φenics: **Vainshtein screening with the finite element method**

Jonathan Braden,[a] **Clare Burrage,**[b] **Benjamin Elder,**[b,c] **and Daniela Saadeh**[b,d]

[a]Canadian Institute for Theoretical Astrophysics, University of Toronto, 60 St. George Street, Toronto, Canada

[b]School of Physics and Astronomy, University of Nottingham, University Park, Nottingham, U.K.

[c]Department of Physics and Astronomy, University of Hawai'i, 2505 Correa Road, Honolulu, USA

[d]Institute of Cosmology and Gravitation, University of Portsmouth, Portsmouth, U.K.

E-mail: daniela.saadeh@port.ac.uk

Abstract. Within the landscape of modified theories of gravity, progress in understanding the behaviour of, and developing tests for, screening mechanisms has been hindered by the complexity of the field equations involved, which are nonlinear in nature and characterised by a large hierarchy of scales. This is especially true of Vainshtein screening, where the fifth force is suppressed by high-order derivative terms which dominate within a radius much larger than the size of the source, known as the Vainshtein radius.

In this work, we present the numerical code φenics, building on the `FEniCS` library, to solve the full equations of motion from two theories of interest for screening: a model containing high-order derivative operators in the equation of motion and one characterised by nonlinear self-interactions in two coupled scalar fields. We also include functionalities that allow the computation of higher-order operators of the scalar fields in post-processing, enabling us to check that the profiles we find are consistent solutions within the effective field theory. These two examples illustrate the different challenges experienced when trying to simulate such theories numerically, and we show how these are addressed within this code. The examples in this paper assume spherical symmetry, but the techniques may be straightforwardly generalised to asymmetric configurations. This article therefore also provides a worked example of how the finite element method can be employed to solve the screened equations of motion. φenics is publicly available and can be adapted to solve other theories of screening.

Contents

1 Introduction

The theory of general relativity is an extremely successful description of gravity on Earth and within the solar system. On cosmological scales the standard cosmological model, ΛCDM, relies on general relativity as the theory of gravity, but is in agreement with observations only after the introduction of a new, dark, matter component, and either fine tuning of the cosmological constant or the introduction of a second new component, called dark energy [1]. Even then, tensions may be starting to arise between differing measurements of the Hubble constant [2–5].

It therefore behoves us to consider the possibility that general relativity may be modified, in particular on large cosmological scales. However, such modifications must still be in agreement with local tests of gravity, which see no deviation from general relativity [6]. This leads to the development of theories with screening mechanisms, where nonlinearities mean that modifications can be large on cosmological scales, but small locally [7, 8].

The simplest and most common modification of gravity is the introduction of an additional scalar degree of freedom in the gravitational sector. The additional scalar mediates a fifth force, which must be suppressed within the solar system via a screening mechanism to be consistent with observations. Such suppression can be achieved if the scalar field displays non-trivial self-interactions, which take the form of nonlinear terms in the Lagrangian. The presence of the nonlinear terms needed for screening should not be surprising, as both general relativity, still our current best theory of gravity, and the Higgs model, the only fundamental scalar to have been detected, require non-trivial self interactions of their fields. Nonlinear terms for the scalar mode occur naturally within, for example, $f(R)$ theories [9], massive gravity [10, 11] and Horndeski theories [12, 13]. In Refs. [14, 15] it has been suggested that a fifth force with screening may be present on galactic scales. Theories of dark energy and modified gravity, including those theories that possess screening, are reviewed in Refs. [7, 8, 16–18].

Of particular interest for this work are Galileon scalar-tensor theories. Galileon theories were introduced in Ref. [19], as higher-order derivative scalar theories that, nevertheless, have second-order equations of motion. In flat space-time, the theories were also required to be invariant, up to total derivatives, under shifts of the vacuum expectation value of the scalar field and shifts in its gradient. As these theories were only required to respect the symmetry up to total derivatives, we will refer to these theories as Wess-Zumino Galileons following Ref. [10]. The presence of the symmetry helps protect the particular form of these theories from quantum corrections.

Wess-Zumino Galileon theories can be generalised in many different ways. We can require that the symmetry be completely respected at the level of the action (not just up to total derivatives), but relax the requirement that the equations of motion be second order. The resulting terms in the equation of motion may be higher order in derivatives, but the resulting Ostrogradski ghost instability appears only at the cut-off of the effective field theory. The discussion in Ref. [10] indicates that the UV behaviour of these theories may differ from Wess-Zumino Galileons. If we do not insist on the symmetry, but do require the equations of motion to remain second order in derivatives, then the Wess-Zumino Galileons generalise to the Horndeski scalar-tensor theory [12, 13], where the equations of motion remain second order even around a curved background. Indeed the Horndeski model can be generalised even

further to beyond-Horndeski [20, 21] and Degenerate Higher-Order Scalar-Tensor (DHOST) theories [22–25].

A final Galileon model of interest here is the UV complete Galileon of Ref. [26]. This is a two scalar field model which has quadratic kinetic terms for the two scalars, but has a nonlinear potential and kinetic mixing between the two scalars. Considered just as a description of the two scalar fields, this theory is complete in the UV. If one field is heavier than the other, and is 'integrated out' of the theory by substituting its equation of motion back into the action, then the resulting theory for the light scalar is a Galileon theory, with higher order derivatives in the equations of motion.

In order to test theories with screening against observations, we need an accurate description of how the screening mechanism operates. As the theory is, by definition, nonlinear, and may contain non-canonical kinetic terms, it is challenging to solve analytically except in very simplified circumstances. The nonlinear structure of the theory can also pose a challenge to numerical simulation. Nevertheless, previous simulations of Vainshtein screening have successfully shed light on the phenomenology of the model. Ref. [27] showed that the Galileon force can be enhanced inside a hole in a planar object. A violation of the equivalence principle has been shown in finite-differencing simulations of a Galileon model for two extended bodies [28], and for two point sources in finite-element simulations of $P(X)$ theories [29]. Ref. [30] solved for the cubic Galileon field in two- and three-body systems on solar system scales, using finite-differencing techniques. Systems with time evolution have also been studied. In particular, a four-dimensional numerical code to study scalar gravitational radiation emitted from binary systems and probe the Vainshtein mechanism in situations that break the assumption of staticity and spherical symmetry was presented in Ref. [31].

Considerable effort has also been devoted to the numerical characterisation of chameleon-like screening. The time evolution of a chameleon and symmetron field around spherically symmetric black holes was studied in Ref. [32] using a combination of pseudo-spectral and Gauss-Legendre methods. In the context of laboratory tests of gravity, the behaviour of the chameleon and symmetron inside an experimental chamber has been studied using finite-differencing and finite-element techniques. This lead to experimental bounds on the chameleon [33, 34] and symmetron models [35, 36]. Ref. [37] studied the symmetron force in Casimir experiments, forecasting the sensitivity of future experiments that are realisable with the current state of the art technology. In Ref. [38], the finite element method was used to study the dependence of chameleon screening on the source shape, assuming cylindrical symmetry. All these studies solved the full equation of motion, without restricting the behaviour to specific approximating regimes.

Much effort has gone into studying the cosmological evolution of these theories numerically. It was shown in Ref. [39] that N-body simulations in theories with Vainshtein screening could be sped up by refining the mesh only in regions where the fifth force is not suppressed. An alternative approach to speeding up such calculations by introducing a screening factor was introduced in Ref. [40], and further attempts to speed up calculations of chameleon screening mechanisms are detailed in Ref. [41]. ISIS, an N-body cosmological code with scalar fields based on RAMSES [42], uses a nonlinear multi-grid solver that can treat a large class of scalar-tensor theories of modified gravity that possess screening mechanisms. This was extended to

include disformal couplings to matter in [43]. The results of a code comparison project for these N-body modified gravity codes were reported in Ref. [44]. In Ref. [45], the ECOSMOG code was used to simulate linear and nonlinear growth of the large-scale structure in Cubic Galileon gravity. Subsequently, it was extended to use adaptive mesh refinement in Ref. [46], and was applied to quartic Galileons in Ref. [47]. Various Boltzmann codes have also been written to simulate cosmological evolution in theories of modified gravity, including ISiTGR [48], MGCAMB [49, 50], EFTCAMB [51, 52] and hi_class [53, 54]. A comparison of these codes can be found in Ref. [55]. Simulations of $f(R)$ theories of modified gravity, including the hydro-dynamical code SHYBONE, are reported in Ref. [56]. A recent review of simulation techniques for modified gravity theories, in particular those with screening, can be found in Ref. [57]. These works have focused on simulating the cosmological solutions to the Galileon equations of motion, and the resulting consequences for observational probes. This has necessarily required making approximations to, or truncations of, the Galileon equations of motion.

In this work we focus on solving the full nonlinear field equations for screened scalar fields around isolated compact objects. To properly test screening mechanisms and obtain robust constraints, one must obtain solutions to the full field equations. In doing so, one immediately encounters a large hierarchy of scales. The simulation grid must be sufficiently large (of order several Compton wavelengths m^{-1}) to accurately capture the Yukawa suppression, while simultaneously having a small enough spacing to capture any features in the source-vacuum transition. At minimum, the source will have a characteristic size r_s, possibly with additional substructure of much smaller width w. This leads to potentially large hierarchies mr_s and mw. Standard numerical approaches using uniformly spaced grids are ill-suited to such problems, as the number of grid-points must exceed the largest of these hierarchies. Incorporating this hierarchy of scales accurately is extremely computationally burdensome even in cases of spherical symmetry, where the equations reduce to a single spatial dimension, and is practically intractable when symmetry assumptions are relaxed and additional spatial dimensions are required. To overcome this challenge, nonuniform meshes are required, which cluster lattice points near narrow features (i.e. boundary layers), while sparsely sampling regions where the solutions are smooth. One particularly natural way to do this, which we explore in this paper, is to use finite-element methods.

In this work, we develop a numerical code – φenics – applying the finite element method to two case studies of interest for Vainshtein screening: a theory with high-order derivative operators in the equation of motion and a second theory with nonlinear potential self-interactions combined with kinetic mixing. The latter of these theories is (a slightly modified form of) the UV-complete Galileon, and the former is the resulting low-energy theory obtained when the heavy field is integrated out. This relationship is explored in more detail in Ref. [58]. However, for the purposes of this paper, these models provide two illustrations of different types of theory that exhibit screening, and so serve as good test beds to demonstrate the versatility of the φenics code and the finite element method.

φenics is developed for the specific purpose of providing prospective users with an easy-to-use framework to test their theories of screening. The finite element method relies on an integral form of the field equations that is better equipped for high-derivative operators, as it helps to lower their order. Additionally, Neumann boundary conditions, typically placed at the origin, are natural in the formulation, without the need to impose them 'by hand'. These advantages are not unique to the finite element method, however, whose key advantage is its

flexibility to deal with complex geometries. As such, the finite element method is employed in a number of fields in physics and engineering. Although, in this work, we restrict to spherical geometries where the solutions found by the code can be compared with analytic results under specific assumptions, the flexibility of the finite element method will allow the code to be straightforwardly extended to deal with less symmetric systems.

This paper is structured as follows. In Sec. 2, we briefly review the Galileon models we use to demonstrate the φenics code, and the density profiles for the massive source adopted for this purpose. Section 3 describes the numerical method used to solve the equation(s) of motion for the scalar field(s), including a description of the meshes used. Section 4 gives examples of the successful working of the code, and of the types of results that can be obtained from it. We conclude in Sec. 5. Further technical details of the code – in particular, tests of numerical convergence – are included as appendices.

Throughout, we work in units with $c = \hbar = 1$, and use the reduced Planck mass $M_{\mathrm{P}} = 1/\sqrt{8\pi G}$. Our metric convention is $(-, +, +, +)$.

φenics is available from https://github.com/scaramouche-00/phienics. It is developed in `Python`, with both `Python 2` and `3` currently supported[1]. The code documentation is hosted at https://phienics.readthedocs.io.

2 Massive Galileons and Vainshtein screening

In this section we briefly present our two case-study theories. They are chosen to demonstrate the ability of φenics to deal with higher-derivative operators and the presence of multiple fields. Furthermore, these theories are related in the manner discussed in Ref. [26, 58]. The equations of motion and boundary conditions required to specify solutions are given in Sec. 2.1, and the source profiles around which we solve for the behaviour of the scalar fields are detailed in Sec. 2.2.

Most field theories must be viewed as low-energy effective field theories (EFTs), and scalar field theories used to explore screening are no exception. To ensure the validity of any solution we have found, it is necessary to compute higher-order operators omitted in the low-energy description and ensure they are smaller than the terms included when defining the equations of motion. The ability to compute these operators accurately is a key ability of φenics. We introduce this aspect in more detail in Sec. 2.3.

2.1 Equations of motion and boundary conditions

We solve for the behaviour of the Galileon fields(s) around a static, spherical compact source in Minkowski space-time, neglecting any space-time curvature coming from the presence of the source, or back-reaction due to the energy density stored in the scalar field profile. At minimum, the compact source is characterised by an overall size r_{s} and mass $M_{\mathrm{S}} = 4\pi \int_0^\infty \mathrm{d}r\, r^2 \rho(r)$, where ρ is the source's energy density.

We consider two models:

---

[1]However, please note `Python 2` support will be discontinued.

1. A theory of a massive scalar field characterised by higher-derivative operators with equation of motion

$$\Box\pi - m_\pi^2\pi - \epsilon\frac{\Box((\Box\pi)^n)}{\Lambda^{3n-1}} = \frac{\rho}{M_{\mathrm{P}}}\,. \tag{2.1}$$

   Here π is a massive scalar field (of mass m_π), Λ is a cut-off scale with units of mass, and ρ is the energy density of the source object. The dimensionless parameter ϵ regulates the strength of the nonlinear term, which is expected to give rise to Vainshtein screening. In the following, this theory will be referred to as the '**single-field theory**'.

2. A theory of two coupled massive scalar fields with nonlinear self-interactions, originally proposed in [26] in the context of massive Galileons, with equations of motion

$$\Box\phi - m_\phi^2\phi - \alpha\Box H = \frac{\rho}{M_{\mathrm{P}}}\,, \tag{2.2a}$$

$$\Box H - m_H^2 H - \alpha\Box\phi - \frac{\lambda}{3!}H^3 = 0\,. \tag{2.2b}$$

   Here, ϕ is a light cosmological scalar field of mass m_ϕ, H is a self-interacting more massive scalar field of mass m_H, ρ is the energy density of the compact source, α is the dimensionless coupling between the two scalar fields, and λ is the dimensionless parameter controlling the strength of the self-coupling of H. Typically, $m_\phi \ll m_H \ll M_{\mathrm{S}}$. Additionally, $m_\phi r_{\mathrm{s}} \ll 1$, i.e. the range of the fifth force is much greater than the characteristic size of the source. The Compton wavelength m_H^{-1} of the heavy field H can be either greater or smaller than the source radius. In the following, this theory will be referred to as the '**two-field theory**'.

In Ref. [26] the two-field model was shown to be UV complete, in the absence of the coupling to matter fields. Including additional couplings to matter necessarily introduces higher order operators, and so may change the UV-complete nature of the theory. Also in Ref. [26], it was shown that integrating-out the heavy field H from the two-field model leaves a single-field higher-derivative theory of the form of Eq. (2.1) with $n = 3$. The single-field higher derivative model of Eq. (2.1) is only valid up to the cut-off Λ.

As the field configurations we consider are static, the d'Alembertian operator, $\Box$, becomes the Laplacian, ∇^2. Further, for spherically symmetric systems considered in this work, the Laplacian becomes $\nabla_r^2 \equiv \frac{1}{r^2}\frac{\partial}{\partial r}\left(r^2\frac{\partial}{\partial r}\right)$. Under these symmetry assumptions, Eq. (2.1) becomes:

$$\nabla_r^2\pi - m_\pi^2\pi - \epsilon\frac{\nabla_r^2((\nabla_r^2\pi)^n)}{\Lambda^{3n-1}} = \frac{\rho}{M_{\mathrm{P}}}\,, \tag{2.3}$$

and similarly Eq. (2.2) becomes

$$\nabla_r^2\phi - m_\phi^2\phi - \alpha\nabla_r^2 H = \frac{\rho}{M_{\mathrm{P}}}\,,$$

$$\nabla_r^2 H - m_H^2 H - \alpha\nabla_r^2\phi - \frac{\lambda}{3!}H^3 = 0\,. \tag{2.4a}$$

These equations must be supplemented by appropriate boundary conditions. Since we are interested in static solutions, far from the source the fields must take values that minimise

their potentials:

$$\pi(\infty) = \phi(\infty) = H(\infty) = 0 \,. \tag{2.5}$$

It is standard to impose a regularity condition at the origin:

$$\frac{\partial}{\partial r}\pi(0) = \frac{\partial}{\partial r}\phi(0) = \frac{\partial}{\partial r}H(0) = 0 \,. \tag{2.6}$$

Even though other choices are, in principle, possible, all known physical theories display this behaviour for extended sources. This choice of Neumann boundary conditions avoids cusps or singularities at the origin, which would cause the energy density of the scalar field(s) to be infinite. Although diverging potentials are present around point sources in Newtonian gravity and electrostatics, these are always regulated when extended sources are considered.

Finally, the single-field theory is order four, and thus requires two additional boundary conditions. As shown in Sec. 3.5 and Sec. 4.1.1, it is convenient to think of $\nabla_r^2\pi$ as an independent dynamical degree of freedom. Far from the source, the fields π, ϕ, and H will be close to their vacuum values. Due to their nonzero masses, the fields decay as $\sim e^{-m_\star r}/r$ far from the source. This also constrains the behaviour of the field derivatives at infinity. In the single-field theory, for example, $\frac{\partial}{\partial r}\pi, \nabla_r^2\pi, \left[\left(\nabla_r^2\pi\right)^n\right]$, and $\frac{\partial}{\partial r}\left[\left(\nabla_r^2\pi\right)^n\right]$ all decay as e^{-mr}/r or e^{-3mr}/r^3 for sufficiently large r. We therefore use

$$\left[\left(\nabla_r^2\pi\right)^n\right](\infty) = 0 \,, \tag{2.7}$$

as a second Dirichlet boundary condition for the single-field theory.

For the final boundary condition, we note the quantities $\nabla_r^2\pi(0)$, $\left[\left(\nabla_r^2\pi\right)^n\right](0)$, and $\frac{\partial}{\partial r}\left[\left(\nabla_r^2\pi\right)^n\right](0)$ are all expected to be finite but nonzero [58]. We therefore demand $\frac{\partial}{\partial r}[(\nabla_r^2\pi)^n](0) < \infty$ as a fourth boundary condition for the single-field theory. These boundary conditions are consistent with the requirements of the UV completion of the theory.

Summarising, the complete set of boundary conditions for the single-field theory is:

$$\left\{\pi(\infty) = 0; \frac{\partial}{\partial r}\pi(0) = 0; \left[\left(\nabla_r^2\pi\right)^n\right](\infty) = 0; \frac{\partial}{\partial r}\left[\left(\nabla_r^2\pi\right)^n\right](0) < \infty\right\} \,, \tag{2.8}$$

and

$$\left\{\phi(\infty) = 0; H(\infty) = 0; \frac{\partial}{\partial r}\phi(0) = 0; \frac{\partial}{\partial r}H(0) = 0\right\} \,, \tag{2.9}$$

for the two-field theory.

2.2 Source profiles

As stated above, we want to find static scalar field profiles around a static, spherically symmetric source of mass M_S and characteristic radius r_s. For a radial source density profile given by a step-function $\rho(r) = (3M_S/4\pi r_{\text{step}}^3)\Theta(r_{\text{step}} - r)$, analytic approximations allow us to estimate the form of the field profile. These analytic results provide a useful test of the code. For density profiles that vary with r, an analytic approximation to the form of the solution is not generally known.

For illustration, we will consider three source profiles: a smoothed top-hat, a truncated cosine, and a 'wedding cake' profile made up of three stacked Gaussians. Each of these profiles has

a natural intrinsic length scale: the location of the boundary layer for the smoothed top-hat, the period of the cosine, or the half-width of the widest Gaussian. For the step profile mentioned above, this intrinsic scale is r_{step}. However, the dynamics we are interested in are sensitive to the total mass of the object, so we define r_{s} as smallest radius enclosing 95% of the mass, rather than these intrinsic scales.

To relate the natural theory scales (denoted here by r_{th}) to our definition of r_{s}, we introduce a rescaling parameter $\bar{t}$ so that $r_{\text{th}} = \bar{t} r_{\text{s}}$. Denoting a profile characterised by theory scale r_{th} as $\rho(r|r_{\text{th}})$, we obtain $\bar{t}$ by solving $4\pi \int_0^{r_{\text{s}}} dr\, r^2 \rho(r|\bar{t} r_{\text{s}}) = 0.95\, M_{\text{S}}$.[2]

The expressions for the three example density profiles, shown in Figure 1, are:

1. A smoothed top-hat profile:

$$\rho(r) = \frac{M_{\text{S}}}{4\pi(-2w^3)\text{Li}_3(-e^{\bar{t} r_{\text{s}}/w})} \frac{1}{\exp\frac{r-\bar{t} r_{\text{s}}}{w} + 1} \tag{2.10}$$

where $\text{Li}_3(x)$ is the polylogarithm function of order 3. In the limit $\frac{w}{r_{\text{s}}} \to 0$ the density profile in Eq. (2.10) becomes a step function. In this limit, we find $\bar{t}^{-3} = 0.95$, so $\rho(r) = \frac{3M_{\text{S}}}{4\pi(\bar{t} r_{\text{s}})^3}\Theta(\bar{t} r_{\text{s}} - r) = 0.95\frac{3M_{\text{S}}}{4\pi r_{\text{s}}^3}\Theta(\bar{t} r_{\text{s}} - r)$.[3] Choosing a profile that is smoother than a step function avoids shocks in the derivatives of the true solution that are difficult to resolve. For this source profile, analytic solutions may be found in the limit of a step function and specific sub-cases.

2. A truncated cosine profile:

$$\rho(r) = \begin{cases} \dfrac{3\pi M_{\text{S}}}{4(\bar{t} r_{\text{s}})^3(\pi^2 - 6)}\left[\cos\left(\pi\dfrac{r}{\bar{t} r_{\text{s}}}\right) + 1\right] & \text{if } r \leq \bar{t} r_{\text{s}} \\ 0 & \text{otherwise} \end{cases} \tag{2.11}$$

describing a more slowly varying source profile. Note that even with this simple source distribution, there is no known analytic solution for the field profile.

3. A 'Gaussian wedding cake' profile, i.e. a linear combination of three Gaussians making a three-layered shape:

$$\rho(r) = \frac{M_{\text{S}}}{4\pi(\bar{t} r_{\text{s}})^3 X}\left[A_1 e^{-\frac{1}{2}\left(\frac{r}{\sigma_1}\right)^2} + A_2 e^{-\frac{1}{2}\left(\frac{r-\mu_2}{\sigma_2}\right)^2} + A_3 e^{-\frac{1}{2}\left(\frac{r-\mu_3}{\sigma_3}\right)^2}\right], \tag{2.12}$$

with the parameters μ_i, σ_i, and A_i controlling to locations, widths, and relative heights of the peaks. The normalisation factor X is set by the constraint $4\pi \int_0^\infty dr r^2 \rho(r) = M_{\text{S}}$. For calculations in this paper, we take $\mu_2 = \bar{t} r_{\text{s}}/3$, $\mu_3 = 2\,\bar{t} r_{\text{s}}/3$, $\sigma_1 = \bar{t} r_{\text{s}}/9$, $\sigma_2 = \bar{t} r_{\text{s}}/7$, and $\sigma_3 = \bar{t} r_{\text{s}}/12$. The dimensionless heights $A_{1,2,3}$ are chosen so $\rho(0) = 3/2 \times \rho(\mu_2) = 3 \times \rho(\mu_3)$ and $A_3 = 1$. The resulting normalisation factor is $X \approx 0.4186008$. The form of this profile is chosen to illustrate the flexibility of the finite element method in solving for non-monotonic source profiles with internal structure, rather than to represent a physical system of interest.

---

[2]We solve this using `Scipy`'s inbuilt Broyden method, with $\bar{t}^{(0)} = 1$ as an initial guess.

[3]With our definition of source radius r_{s} (i.e. the radius enclosing 95% of the source mass) the transition of the limiting top hat profile does not occur at r_{s}.
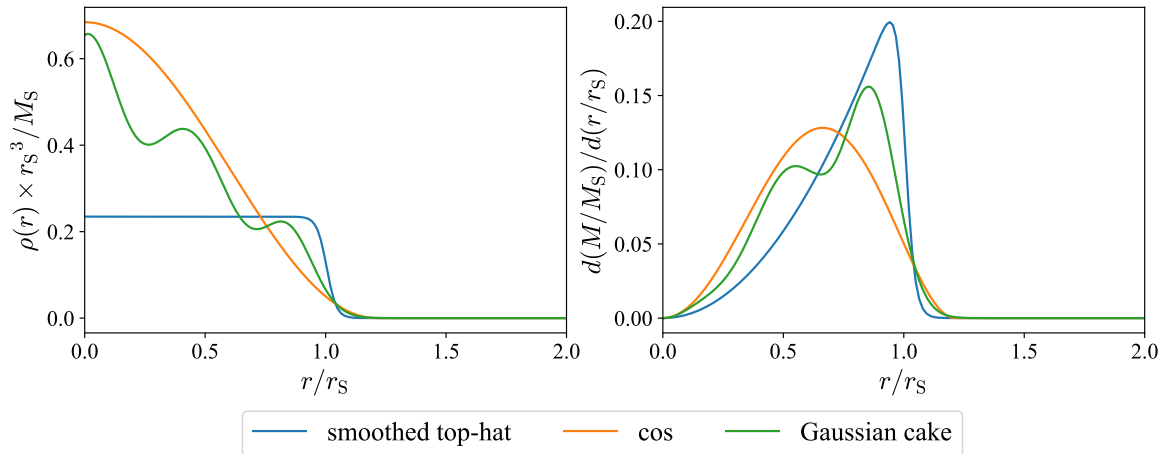
Figure 1. *Left*: The dimensionless source profiles in Eqs. (2.10), (2.11), and (2.12). *Right*: The dimensionless radial mass density $\left(\frac{r_{\mathrm{s}}}{M_{\mathrm{S}}}\right) \frac{dM}{dr} = 4\pi\rho \left(\frac{r}{r_{\mathrm{s}}}\right)^2 \frac{r_{\mathrm{s}}^3}{M_{\mathrm{S}}}$ (i.e. the mass of a spherical shell of radius r) for the same source profiles. In both cases the mass and radial distance are normalised to the source mass and radius, respectively.

In Eqs. (2.10), (2.11), and (2.12), the symbol $\pi = 3.14159...$ is a normalisation parameter and does not represent the field π in Eq. (2.1).

Numerically, the smoothed-top hat profile is the most difficult to solve for, as it has a boundary layer, introducing an additional hierarchy. The truncated cosine has a discontinuous second derivative, which can cause numerical challenges. The 'Gaussian wedding cake' profile is the easiest to obtain numerically, as it does not present discontinuities or internal structures on very different scales than its overall radius. However, it is the most challenging profile to understand analytically.

The user can easily implement their favourite source profile by adding a subclass in the `source` module using the existing examples as a template. More details are given in Appendix C and the φ`enics` documentation.

2.3 Screening and higher-order operators

To compare the relative strengths of the scalar and gravitational forces – which in turn allows one to determine whether the theory displays Vainshtein screening – φ`enics` computes the Newtonian potential from the Poisson equation

$$\nabla_r^2 \Phi_{\mathrm{N}} = \frac{\rho}{2M_{\mathrm{P}}^2} \,. \tag{2.13}$$

To make a fair numerical comparison, we compute Φ_{N} using the same mesh (see Sec. 3.2 for a definition) and source density profile used to compute the scalar field profile in the theory we are comparing to. Note that when the scalar field couples to matter with a coupling strength $M = M_{\mathrm{P}}$, the ratio of the scalar fifth force to the Newtonian gravitational force around a compact object is two.

When considering a theory as a low-energy effective field theory, as we are here for the single-field model, one should check whether the assumption that higher-order operators can be neglected when computing field profiles is indeed valid. This is an important test of the validity of any solution. The specific scheme used in φenics is especially optimised for the computation of such operators:

$$Q_n^{(\text{single}-\text{field})} \propto \frac{\epsilon^n}{\Lambda^{3n-1}} \nabla_r^2((\nabla_r^2\pi)^n) \tag{2.14}$$

and

$$Q_n^{(\text{two}-\text{field})} \propto \frac{\alpha^n}{m_H^{3n-1}} \nabla_r^2((\nabla_r^2\phi)^n) \tag{2.15}$$

In Eq. (2.14), the dimensionless parameter ϵ controls the strength of the nonlinear terms in π. In Eq. (2.15), the dimensionless parameter α controls the effect that the heavy field H and its associated nonlinearities has on the light field ϕ. Computing the operators in Eq. (2.14) and (2.15) across the radial domain accurately requires considerably more care than the field profiles or gradients alone. We discuss this aspect in Sec. 3.5, where we describe how these operators are computed within the φenics code.

3   Numerical method

In this section, we describe our numerical approach to solve the equations of motion Eq. (2.3) and Eq. (2.4), with boundary conditions Eq. (2.8) and Eq. (2.9), respectively. As analytic solutions are not generally known, numerical methods are essential. A number of difficulties arise that any numerical method must address:

1. we are solving a nonlinear boundary value problem, requiring a high dimensional nonlinear solver;

2. there are large hierarchies in the problem, in particular:

   – the ratio of the overall source size to the Compton wavelength of the fields; and

   – the ratio of the overall source size to the size to the scales of its internal structure;

3. for the single-field model (Eq. (2.3)) and in the computation of the Q_n operators (Eq. (2.14) and Eq. (2.15)), we must deal with numerically ill-behaved high-order derivative operators.

In the one-dimensional case considered here, a large hierarchy between the overall source radius and typical scales of its internal structure may lead to a similarly large hierarchy of scales in the structure of the sourced solutions. For example, if the source term is constant except in a narrow transition region, we expect the field solutions to possess a boundary layer. This is a manifestation of complex solution geometry derived from complexity in the source itself. For the spherically symmetric case considered here, the possible geometric complexity is limited by the lack of angular degrees of freedom. Absent these simplifying symmetries, solutions can be substantially more complex. In these cases, the numerical grid used to compute the solution must be sufficiently flexible to capture the full structure of the field solutions. A natural approach, which we will use here and outline below, is the finite element

method. For the simple spherically symmetric case studied here, more powerful techniques are available, such as globally based pseudospectral methods [59, 60]. However, unlike finite element methods, these global methods can be difficult to generalise to more complex source geometries. Therefore, in this paper we pursue a finite element approach, providing a simple test case in which to illustrate the utility of the methods, and providing a foundation on which to base further studies where the assumption of spherical symmetry is relaxed.

In the remainder of this section, we outline our methods to resolve each of these difficulties. First, in Sec. 3.1 we describe the Newton method for solving the equations of motion, which addresses their nonlinear nature. Next, in Sec. 3.2 we review our spatial discretisation scheme—the finite element method, briefly mentioning key features that allow us to address the issues of both large hierarchies and high-order derivatives. We continue by introducing the weak form of the differential equations in Sec. 3.3. In the final two subsection, we outline our approach to deal with large hierarchies and high-order derivatives. To tackle the large hierarchies, we apply nonlinear transformations of the integration domain, described in Sec. 3.4. The finite element method provides a natural way to implement such transformations. Finally, we introduce a modified form of the equations of motion which is better suited to dealing with high derivative operators, presented in Sec. 3.5.

3.1 Nonlinear Solver: Newton Method

In this subsection, we discuss how we approach the nonlinearity of the equations of motion Eq. (2.3) and Eq. (2.4). Of course, we cannot solve the continuum problem on a finite computer, so we must solve a discrete version of the equations. We use an (appropriately) discretised version of the Newton method applied to differential equations, which we briefly outline here. We restrict to a single differential equation, since the generalisation to multiple coupled differential equations is straightforward. Conceptually, we want to solve for the equation of motion at each lattice site. However, in practice we will seek to find a solution of the equation of motion projected into a convenient discrete basis to be outlined in Sec. 3.2. Our nonlinear solver is thus a high dimensional root finding problem, seeking the coefficients of the expansion of the solution in our discrete basis.

We begin by expressing our differential equation in the form $\mathcal{F}[u] = 0$. We now assume we have an approximate solution $u^{(k)}$ and wish to find the true solution u_{true}. If our guess is sufficiently close, we are justified to expand in the difference between the current guess and the true solution

$$0 = \mathcal{F}[(u_{\text{true}} - u^{(k)}) + u^{(k)}] \approx \mathcal{F}[u^{(k)}] + \int \mathrm{d}\mathbf{x} \frac{\delta \mathcal{F}}{\delta u(\mathbf{x})}[u^{(k)}] \left(u_{\text{true}} - u^{(k)}\right) . \qquad (3.1)$$

This gives us a linear equation for the difference $\delta u^{(k)} \equiv u_{\text{true}} - u^{(k)}$, which is sourced by the violation of the equation of motion $-\mathcal{F}[u^{(k)}]$. We then construct the next iterative approximation

$$u^{(k+1)} = u^{(k)} + \delta u^{(k)} , \qquad (3.2)$$

where $\delta u^{(k)}$ is the solution to Eq. (3.1). Finally, integrating against $\frac{\delta \mathcal{F}}{\delta u(\mathbf{x})}[u^{(k)}]$ and substituting (3.1), we obtain an equation for $u^{(k+1)}$ directly in terms of $u^{(k)}$

$$\int \mathrm{d}\mathbf{x} \frac{\delta \mathcal{F}}{\delta u(\mathbf{x})}[u^{(k)}]u^{(k+1)} = -\mathcal{F}[u^{(k)}] + \int \mathrm{d}\mathbf{x} \frac{\delta \mathcal{F}}{\delta u(\mathbf{x})}[u^{(k)}]u^{(k)} . \qquad (3.3)$$

If the initial guess is sufficiently close to the correct solution, repeated iterations of this procedure will converge to a solution to the equations of motion. Our discretisation procedure consists of first expressing these Newton iterations in weak form, and then discretising them using the finite element method. These aspects are discussed in Sec. 3.2 and Sec. 3.3.

Since we want a convergent solution to the equations of motion, there are two natural stopping criteria for our iterations:

1. Residual: Since a true solution to the equations of motion satisfies $\mathcal{F}[u_{\text{true}}] = 0$, we require that the *residual* $\mathcal{F}[u^{(k)}]$ satisfies

$$\left\| \mathcal{F}[u^{(k)}] \right\| \leq \epsilon_{\text{rel}}^{(\text{R})} \left\| \mathcal{F}[u^{(0)}] \right\| + \epsilon_{\text{abs}}^{(\text{R})}, \tag{3.4}$$

2. Change in the solution: As the sequence of solutions $u^{(k)}$ convergences, the change in the solution $\left\| u^{(k)} - u^{(k-1)} \right\|$ tends to 0. A convergence criterion can then be written as:

$$\left\| u^{(k)} - u^{(k-1)} \right\| \leq \epsilon_{\text{rel}}^{(\text{S})} \left\| u^{(0)} \right\| + \epsilon_{\text{abs}}^{(\text{S})}. \tag{3.5}$$

In the above, $\epsilon_{\text{rel}}^{(\cdot)}$ and $\epsilon_{\text{abs}}^{(\cdot)}$ are user set relative and absolute tolerances, $u^{(0)}$ is the initial guess, and $\|\cdot\|$ is some norm on the solution space. Note that the tolerances $\epsilon_{\text{rel}}^{(\cdot)}$ and $\epsilon_{\text{abs}}^{(\cdot)}$ need not be the same for both criteria. Relative criteria have the advantage that they do not depend on the characteristic size of u (for the change in solution criterion) or that of the terms in the equations of motion (for the residual criterion). However, if the initial guess is very close to the actual solution, they may enforce an unreasonable number of iterations on the solver, potentially beyond the accuracy the scheme is actually capable of achieving. In this case, absolute criteria become helpful because they stop the iterative process once an absolute tolerance has been met. In the formulae shown, relative and absolute criteria have been combined to retain both advantages. The specific norm used in φenics will be explained in Sec. 4.1.

As is clear from the above derivation, the key ingredient for success of this method is to start with a sufficiently accurate initial guess. Often, this is the most difficult step in the numerical solution of a nonlinear boundary value problem. Our approach for the screening problems studied in this paper is discussed in Sec. 4.1.1-4.1.2. For the problems considered here, we find these guesses to be sufficient to guarantee convergence. However, for other problems, additional tricks may be needed. This may include expanding the basin of convergence of the nonlinear solver by combining a gradient descent with Newton iterations or allowing variable Newton steps. It may also include improved initial guesses by considering a family of theories with an adjustable parameter that allows us to deform from some theory we can solve (either numerically or analytically) to the theory of interest. Starting from the solvable theory, we can adjust the parameter in a series of steps, using the solution at each intermediate step as the initial guess of the next step. Since these improvements were not needed for the problems studied here, we leave their implementation to future work.

3.2   Spatial Discretisation: The Finite Element Method

In order to obtain a discrete approximation to the Newton iterations in Eq. (3.1), we use the finite element method, which we briefly outline here. For an accessible reference on the finite

element method, see Ref. [61]. As will be clear from the presentation, many implementation choices are possible, and we make use of the publicly available `FEniCS` library[4] [62, 63] for the required functionality[5]. To avoid unnecessary technical complications, here we restrict the discussion to implementation in one spatial dimension.

Our first step is to replace the infinite radial domain $[0, \infty)$ by a large but finite domain $[0, r_{\max}]$. In doing so, we must ensure that imposing the Dirichlet boundary condition at a finite radius (instead of infinity) does not distort the numerical solution. In other words, we must ensure that the systematic error made by applying the Dirichlet boundary condition at a finite radius is small.[6] For the problems of interest in this paper, the Compton wavelength is the largest scale and therefore controls the fields' behaviour at infinity. This means $r_{\max}$ must be orders of magnitude larger than the Compton wavelength of all the fields involved. To see this explicitly, we can approximate the profile ψ of a field of mass m far away from the source as $\psi(r_{\rm s}) \left(\frac{r_{\rm s}}{r}\right) e^{-m(r-r_{\rm s})}$, where $\psi(r_{\rm s})$ is the value at the source radius. In order for the true value of the field at $r_{\max}$, i.e. $\psi_{\rm true}(r_{\max})$, to be approximated by 0, i.e. the boundary value, this difference must be small. A good starting point is requiring $\psi_{\rm true}(r_{\max})/\psi(r_{\rm s}) \ll 1$, from which it follows:

$$\frac{r_{\max}}{r_{\rm s}} + \frac{1}{mr_{\rm s}} \ln\left(\frac{r_{\max}}{r_{\rm s}}\right) \gtrsim 1 + \frac{1}{mr_{\rm s}} \ln\left|\frac{\psi(r_{\rm s})}{\psi_{\rm true}(r_{\max})}\right| . \tag{3.6}$$

This estimate can be improved by applying a similar condition to the derivatives appearing on the equation of motion. In practice, one needs to check that changing the specific value of $r_{\max}$ does not change the numerical solution. For the values of $r_{\max}$ used in production runs, we verified that the results were insensitive to the precise numerical value.

The next step, from which the finite element scheme derives its name, is to divide this large (but finite) interval into a collection of N "elements", each storing a local representation of the functions defined on the interval. These elements are typically of varying sizes, and in two or more dimensions can even have varying geometries. These local function spaces are then stitched together to give a function defined on the whole domain. This is particularly powerful in more than one spatial dimension, where the cells can be joined together in highly nontrivial ways, allowing for an accurate representation of extremely complex geometries. In the one-dimensional case considered here, the adaptive nature of these elements provides a natural way to resolve the various hierarchies mentioned above. To fully specify our discretisation, we are now left with a number of choices:

- How do we choose the sizes of each element and distribute them through the domain?

- How do we represent a function (and its derivatives) within each element?

- How do we deal with the boundaries joining the elements together?

Our first step is to size our individual elements and position them in the domain. For the one-dimensional case considered here, this amounts to deciding the location of the left ($r_{l,\min}$) and right ($r_{l,\max}$) endpoints of the lth element. Equivalently, we may specify the length of each element and the position of its center. A convenient approach, which we use here, is

---

[4]https://fenicsproject.org/

[5]The `FEniCS` library was also used to study the phenomenology of screening in Refs. [29, 38].

[6]Another option is to expand in basis functions that are themselves defined on the semi-infinite interval, in which case the Dirichlet boundary condition can be imposed at infinity.
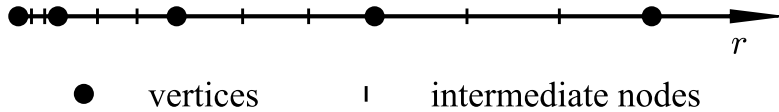
Figure 2. An illustration of a finite element discretisation with 4 elements. In this example, the degree of interpolating polynomials is 3.

to use a monotonically increasing mapping function $T(x)$ from an underlying variable x to the (dimensionless) radial variable r appearing in our equations. In other words, the radial element vertices r_i are related to the element vertices x_i in the original space by

$$r_i = r_{\mathrm{s}}T(x) \,. \tag{3.7}$$

If we uniformly size the elements in the x space, then the size of the element Δr at radial position r will be approximately

$$\Delta r(r) \approx \Delta x \frac{\partial T}{\partial x}(T^{-1}(r)) \,. \tag{3.8}$$

More details of the specific mapping functions are provided in Sec. 3.4.

Next we choose our local representation for functions within each element. A plethora of choices are available, and we follow the prescription used in `FEniCS`, which we outline below. This gives us access to a powerful library for the functionalities of the finite element method. Within each element, we choose a set of basis functions B_i and locally expand the function

$$f_l(r) = \sum_{i=0}^{O} c_i B_i(r) \,, \qquad r \in [r_{l,\min}, r_{l,\max}] \tag{3.9}$$

where the index l labels the element, and we have expanded in a basis of order O. $f_l(r)$ indicates the restriction of the function f to the element l. From this expansion, numerical derivatives are defined by

$$\frac{\mathrm{d}^n}{\mathrm{d}r^n}f_l(r) = \sum_{i=0}^{O} c_i \frac{\mathrm{d}^n}{\mathrm{d}r^n}B_i(r) \,. \tag{3.10}$$

Note, however, that if the expansion is truncated at low-order, higher-order numerical derivatives may poorly approximate the true derivative of the function. To choose a set of basis functions B_i and interpolation order O, we use Lagrange polynomials, which are specified by choosing a set of $O+1$ collocation points r_i for $i = 0, \ldots O$. In the `FEniCS` library, two points are given by the cell extrema, whereas the other $O-1$ points are taken to be uniformly spaced within each element. A simple collocation grid is illustrated in Figure 3.2.

To obtain a global approximation for a function we must join neighbouring elements together. For a connected domain in one-dimension, we take

$$r_{l+1,\min} = r_{l,\max} \,, \tag{3.11}$$

for $l = 1, \ldots N - 1$ where N is the number of elements, and

$$r_{1,\min} = 0, \quad r_{N,\max} = r_{\max}, \tag{3.12}$$

for the boundary elements. At this point in the construction, the functions are piecewise continuous, but discontinuities can arise at the element boundaries. If we wish to work in a continuous function space, we make use of the fact that the element boundaries are collocation points, and impose an additional global constraint that function values at the cell boundaries are equal in neighbouring cells. For a discretisation using N elements, each with an order p spectral approximation, there are then $N(p + 1)$ independent degrees of freedom for the discontinuous function space, and $Np + 1$ independent degrees of freedom for the piece-wise continuous function space. As explained in Sec. 4.1.3, at various points we will find it advantageous to either impose global continuity or to work within a piece-wise continuous function space.

3.3 Weak Form of Differential Equations

The finite element method naturally lends itself to the integral (weak) form of the field equations. Consider a differential operator $\mathcal{D}[u]$ acting on a function u, and the corresponding equation $\mathcal{D}[u] = 0$. In our case, this equation is a Newton iteration as in Eq. (3.1). We consider the inner product $\langle \mathcal{D}, v \rangle$ between $\mathcal{D}$ and a test function v selected from some test function space. A weak solution to the differential equation is one for which this inner product vanishes for *any* function in the test space. This is, of course, a weaker condition than demanding that the equation $\mathcal{D}[u(x)] = 0$ holds at every point x (i.e. the strong form of the equation). As a result, weak solutions can accommodate the presence of shocks and point-wise discontinuities in solutions. As test space, we use the Sobolev space H^1 of square-integrable functions in $\mathbb{R}^3$ whose gradient is also square-integrable in $\mathbb{R}^3$. Numerically, we approximate this (infinite dimensional) test space by the same basis of finite elements used to expand our solution. Systems of equations can be expressed as components of a single vector equation, with the scalar product taken with respect to a vector of test functions.

Since we work with spherically symmetric solutions in three spatial dimensions, it is reasonable to use the inner product

$$\langle f, g \rangle = \int_0^\infty \mathrm{d}r\, r^2 f g\,. \tag{3.13}$$

Note that by introducing the r^2 weight into our inner product, we are effectively looking for weak solutions to $r^2\mathcal{D}$. Without loss of generality, we can assume our test functions vanish at any points where we have specified Dirichlet boundary conditions. Inhomogeneous boundary conditions can be converted into homogeneous boundary conditions by writing $u = u_{\mathrm{bc}} + \Delta u$, where u_{bc} is some function satisfying the inhomogeneous boundary conditions. We then solve for Δu instead, which lives in a function space satisfying homogeneous boundary conditions. The integral projections appearing in the weak formulation must be estimated numerically. These are approximated internally by `FEniCS` using the `assemble` method, which implements accurate quadrature or tensor methods to numerical integration.

We can also use the weak form of the differential equation to reduce the order of differential operators. Consider, for example, the radial Laplacian of a function

$$\langle v, \nabla_r^2 f \rangle = \int_0^\infty \mathrm{d}r\, r^2 v \frac{1}{r^2} \frac{\partial}{\partial r} \left(r^2 \frac{\partial f}{\partial r} \right) = -\langle \nabla_r v, \nabla_r f \rangle + \left[r^2 v \frac{\partial f}{\partial r} \right]_0^\infty\,. \tag{3.14}$$

Therefore, up to a boundary term, we can swap projections of $\nabla_r^2 f$ onto the test function v with projections of $\nabla_r f$ with the gradient $\nabla_r v$. For our set of boundary conditions Eq. (2.8) and Eq. (2.9), the boundary term vanishes. We will make use of this several times in Sec. 4.1.

3.4 Mesh Construction: Distributing and Sizing the Elements

To accurately capture the large hierarchies present in theories of screening, we take advantage of the flexibility of the finite element method to specify non-uniform spatial discretisations. As explained above, for an N element discretisation on a connected one-dimensional domain, this is accomplished by specifying a collection of $N+1$ element vertices. In what follows, we denote this collection of vertices (which define the lengths and locations of the elements) a mesh. To generate a radial mesh, we begin with a uniform mesh on the interval $[x_{\min}, x_{\max}] \subset \mathbb{R}$. We then introduce a nonlinear map from this uniform mesh into our radial coordinates

$$r_i = T(x_i) \text{ for } x_i = x_{\min} + i \left(\frac{x_{\max} - x_{\min}}{N} \right) , \ i = 0, \dots, N . \tag{3.15}$$

Finally, the left and right boundaries of element l are identified as $r_{l,\min} = r_{l-1}$ and $r_{l,\max} = r_l$, where $l = 1, \dots, N$ labels the elements. A subtle point (discussed in Sec. 3.4.1, Sec. 3.4.2, and Sec. 3.4.3 below), is that the boundaries of the preimage uniform mesh can depend on the map T and the outer radial boundary of the final mesh $r_{\max}$.

Our goal is to create a mesh that is finest around the source-vacuum transition (where the sharpest variation in the field profiles are expected) and very coarse far from the source (where the solutions slowly decay with a characteristic scale given by the Compton wavelength). This allows us to represent a large simulation box and multiple large hierarchies with a reasonable number of points.

The mapping functions we introduce below have a non-monotonic radial vertex density, with the maximal point density (and thus the greatest spatial resolution) occurring at some radius r_{mpd} in the interior of the interval. It is often convenient to fix the radial position of this maximal vertex density to the location where either the source or solution is most rapidly varying. A common example occurs when the source profile includes sharp features, such as the top hat, which manifest as corresponding sharp features in the field profiles. Because our inner radial vertex is zero, one strategy to reposition r_{mpd} is to simply rescale all of the radial points using a new transform $T_{\mathrm{scl}}(x) = T(x)/c$. The scale c is obtained by solving $T''(x_{\mathrm{mpd}}) = 0$ for x_{mpd} and then setting $c = T(x_{\mathrm{mpd}})/r_{\mathrm{mpd}}$. The extrema of the original x mesh are then calculated to ensure the size of the radial domain is unchanged. For convenience, this procedure is included in φenics[7].

Additionally, we want to avoid clustering too many points near the origin, which can lead to either additional numerical noise or slower convergence as the number of elements is increased. We find enforcing linear scaling of the map near the origin is sufficient to avoid these numerical issues.

---

[7]Our implementation in φenics assumes that $r_{\mathrm{mpd}} = r_{\mathrm{s}}$. We find this is useful for the smoothed top-hat function, but not the Gaussian wedding cake or truncated cosine profiles. However, in more general problems it may be convenient to shift the location of maximal resolution, so we have introduced r_{mpd} as a free parameter in the above discussion.

After building a nonlinear mesh, it is sometimes necessary to either increase or decrease the density of mesh vertices in some regions. φenics includes built-in functionality for each of these tasks. In the former case, targeted refinement of the mesh can be achieved by splitting elements in half within specified radial intervals, thus increasing the number of elements. For the former case, φenics provides functionality to repel (or decluster) element vertices from a specific radius. Unlike the refinement algorithm, the declustering algorithm does not alter the number of mesh vertices. This algorithm is described in sub-section 3.4.3. After producing the mesh, we always check that mesh vertices are sufficiently spaced to be resolved well within machine precision (i.e. 10^4 times the FEniCS machine precision of 3×10^{-16}).

We now describe the two parametrised nonlinear transformations $T(x)$ used in this paper, as well as additional modifications that may be useful in specific problems. Note that parameter choices that lead to small residuals when solving the equations of motion may lead to noisy calculations of the operators Q_n in post-processing. This is likely because when more nodes are placed than are necessary to capture the relevant modes in a function, the only effect is to enhance round-off error by summing additional modes that are unresolved by machine-precision. The interested reader can find more details in Appendix B.

Users can implement their own choice of transformation, which may make use of φenics's linear refinement and point removal features. This is done by adding a custom mesh subclass, following the examples of the `ArcTanExp` and `ArcTanPowerLaw` subclasses included with φenics. These two transformations are outlined in this section. Further details are given in Appendix C.

3.4.1   `ArcTanExp` mesh

This mesh is generated by the nonlinear transformation:

$$r = T(x) = r_{\rm s} \frac{2}{\pi} \arctan{(kx)} \exp{\left(ax^3 + bx\right)} ,\qquad(3.16)$$

where $k > 0, a, b \geq 0$, and a and b are not simultaneously zero. Figure 3 gives an example of the resulting element vertices.

The arctan factor regulates vertex density around the origin while maintaining linear scaling. Meanwhile, the cubic exponential factor clusters vertices around the source radius $r_{\rm s}$, while spreading them out at large radii. The specific point density at $r_{\rm s}$ can be tuned using the parameters a and b. The flatter the transformation at $T^{-1}(r_{\rm s})$, the denser the number of points at $r_{\rm s}$, so that higher values of b/a make it less dense. Note that b and a also control the spacing at large radii, which can lead to a poor sampling of the Yukawa suppressed regime for extreme parameter values.

The extrema of the starting linear mesh $x_{\rm min} = T^{-1}(r_{\rm min})$ and $x_{\rm max} = T^{-1}(r_{\rm max})$ are obtained by solving $T(x_*) - r_* = 0$ numerically, with $* = \min, \max$, using Scipy's implementation of the Newton method. To obtain an initial guess for the algorithm, we approximate the inverse transformation T^{-1} as linear around the origin, and take

$$x_{\rm min}^{(0)} = \frac{r_{\rm min}}{T'(0)} .\qquad(3.17)$$

For the outer boundary, we approximate $T(x) \approx r_{\rm s} e^{ax + bx^3}$. If $a \neq 0$, we obtain a guess for $x_{\rm max}^{(0)}$ using Cardano's formula for cubic equations. We define $p \equiv b/a$, $q(r) = \frac{1}{a} \log{\left(\frac{r_{\rm s}}{r}\right)}$, and
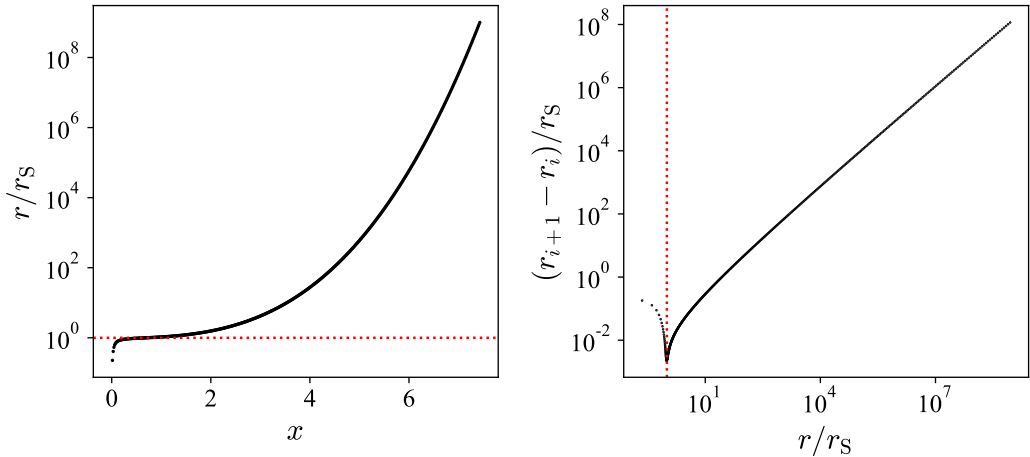
Figure 3. *Left:* A sample distribution of element vertices r_i for the `ArcTanExp` mesh (Eq. (3.16)) on a truncated domain with $r_{max} = 10^9 r_s$ and 500 elements. The mapping parameters are $k = 25$, $a = 5 \times 10^{-2}$, and $b = 3 \times 10^{-2}$. *Right:* The corresponding lengths of each element. In both panels, the red dotted line indicates the source radius.

$D(r) = (q(r)/2)^2 + (p/3)^3$. Our initial guess for x_{max} is then:

$$x_{max}^{(0)} = \sqrt[3]{-\frac{q(r_{max})}{2} + \sqrt{D(r_{max})}} + \sqrt[3]{-\frac{q(r_{max})}{2} - \sqrt{D(r_{max})}}. \tag{3.18}$$

Cardano's formula is justified because the exponential term in the transformation is monotonic, so there is only one real root to $T(x) - r = 0$. When $b \neq 0$ and $a = 0$, we instead use $x_{max}^{(0)} = \frac{1}{b} \log\left(\frac{r_{max}}{r_s}\right)$.

3.4.2 `ArcTanPowerLaw` mesh

This mesh is generated by the transformation:

$$T(x) = r_s \left(\frac{2}{\pi} \arctan(kx) + x^\gamma\right) , \tag{3.19}$$

where $k > 0$ and $\gamma \geq 1$. This transformation is visualised in Figure 4.

As in Eq. (3.16), the arctan term allows us to enforce linear scaling at the origin, which avoids placing an unnecessary number of points at $r = 0$. Compared to the previous mesh, this transformation distributes points more uniformly across the box, therefore placing a larger number of points at larger radii.

As in the previous case, we obtain the extrema of the starting mesh by using the Newton's method. At small radii, it is still a good approximation to describe the inverse transformation as linear around the origin, whereas at large r, where the power-law term dominates, $x \approx (r/r_s - 1)^{1/\gamma}$. Therefore we take

$$x_{min}^{(0)} = \frac{r_{min}}{T'(0)} , \tag{3.20}$$

and

$$x_{\max}^{(0)} = \left(\frac{r_{\max}}{r_{\mathrm{s}}} - 1 \right)^{1/\gamma} , \tag{3.21}$$
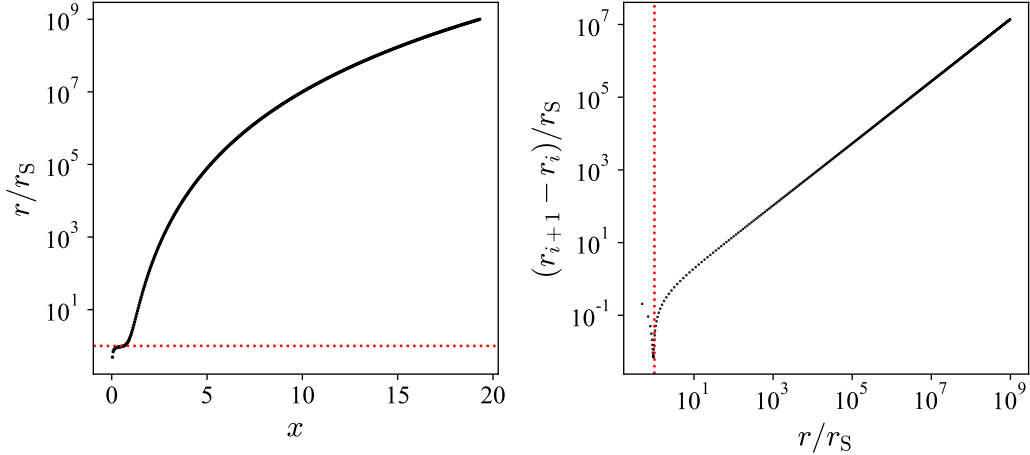
as initial guesses for the algorithm.

Figure 4. *Left*: A sample distribution of element vertices r_i for the `ArcTanPowerLaw` mesh (Eq. (3.19)), with mapping parameters $k = 25, \gamma = 7$, and $r_{\max} = 10^9 r_{\mathrm{s}}$. The finite element decomposition consists of 500 elements. *Right*: The corresponding lengths of each element. In both panels, the red dotted line indicates the source radius.

3.4.3   Linear Refinement and Vertex Declustering

For all choices of mesh, φenics allows the mesh to be linearly refined in a specified radial domain, which allows the user to further resolve specific features. This is obtained by using FEniCS's inbuilt `refine` method, which halves cells in a specified interval $[r_{\mathrm{start}}, r_{\mathrm{stop}}]$. This procedure is repeated a user-set number of times.

Occasionally, numerical noise in specific parts of the domain, particularly at the Vainshtein radius where the nonlinear operators briefly go to 0, may propagate into surrounding regions due to the nonlinear nature of the equations of motion. This leakage results in numerical inaccuracies in the computed solution. To alleviate this problem, φenics includes an option to decluster points at some radius of interest r_{rm}.

The point declustering is implemented as follows.[8] The user provides a radial declustering point r_{rm}, declustering parameter k_{rm}, and radial extent A_{rm}. Given a nonlinear transformation T, such as Eq. (3.16) or Eq. (3.19), we define a new transformation T_*:

$$T_*(x) \equiv \frac{T(x)}{c} T_{\mathrm{rm}}(x), \tag{3.22}$$

where

$$T_{\mathrm{rm}}(x) = \frac{A_{\mathrm{rm}}}{\pi} \arctan \left(k_{\mathrm{rm}} \left(x - x_{\mathrm{rm}} \right) \right) + 1 + \frac{A_{\mathrm{rm}}}{2} . \tag{3.23}$$

---

[8]We found that this algorithm gave better results than fusing cells together.

In this expression, $x_{\rm rm}$ is a derived parameter determined from the user-supplied declustering radius via $r_{\rm rm} = T_*(x_{\rm rm})$.

The new function $T_{\rm rm}$ adjusts the slope of the transformation T_* at $r_{\rm rm}$. For the `ArcTanExp` transformation Eq. (3.16), choosing $k_{\rm rm} \gtrsim k_c \equiv 3\sqrt[3]{a}\log{(r_{\rm rm}/r_{\rm s})}^{2/3} + b$ decreases the slope, whereas $k_{\rm rm} \lesssim -k_c$ increases it; for the `ArcTanPowerLaw` transformation Eq. (3.19), $k_{\rm rm} \gtrsim k_c \equiv \left(\frac{1+A_{\rm rm}/2}{r_{\rm rm}/r_{\rm s}}\right)^{1/\gamma}\sqrt{\gamma^2 - 1}$ decreases the slope, whereas $k_{\rm rm} \lesssim -k_c$ increases it.[9] In the former case, points are repelled (i.e. declustered) from the location $x_{\rm rm}$ leading to decreased resolution, while in the latter they are attracted (i.e. clustered) leading to increased resolution. The user-supplied parameter $A_{\rm rm}$ sets the radial extent of the (de)clustering effect. Although clustering and declustering are both possible with this algorithm, we note that it was designed with declustering at the Vainshtein radius in mind, which is its main use.

Introducing the factor $T_{\rm rm}$ generically shifts the radial location with the densest clustering of points. We counteract this unwanted effect by introducing the scaling parameter c in Eq. (3.22), as discussed earlier in this section. We want to ensure that the maximal vertex density occurs at the source radius while simultaneously keeping the declustering/clustering at $r_{\rm rm}$. This gives us the following three conditions: $T_*''(x_{\rm s}) = 0$ at $x_{\rm s} = T_*^{-1}(r_{\rm s})$, and $T_*(x_{\rm rm}) = r_{\rm rm}$ for the unknown variables $\{x_{\rm s}, x_{\rm rm}, c\}$. We can solve explicitly for $c = T(x_{\rm rm})\left(1 + \frac{A_{\rm rm}}{2}\right)/r_{\rm rm}$, leaving two equations:

$$T_*(x_{\rm s}) = r_{\rm s} \,, \tag{3.24a}$$
$$T_*''(x_{\rm s}) = 0 \,, \tag{3.24b}$$

in two unknowns $\{x_{\rm s}, x_{\rm rm}\}$. We solve Eqns. (3.24) by using `Scipy`'s built-in `root` method, which applies a modification of the Powell hybrid method [64].

If the user implements clustering at $r_{\rm rm}$ (i.e. negative $k_{\rm rm}$), it is important to check that T_*' remains positive around $x_{\rm rm}$, so the transformation T_* is monotonic. We automate this check in φ`enics` by resampling T_*' in the vicinity of $x_{\rm rm}$ and explicitly checking that it remains positive.[10]

The remaining task is to update the extrema of the starting mesh, which will generally shift. After updating the derivative of the transformation as $T_*'(x) = T'(x)T_{\rm rm}(x)/c + T(x)T_{\rm rm}'(x)/c$, we repeat the procedure of solving $T_*(x_{\rm min,max}) = r_{\rm min,max}$ using the Newton method.

3.5 Treatment of High-Order Derivative Operators

A key challenge in the single-field theory (Eq. (2.1)) is to compute the highly nonlinear term $\nabla_r^2((\nabla_r^2\pi)^n)$ accurately. A similar challenge is encountered in the computation of the Q_n operators Eq. (2.14) and Eq. (2.15), which encode corrections to the effective IR theory.

---

[9]Please note these are rules of thumb only. Also, for $|k| \lesssim k_c$, the slope is changed, but not enough to place a new local maximum or minimum in the vertex distance.

[10]To resample, we first define new coordinates y via $x(y) = \sqrt{1 + A_{\rm rm}/2}\tan{(y/10)}/|k_{\rm rm}| + x_{\rm rm}$. We then select eleven resampling points y_i uniformly distributed in the interval $y \in [-10\arctan{(1)}, -10\arctan{(1)}]$.

To address the former challenge, we write the equations of motion as a system of three lower-order equations

$$\begin{pmatrix} Y - m_\pi^2 \pi - \epsilon \nabla_r^2 W \\ Y \\ W \end{pmatrix} = \begin{pmatrix} \frac{\rho}{M_P} \\ \nabla_r^2 \pi \\ Y^n \end{pmatrix} , \tag{3.25}$$

through the introduction of the auxiliary variables W and Y.

Similarly, we re-express the system of equations of the two-field theory

$$\begin{pmatrix} Y - m_\phi^2 \phi - \alpha Z \\ Z - m_H H - \alpha Y - \frac{\lambda}{6} H^3 \\ Y \\ Z \end{pmatrix} = \begin{pmatrix} \frac{\rho}{M_P} \\ 0 \\ \nabla_r^2 \phi \\ \nabla_r^2 H \end{pmatrix} , \tag{3.26}$$

through the introduction of the auxiliary variables Y and Z. For the two-field theory, this formulation is not needed to compute the field profiles. However, we find that introducing Y as an independent variable improves the calculation of the operators Q_n in post-processing. This is most likely because we have introduced additional degrees of freedom to resolve structure in the Laplacian, which are subsequently constrained by the equations of motion. The dimensionless units used in our numerics are presented in Appendix A. However, in this section we express the equations in their native form for readability.

In the single-field theory, the introduction of the the auxiliary variable is sufficient to accurately compute the nonlinear term in the equations of motion. Meanwhile, the Q_n operators obtained in post-processing require some more work, as they are not constrained by the equations of motion, with the exception of the $Q_3^{(\text{single}-\text{field})} \propto \nabla_r^2 W$ term in the single-field theory.

We begin by expanding

$$Q_n \propto \nabla_r^2(Y^n) = nY^{n-1}\nabla_r^2 Y + n(n-1)Y^{n-2}\nabla_r Y \cdot \nabla_r Y , \tag{3.27}$$

and we use the right-hand side of this expression for numerical evaluations. Furthermore, we project this expression onto the same function space as Y. Using some simple test profiles for Y, we verified that this yielded more accurate results than directly evaluating the left-hand side. We also did heuristic checks of accuracy by comparing numerical calculations of the left- and right-sides of Eq. (3.27). The most likely reason for the improved accuracy is that the Laplacian upweights the short-wavelength (i.e. high-order polynomial) modes, which can amplify both short-distance numerical noise and aliased short-wavelength power generated by the high degree of nonlinearity.

High orders of nonlinearity introduce one additional subtlety. The most natural set of units to evaluate the equations of motion do not necessarily lead to dimensionless Y's that are order one. As a result, high-order powers may either underflow or overflow floating point precision. Moreover, numerical aliasing may become more pronounced as Y increases due to the nonlinearity. Therefore, for practical computations it is convenient to first compute the rescaled operator

$$Q_n^K = \nabla_r^2 \left(\left(K \nabla_r^2 \phi \right)^n \right) , \tag{3.28}$$

for some K. From this we obtain

$$Q_n^{\text{num}} = Q_n^K / K^n. \tag{3.29}$$

Note that this is equivalent to working with a different set of of dimensionless units.

4 Example Solutions

In this section, we present results from φenics for three representative models — one for the single-field theory and two for the two-field theory. First, we derive the weak form of the Newton iterations in Sec. 4.1. This serves as a worked example for future φenics users. We also outline some additional numerical details specific to each example. Sec. 4.2 presents the numerical solutions computed by φenics. Finally, we compare our numerical results to theoretical predictions in Sec. 4.3. For the interested reader, numerical convergence tests are presented in Appendix B. These tests provide an important test of our numerics, and are meant as a basic template for future φenics users to ensure accuracy of their results.

In our single-field example, we consider the field behaviour around a source of large mass similar to the star Antares. This example allows us to explore the behaviour of a nonlinear, high-derivative theory in the presence of a realistic source. We choose m_π so that the Compton wavelength is much longer than the Vainshtein radius, which is typically the case of physical interest.

For the two-field theory, we consider two cases. In the first, the Compton wavelengths of both scalar fields are comparable to the size of the source. We expect features to be excited by the source-vacuum transition, with a characteristic size dependent on the width of the transition region. Characterising these features analytically, or even numerically, is challenging, so this model is a good testing ground for our numerical method.

In the second two-field model, we instead consider a different hierarchy of scales: $m_\phi r_{\text{s}} \ll m_H r_{\text{s}} \ll 1$. This allows us to test the regime in which the Compton wavelength of the heavy field is larger than the source (which results in different numerical and phenomenological behaviour). This hierarchy also allows us to demonstrate the behaviour of extremely light fields ϕ, arguably the regime of greatest interest for these theories.

The parameters specifying our three models are:

(M1) Single-field theory; the model parameters are $m_\pi = 10^{-50} M_{\text{P}}$, $\epsilon = 3 \times 10^{-3}$, $\Lambda = 10^{-39} M_{\text{P}}$, and $n = 3$. The source has mass $M_{\text{S}} = 5 \times 10^{39} M_{\text{P}}$ and radius $r_{\text{s}} = 7 \times 10^{45} M_{\text{P}}^{-1}$.

(M2) Two-field theory; the model parameters are $m_\phi = 10^{-48} M_{\text{P}}$, $m_H = 10^{-46} M_{\text{P}}$, $\alpha = 0.4$, and $\lambda = 0.7$. The source has mass $M_{\text{S}} = 10^{10} M_{\text{P}}$ and radius $r_{\text{s}} = 10^{47} M_{\text{P}}^{-1}$.

(M3) Two-field theory; the model parameters are $m_\phi = 10^{-51} M_{\text{P}}$, $m_H = 10^{-48} M_{\text{P}}$, $\alpha = 0.4$; we consider both $\lambda = 0$ and $\lambda = 0.7$. The source has mass $M_{\text{S}} = 10^{10} M_{\text{P}}$ and radius $r_{\text{s}} = 10^{47} M_{\text{P}}^{-1}$.

Our results for models (M1) and (M2) are given in Sec. 4.2, where we present numerically computed field profiles and derived quantities for the source profiles defined in Eqs. (2.10)-(2.12). We will instead use Model (M3) to illustrate numerical convergence tests in Appendix B,

as well as to compare against theoretical predictions for the special linear subcase $\lambda = 0$ in the UV theory (for which expressions were given in [65]). Unless otherwise stated, for the smoothed top-hat profile Eq. (2.10), we use a width $w = 0.02\,r_{\mathrm{s}}$.

4.1 Weak form of the Newton iterations

In this section, we derive the weak form of the Newton iterations as discussed in Sec. 3.3. The dimensionless units used in φenics are presented in Appendix A. However, in this section we will use dimensionful variables for readability. We also provide some additional technical details about the choice of continuous or discontinuous function spaces, and the stopping criteria for our Newton iterations. Throughout, we denote our current estimate of variable u by $u^{(k)}$, with the subsequent iteration denoted $u^{(k+1)}$.

4.1.1 Single-field, higher-derivative theory

In Eq. (3.25), we expressed the original equation of motion of the single-field theory as a system of equations that allows us to lower the order of the nonlinear, high-derivative operator. In order to be solved iteratively using the Newton method, Eq. (3.25) must be linearised as shown in Eq. (3.1). For the single-field theory, the linearised system of equations is:

$$
\begin{pmatrix} -m_\pi^2 & 1 & -\frac{\epsilon}{\Lambda^{3n-1}}\nabla_r^2 \\ -\nabla_r^2 & 1 & 0 \\ 0 & -n(Y^{(k)})^{n-1} & 1 \end{pmatrix} \begin{pmatrix} \pi^{(k+1)} \\ Y^{(k+1)} \\ W^{(k+1)} \end{pmatrix} =
$$
$$
\begin{pmatrix} Y^{(k+1)} - m_\pi^2 \pi^{(k+1)} - \frac{\epsilon}{\Lambda^{3n-1}}\nabla_r^2 W^{(k+1)} \\ Y^{(k+1)} - \nabla_r^2 \pi^{(k+1)} \\ W^{(k+1)} - n(Y^{(k)})^{n-1}Y^{(k+1)} \end{pmatrix} = \begin{pmatrix} \frac{\rho}{M_{\mathrm{P}}} \\ 0 \\ (1-n)(Y^{(k)})^n \end{pmatrix} , \qquad (4.1)
$$

with boundary conditions

$$
\left\{ \pi^{(k+1)}(r_{\max}) = 0; \nabla_r \pi^{(k+1)}(0) = 0; W^{(k+1)}(r_{\max}) = 0; \nabla_r W^{(k+1)}(0) = 0 \right\} . \qquad (4.2)
$$

We interpret this system of equations as a single vector equation and integrate each side against a vector test function (v_1, v_2, v_3)

$$
\left\langle \begin{pmatrix} Y^{(k+1)} - m_\pi^2 \pi^{(k+1)} - \epsilon\frac{\nabla_r^2 W^{(k+1)}}{\Lambda^{3n-1}} \\ \nabla_r^2 \pi^{(k+1)} - Y^{(k+1)} \\ W^{(k+1)} - n(Y^{(k)})^{n-1}Y^{(k+1)} \end{pmatrix}, \begin{pmatrix} v_1 \\ v_2 \\ v_3 \end{pmatrix} \right\rangle = \left\langle \begin{pmatrix} \frac{\rho}{M_{\mathrm{P}}} \\ 0 \\ (1-n)(Y^{(k)})^n \end{pmatrix}, \begin{pmatrix} v_1 \\ v_2 \\ v_3 \end{pmatrix} \right\rangle . \quad (4.3)
$$

The weak form is therefore:

$$
\int Y^{(k+1)}v_1 r^2 \mathrm{d}r - \int m_\pi^2 \pi^{(k+1)}v_1 r^2 \mathrm{d}r - \frac{\epsilon}{\Lambda^{3n-1}}\int \nabla_r^2 W^{(k+1)}v_1 r^2 \mathrm{d}r
$$
$$
+ \int \nabla_r^2 \pi^{(k+1)}v_2 r^2 \mathrm{d}r - \int Y^{(k+1)}v_2 r^2 \mathrm{d}r
$$
$$
+ \int W^{(k+1)}v_3 r^2 \mathrm{d}r - n\int (Y^{(k)})^{n-1}Y^{(k+1)}v_3 r^2 \mathrm{d}r
$$
$$
= \int \frac{\rho}{M_{\mathrm{P}}}v_1 r^2 \mathrm{d}r + (1-n)\int (Y^{(k)})^n v_3 r^2 \mathrm{d}r .
$$

$$(4.4)$$

Following Eq. (3.14), we can express all the terms that contain a Laplacian as products of gradients by integrating by parts. For example,

$$
\begin{aligned}
\int \nabla_r^2 \pi^{(k+1)} v_2 r^2 \mathrm{d}r &= \left[\frac{\partial \pi^{(k+1)}}{\partial r} v_2 r^2 \right]_0^\infty - \int_0^\infty \nabla_r \pi^{(k+1)} \cdot \nabla_r v_2 r^2 \mathrm{d}r \ , \\
&= - \int_0^\infty \nabla_r \pi^{(k+1)} \cdot \nabla_r v_2 r^2 \mathrm{d}r \ ,
\end{aligned}
\tag{4.5}
$$

since the boundary term is zero for our choice of boundary conditions Eq. (2.8). This gives us the weak form of the equations used in φenics

$$
\begin{aligned}
&\int \mathrm{d}r\, r^2 \left(Y^{(k+1)} - m_\pi^2 \pi^{(k+1)} + \frac{\epsilon}{\Lambda^{3n-1}} \nabla_r W^{(k+1)} \nabla_r \right) v_1 \\
&- \int \mathrm{d}r\, r^2 \left(\nabla_r \pi^{(k+1)} \nabla_r + Y^{(k+1)} \right) v_2 \\
&+ \int \mathrm{d}r\, r^2 \left(W^{(k+1)} - n(Y^{(k)})^{n-1} Y^{(k+1)} \right) v_3 \\
&= \int \mathrm{d}r\, r^2 \frac{\rho}{M_{\mathrm{P}}} v_1 + (1-n) \int \mathrm{d}r\, r^2 (Y^{(k)})^n v_3 \ .
\end{aligned}
\tag{4.6}
$$

As a final step, we require an initial guess as the starting point for Newton's method. This is provided by the limit in which the nonlinear term is dominant. In this case, we solve the Poisson equation

$$
-\nabla_r^2 W^{(0)} = \frac{\Lambda^{3n-1}}{\epsilon} \frac{\rho}{M_{\mathrm{P}}} \ ,
\tag{4.7}
$$

for $W^{(0)}$. From this, we compute $Y^{(0)} = \sqrt[n]{W^{(0)}}$, and finally we obtain $\pi^{(0)}$ from the Poisson equation $\nabla_r^2 \pi^{(0)} = Y^{(0)}$. In principle, another initial guess could be found by setting $\epsilon = 0$ to remove the nonlinear term. However, this almost never leads to convergence, except for extremely small values of the nonlinearity parameter ϵ, hinting that the nonlinearities are indeed very important, as expected.[11]

4.1.2 Two-field, second-derivative theory

The system of equations for the UV theory in Eq. (3.26) is linearised as:

$$
\begin{pmatrix}
Y^{(k+1)} - m_\phi^2 \phi^{(k+1)} - \alpha Z^{(k+1)} \\
Z^{(k+1)} - m_H^2 H^{(k+1)} - \alpha Y^{(k+1)} - \frac{\lambda}{2}(H^{(k)})^2 H^{(k+1)} \\
Y^{(k+1)} - \nabla_r^2 \phi^{(k+1)} \\
Z^{(k+1)} - \nabla_r^2 H^{(k+1)}
\end{pmatrix}
=
\begin{pmatrix}
\frac{\rho}{M_{\mathrm{P}}} \\
-\frac{\lambda}{3}(H^{(k)})^3 \\
0 \\
0
\end{pmatrix} ,
\tag{4.8}
$$

which is the strong form of the Newton iterations. These are supplemented by the boundary conditions

$$
\left\{ \phi^{(k+1)}(r_{\max}) = 0; H^{(k+1)}(r_{\max}) = 0; \nabla_r \phi^{(k+1)}(0) = 0; \nabla_r H^{(k+1)}(0) = 0 \right\} .
\tag{4.9}
$$

---

[11]Setting $\epsilon = 0$ gives a good guess far from the source, where the linear theory is recovered, but nonlinearities generally dominate closer to the source.

As with the single-field theory, we interpret the system of equations as a single vector equation and take the scalar product of both sides with a vector test function (v_1, v_2, v_3, v_4). We then integrate by parts to turn integrals of Laplacians into products of gradients, as in Eq. (3.14). This gives us the final form of our equations

$$
\int \mathrm{d}r\, r^2 \left(Y^{(k+1)} - m_\phi^2 \phi^{(k+1)} - \alpha Z^{(k+1)} \right) v_1
$$
$$
+ \int \mathrm{d}r\, r^2 \left(Z^{(k+1)} - m_H^2 H^{(k+1)} - \alpha Y^{(k+1)} - \frac{\lambda}{2}(H^{(k)})^2 H^{(k+1)} \right) v_2
$$
$$
+ \int \mathrm{d}r\, r^2 \left(\nabla_r \phi^{(k+1)} \nabla_r + Y^{(k+1)} \right) v_3 - \int \mathrm{d}r\, r^2 \left(\nabla_r H^{(k+1)} \nabla_r + Z^{(k+1)} \right) v_4 \qquad (4.10)
$$
$$
= \int \mathrm{d}r\, r^2 \frac{\rho}{M_\mathrm{P}} v_1 + \int \mathrm{d}r\, r^2 \frac{\lambda}{3}(H^{(k)})^3 v_2
$$

For our initial guess, we use the solution for $\lambda = 0$ as a "linear guess". Unlike the single-field theory, in all cases examined this choice led to convergence. This suggests very different behaviour in the UV-complete two-field theory compared to the one-field effective IR theory.

4.1.3   Numerical Implementation

In the following, we will use piecewise continuous functions to characterise quantities we integrate directly when solving the equations of motion, specifically the field profiles and the auxiliary variables Y, Z, and W. In fact, when integrating the equations of motion, we found that using the same mathematical representation for all physical terms gave the lowest residuals to the strong form of the equations. Intuitively, in this case the auxiliary variables are both constrained to be Laplacians of some function and to satisfy the equations of motion. This double information can be used to stitch together the values of the functions at the vertices correctly.

The situation is different when derivatives are computed in post-processing, without supplemental information from the physics. This is useful, for instance, when computing the fifth force arising from a scalar field, which is proportional to its gradient, or in the computation of the Q_n operators. In this case, the gradient is computed as in Eq. (3.10). To preserve the most mathematical information, we use a discontinuous function space in this case. We have found in the tests described in Appendix B that this choice gives the most accurate results. This is because the approximating polynomials are finite order, so there will be discontinuities in some derivatives of the function. For instance, if linear (i.e. degree 1) Lagrange polynomials are chosen to approximate a field profile, its gradient will be piece-wise constant (and discontinuous).

As the stopping criterion for the Newton iterations, we employ a residual criterion as in Eq. (3.4), together with the norm

$$
\| \mathcal{F} \cdot B_{l,o} \|_\infty \equiv \max \left\{ |\langle \mathcal{F}, B_{l,o} \rangle| : l = 1, \dots, N, o = 0, \dots, p \right\}, \qquad (4.11)
$$

where the $B_{l,o}$'s are the o-th basis function in the l-th element as described in Sec. 3.2. Here N is the number of elements, and p is the order of the interpolating polynomials. This norm can be thought of as the infinity norm $\|V\|_\infty \equiv \max\limits_{i=1,\cdots,Np+1} \{|V_i|\}$ (i.e. maximum absolute component) of a vector whose $(Np+1)$ entries are the scalar products[12] $\langle \mathcal{F}, B_{l,o} \rangle =$

---

[12] The quantity $\mathcal{F} = \int \mathcal{F} v r^2 \mathrm{d}r$ is called the *weak residual*.

$\int \mathcal{F} B_{l,o} r^2 \mathrm{d}r$.

When presenting results, we will also find it helpful to present them in terms of the infinity norm of a (discretised) function f, which is defined to be:

$$\|f\|_\infty \equiv \max_{i=1,\cdots,Np+1}\{|f(r_i)|\} , \tag{4.12}$$

i.e. infinity norm of the vector made up by the function values at mesh nodes.

4.2 Numerical solutions

In Fig. 5, we show φenics's output for model (M1). The first two panels show the π profile near the source (left panel) and throughout the simulation volume (middle panel). The right panel shows the gradient of the scalar field $\nabla_r \pi$, which is proportional to the fifth force. The grey band in Fig. 5 indicates the scale at which the behaviour of the numerical solutions changes, and therefore represents our numerical estimate of the Vainshtein radius of this system. The source radius r_s and Compton wavelength m_π^{-1} are indicated by a dashed red and dashed green line, respectively. As expected, the field reaches its minimum (i.e. maximum excursion) in the middle of the source, then damps slowly within the Vainshtein radius. Outside of the Vainshtein radius, we recover the expected $\sim e^{-m_\pi r}/r$ decay of a massive scalar field. The large Compton wavelength in this model (relative to both the Vainshtein and source radii) mean structures within the source are not detected. Therefore, the field profiles and gradients are very similar for all three source profiles in this case. As shown in Appendix A, these results are applicable for other sources which keep the effective parameters $\bar{m}_\pi \equiv m_\pi r_s$ and $\bar{\epsilon} \equiv \epsilon \left(\frac{M_S}{M_P}\right)^{n-1} (r_s \Lambda)^{1-3n}$ unchanged.
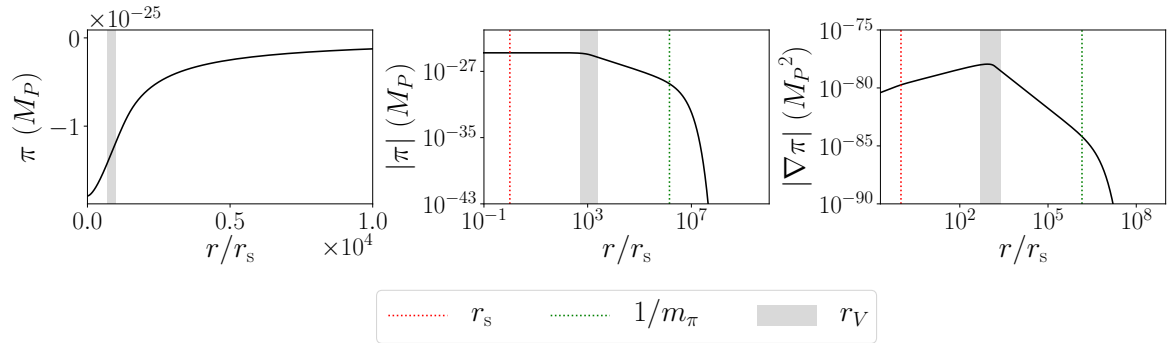

Figure 5. φenics's output for the single-field theory, example model (M1). *Left and middle*: The field profile around the source in linear scale (*left*) and across the box in log-log scale (*middle*). *Right*: The gradient of the field π, which is proportional to the fifth force. The numerical profile was computed using an `ArcTanPowerLaw` mesh with 2300 elements, $k = 14$, $\gamma = 8$, and $r_{\max} = 10^{13} r_s$. The interpolating polynomials were degree 7, and the overall field scale was taken to be $\mu_\pi = 10^{-35} M_P$ (see Appendix A for the definition). The source profile was a smoothed top-hat. For this choice of mass/Compton wavelength $1/m_\pi$, the details of the field profile and gradient are similar to those obtained from the other two source profiles.

In Fig. 6, we show the output for model (M2). As expected, in this case the impact of the source structure is dramatic. In particular, the heavy field profile and its gradient develop complicated features thanks to the smaller Compton wavelength. The oscillations in the field
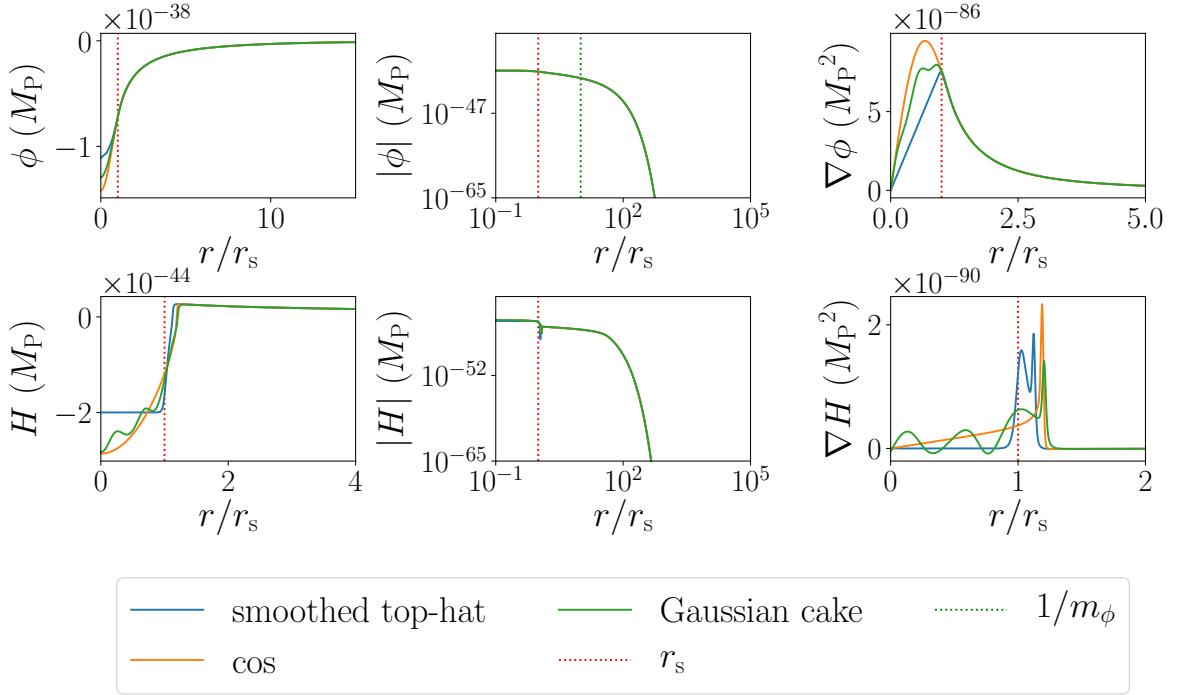
Figure 6. φenics's output for the two-field theory example model (M2), in the presence of a smoothed top-hat, cosine or Gaussian cake source profile. Top panels: ϕ's field profile around the source (left panel) and across the simulation box (middle panel), π's gradient around the source. Bottom panels: Same for field H. These plots have been obtained with the following settings: for all source profiles, `ArcTanExp` mesh with 250 elements before refinement, $a = 5 \times 10^{-2}$, $b = 10^{-2}$, $r_{\max} = 10^9 \, r_{\mathrm{s}}$; degree of interpolating polynomials: 5; $\mu_\phi = 10^{13} M_{\mathrm{P}}$, $\mu_H = 10^{12} M_{\mathrm{P}}$ (see Appendix A for the definition of the numerical parameters μ_ϕ, μ_H). For the smoothed top-hat profile: $k = 20$, additionally twice linearly refined between $r/r_{\mathrm{s}} = 1.05$ and 1.2; for the cos profile: $k = 1$, additionally thrice linearly refined between $r/r_{\mathrm{s}} = 1.1$ and 1.25; for the Gaussian cake profile: $k = 1$, additionally thrice linearly refined between $r/r_{\mathrm{s}} = 1.1$ and 1.3.

profile near the surface of the source arise because H and ϕ are not eigenmodes for linear fluctuations about the vacuum. This mode mixing, combined with the nonlinearity of the system, gives rise to the oscillations observed near the edge of the source.

To better understand these features, we study them using a smoothed top-hat source profile of varying width — the results are in Fig. 7. The features we see in the field profile at distances close to the source radius are induced by the sharp change in density at the surface of the source mass. As seen in Fig. 7, the amplitude of the oscillatory features decreases as the steepness of the density profile is decreased. The characteristic structure of these features evident in Fig. 7, with two peaks in $\nabla_r H$ close to the source radius, is a result of the field H not being an eigenmode of the linear system, and instead a mix of the light and heavy modes. Both modes are excited by the change in density at the source-vacuum transition, and so two fluctuations with different characteristic scales are seen in H. Although not detailed here, we verified that φenics resolves these complicated features accurately using methods analogous to those in Appendix B.

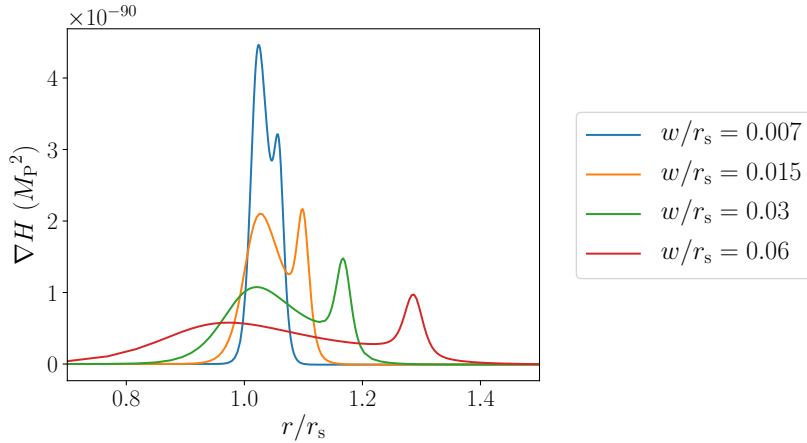For both theories, the behaviour of the terms in the equations of motion Eqs. (2.1)-(2.2)

Figure 7. Behaviour of the heavy field gradient $\nabla_r H$ with increasing source steepness, for a smoothed top-hat profile of varying width w in model (M2). The gradient exhibits features that are a consequence of H not being an eigenstate of the linear operator. This plot has been obtained with the following settings: `ArcTanExp` mesh with 250 elements, $k = 20$, $a = 5 \times 10^{-2}$, $b = 10^{-2}$, $r_{\max} = 10^9 \, r_{\mathrm{s}}$; degree of interpolating polynomials: 5; $\mu_\phi = 10^{13} M_{\mathrm{P}}$, $\mu_H = 10^{12} M_{\mathrm{P}}$ for all models. For $w = 0.007 \, r_{\mathrm{s}}$: additionally thrice linearly refined between $r/r_{\mathrm{s}} = 1$ and 1.1. For $w = 0.015 \, r_{\mathrm{s}}$: additionally twice linearly refined between $r/r_{\mathrm{s}} = 1.05$ and 1.15. For $w = 0.03 \, r_{\mathrm{s}}$: additionally twice linearly refined between $r/r_{\mathrm{s}} = 1.13$ and 1.2. For $w = 0.06 \, r_{\mathrm{s}}$: additionally twice linearly refined between $r/r_{\mathrm{s}} = 1.2$ and 1.35.
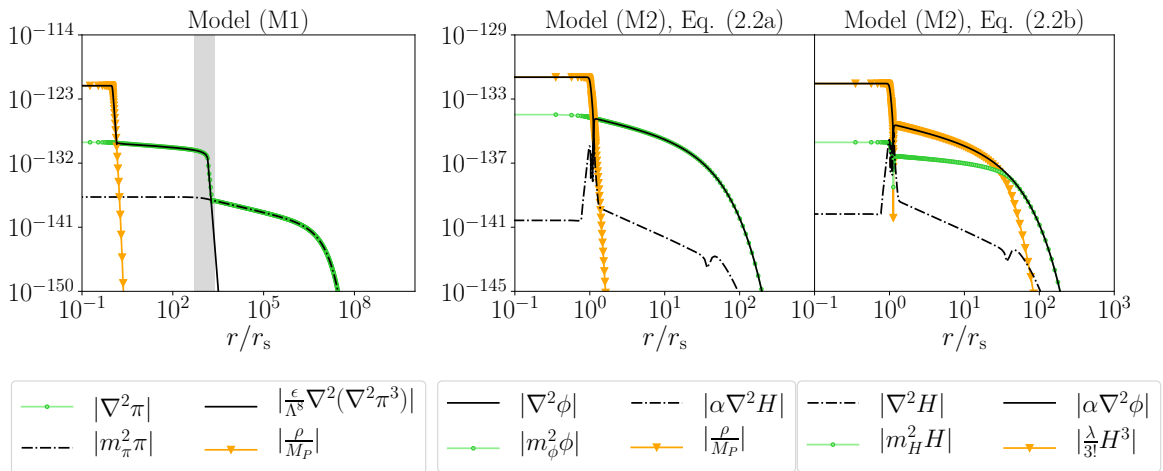
Figure 8. Left panel: terms in the equation of motion Eq. (2.1) for the single-field theory, model (M1). The grey shaded area indicates the Vainshtein radius. Middle and right panels: terms in the equations of motion Eq. (2.2) for the two-field theory, model (M2). For all panels, the vertical axis is in units M_{P}^3. This plot has been obtained with the same settings as those in Figure 5 and 6 for a smoothed top-hat source of width $w = 0.02 \, r_{\mathrm{s}}$.

across the radial domain can be seen in Figure 8. For readability, we only plot them for a top-hat source of width $w = 0.02 \, r_{\mathrm{s}}$. The qualitative behaviour of the terms outside of the source does not change for the other source profiles.

4.3 Comparison with theoretical expectations

We now compare the results obtained with φenics against theoretical expectations for the behaviour of these theories. As discussed earlier the analytic solutions to these theories' equations of motion are not known in general, but in certain limits and for certain parameter values a useful comparison is possible.

4.3.1 Comparison of profiles and equations of motion

When considering the equations of motion for the theories of interest, it is easiest to understand the resulting behaviour of the field(s) if certain terms in the equation of motion dominate over others. For example, in the single-field theory (Eq. (2.1)) if the dominant terms are the Laplacian of π and the mass term $m_\pi^2 \pi$, then we can expect that the field π has either exponential growth or decay. Our analytic intuition for the behaviour of the solutions to the equations of motion is also informed by what we know about the boundary conditions for the theory. In particular, the field evolves to a constant value, its vacuum expectation value, far away from the compact source. From this we can conclude that there is a region of space sufficiently far away from the source, where 'sufficiently' remains to be determined, within which the field value is close to its vacuum expectation value and gradients are small. In this regime, in the single-field theory, the dominant terms in the equation of motion are the Laplacian and mass terms, and because we know the boundary condition at infinity we can conclude that far from the source we have

$$\pi \sim \frac{1}{r} e^{-m_\pi r} \,. \tag{4.13}$$

We can further anticipate the existence of other regimes with different behaviour. As the gradients of the field grow as we approach the source, the term in the equation of motion containing higher powers of derivatives could dominate over the Laplacian of the field π. Inside the source, the term proportional to the density of the source will be significant.

For the single-field theory and model (M1), Fig. 8 shows the different terms in the equation of motion across the radial domain. We observe three regimes, which are separated by the source radius r_s and the Vainshtein radius $r_V \sim 8.6 \times 10^2 \, r_s$. Inside the source, the nonlinear (solid black line) and source (orange line with triangles) terms dominate, tracking each other. Outside of the source but well inside the Vainshtein radius, the source term decays rapidly and the kinetic terms become dominant, as can be seen by the nonlinear (solid black line) and Laplacian (green line with circles) terms. Outside of the Vainshtein radius, the Laplacian (green line with circles) tracks the mass term (black dot-dashed line), and the field's behaviour is described by Eq. (4.13). As a result a further change in the behaviour of the solution is seen at the Compton wavelength m_π^{-1}. For $r < m_\pi^{-1}$ we see a power-law decay of the field, and for $r > m_\pi^{-1}$ we see exponential decay.

This behaviour can be understood as follows. Inside the Vainshtein radius, we can neglect the mass term $-m_\pi^2 \pi$, so that the equation of motion Eq. (2.3) can be approximated as $\nabla_r^2 \pi - \frac{\epsilon}{\Lambda^{3n-1}} \nabla_r^2 (\nabla_r^2 \pi^n) \sim \frac{\rho}{M_P}$ and integrated once. Within the Vainshtein radius we also expect that the nonlinear term will be the dominant term driving the evolution of the field. Inside the source, we therefore neglect the Laplacian of π, which reduces the equation of motion to the Poisson-like equation $-\frac{\epsilon}{\Lambda^{3n-1}} \nabla_r^2 (\nabla_r^2 \pi^n) \sim \frac{\rho}{M_P}$ for $r \lesssim r_V$. This is in agreement

with Fig. 8, where the nonlinear and source terms overlap within the source. In analogy with gravitation (except for a sign), we expect $(\nabla_r^2 \pi)^3$ to reach its maximum at the origin and then decay as $(\nabla_r^2 \pi)^3 \sim 1/r$ outside of the source, as long as we remain inside the Vainshtein radius. We indeed see in Fig. 8 that the Laplacian scales as $\nabla_r^2 \pi \sim 1/\sqrt[3]{r}$ in the region $r_S < r < r_V$. Finally, outside the Vainshtein radius, the equation of motion takes the linear form, and we expect the field to decay as $\sim e^{-m_\pi r}/r$.

For the two-field theory with a step-function source, an exact solution for the linear (i.e. $\lambda = 0$) case was given in Ref. [65] (see Eqns. (4.23)-(4.26) within). This provides a useful test case for our numerical solution obtained with φenics. Although this comparison does not test the nonlinear solver — as it requires only a single iteration of a linear solver — we can make consistency checks. Importantly, it allows us to check that our grid can capture sharp features. In particular, numerical solutions obtained for a smoothed top-hat distribution Eq. (2.10) should converge to those corresponding to a step-function source as the width of the source-vacuum transition relative to r_s is taken to zero.[13]

For this comparison we will use model (M3) with $\lambda = 0$. For every choice of r_{max}, there is a radius after which the solution will be sensitive to the finiteness of the box. For the purposes of this test we use $r_{max} = 10^9 r_s$ and compare the value of our numerical solutions for radii $r < 10^5 r_s$.

Figure 9 shows the maximum relative difference between φenics's solution and the analytical result, for both fields. We have used the sequence of source widths:

$$wr_s^{-1} = \{1, 1.6, 2.5, 4, 6.2, 10, 15.3, 24.2, 38, 60\} \times 10^{-3} \,, \tag{4.14}$$

which were chosen to be approximately equispaced logarithmically. As expected, we see the numerical solution approach the analytic solution (with $w = 0$) as we decrease the width of the boundary layer connecting the source to the vacuum. For all models, we observe that the maximum error is localised at the origin, where the field reaches its minimum, and sharply decreases outside of the source.

4.3.2   Operators Q_n

Above we demonstrated the existence of various dynamical regimes, each characterised by a different combination of dominant terms in the equations of motion. In the single-field case, in particular, the nonlinear term was important within the Vainshtein radius r_V. To verify that the low-energy effective theory is valid in this nonlinear regime, we must evaluate further higher-order corrections on these solutions, which have the form of the Q_n's in Eq. (2.14)-(2.15). If the size of these terms is comparable to those we have kept in the equation of motion, then our effective field theory breaks down.[14] For the case $n = 3$ considered here, only $Q_3^{(single-field)}$ is directly evaluated by the integrator. The others are obtained in post-processing. Given the highly nonlinear nature of these operators, and the shift of power to smaller scales induced by the gradients, it is important to verify that this post-processing produces accurate results. Here we will verify that our post-processed Q_n's agree

---

[13]Note that this test would not allow checking the accuracy of the features displayed in Figure 7 as they are not present for $\lambda = 0$.

[14]It should be noted that the leading Q_n operators being small, for n finite, is a necessary though not sufficient condition for the operator series to converge, and therefore for the theory to be predictive.
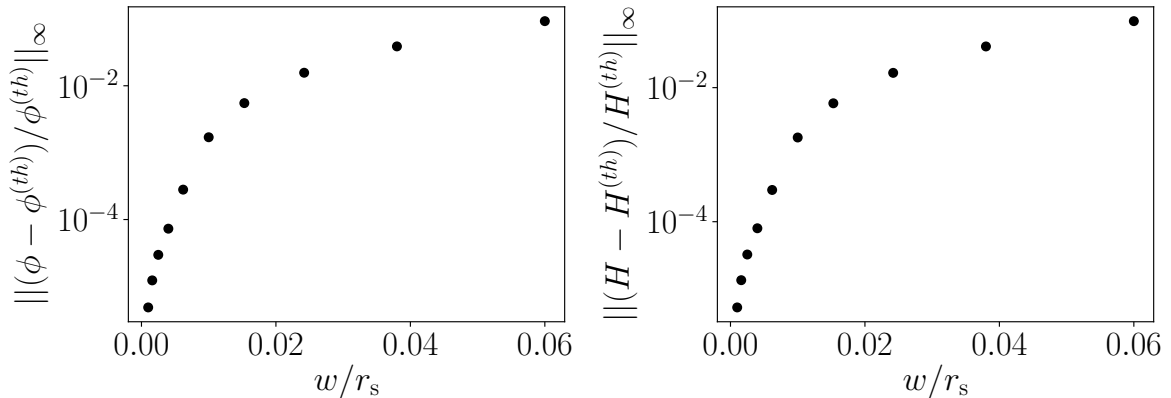
Figure 9. Infinity norm (i.e. maximum absolute value at mesh nodes, whether mesh vertices or intermediate collocation points) of the relative error between the solution obtained with φenics for the model (M3) ($\lambda = 0$), and the analytical result of Ref. [65], Eqns. (4.23)-(4.26). The left and right panels show, respectively, the results for ϕ and H. For all solutions, we used the following settings: `ArcTanExp` mesh with 800 elements, $k = 7$, $a = 10^{-4}$, $b = 0$, $r_{\max} = 10^9 \, r_{\mathrm{s}}$; degree of interpolating polynomials: 5.

with approximate analytic calculations motivated from the dynamical behaviour observed above.

In particular, we will consider the following operators:

$$O_p^{(\text{single}-\text{field})} \equiv (-1)^{p+1} \binom{3p}{p} \frac{\epsilon^p}{2p+1} \nabla_r^2 \left((\nabla_r^2 \pi)^{2p+1} \right) \Lambda^{-6p-2} \,, \qquad (4.15)$$

for the single-field theory and

$$O_p^{(\text{two}-\text{field})} \equiv (-1)^{p+1} \binom{3p}{p} \frac{\alpha^{2p+2}}{2p+1} \left(\frac{\lambda}{3!} \right)^p \nabla_r^2 \left((\nabla_r^2 \phi)^{2p+1} \right) m_H^{-6p-2} \,, \qquad (4.16)$$

for the two-field theory. These two types of operator have the same functional form, and they differ only in the dependence of the constant prefactor on the parameters of the two theories. These operators are motivated by the connection between the single-field theory with $n = 3$ and the low-energy ($E \ll m_H$) limit of the two-field theory, where higher-order operators of this type may be expected to arise [58]. Note, in particular, the correction O_p is related to our original operators Q via $O_p \propto Q_{2p+1}$.

Single-Field Theory: We first derive predictions for the single-field theory based on the dynamical behaviour outlined above. For the single-field theory, we distinguish between the nonlinear regime ('NL' in the following) at scales of $0 < r < r_V$ (combining the nonlinear behaviour of the field inside and outside the source) and Yukawa suppression ('YS') for $r > r_V$. As discussed above, the term nonlinear in π dominates the evolution of the field inside the Vainshtein radius, which yields the following result for $W \equiv (\nabla_r^2 \pi)^n$:

$$W_{\mathrm{NL}} = \frac{1}{4\pi\epsilon} \frac{\Lambda^{3n-1} M_{\mathrm{S}}}{\bar{t} r_{\mathrm{s}} M_{\mathrm{P}}} \begin{cases} \dfrac{1}{2} \left(3 - \dfrac{r^2}{\bar{t}^2 r_{\mathrm{s}}^2} \right), & \text{for } 0 < r \leq \bar{t} r_{\mathrm{s}} \\[4mm] \dfrac{\bar{t} r_{\mathrm{s}}}{r}, & \text{for } \bar{t} r_{\mathrm{s}} \leq r \leq r_{\mathrm{V}} \end{cases} \,. \qquad (4.17)$$

From this result, we obtain the Laplacian $(\nabla_r^2 \pi)_{NL} \equiv W_{NL}^{1/n}$, which is subsequently used to evaluate O_p:

$$O_{p,\mathrm{NL}} = \frac{\Lambda^{1-\beta}}{\bar{t}^2 r_{\mathrm{s}}^2} 2 c_p \epsilon^{p-\beta} \beta \left(\frac{M_{\mathrm{S}}}{8\pi \bar{t} r_{\mathrm{s}} M_{\mathrm{P}}}\right)^\beta \begin{cases} \left(3 - \dfrac{r^2}{\bar{t}^2 r_{\mathrm{s}}^2}\right)^{\beta-2} \left((1+2\beta)\dfrac{r^2}{\bar{t}^2 r_{\mathrm{s}}^2} - 9\right) & \text{for } 0 < r \le \bar{t} r_{\mathrm{s}}\,, \\[12pt] 2^{\beta-1}(\beta-1)\left(\dfrac{\bar{t} r_{\mathrm{s}}}{r}\right)^{\beta+2} & \text{for } \bar{t} r_{\mathrm{s}} \le r \le r_V\,, \end{cases}$$
$$(4.18)$$

where we have defined $\beta \equiv (2p+1)/n$ and $c_p \equiv (-1)^{p+1}\binom{3p}{p}\frac{1}{2p+1}$. Outside of the Vainshtein radius, the field decays as $\pi \approx \pi(r_{\mathrm{piv}})\frac{r_{\mathrm{piv}}}{r}e^{-m_\pi(r-r_{\mathrm{piv}})}$, where $r_{\mathrm{piv}} > r_V$ is some pivot scale. The undetermined normalisation is calibrated to simulations by computing $a_{\mathrm{piv}} \equiv \pi(r_{\mathrm{piv}})r_{\mathrm{piv}}e^{m_\pi r_{\mathrm{piv}}}$, which is nearly independent of the choice of r_{piv}, provided $r_{\mathrm{piv}} > r_V$. We obtain:

$$O_{p,\mathrm{YS}} = \Lambda m_\pi^2 c_p (2p+1)\epsilon^p \left(a_{\mathrm{piv}}\frac{m_\pi^3}{\Lambda^3}\frac{e^{-m_\pi r}}{m_\pi r}\right)^{2p+1} \left(2p+1 + \frac{4p}{m_\pi r} + \frac{2p}{m_\pi^2 r^2}\right). \qquad (4.19)$$

We compare these analytic single-field predictions to the post-processing calculations of φenics in the left panel of Fig. 10. We see excellent agreement with the nonlinear solution for $r < r_{\mathrm{v}}$, and with the Yukawa suppressed solution for $r > r_V$. The Vainshtein radius r_V provides the radius at which the two solutions match onto each other, which is a second parameter that must be obtained from simulations. The apparent discrepancy at $r = r_{\mathrm{s}}$ is due to our source profile being a smoothed top-hat profile instead of a step function.

Two-Field Theory: For the two-field theory, similarly detailed predictions are not available in all regimes. Instead, we will make consistency checks by comparing the numerical value of the post-processed operators with analytic approximations based on the dynamical field behaviour observed in Fig. 8. The middle panel of that figure shows that the Laplacian $\nabla_r^2 \phi$ tracks either the source (for $r \lesssim r_{\mathrm{s}}$) or mass term (for $r \gtrsim r_{\mathrm{s}}$) of the light field almost everywhere in the simulation volume. The exception is a small region just outside of the source radius ($r \sim r_{\mathrm{s}}$), where the heavy-field term $\alpha \nabla_r^2 H$ briefly dominates. Based on this observation, we expect the operator O_3 to also demonstrate these different regimes of behaviour. To fix terminology, we refer to these two regimes as 'light-field dominated' and 'heavy-field dominated', respectively.

In the light-field dominated regime, we distinguish three cases:

- *source dominated*: Well inside the source, the Laplacian $\nabla_r^2 \phi$ tracks the source term, which is approximately constant for our smoothed top-hat profile. In the same region, the heavy-field term $\alpha \nabla_r^2 H$ is negligible. The mass term $m_\phi^2 \phi$ is small compared to the source term, but varies more rapidly, so it cannot be neglected. We can then write $\nabla_r^2 (\nabla_r^2 \phi)^{2n+1} \sim \nabla_r^2((\frac{\rho}{M_{\mathrm{P}}} + m_\phi^2 \phi)^{2n+1})$. Because the source density, $\rho \approx \rho_0 \equiv \rho(r=0)$, is approximately constant around the origin, we can write:

$$\nabla_r^2\left(\left(\frac{\rho}{M_{\mathrm{P}}} + m_\phi^2 \phi\right)\right)^{2n+1} \approx \sum_{q=0}^{2n+1} \binom{2n+1}{q}\left(\frac{\rho_0}{M_{\mathrm{P}}}\right)^{2n+1-q} m_\phi^{2q}\,\nabla_r^2(\phi^q)\,. \qquad (4.20)$$
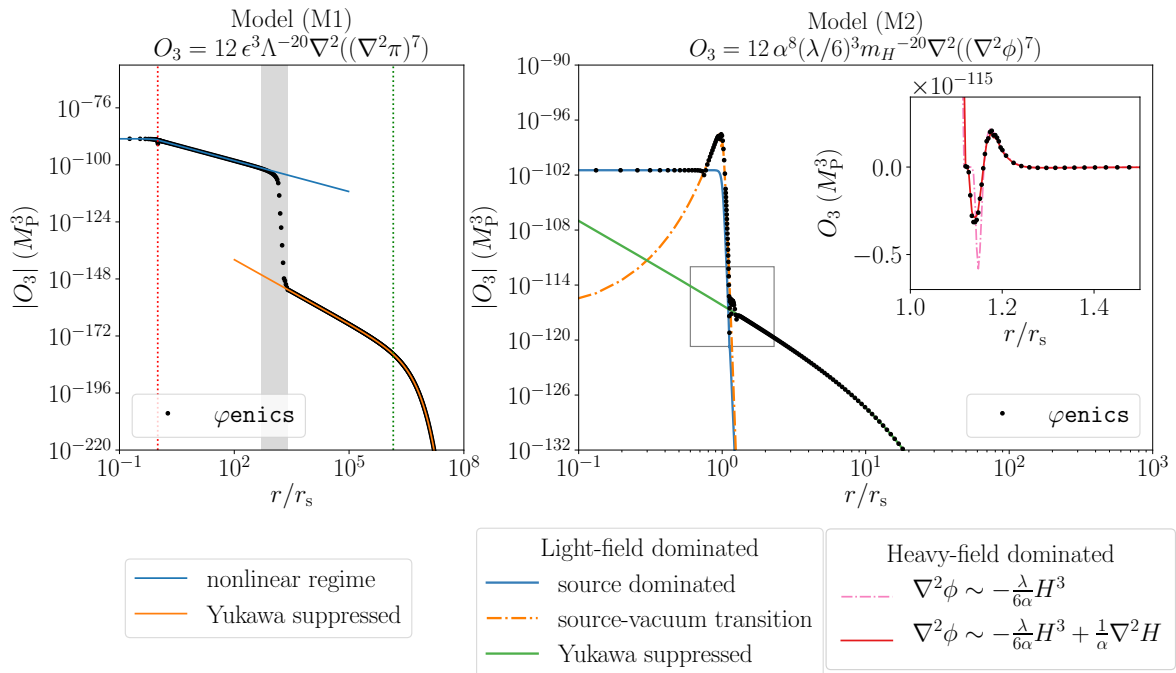
Figure 10. Comparison of φenics's prediction for the operator O_3 defined in Eq. (4.15)-(4.16) for the single-field (left) and two-field (right) theory example models (M1) and (M2). In the single-field theory, we recover the nonlinear regime and the Yukawa suppression, matching between them. The red and green vertical lines indicate, respectively, the source radius and the Compton wavelength of the field π, whereas the grey band highlights the Vainshtein radius. The rich phenomenology of the two-field theory is reflected in the many regimes that characterise it. The light field ϕ dominates across most of the radial domain, where we can compare against three regimes (defined in the text): source domination, source-vacuum transition and Yukawa suppression. In a small region just outside the source radius, indicated with the grey box, the heavy field dominates. We show this (in linear scale) in the inset, where we display the leading term contributing to the operator (i.e. the nonlinear term $\lambda H^3/(6\alpha)$) as well as the next-order correction, i.e. the Laplacian $\alpha^{-1}\nabla_r^2 H$. This plot was obtained with the following settings: `ArcTanExp` mesh with 250 elements before refinement, $k = 3, a = 5 \times 10^{-2}, b = 10^{-2}, r_{\max} = 10^9 \, r_{\mathrm{s}}$, twice linearly refined between $r/r_{\mathrm{s}} = 1.05$ and $r/r_{\mathrm{s}} = 1.2$; degree of interpolating polynomials: 5; $\mu_\phi = 10^{13} M_{\mathrm{P}}$, $\mu_H = 10^{12} M_{\mathrm{P}}$. Also for the single-field theory, the normalisation constant a_{piv} is computed at a pivot radius $r_{\mathrm{piv}} = 10^4 \, r_{\mathrm{s}}$, and for the two-field theory at $r_{\mathrm{piv}} = 10 \, r_{\mathrm{s}}$.

For the term $\nabla_r^2(\phi^q)$ in this expression, we obtain $q\phi^{q-2}(\phi\nabla_r^2\phi + (q-1)(\nabla_r\phi)^2)$. Knowing a priori which of the last terms dominates in $\nabla_r^2(\phi^q)$ is challenging. We evaluate them with φenics using the dimensionless units of Appendix A, which are better suited to this kind of comparison. Due to the Neumann boundary condition, we expect $|\phi\nabla_r^2\phi| \gg (\nabla_r\phi)^2$ around the origin and numerically we see that this is true. We can therefore rewrite the right-hand side of Eq. (4.20) as $\left(\frac{\rho}{M_{\mathrm{P}}}\right)^{2n+1} \frac{\nabla^2\phi}{\phi} \sum q\binom{2n+1}{q} \left(\frac{m_\phi^2\phi}{\rho/M_{\mathrm{P}}}\right)^q$. We can see from Fig. 8 that, for model (M2) and around the origin, the mass term is much smaller than the source term, i.e. $m_\phi^2\phi \ll \rho/M_{\mathrm{P}}$. Therefore, we expect the $q = 1$

term to dominate, and we have:

$$O_p \approx c_p(2p+1)\alpha^{2p+2}\left(\frac{\lambda}{6}\right)^p \left(\frac{\rho_0}{M_{\mathrm{P}}m_H^3}\right)^{2p}\left(\frac{m_\phi}{m_H}\right)^2 \nabla_r^2\phi \ . \tag{4.21}$$

- *source-vacuum transition*: At the source-vacuum transition r_{s}, the operator O_p is instead dominated by the rapid variation of the source profile. This contribution can be computed exactly from Eq. (2.10), leading to the approximation:

$$O_p \approx \frac{m_H}{w^2}c_p(2p+1)\alpha^{2p+2}\left(\frac{\lambda}{6}\right)^p\left(\frac{\rho(r)}{M_{\mathrm{P}}m_H^3}\right)^{2p+1}\frac{e^{-\frac{r-\bar{t}r_{\mathrm{s}}}{w}}}{\left(1+e^{-\frac{r-\bar{t}r_{\mathrm{s}}}{w}}\right)^2}\times$$
$$\times \left(e^{-\frac{r-\bar{t}r_{\mathrm{s}}}{w}}\left(1+2p-\frac{2w}{r}\right)-1-\frac{2w}{r}\right) \ . \tag{4.22}$$

- *Yukawa suppressed*: Similarly to the single-field theory, by evaluating the normalisation $a_{\mathrm{piv}} \equiv \phi(r_{\mathrm{piv}})r_{\mathrm{piv}}e^{m_\phi r_{\mathrm{piv}}}$ at some pivot scale r_{piv} in the regime where ϕ obeys the linearised equation of motion, we obtain:

$$O_p = c_p(2p+1)\left(\frac{\lambda}{6}\right)^p\alpha^{2p+2}m_H m_\phi^2\left(a_{\mathrm{piv}}\frac{m_\phi^3}{m_H^3}\frac{e^{-m_\phi r}}{m_\phi r}\right)^{2p+1}$$
$$\times \left((2p+1)+\frac{4p}{m_\phi r}+\frac{2p}{m_\phi^2 r^2}\right) \ . \tag{4.23}$$

In the heavy-field dominated regime, we perform consistency checks between different quantities computed by φenics. In this regime, it is helpful to express the Laplacian $\nabla_r^2\phi$ in terms of the equation of motion of the heavy field: $\nabla_r^2 H = -\frac{\lambda}{6\alpha}H^3 + \frac{1}{\alpha}\nabla_r^2 H - m_H^2 H$. We can see in Fig. 8 that, at $r \sim r_{\mathrm{s}}$, the Laplacian $\nabla_r^2\phi$ tracks the nonlinear term $\frac{\lambda H^3}{6\alpha}$, which is therefore expected to have a large contribution. Although the Laplacian $\nabla_r^2 H/\alpha$ is strongly subdominant in that regime, it has a large gradient, so we also expect it to have a large contribution on the operator O_n. The mass term instead does not show large gradients and is subdominant compared to the nonlinear term, so that we can expect it to generate the smallest correction. A consistency check can be therefore obtained by comparing the φenics numerical result for O_n against the increasingly accurate approximations $\nabla_r^2\phi \approx -\frac{\lambda}{6\alpha}H^3$, $\nabla_r^2\phi \approx -\frac{\lambda}{6\alpha}H^3 + \frac{1}{\alpha}\nabla_r^2 H$ and $\nabla_r^2\phi = -\frac{\lambda}{6\alpha}H^3 + \frac{1}{\alpha}\nabla_r^2 H - m_H^2 H$. This is obtained by evaluating Eq. (3.27) after having set Y to coincide with the desired function (e.g. $Y = \frac{\lambda}{6\alpha}H^3$), also using φenics. The inset of Fig. 10 compares the φenics postprocessing result for O_p against the leading contributions coming from the heavy-field dominated regime: the nonlinear term and the Laplacian $\nabla_r^2 H$. The mass term induces a further, very small correction that we do not show for simplicity.

5 Summary and Outlook

In this work we have introduced a new finite-element code, φenics, which computes the behaviour of nonlinear scalar field theories around isolated, spherical compact objects. The

motivation for introducing these fields comes from attempts to modify the theory of gravity, and understand the nature of dark energy, whilst still being compatible with local tests of gravity. In order to robustly test these theories against observations, precision predictions for their phenomenology are needed. It is to facilitate this that we present the φenics code.

As scalar field theories with screening are, by construction, nonlinear theories, numerical simulation can be challenging. Additional difficulties come from the large hierarchies of scale in the problems we wish to solve and the presence of higher-derivative operators in the equations of motion. In this work, we have demonstrated the power of the finite element method to address these problems. We show explicitly how the code works for two very different nonlinear scalar field theories: one with multiple fields, and the other with a single field, but relevant higher order derivatives. We have also shown that φenics can solve for the behaviour of the fields around different source density profiles. This includes profiles that approximate a step function, and profiles where the density does not scale monotonically with distance.

For both example theories, we find that we can accurately simulate the behaviour of the scalar field across all dynamical regimes, including deep inside the source object, regions where nonlinear and higher derivative terms dominate, as well as the Yukawa suppression far from the source due to the Compton wavelength of the scalar field. Where analytic predictions are available, we confirm that φenics reproduces the expected scalar field profile. We also found cases where, despite the lack of a full analytic solution for the profile, we were able to analytically estimate the scaling behaviour of the solution across the various regimes of behaviour, using our numerical results as a guide. We again found that φenics agrees with all expectations. An illustration of this excellent agreement is shown in Fig. 10. We have also been able to identify new behaviours, such as oscillations in the scalar field profile at the surface of the source due to a rapid change in the source density.

This paper lays the ground work for future applications of the finite element method to more complex geometries, which will allow us to explore the behaviour of screened scalar fields around asymmetrical objects. Currently φenics assumes the sources are static and spherically symmetric. In future work we will aim to relax these two assumptions, allowing for arbitrary spatial source shapes and time-dependence, with the corresponding aspherical and time-dependent field profiles. This will allow us to study the full spectrum of behaviour for screened scalar theories, providing new opportunities to detect or constrain these interesting cosmological theories.

φenics is publicly available from https://github.com/scaramouche-00/phienics and can be easily adapted for other theories, by modifying the relevant module as described in Appendix C.

Acknowledgments

Appendices

A   Dimensionless equations

In this appendix we present the dimensionless variables and corresponding equations of motion used in our code. Throughout, we work in units with $\hbar = c = 1$ and assume three spatial dimensions, so that time and space have units of inverse mass, and fields have units of mass. To fix notation, we denote dimensionless units by an overbar $\bar{\cdot}$. We introduce a length scale κ to define the dimensionless radius $\bar{r} = \kappa^{-1} r$, and mass scales μ_i to define dimensionless fields $\bar{\pi} = \mu_\pi^{-1} \pi$, $\bar{H} = \mu_H^{-1} H$, and $\bar{\phi} = \mu_\phi^{-1} \phi$. It is further convenient to express the source energy density $\rho = \frac{M_{\rm S}}{r_{\rm s}^3} f\left(\frac{r}{r_{\rm s}}\right)$ in terms of a dimensionless $\mathcal{O}(1)$ function $f(x)$ which vanishes for $x \gtrsim 1$. For the profiles used in this paper, the function $f(x)$ is shown in Fig. 1. For generality, we keep these choices arbitrary for the moment, with the specific choices made in φ`enics` indicated below.

In dimensionless units, the equations of motion are:

1. Single-field theory:

$$\bar{\nabla}_r^{\,2} \bar{\pi} - (m_\pi \kappa)^2\, \bar{\pi} - \epsilon \left(\mu_\pi \kappa\right)^{n-1} \left(\kappa \Lambda\right)^{1-3n} \bar{\nabla}_r^{\,2} \left[\left(\bar{\nabla}_r^{\,2} \bar{\pi}\right)^n\right] = \frac{M_{\rm S}}{M_{\rm P}} \left(\frac{\kappa}{r_{\rm s}}\right)^3 \frac{1}{\mu_\pi \kappa} f\left(\frac{\kappa}{r_{\rm s}} \bar{r}\right).$$

$$\tag{A.1}$$

   The single-field theory displays highly nonlinear behaviour: as a guide to obtaining optimum choice of units, it is then helpful to consider the deep nonlinear regime of Eq. (4.7), which gives the initial guess for the Newton algorithm. In dimensionless units, it is:

$$-\epsilon \left(\mu_\pi \kappa\right)^{n-1} \left(\kappa \Lambda\right)^{1-3n} \bar{\nabla}_r^{\,2} \left[\left(\bar{\nabla}_r^{\,2} \bar{\pi}\right)^n\right] = \frac{M_{\rm S}}{M_{\rm P}} \left(\frac{\kappa}{r_{\rm s}}\right)^3 \frac{1}{\mu_\pi \kappa} f\left(\frac{\kappa}{r_{\rm s}} \bar{r}\right).$$

$$\tag{A.2}$$

   By setting $\kappa = r_{\rm s}$ and $\mu_\pi \kappa = (\kappa \Lambda)^{(3n-1)/n} (M_{\rm S}/M_{\rm P})^{1/n}$, one obtains the Poisson equation $-\epsilon \bar{\nabla}_r^{\,2} \left[\left(\bar{\nabla}_r^{\,2} \bar{\pi}\right)^n\right] = f(\bar{r})$, which is particularly suited to solve for $\bar{W} \equiv \left(\bar{\nabla}_r^{\,2} \bar{\pi}\right)^n$.

However, to better match between the nonlinear and linear regimes of the single-field theory, it is helpful to include information on the Galileon mass m_π in the rescaling μ_π. This is the default choice in φenics, where we set:

$$\mu_\pi = (\kappa\Lambda)^{(3n-1)/n}\,(M_{\rm S}/M_{\rm P})^{1/n}m_\pi\,, \tag{A.3}$$

together with $\kappa = r_{\rm s}$. With this choice of units, Eq. (A.1) can be now expressed as:

$$\bar{\nabla}_r{}^2\bar{\pi} - (m_\pi r_{\rm s})^2\,\bar{\pi} - \tilde{\epsilon}\bar{M}\bar{\nabla}_r{}^2\left[\left(\bar{\nabla}_r{}^2\bar{\pi}\right)^n\right] = \bar{M}f\left(\bar{r}\right)\,, \tag{A.4}$$

where we have defined $\tilde{\epsilon} = \epsilon(m_\pi r_{\rm s})^n$ and $\bar{M} \equiv \left(\frac{M_{\rm S}}{M_{\rm P}}\right)^{1-\frac{1}{n}}(r_{\rm s}\Lambda)^{-3+\frac{1}{n}}(m_\pi r_{\rm s})^{-1}$.

For a given source profile f, the solution is determined (up to global scalings) by the two dimensionless parameters $\bar{m}_\pi \equiv m_\pi r_{\rm s}$ and $\bar{\epsilon} \equiv \epsilon\left(\frac{M_{\rm S}}{M_{\rm P}}\right)^{n-1}(r_{\rm s}\Lambda)^{1-3n}$, which we see by choosing $\kappa = r_{\rm s}$ and $\mu_\pi\kappa = M_{\rm S}/M_{\rm P}$, resulting in

$$\bar{\nabla}_r{}^2\bar{\pi} - \bar{m}_\pi^2\bar{\pi} - \bar{\epsilon}\bar{\nabla}_r{}^2\left[\left(\bar{\nabla}_r{}^2\bar{\pi}\right)^n\right] = f\left(\bar{r}\right)\,. \tag{A.5}$$

However, this form of the equation is normally numerically less-suited to the computation of the solution, as $\bar{\epsilon}$ can become very large.

2. Two-field theory:

$$\begin{cases} \bar{\nabla}_r{}^2\bar{\phi} - (m_\phi\kappa)^2\,\bar{\phi} - \alpha\left(\frac{\mu_H}{\mu_\phi}\right)\bar{\nabla}_r{}^2\bar{H} = \dfrac{M_{\rm S}}{M_{\rm P}}\left(\dfrac{\kappa}{r_{\rm s}}\right)^3\dfrac{1}{\mu_\phi\kappa}f\left(\dfrac{\kappa}{r_{\rm s}}\bar{r}\right)\,, & \text{(A.6a)} \\[2mm] \bar{\nabla}_r{}^2\bar{H} - (m_H\kappa)^2\,\bar{H} - \alpha\left(\frac{\mu_\phi}{\mu_H}\right)\bar{\nabla}_r{}^2\bar{\phi} - \dfrac{\lambda}{3!}\left(\mu_H\kappa\right)^2\bar{H}^3 = 0\,. & \text{(A.6b)} \end{cases}$$

To normalise to the source profile, it is convenient to choose $\kappa = r_{\rm s}$, $\mu_\phi\kappa = \frac{M_{\rm S}}{M_{\rm P}}$, and $\mu_H = |\alpha|\mu_\phi$,[15] resulting in

$$\begin{cases} \bar{\nabla}_r{}^2\bar{\phi} - (m_\phi r_{\rm s})^2\,\bar{\phi} - \text{sgn}(\alpha)\alpha^2\bar{\nabla}_r{}^2\bar{H} = f\left(\bar{r}\right) & \text{(A.7a)} \\[2mm] \bar{\nabla}_r{}^2\bar{H} - (m_H r_{\rm s})^2\,\bar{H} - \text{sgn}(\alpha)\bar{\nabla}_r{}^2\bar{\phi} - \dfrac{\lambda}{3!}\alpha^2\left(\dfrac{M_{\rm S}}{M_{\rm P}}\right)^2\bar{H}^3 = 0\,. & \text{(A.7b)} \end{cases}$$

This time, for a given source profile f, the solution is determined (up to global scalings) by four dimensionless parameters: $m_\phi r_{\rm s}$, $m_H r_{\rm s}$, α, and $\frac{\lambda}{3!}\alpha^2\frac{M_{\rm S}^2}{M_{\rm P}^2}$.

For both theories, φenics's default choices are a helpful guide to obtaining optimal units that allow for convergence in a physical model of interest. However, for every specific model, setting the quantities μ_π, μ_H, μ_ϕ 'by hand' may give the best results.

Before being solved, the rescaled equations of motion are expressed in the weak form detailed in Sec. 3.3 and 4.1.

---

[15]The choice $\mu_H = |\alpha|\mu_\phi$ ensures the dominant term in the second order equations of motion have similar scales, which makes residuals easier to interpret.

B Numerical Convergence Tests

In this appendix, we demonstrate the numerical convergence properties of our solver. In particular, we provide independent tests of both our iterative Newtonian solver and the convergence of the solution with varying spatial discretisations. For concreteness, we focus on the UV complete model (M3). In the following, we will solve the field profiles around a smoothed top-hat source profile of width $w = 0.02\,r_{\mathrm{s}}$.

B.1 Convergence of Newton iterations

We start with convergence of our Newton iterations. In Fig. 11, we show the behaviour of the error in the solution at every iteration: we quantify the error both in terms of the change in the solution (left panel, also see Eq. 3.5) and residual (right panel and Eq. 3.4). In this run, tolerances have been adjusted to exceedingly small values so that the full behaviour of the error could be observed. The norms in Fig. 11 are as defined in Eqns. (4.11)-(4.12).

As expected, the residual decreases with each iteration (in this case exponentially), eventually reaching a plateau when the scheme reaches its accuracy limits. Interestingly, in this case we find an intermediate regime where the required field perturbation temporarily increases with iteration, while the residual always decreases. We typically find a larger number of iterations is required as λ is increased. This is easy to understand when considering that, for the two-field model, we take the $\lambda = 0$ solution as our initial guess. As a result, the "distance" between the initial guess and the converged solution for $\lambda \neq 0$ increases as the nonlinear parameter λ is turned on. For the single-field model, on the other hand, our initial guess is instead based on assuming the nonlinear term dominates, and as a result it typically requires fewer iterations as the parameter ϵ is increased.
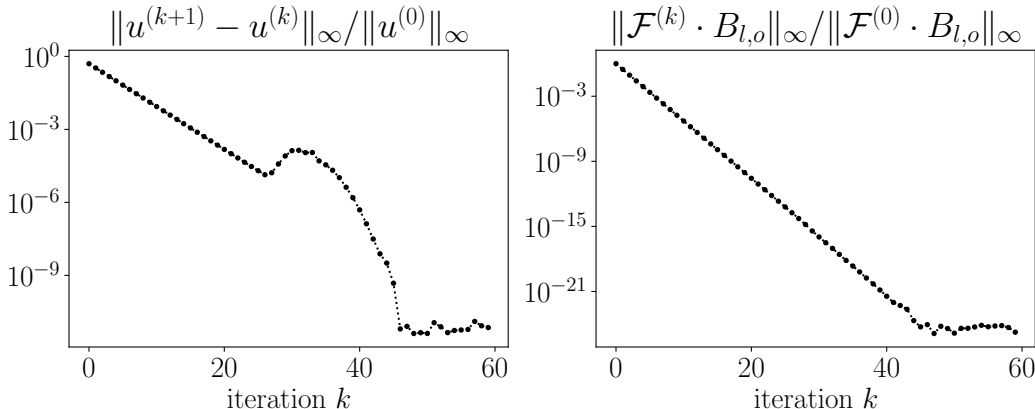

Figure 11. The error in the solution as a function of iteration of the Newton solver, expressed as change in the solution (*left panel*) and residuals of the equation (*right panel*), for model (M3). The norms used for this test are defined in Eqns. (4.11)-(4.12). This plot has been obtained with the following settings: `ArcTanExp` mesh with 250 elements, $k = 25, a = 0.5, b = 6, r_{\mathrm{max}} = 10^9 r_{\mathrm{s}}$; degree of interpolating polynomials: 5; $\mu_\phi = 10^{13} M_{\mathrm{P}}$, $\mu_H = 10^{12} M_{\mathrm{P}}$.

B.2 Spatial Discretisation Convergence

We shall now present a test that the obtained solution is a good approximation of the true continuous solution. If that is the case, the properties of the solution should be independent of the particular chosen discretisation. In Fig. 12, we show the field profiles as obtained for two meshes characterised by the same number of elements (500) but with considerable difference in their distributions:

- `ArcTanExp` mesh, $k = 15, a = 5 \times 10^{-2}, b = 3 \times 10^{-2}, r_{\max} = 10^9 \, r_{\rm s}$,

- `ArcTanPowerLaw` mesh, $k = 40, \gamma = 20, r_{\max} = 10^9 \, r_{\rm s}$.

In both cases, we use interpolating polynomials of degree 5.

Under every plot in Fig. 12, we show the relative difference between the two obtained solutions, which can be taken to be a measure of the error in the computation of the true, unknown continuous profile. We can see that, for $r \lesssim 10^4 \, r_{\rm s}$, we have a precision of 10^{-9} or better in the characterisation of ϕ's profile, and one of 10^{-5} or better in the characterisation of H's profile. As the radius grows larger, the relative difference grows: this is due to the `ArcTanExp` mesh increasing its sizing exponentially at large radii, which is less suited to describe exponential decay compared to the `ArcTanPowerLaw`, which increases element sizes polynomially. It should be noted, however, that Yukawa suppression, unlike other features in the field profiles, is completely captured by linear physics: as such, it is less of a focus for numerical codes, like φ`enics`, tailored to the description of nonlinear physics inaccessible to analytical understanding. In particular, the interior solution is insensitive to any numerical errors appearing in the Yukawa-suppressed tail. Finally, note that the large increase in the relative error of H around $r \sim 2 - 3 \, r_{\rm s}$ is due to the field going to 0, and does not reflect a correspondingly large increase in the absolute difference between the two computed solutions.

We now consider the spatial convergence properties of our solution with varying resolution (i.e. number of elements). As we refine our spatial discretisation, we expect our discrete solutions to approach the continuum limit. We can determine the rate at which we approach the continuum by comparing solutions as the resolution is varied. This standard convergence testing is illustrated in Figure 13 for a collection of `ArcTanExp` meshes with mapping parameters $k = 7$, $a = 5 \times 10^{-2}$, and $b = 5 \times 10^{-2}$. We consider a series of eight of these meshes, each constructed from twice as many elements as the previous one: i.e. 25, 50, 100, 200, 400, 800, 1600, and 3200 elements. For each mesh, $r_{\max}$ is fixed to $r_{\max} = 10^9 \, r_{\rm s}$, and we employ interpolating polynomials of order 5. Due to the nonlinear nature of the transformation Eq. (3.16), the sequence of meshes are non-nested. To facilitate a direct comparison, we thus introduce a separate "global" mesh consisting of the union of the element vertices from each of the individual meshes. We then interpolate each of the individual solutions onto this global mesh for comparison. For notational convenience, we denote the resulting interpolated profile arising from a mesh with N elements by ϕ_N (and similarly for H, Z, and W). Figure 13 shows the maximal pointwise difference between solutions for neighbouring resolutions. For this test we will calculate the norm $\|\cdot\|_\infty$ on the "global" mesh vertices only, without including intermediate nodes. As expected, the maximal pointwise difference decreases with increasing resolution, until it saturates due to roundoff errors inherent in finite precision arithmetic. Continuing to enhance the resolution beyond this point is not only computationally wasteful, but can also be detrimental as the additional modes will simply be populated with random numerical noise. This is especially harmful when computing the
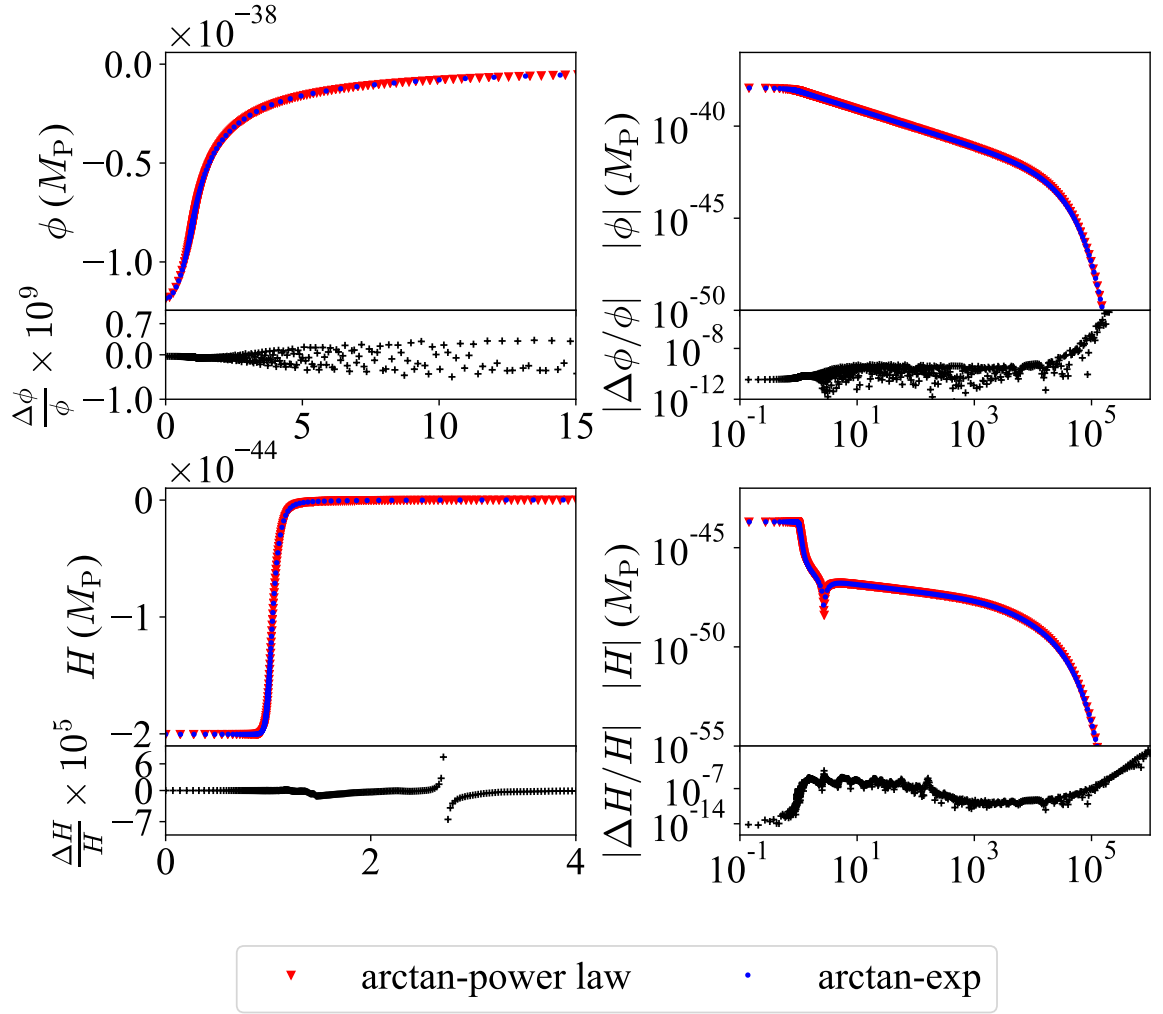
Figure 12. φenics computation of the field profiles for model (M3), as obtained with an `ArcTanExp` mesh (*blue circles*) with 500 elements and $k = 15, a = 5 \times 10^{-2}, b = 3 \times 10^{-2}, r_{\max} = 10^9 r_{\rm s}$, and an `ArcTanPowerLaw` mesh (*red triangles*) with the same number of elements and $k = 40, \gamma = 20, r_{\max} = 10^9 r_{\rm s}$. In both cases, we use order 5 interpolating polynomials. Underneath the profiles, the relative difference in the solutions are shown. Please refer to the the text for discussion of the features of these plots.

Q_n operators in post-processing, since derivatives enhance these noise-dominated short-scale modes. In Fig. 13, we can see that the relative error for the auxiliary variable Z is significantly larger than for the other fields. Examining Fig. 8, we see this occurs because the magnitude of Z is considerably smaller than the other variables, so that the precision in its characterisation will be limited by the roundoff errors in the remaining variables.

A similar test of accuracy can be made by studying the strong residuals of the system of equations in Eq. (3.26) for varying numbers of mesh elements. We make this comparison in Fig. 14, for the same meshes considered in the previous test and in Fig. 13. In particular, we focus on the consistency relation $Y - \nabla_r^2 \phi$ (third equation), which is the dominant residual for the example model considered — this suggests the largest source of error comes from
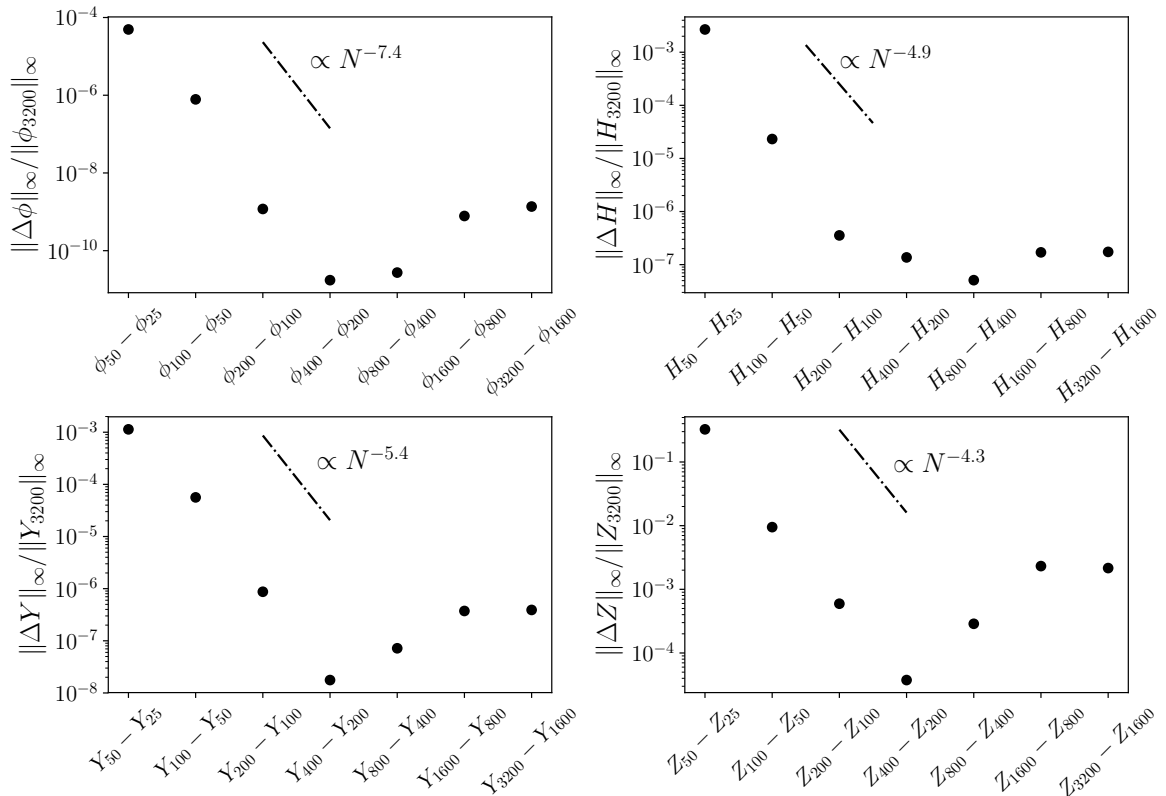
Figure 13. Convergence of our discrete solution to the continuum limit. We consider a family of solutions evaluated on an `ArcTanExp` mesh ($k = 7$, $a = 5 \times 10^{-2}$, and $b = 5 \times 10^{-2}$), with each solution indexed by the number of elements varying by factors of 2 from 25 to 3200. To compare solutions, we define an interpolation grid comprised of the union of all mesh vertices from the meshes of different resolutions. We then compute the maximal pointwise difference on this interpolation grid (denoted here by $\|\cdot\|_\infty$) between solutions on grids of neighbouring resolutions. This provides an estimate of the difference between the continuum and discrete solution. As the number of elements increases, the error in the solution decreases as a power of N, until it reaches the final accuracy in the scheme and no further improvement is obtained. For this model and choice of mesh, we find that the error saturates with around $N = 200$ elements, corresponding to roughly 1000 degrees of freedom.

inaccuracies in the discretisation of the product rule, an effect we will further investigate below. To make this comparison, we compute (in post-processing) the Laplacian $\nabla_r^2 \phi$ as the sum of derivatives $\left(\frac{\partial^2}{\partial r^2} + \frac{2}{r} \frac{\partial}{\partial r} \right) \phi$, projecting this expression onto a discontinuous function space of Lagrange elements of degree 5 (i.e. the same degree as used when solving the system of equations). We can again see that increasing the number of elements initially improves the error in the solution, until an excessive number of them just enhances the effects of numerical noise without further improvement. Between $r \sim 10^1$ and $r \sim 10^5$ the residuals decrease polynomially with r; their magnitude further decreases with the number of elements as $\sim N^{-3.7}$.

As discussed above, the dominant source of error in our numerical solution to model (M3) appears to stem from the breakdown of the product rule upon discretisation. When computing the weak form of the equations of motion, the Laplacian is expressed as a product
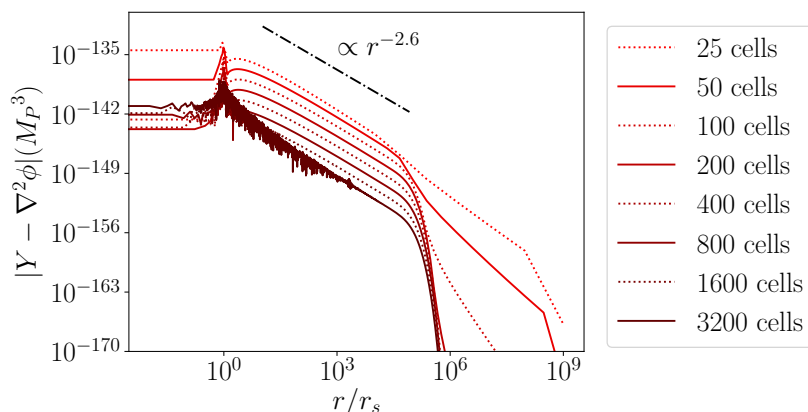
Figure 14. This figure shows the strong residual of the consistency relation $Y - \nabla^2_r\phi$, for model (M3) and increasing numbers of elements. In this plot, we consider the same family of `ArcTanExp` meshes as in Fig. 13: mapping parameters $k = 7$, $a = 5 \times 10^{-2}$, $b = 5 \times 10^{-2}$, $r_{\max} = 10^9 r_s$ and resolution increasing from 25 to 3200 elements doubling each time. Similarly to Fig. 13, we can see that a higher number of elements initially decreases the error at mesh vertices, until no additional improvement is obtained. Once the element number become excessive and an overly high number of modes is used to described the target function, the effects of round-off errors dominate.

of gradients via integration by parts (see Sec. 4.1.1-4.1.2). For continuous functions and the given boundary conditions, the equality $\int \nabla^2_r f v r^2 dr = \int \nabla_r f \cdot \nabla_r v r^2 dr$ is exact. However, for discretised functions, it is approximate: the result is an extra, artificial, boundary term.[16]

In order to study this effect more closely, we now examine the residuals for numerical solutions to the Poisson equation $\nabla^2_r \eta = \frac{1}{2} \frac{\rho}{M_p{}^2}$, for the same source profile as in model (M3). For the Poisson equation, the source term $\rho/w M_{\mathrm{P}}^2$ can be evaluated within machine precision at every vertex, so that the residuals are ascribable to the Laplacian alone: this setup is then particularly convenient to explore the breakdown of the product rule in the discretum. We solve the Poisson equation for a collection of seven `ArcTanExp` meshes with the same transformation parameters as in Figs. 13-14 (i.e. $k = 7$, $a = 5 \times 10^{-2}$, $b = 5 \times 10^{-2}$ and $r_{\max} = 10^9 r_s$), as well as a collection of `ArcTanPowerLaw` meshes with transformation parameters $k = 15, \gamma = 6.5, r_{\max} = 10^9 r_s$. For both meshes, we use a resolution of $50, 100, 200, 400, 800, 1600,$ and 3200 elements. To explore the effect of element degree on the accuracy of the product rule, we also vary the degree of interpolating polynomials from 2 to 9. The results are shown in Fig .15. We can see (left panels) that the error in the product rule decreases as the number of elements increases (as approximately $N^{-3.9}$ for the `ArcTanExp` meshes and $N^{-3.5}$ for the `ArcTanPowerLaw` meshes) until the effects of numerical noise are felt again. Similarly, the residuals decrease with increasing order of interpolating polynomials (right panels), until the number of nodes similarly becomes too large and round-off errors dominate.

---

[16]Note that it is, in principle, possible to discretise the domain in such a way that the product rule is preserved in the discretum.
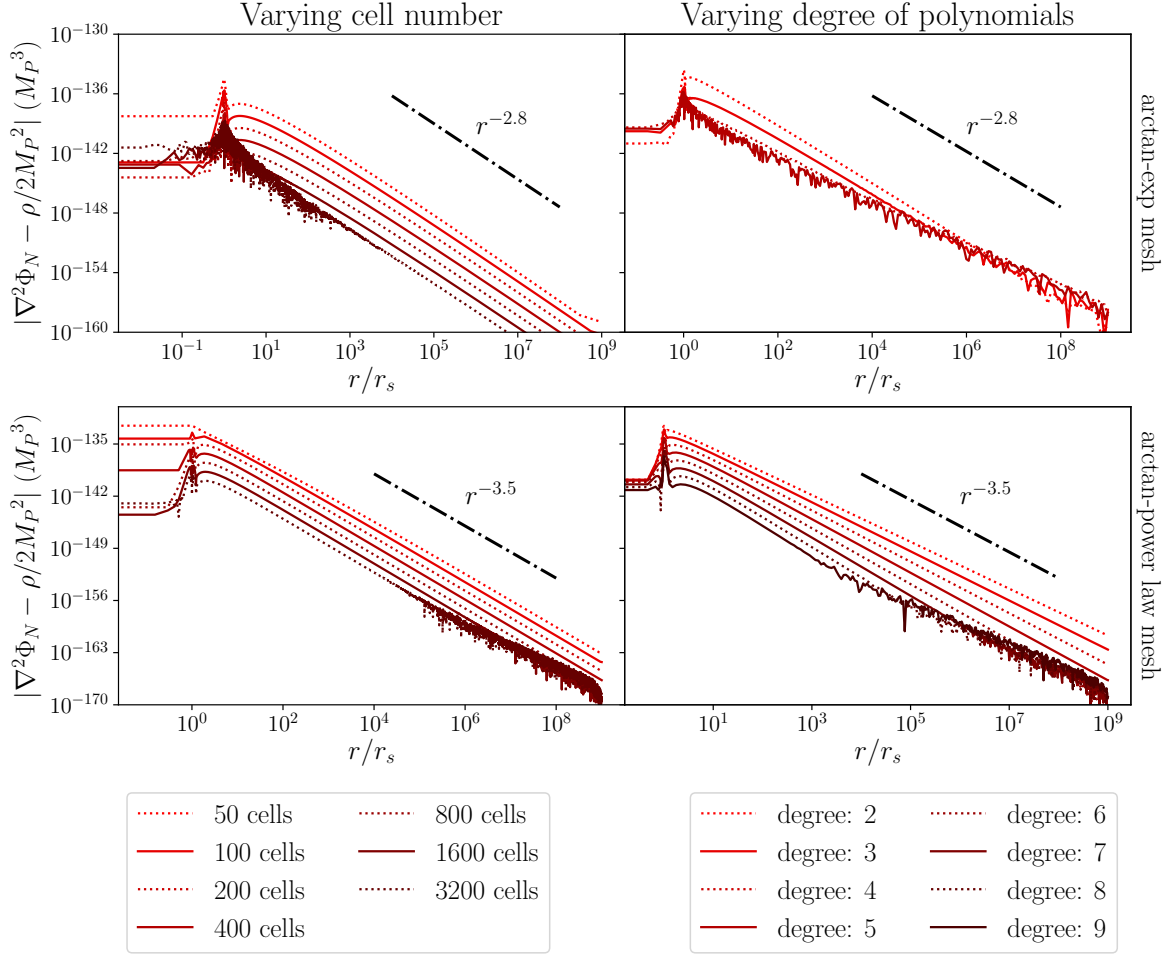
Figure 15. Residuals of the Poisson equation, for the same source profile as in model (M3). For the Poisson equation, the residuals are ascribable to the discretisation of the Laplacian alone: as such, they are a useful measure of the breakdown of the product rule in the discretum. In the left panels, we show the residuals for varying element numbers, from 50 to 3200 doubling each time, whilst keeping the order of interpolating polynomials fixed to 5. In the right panels, we instead vary the order of interpolating polynomials from 2 to 9, whilst keeping the number of elements fixed to 400. We repeat the experiment for both of the meshes described in Sec. 3.4: the results for the `ArcTanExp` meshes are shown in the top panels, and in the bottom panels for the `ArcTanPowerLaw` meshes. The transformation parameters for the family of `ArcTanExp` meshes are $k = 7$, $a = 5 \times 10^{-2}$, $b = 5 \times 10^{-2}$, and for the `ArcTanPowerLaw` meshes $k = 15$, $\gamma = 6.5$. In both cases, we use $r_{\max} = 10^9 \, r_{\rm s}$. The residuals of the Poisson equation improve with increasing number of elements and degree of interpolating polynomials, until numerical noise resulting from round-off errors dominates, similarly to what is shown in Figs. 13-14.

C   Code structure

In this appendix, we briefly describe φenics's structure, and outline the steps that prospective users will need to take if they wish to adapt the code to other screening models. Further details on the listed code elements may be found in φenics's documentation; additionally, the Jupyter notebooks provided with the code (`IR_main.ipynb` and `UV_main.ipynb`) show a

suggested workflow that generalises to any static and spherically symmetric screening problem.

C.1 Discretisation

Any prospective user of φenics should start by discretising the domain. For the example theories considered in this paper, the `ArcTanExpMesh` and `ArcTanPowerLawMesh` meshes presented in Sec. 3.4 were implemented with the provided mesh classes of the same name. We expect them to be useful for other problems of screening and not to be specific of our models. However, prospective users may wish to implement different transformations based on specific knowledge of their problem and thus generate different meshes. φenics's `ArcTanExpMesh` and `ArcTanPowerLawMesh` were implemented as subclasses of a `Mesh` base class to allow this while retaining other useful functionalities, including linear refinement, point declustering, and checks that mesh vertices are not closer than machine precision allows or that the derivative of the transformation is negative (for declustering and k_{rm} negative). These general functionalities, useful for any mesh, can therefore be performed without any specific knowledge of the nonlinear transformation used.

To add a new mesh to φenics, we encourage users to implement a new `Mesh` subclass modelled on the `ArcTanExpMesh` and `ArcTanPowerLawMesh` subclasses provided. Any new subclass should contain a definition of the transformation to be applied, together with its first and second derivatives, plus either an exact or approximate inverse transformation at small and large radii.

Further finite-element settings are given in the `fem` module: they are expected to work 'as is' for generic models of screening. In particular, they define continuous and discontinuous Lagrange elements to represent functions as described in Sec. 3.2.

C.2 Source

The `source` modules contain definitions of static and spherically symmetric source profiles of interest for general models of screening. They are: smoothed top-hat (`TopHatSource`, Eq. 2.10), step function (`StepSource`), truncated cosine (`CosSource`, Eq. 2.11), Gaussian (`GaussianSource`), and 'Gaussian cake' (`GCakeSource`, Eq. 2.12). Any user wishing to implement a different source profile can add a corresponding subclass of `Source` following the examples provided.

C.3 Solver

φenics contains a base solver class (`Solver`) that takes care of functionalities that are general to any nonlinear solver, and not restricted to specific models of screening. These functionalities include: advancing the Newton iterations, computing the error at every iteration (given a weak residual form and/or the change in the solution at that iteration) and implementing the stopping criteria.

Similarly to the way in which the mesh is built, model-specific tasks for the single- and two-field theories are implemented in the `Solver` subclasses `IRSolver` (IR module) and `UVSolver` (UV module). Users wishing to apply φenics to different theories are encouraged to write

a model-specific subclass of `Solver` following the examples provided, and use the base-class functionalities for non-model specific tasks. Information to be supplied in the custom subclass includes:

- a linear solver for the Newton iteration (Eq. 3.1), expressed in weak form. For our example theories, the analytic formulae for this step have been written down explicitly in Eq. (4.6) and Eq. (4.10). Note that the weak form must correctly implement the Neumann boundary conditions;

- the Dirichlet boundary conditions for the theory;

- the weak and strong form of the residuals;

- where the equation of motion is expanded into a system of equations (or is natively a system of equations), the definition of the composite function space for the vector equation. The `IRSolver` and `UVSolver` subclasses perform this task and can be taken as examples.

References

[1] N. Aghanim et al. Planck 2018 results. VI. Cosmological parameters. 2018.

[2] Jose Luis Bernal, Licia Verde, and Adam G. Riess. The trouble with H_0. *JCAP*, 1610(10):019, 2016.

[3] Wendy L. Freedman. Cosmology at a Crossroads. *Nat. Astron.*, 1:0121, 2017.

[4] Adam G. Riess, Stefano Casertano, Wenlong Yuan, Lucas M. Macri, and Dan Scolnic. Large Magellanic Cloud Cepheid Standards Provide a 1% Foundation for the Determination of the Hubble Constant and Stronger Evidence for Physics beyond ΛCDM. *Astrophys. J.*, 876(1):85, 2019.

[5] Kenneth C. Wong et al. H0LiCOW XIII. A 2.4% measurement of H_0 from lensed quasars: 5.3σ tension between early and late-Universe probes. 2019.

[6] Clifford M. Will. Was Einstein Right? A Centenary Assessment. 2014.

[7] Austin Joyce, Bhuvnesh Jain, Justin Khoury, and Mark Trodden. Beyond the Cosmological Standard Model. *Phys. Rept.*, 568:1–98, 2015.

[8] Mustapha Ishak. Testing General Relativity in Cosmology. *Living Rev. Rel.*, 22(1):1, 2019.

[9] Thomas P. Sotiriou and Valerio Faraoni. f(R) Theories Of Gravity. *Rev. Mod. Phys.*, 82:451–497, 2010.

[10] Kurt Hinterbichler. Theoretical Aspects of Massive Gravity. *Rev. Mod. Phys.*, 84:671–710, 2012.

[11] Claudia de Rham, Gregory Gabadadze, and Andrew J. Tolley. Resummation of Massive Gravity. *Phys. Rev. Lett.*, 106:231101, 2011.

[12] Gregory Walter Horndeski. Second-order scalar-tensor field equations in a four-dimensional space. *Int. J. Theor. Phys.*, 10:363–384, 1974.

[13] C. Deffayet, S. Deser, and G. Esposito-Farese. Generalized Galileons: All scalar models whose curved background extensions maintain second-order field equations and stress-tensors. *Phys. Rev.*, D80:064015, 2009.

[14] Harry Desmond, Pedro G. Ferreira, Guilhem Lavaux, and Jens Jasche. Fifth force constraints from the separation of galaxy mass components. *Phys. Rev.*, D98(6):064015, 2018.

[15] Harry Desmond, Pedro G. Ferreira, Guilhem Lavaux, and Jens Jasche. Fifth force constraints from galaxy warps. *Phys. Rev.*, D98(8):083010, 2018.

[16] Philip Bull et al. Beyond ΛCDM: Problems, solutions, and the road ahead. *Phys. Dark Univ.*, 12:56–99, 2016.

[17] Kazuya Koyama. Cosmological Tests of Modified Gravity. *Rept. Prog. Phys.*, 79(4):046902, 2016.

[18] Anže Slosar et al. Dark Energy and Modified Gravity. 2019.

[19] Alberto Nicolis, Riccardo Rattazzi, and Enrico Trincherini. The Galileon as a local modification of gravity. *Phys. Rev.*, D79:064036, 2009.

[20] Jérôme Gleyzes, David Langlois, Federico Piazza, and Filippo Vernizzi. Healthy theories beyond Horndeski. *Phys. Rev. Lett.*, 114(21):211101, 2015.

[21] Jérôme Gleyzes, David Langlois, Federico Piazza, and Filippo Vernizzi. Exploring gravitational theories beyond Horndeski. *JCAP*, 1502:018, 2015.

[22] David Langlois and Karim Noui. Degenerate higher derivative theories beyond Horndeski: evading the Ostrogradski instability. *JCAP*, 1602(02):034, 2016.

[23] David Langlois and Karim Noui. Hamiltonian analysis of higher derivative scalar-tensor theories. *JCAP*, 1607(07):016, 2016.

[24] Jibril Ben Achour, David Langlois, and Karim Noui. Degenerate higher order scalar-tensor theories beyond Horndeski and disformal transformations. *Phys. Rev.*, D93(12):124005, 2016.

[25] Marco Crisostomi, Kazuya Koyama, and Gianmassimo Tasinato. Extended Scalar-Tensor Theories of Gravity. *JCAP*, 1604(04):044, 2016.

[26] Claudia de Rham, Scott Melville, Andrew J. Tolley, and Shuang-Yong Zhou. Massive Galileon Positivity Bounds. *JHEP*, 09:072, 2017.

[27] Hiromu Ogawa, Takashi Hiramatsu, and Tsutomu Kobayashi. Anti-screening of the Galileon force around a disk center hole. *Mod. Phys. Lett.*, A34(02):1950013, 2018.

[28] Takashi Hiramatsu, Wayne Hu, Kazuya Koyama, and Fabian Schmidt. Equivalence Principle Violation in Vainshtein Screened Two-Body Systems. *Phys. Rev.*, D87(6):063525, 2013.

[29] Adrien Kuntz. Two-body potential of Vainshtein screened theories. *Physical Review D*, 100(2):024024, July 2019.

[30] Nicholas C. White, Sandra M. Troian, Jeffrey B. Jewell, Curt J. Cutler, Sheng-wey Chiow, and Nan Yu. Robust numerical computation of the 3D scalar potential field of the cubic Galileon gravity model at solar system scales. *Physical Review D*, 102(2):024033, July 2020.

[31] Furqan Dar, Claudia De Rham, J. Tate Deskins, John T. Giblin, and Andrew J. Tolley. Scalar Gravitational Radiation from Binaries: Vainshtein Mechanism in Time-dependent Systems. *Class. Quant. Grav.*, 36(2):025008, 2019.

[32] Andrei V. Frolov, José T. Gálvez Ghersi, and Alex Zucca. Unscreening scalarons with a black hole. *Physical Review D*, 95(10), May 2017.

[33] Amol Upadhye, Steven S. Gubser, and Justin Khoury. Unveiling chameleon fields in tests of the gravitational inverse-square law. *Physical Review D*, 74(10), Nov 2006.

[34] Benjamin Elder, Justin Khoury, Philipp Haslinger, Matt Jaffe, Holger Müller, and Paul Hamilton. Chameleon dark energy and atom interferometry. *Physical Review D*, 94(4), Aug 2016.

[35] Matt Jaffe, Philipp Haslinger, Victoria Xu, Paul Hamilton, Amol Upadhye, Benjamin Elder, Justin Khoury, and Holger Müller. Testing sub-gravitational forces on atoms from a miniature in-vacuum source mass. *Nature Physics*, 13(10):938–942, Jul 2017.

[36] Philippe Brax, Anne-Christine Davis, Benjamin Elder, and Leong Khim Wong. Constraining screened fifth forces with the electron magnetic moment. *Phys. Rev. D*, 97(8):084050, 2018.

[37] Benjamin Elder, Valeri Vardanyan, Yashar Akrami, Philippe Brax, Anne-Christine Davis, and Ricardo S. Decca. Classical symmetron force in casimir experiments. *Physical Review D*, 101(6), Mar 2020.

[38] Clare Burrage, Edmund J. Copeland, Adam Moss, and James A. Stevenson. The shape dependence of chameleon screening. *JCAP*, 1801(01):056, 2018.

[39] Alexandre Barreira, Sownak Bose, and Baojiu Li. Speeding up N-body simulations of modified gravity: Vainshtein screening models. *JCAP*, 1512(12):059, 2015.

[40] Hans A. Winther and Pedro G. Ferreira. Fast route to nonlinear clustering statistics in modified gravity theories. *Phys. Rev.*, D91(12):123507, 2015.

[41] Sownak Bose, Baojiu Li, Alexandre Barreira, Jian-hua He, Wojciech A. Hellwing, Kazuya Koyama, Claudio Llinares, and Gong-Bo Zhao. Speeding up N-body simulations of modified gravity: Chameleon screening models. *JCAP*, 1702(02):050, 2017.

[42] Claudio Llinares, David F. Mota, and Hans A. Winther. ISIS: a new N-body cosmological code with scalar fields based on RAMSES. Code presentation and application to the shapes of clusters. *Astron. Astrophys.*, 562:A78, 2014.

[43] R. Hagala, C. Llinares, and D. F. Mota. Cosmological simulations with disformally coupled symmetron fields. *Astron. Astrophys.*, 585:A37, 2016.

[44] Hans A. Winther et al. Modified Gravity N-body Code Comparison Project. *Mon. Not. Roy. Astron. Soc.*, 454(4):4208–4234, 2015.

[45] Alexandre Barreira, Baojiu Li, Wojciech A. Hellwing, Carlton M. Baugh, and Silvia Pascoli. Nonlinear structure formation in the Cubic Galileon gravity model. *JCAP*, 1310:027, 2013.

[46] Baojiu Li, Gong-Bo Zhao, and Kazuya Koyama. Exploring Vainshtein mechanism on adaptively refined meshes. *JCAP*, 1305:023, 2013.

[47] Baojiu Li, Alexandre Barreira, Carlton M. Baugh, Wojciech A. Hellwing, Kazuya Koyama, Silvia Pascoli, and Gong-Bo Zhao. Simulating the quartic Galileon gravity model on adaptively refined meshes. *JCAP*, 1311:012, 2013.

[48] Jason N. Dossett, Mustapha Ishak, and Jacob Moldenhauer. Testing General Relativity at Cosmological Scales: Implementation and Parameter Correlations. *Phys. Rev.*, D84:123001, 2011.

[49] Alireza Hojjati, Levon Pogosian, and Gong-Bo Zhao. Testing gravity with CAMB and CosmoMC. *JCAP*, 1108:005, 2011.

[50] Alex Zucca, Levon Pogosian, Alessandra Silvestri, and Gong-Bo Zhao. MGCAMB with massive neutrinos and dynamical dark energy. *JCAP*, 1905(05):001, 2019.

[51] Bin Hu, Marco Raveri, Noemi Frusciante, and Alessandra Silvestri. Effective Field Theory of Cosmic Acceleration: an implementation in CAMB. *Phys. Rev.*, D89(10):103530, 2014.

[52] Bin Hu, Marco Raveri, Noemi Frusciante, and Alessandra Silvestri. EFTCAMB/EFTCosmoMC: Numerical Notes v3.0. 2014.

[53] Miguel Zumalacárregui, Emilio Bellini, Ignacy Sawicki, Julien Lesgourgues, and Pedro G. Ferreira. hi_class: Horndeski in the Cosmic Linear Anisotropy Solving System. *JCAP*, 1708(08):019, 2017.

[54] Emilio Bellini, Ignacy Sawicki, and Miguel Zumalacárregui. hi_class: Background Evolution, Initial Conditions and Approximation Schemes. *JCAP*, 2002(02):008, 2020.

[55] E. Bellini et al. Comparison of Einstein-Boltzmann solvers for testing general relativity. *Phys. Rev.*, D97(2):023520, 2018.

[56] Christian Arnold and Baojiu Li. Simulating galaxy formation in f(R) modified gravity: Matter, halo, and galaxy-statistics. 2019.

[57] Claudio Llinares. Simulation techniques for modified gravity. *Int. J. Mod. Phys.*, D27(15):1848003, 2018.

[58] Clare Burrage, Ben Coltman, Antonio Padilla, Daniela Saadeh, and Toby Wilson. Massive Galileons and Vainshtein screening. *JCAP*, 2021(2):050, February 2021.

[59] John P. Boyd. *Chebyshev and Fourier Spectral Methods.* Dover Books on Mathematics. Dover Publications, Mineola, NY, second edition, 2001.

[60] J. Richard Bond, Jonathan Braden, and Laura Mersini-Houghton. Cosmic bubble and domain wall instabilities III: The role of oscillons in three-dimensional bubble collisions. *JCAP*, 09:004, 2015.

[61] H.P. Langtangen and K.A. Mardal. *Introduction to Numerical Methods for Variational Problems.* Texts in Computational Science and Engineering. Springer International Publishing, 2019.

[62] Martin S. Alnæs, Jan Blechta, Johan Hake, August Johansson, Benjamin Kehlet, Anders Logg, Chris Richardson, Johannes Ring, Marie E. Rognes, and Garth N. Wells. The fenics project version 1.5. *Archive of Numerical Software*, 3(100), 2015.

[63] Anders Logg, Kent-Andre Mardal, Garth N. Wells, et al. *Automated Solution of Differential Equations by the Finite Element Method.* Springer, 2012.

[64] Jorge J. Moré, Burton S. Garbow, and Kenneth E. Hillstrom. *User Guide for MINPACK-1.* Argonne National Laboratory and United States. Department of Energy. Office of Scientific and Technical Information, 1980.

[65] Toby Wilson. Scalar fields and gravity, July 2018.

[66] Hans Petter Langtangen and Anders Logg. *Solving PDEs in Python.* Springer, 2017.