. We note that the FBP reconstruction forms an imperfect baseline: it exhibits noise, "cupping" artefacts due to beam-hardening, and motion artefacts where Haines jumps have occurred during drainage. By quantifying these imperfections in the FBP reference reconstruction, we can establish a measure of acceptable error in our dynamic tomography reconstructions.

To quantify the errors in the FBP reference reconstruction we segment out the static rock scaffold from the initial static scan μ_s, and assess how well this unchanging feature has been reconstructed by FBP. We calculate the error (in the L2 norm) in the FBP reconstruction of the rock scaffold, using the initial static scan of the sample as a motion-artefact-free baseline. Normalising this error by the number of voxels in the rock scaffold, we obtain an estimate of noise and motion artefacts in the reference FBP reconstruction. This error measurement does not account for beam hardening, as it is equally present in the FBP and static reconstructions. We use this estimate of experimental error to define arbitrary units for the L2 error-per-voxel in our reconstructions: the rock scaffold in the FBP reconstruction has an error of 1.0 a.u., so anything below this may be regarded as being within the limits of experimental error.

The empirically-derived SIRT-based dynamic tomography algorithm [4] produces a reconstruction with an error of 1.77 a.u.; the statistical SIRT-based algorithm [17] produces better

Fig. 2. Reconstructed 2D (right) and 3D (left) slices through a Bentheimer sandstone during drainage, calculated as a difference image between the initial and current reconstructed states of the sample [i.e. $\mu(\mathbf{x},t)$]. 2D slices are a 480×400 pixels (6.7×5.6mm), taken normal to the rotation axis. *Middle row:* a reference reconstruction from full data, using FBP. In the upper right is a pore which displays partial volume effects. This violates our assumption that the rock is present at only a single grey level. *Top row:* reconstruction from every 10th viewing angle, using an empirical SIRT-based 2-phase flow algorithm [4]. *Bottom row:* reconstruction from every 10th viewing angle, using MLTR-like algorithm [Eq. (5)]. Note successful reconstruction of partial volume effects.
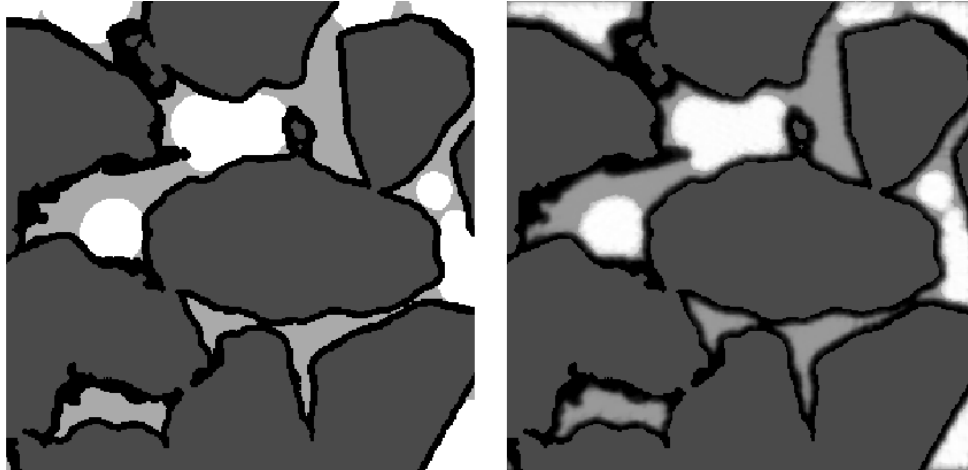
Fig. 3. *Left:* a 292×292 pixel slice through the synthetic 420-pixel diameter, 4-material object, taken normal to the rotation axis when the time-evolution was roughly 50% complete. The immobile rock phase is in dark grey, and the three mobile fluid phases are black, light grey and white respectively. *Right:* a time-resolved reconstruction from simulated data, using 72 radiographs per time-step. The majority of thin-film structures are correctly reconstructed.

results with an error of 1.24 a.u., though still slightly less accurate than the FBP algorithm. Of the three dynamic tomography algorithms, only the MLTR-like algorithm presented in this paper [Eq. (5)] produces a reconstruction within experimental error (0.81 a.u.) of the under-sampled data.

6. Synthetic 4-material data

The algorithm presented in Eq. (5) is, in theory, applicable to multi-phase flow. This adaptability was not present in our previous approaches to dynamic tomography, which did not generalise beyond two-phase fluid flow in porous materials [4]. In order to test this algorithm we simulated CT imaging of a sample containing three mobile phases.

We use a material dictionary consisting of an impermeable scaffold and three mobile fluids. As discussed in the introduction, it is quite difficult to construct physically-realistic simulations of three-phase fluid flow. Having segmented a pore-space distribution from a tomogram of sandstone, a physically un-realistic time-evolving distribution of fluids was created by using several different region-growing transforms to slowly fill the pore space and overlaying the results (see Fig. 3 and Fig. 4, left). Imaging was simulated using a geometry identical to that described in section 5: a 512^2 detector, with cone-angle 52 degrees. We simulated 1512 radiographs with an angular separation of 5 degrees.

Figures 3 and 4 show the reconstruction from synthetic data. As in section 5, we used 72 radiographs per volume time-frame. The time-evolving sample contained several film-like structures. Thin films are of particular interest in multi-phase flow, so it is encouraging to see that the majority of these structures were successfully reconstructed. We repeated our quantitative error analysis from the previous section, generating an FBP reconstruction from 720 noiseless radiographs per volume time-frame. This reference dataset was compared to the original numerical phantom, and the error-per-pixel in the L2 norm was again used to calibrate an "arbitrary unit" of error. The reconstruction from under-sampled data generated by MLTR-like algorithm in Eq.
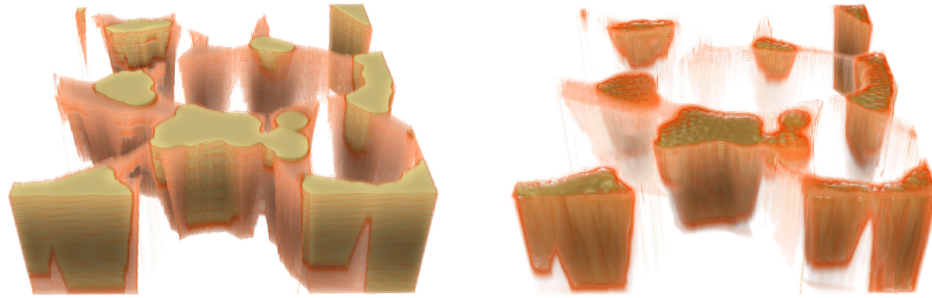
Fig. 4. 3D renderings showing the time-evolution of the slice shown in figure 3. Time evolution occurs along the vertical axis, proceeding from bottom to top. The two non-air fluid phases are rendered in brown and red respectively. *Left:* the numerical phantom used to generate the data. *Right:* a time-resolved reconstruction from simulated data, using 72 radiographs per time-step.

(2) had an error of approximately 1.3 a.u., though we are wary of directly comparing these arbitrary units to those used in the previous section. This is due to the lack of detector noise in the reference reconstruction, and the fact that in this case we have access to the original numerical phantom and are thus able to exactly calculate motion artefacts in the FBP reconstruction. The un-realistic nature of the synthetic fluid distribution implies that future work should focus on the collection of real, experimental 3-phase flow data.

7. Conclusion

We have presented a modular framework for time-resolved tomography of complex, highly constrained samples, using *maximum a posteriori* estimation. This framework has demonstrated both theoretical and practical value: its flexibility allows us to place previous results in a wider mathematical context, and quantitatively improves reconstruction performance when real samples display inevitable, small deviations from our assumptions. Finally, we have demonstrated the potential for applying this framework to the imaging of three-phase flow in porous materials, by applying it to synthetic data.

Acknowledgments