

Parallel inexact Newton-Krylov and quasi-Newton solvers for nonlinear elasticity

Nicolás A. Barnafi*

Luca F. Pavarino*

Simone Scacchi†

Abstract

In this work, we address the implementation and performance of inexact Newton-Krylov and quasi-Newton algorithms, more specifically the BFGS method, for the solution of the nonlinear elasticity equations, and compare them to a standard Newton-Krylov method. This is done through a systematic analysis of the performance of the solvers with respect to the problem size, the magnitude of the data and the number of processors in both almost incompressible and incompressible mechanics. We consider three test cases: Cook’s membrane (static, almost incompressible), a twist test (static, incompressible) and a cardiac model (complex material, time dependent, almost incompressible). Our results suggest that quasi-Newton methods should be preferred for compressible mechanics, whereas inexact Newton-Krylov methods should be preferred for incompressible problems. We show that these claims are also backed up by the convergence analysis of the methods. In any case, all methods present adequate performance, and provide a significant speed-up over the standard Newton-Krylov method, with a CPU time reduction exceeding 50% in the best cases.

1 Introduction

Nonlinear elasticity is a continuum framework for modeling the deformation of an elastic body, and has a large spectrum of applications, ranging from industrial to academic. The flexibility of continuum mechanics resides in its abstract formulation starting from fundamental principles, which allows for an accurate representation of many different materials through constitutive modeling and complex boundary conditions [Hol02]. After a precise physical phenomenon has been devised and modeled, the resulting equations yield a complex system of nonlinear Partial Differential Equations (PDEs). Limited progress has been achieved in terms of its numerical analysis [CD04, ADVL⁺13], so it is difficult to know a-priori efficient strategies for its numerical approximation. Indeed, the problem is not convex but only polyconvex [Cia21], which deteriorates the performance of many well established numerical methods. The gold standard is Newton’s method [ZTNZ77], consisting in a second order approximation of the associated variational principle (or a first order approximation of the Euler-Lagrange equations, referred to as the Newton-Raphson method). This method is popular mainly due to its robustness and low iteration count, as it converges quadratically whenever a good initial guess of the solution is considered [WN99]. Its main drawback is that it requires the repeated assembly of the Jacobian matrix, associated to the tangent problem, which usually becomes a bottleneck of the solution process (see [JHN11] for an example in CFD). This can be more evidently appreciated when using higher order approximations, required for avoiding numerical locking effects [EG13].

In general, not much attention has been given in the community to other methods, such as linearly convergent (gradient descent [Rud16], Richardson [Saa03]) and superlinearly convergent (quasi-Newton [WN99], inexact Newton [DES82]) ones. These methods yield a higher iteration count due to their reduced convergence rate, but can potentially present a drastically overall reduced computational complexity. Indeed, the use of different nonlinear solvers has shown enormous speed-ups in other problems [SHOL06, BRKN17], where the

*Department of Mathematics, Università di Pavia, Via Ferrata 1, 27100 Pavia, Italy. {nicolas.barnafi, luca.pavarino}@unipv.it

†Department of Mathematics, Università di Milano, Via Saldini 50, 20133 Milano, Italy. simone.scacchi@unimi.it

properties of the equations have been exploited to devise a more adequate nonlinear solution method. Some recent works have numerically explored quasi-Newton methods for the mechanics [LLS05, GP88, LBK17], with their main focus being on the nonlinear iterations and CPU time. This type of study does not take into account the linear system involved in each iteration of the method, and thus is not conclusive with respect to the applicability of these methods in an HPC setting, where problems with millions of degrees of freedom need to be solved, possibly for many time steps. A numerical method can be deemed adequate in such case if at least it satisfies the following requirements: i) the nonlinear and linear iterations incurred during the solution procedure are independent of the mesh size used, ii) the method is robust with respect to the model parameters, at least in the scenarios of interest and iii) the method is (strongly) scalable, meaning that the overall solution time improves when more processors are used. We finally highlight [WDE07], where a novel conjugate Newton method was proposed and analysed for nonlinear elasticity.

The scope of this work is to provide a first step in the adequate usage of alternatives to Newton’s method, where we focus in superlinearly convergent methods as they provide an excellent overall computational complexity [WN99]. We stress that all gradient descent algorithms we tested on a preliminary phase failed, and this is indeed consistent with the literature: there are no works, up to our knowledge, of descent nor fixed point algorithms for nonlinear mechanics. We focus on three tests: i) Cook’s membrane test, an almost incompressible problem, ii) a twist problem, where we test an incompressible material and iii) a cardiac modeling test, where we use an idealized left ventricle geometry to model a human heartbeat. The work is structured as follows: in Section 2 we review the mechanics problem and fix some notation, in Section 3 we describe the nonlinear solvers to be used throughout this study, in Section 4 we define the tests in which the methods will be tested together with a thorough description of the linear solver involved, in Section 5 we recall the convergence results of the methods and discuss what their expected performance should be, in Section 6 we show and comment the results obtained and we conclude with a discussion of the results in Section 7.

2 The hyperelasticity problem

Consider a connected domain $\Omega \subset \mathbb{R}^{d \in \{2,3\}}$ that represents the geometry to be deformed, and define its Dirichlet and Neumann boundaries as $\partial\Omega_D$ and $\partial\Omega_N$ respectively, such that $\overline{\partial\Omega} = \overline{\partial\Omega_D} \cup \overline{\partial\Omega_N}$. We look for a displacement $\mathbf{d} : [0, T] \times \Omega \rightarrow \mathbb{R}^d$ such that it solves the momentum conservation equation

$$\begin{aligned} \ddot{\mathbf{d}} - \operatorname{div} \mathbf{P} &= \mathbf{f} && \text{in } (0, T) \times \Omega, \\ \mathbf{d} &= \mathbf{d}_0 && \text{in } \{0\} \times \Omega, \\ \dot{\mathbf{d}} &= \mathbf{v}_0 && \text{in } \{0\} \times \Omega, \\ \mathbf{P}n &= \mathbf{t} && \text{on } [0, T] \times \partial\Omega_N, \\ \mathbf{d} &= \mathbf{d}_D && \text{on } [0, T] \times \partial\Omega_D, \end{aligned} \tag{1}$$

where \mathbf{d}_0 is the initial displacement, \mathbf{v}_0 is the initial velocity, \mathbf{t} is the surface traction, \mathbf{d}_D is a distributed load, $\dot{(\)}$ denotes a time derivative, and \mathbf{P} is known as the Piola stress tensor, usually obtained from a predefined Helmholtz potential Ψ such that

$$\mathbf{P}(\mathbf{F}) = \frac{\partial \Psi}{\partial \mathbf{F}}(\mathbf{F}),$$

where $\mathbf{F} = \mathbf{I} + \nabla \mathbf{d}$. The existence of such a potential is actually a hypothesis, and whenever it exists we refer to the material as a *hyperelastic material*. Approximating this problem by neglecting the inertial term $\ddot{\mathbf{d}}$ results in the well-known *static* mechanics, or quasi-static whenever one of the problem’s data is time-dependent. If we require the volume to be conserved through $J := \det \mathbf{F} = 1$, we obtain an *incompressible* elasticity problem. This can be instead approximated by further penalizing the deviation of $J - 1$ from 0, which results in a modified potential

$$\Psi(\mathbf{F}) = \Psi_{\text{sol}}(\bar{\mathbf{F}}) + \Psi_{\text{vol}}(J),$$

where the isochoric component $\bar{\mathbf{F}} = J^{-1/d}\mathbf{F}$ is such that $\det \bar{\mathbf{F}} = 1$. This separates the energetic contribution of volumetric deformation, and is referred to as almost (or quasi-) incompressibility whenever the term Ψ_{vol} yields $J \approx 1$. More details can be found in [Hol02].

3 Numerical solvers

In this section we briefly review the nonlinear methods we consider, namely Newton-Krylov and quasi-Newton methods. Details regarding their geometric motivation can be found in [WN99], whereas the convergence properties are detailed in Section 5. Our study is mainly motivated by the use of these methods in an HPC infrastructure, so the use of iterative solvers within the nonlinear method is not only inevitable, but also desirable due to their well-established parallel performance. Inspired by this, we have classified our methods according to the degree of exactness of the linear solver it uses, so that a fully inexact method will simply consider the action of the associated preconditioner instead of solving a linear system. Instead, a quasi-exact¹ method will consider an accurate linear solver given by small tolerances, i.e. 10^{-10} and 10^{-8} absolute and relative tolerances, respectively. Now we provide more details regarding the methods, and in particular emphasize where the linear solver is used.

Newton-Krylov methods. These methods can be seen as either a minimization procedure or as a root finding algorithm, the latter often referred to as Newton-Raphson method. We present it as a root finding method. Consider Equation (1) in quasi-static, residual form

$$\mathbf{R}(\mathbf{d}) := \mathbf{f} - \text{div } \mathbf{P}(\mathbf{F}) = \mathbf{0},$$

with $\mathbf{P} = \frac{\partial \Psi}{\partial \mathbf{F}}$ for a given Helmholtz potential Ψ , and consider an initial iterate \mathbf{d}^0 . Then, given a previous iteration \mathbf{d}^{n-1} , the next one is given as the solution $\delta \mathbf{d}^n$ of the linearized problem

$$\partial_{\mathbf{d}} \mathbf{R}(\mathbf{d}^{n-1})[\delta \mathbf{d}^n] = -\mathbf{R}(\mathbf{d}^{n-1}), \quad (2)$$

where $\partial_{\mathbf{d}}$ stands for the Frechét derivative with respect to \mathbf{d} , and the update is then given by $\mathbf{d}^n = \mathbf{d}^{n-1} + \delta \mathbf{d}^n$. Equation (2) gives rise to a linear system of equations often referred to as the tangent problem or simply the linearized problem, which is solved by means of an iterative Krylov space method, accelerated by a preconditioner, see e.g. [CGKT94, CGK⁺98]. In practice, the most used one is referred to as Newton-MG, meaning that problem (2) is solved with an iterative method, preconditioned by an algebraic multigrid preconditioner (AMG) [Hac13, CFPS15].

Inexact Newton-Krylov methods. Problem (2) is an approximation of the actual equation $\mathbf{R}(\mathbf{d}) = 0$, and thus it might not be true that an accurate solution of the tangent problem will also yield an accurate solution of the original equation. In fact, this usually gives rise to over-solving the problem [EW94] in the first iterations, which motivates the use of inexact solvers for (2), i.e. to use large tolerances for the solution of the tangent problem. In addition, the quality of the linearization improves as the iterates are closer to the solution, which motivates the use of adaptive (relative) tolerances. In particular, we use the Einstat-Walker strategy [EW96], given by

$$\text{tol}^n = \frac{\|\|\partial_{\mathbf{d}} \mathbf{R}(\mathbf{d}^{n-1})[\delta \mathbf{d}^n] - \mathbf{R}(\mathbf{d}^{n-1})\| - \|\mathbf{R}(\mathbf{d}^{n-1})\|\|}{\|\mathbf{R}(\mathbf{d}^{n-1})\|}.$$

The choice of the norm $\|\cdot\|$ is arbitrary, so we consider in what follows the ℓ^2 norm. This of course makes sense only in the discrete setting, meaning that for each residual vector \mathbf{r} we consider the norm $\|\mathbf{r}\| = (\sum_i r_i^2)^{1/2}$. The resulting scheme guarantees superlinear convergence [DES82], which is of course worse than the quadratic convergence of a standard Newton-Krylov scheme, but gives an overall reduced complexity as it avoids oversolving the linearized problem.

¹We emphasize that all methods are inexact given their iterative nature. Still, a sufficiently low tolerance can be regarded as a sufficiently exact solver, so we refer to this as quasi-exact.

BFGS method. This is a Quasi-Newton method, and although it was initially devised as a minimization procedure [WN99], it can also be adapted to be used as a root-finding algorithms [DS96]. For this, we look at the minimization principle from which (1) in quasi-static form is obtained:

$$\min_{\mathbf{d}} \Pi(\mathbf{d}) := \int_{\Omega} \Psi(\mathbf{F}) - \mathbf{f} \cdot \mathbf{d} \, dx. \quad (3)$$

A quadratic approximation of this problem yields

$$\Pi(\mathbf{d}) \approx \Pi(\bar{\mathbf{d}}) + \partial_{\mathbf{d}}\Pi(\bar{\mathbf{d}})[\mathbf{d} - \bar{\mathbf{d}}] + \frac{1}{2}\partial_{\mathbf{d}}^2\Pi(\bar{\mathbf{d}})[\mathbf{d} - \bar{\mathbf{d}}],$$

and such scheme is indeed equivalent to (2) when applied as an iterative procedure with an exact Hessian. This method requires an initial approximation of the Hessian $\mathbf{B}^{-1,0} \approx [\partial_{\mathbf{d}}^2\Pi(\mathbf{d}^0)]^{-1}$, which is then enriched at each iteration with a rank two perturbation that includes curvature information given by

$$\mathbf{B}^{-1,k+1} = (\mathbf{I} - \rho_k \mathbf{s}^k \otimes \mathbf{y}^k) \mathbf{B}^{-1,k} (\mathbf{I} - \rho_k \mathbf{y}^k \otimes \mathbf{s}^k) + \rho_k \mathbf{s}^k \otimes \mathbf{s}^k,$$

where $\mathbf{s}^k = \mathbf{x}^{k+1} - \mathbf{x}^k$, $\mathbf{y}^k = \mathbf{F}(\mathbf{x}^{k+1}) - \mathbf{F}(\mathbf{x}^k)$, and $\rho^k = 1/\langle \mathbf{s}^k, \mathbf{y}^k \rangle$ with $\langle \mathbf{a}, \mathbf{b} \rangle = \sum_i a_i b_i$. The approximate Hessian is then used to compute the next iteration as in Equation (2):

$$\delta \mathbf{d}^k = -\mathbf{B}^{-1,k+1} \partial_{\mathbf{d}}\Pi(\mathbf{d}^{k-1}).$$

The action of $\mathbf{B}^{-1,k+1}$ is implemented by means of a two-level recursion that allows for a limited memory implementation [Noc80]. This method, as the inexact Newton-Krylov, yields superlinear convergence, and has the additional advantage of requiring only one assembly of the Hessian matrix (or any other initial matrix). The initial approximation of the Hessian is critical for the convergence and performance of this method, and we have indeed observed that simpler approaches do not yield satisfactory results for this method in the context of nonlinear mechanics. Motivated by this, we leverage the preconditioners obtained from the initial matrix $\partial_{\mathbf{d}}^2\Pi(\mathbf{d}^0)$ to obtain better approximations of the Hessian. If we denote by \mathbf{P} the preconditioner arising from the initial Hessian, we consider the following approximations:

- $\mathbf{B}^{-1,0}$ = action of \mathbf{P} obtained from Hessian matrix (BFGS-preonly)
- $\mathbf{B}^{-1,0}$ = inexact Krylov solver, large ($\approx 10^{-2}$) relative tolerance (inexact BFGS)
- $\mathbf{B}^{-1,0}$ = quasi-exact Krylov solver, small ($\approx 10^{-6}$) relative tolerance (quasi-exact BFGS)

We have added on the right the names corresponding to Figure 1. We note that the quasi-exact and inexact versions are not covered by the theory, as the action of $\mathbf{B}^{-1,0}$ changes at each iteration. This happens because the number of linear iterations required at each nonlinear iteration varies according to the tolerance. We highlight that the use of a fixed tolerance in the inexact scenario is considered for simplicity, as an equivalent Einstat-Walker type of adaptive tolerance could be considered. This is an interesting alternative and by no means a trivial one to set up, so we leave this for future work.

We show a classification of these methods according to the exactness of the linear solver in Figure 1, where we have fixed three degrees of exactness: preconditioner only (preonly²), inexact and quasi-exact. The two families present different types of performance: the Newton family relies on a low iteration count, where the level of inexactness relaxes the linear solver at each iteration, at the cost of some additional Newton steps. The quasi-Newton family instead yields a larger number of iterations, but each iteration is much cheaper than a Newton iteration, as it does not require the Jacobian matrix to be reassembled. The level of exactness in this case reduces the iteration count, at the cost of making the iterations more expensive. The extreme cases Newton-preonly and quasi-exact BFGS have shown to be noncompetitive in our experiments, so we do not consider them in what follows. Throughout the remaining parts of this work, whenever no confusion

²The name is the same as the PETSc option to by-pass the iterative solver.

arises, we may refer to Newton-Krylov and inexact Newton-Krylov methods as simply Newton methods. Throughout this manuscript, we will denote the Newton-Krylov, inexact Newton-Krylov, BFGS-preonly (or simply BFGS) and inexact-BFGS methods as **NK**, **iNK**, **B** and **iB** respectively, as shown in Figure 1. We also show in Appendix A the commands used to run each of the four methods under consideration.

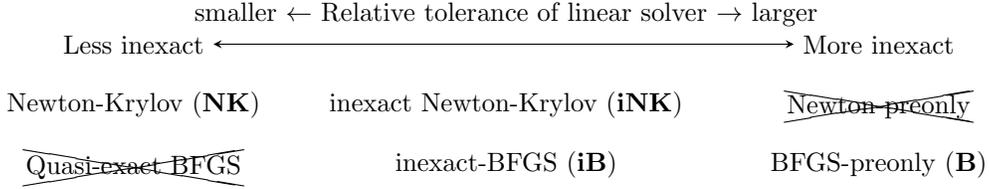


Figure 1: Description of exactness level in the linear solver and nonlinear solvers used in the tests, where we denote "preconditioner only" as "preonly". Newton-preonly and quasi-exact BFGS have been discarded from our experiments because they are not competitive: the first requires too many matrix assembles, the second spends too much time at each iteration.

4 Test problems and solvers details

In this section, we describe the test cases we use to compare the performance of each method. The scope of the heartbeat test is that of providing a more realistic scenario in which the insight acquired in the other two tests will be transferred. Details on the implementation will be given on each problem.

Cook test. This is a static almost incompressible mechanics benchmark [PM16]. We use homogeneous Dirichlet conditions on $\{x = 0\}$, a vertical traction given by $\mathbf{t} = (0, \tau, 0)$ with $\tau = 10^6$ in deformed configuration on $\{x = L\}$, a homogeneous distributed load $\mathbf{f} = \mathbf{0}$, and homogeneous Neumann conditions on the remaining parts of the boundary. The constitutive modeling is given by a almost incompressible Neo-Hookean material:

$$\Psi_{\text{Cook}}(\mathbf{F}) = C_1 (\text{tr}(\bar{\mathbf{F}}^T \bar{\mathbf{F}}) - 3) + k ([\det(\mathbf{F})]^2 - 1 - 2 \log(\det(\mathbf{F}))),$$

where $C_1 = \frac{1}{2}\mu$, $\mu = 8.194 \cdot 10^7$, $k = \lambda + \frac{2}{3}\mu$, $\lambda = 2\mu\nu/(1 - 2\nu)$, and $\nu = 0.3$. The Krylov solver was configured with a GMRES linear solver without restart and with a modified Gram-Schmidt procedure which is more robust [GVL96], preconditioned with HYPRE-BoomerAMG [FY02], using its default configuration and setting the matrix block size to 3 in PETSc to improve the efficiency of the preconditioner. The absolute tolerance was set to 10^{-8} for the exact solver, and to 0 for the inexact ones, whereas the relative tolerance was set to 10^{-6} for the exact solver and to 10^{-2} for the inexact ones. We show the solution in Figure 2 (a).

Twist test. This is a static incompressible mechanics problem [BGO15], where we use the inf-sup stable finite elements $\mathbb{P}_2 - \mathbb{P}_0$ for the approximation of the displacement and pressure. Boundary conditions are given by a homogeneous on $\{z = 0\}$, a $\pi/6$ rotation in the x and y components on $\{z = L_z\}$, homogeneous data $\mathbf{f} = \mathbf{t} = \mathbf{0}$, and homogeneous Neumann conditions on the remaining parts of the boundary. The constitutive modelling is given by the following polyconvex potential:

$$\Psi_{\text{Twist}}(\mathbf{F}, p) = \alpha_p(\mathbf{F} : \mathbf{F} - 3) + \beta_p(\text{cof } \mathbf{F} : \text{cof } \mathbf{F} - 3) - (4\beta_s + 2\alpha_s) \log(\mathbf{F}) - p(\det(\mathbf{F}) - 1),$$

where $\alpha_p = \beta_p = \alpha_s = \beta_s = 9000$ and $\text{cof } \mathbf{F} := \det(\mathbf{F})\mathbf{F}^{-T}$. The logarithmic term with $\det(\mathbf{F})$ is required because of the weak imposition of the incompressibility constraint [ADVL⁺13]. The preconditioning of a saddle point problem is more challenging, so to obtain a more robust performance we used a lower Schur complement preconditioner based on a (\mathbf{d}, p) field split [BKM⁺12], we provide details on the Schur complement preconditioner in Appendix B. The displacement block uses an AMG preconditioner, the Schur complement block instead uses the SIMPLE preconditioner from fluid mechanics

[EHS⁺06], meaning that the inverse of the displacement block is approximated by the inverse of its diagonal, and for the resulting block we used a block Jacobi preconditioner, which we observed to suffice. Two additional comments are in place: (i) to calibrate the effectiveness of each solver in the Schur complement, we started from using a direct solver in each block to guarantee convergence in at most two iterations [Man90], and started relaxing the blocks from there. The difficulty of the problem is predominantly the displacement block, which can be appreciated by the simple preconditioner used for the Schur complement block; (ii) the pressure block presents much less degrees of freedom (DoFs), so we used a telescopic approach that creates an MPI subcommunicator with a fixed 25% of the original processes to avoid excessive communication. We show the solution in Figure 2 (b).

Heartbeat test. This problem represents a real-case study and is indeed our case of interest. It models the contraction of a human left ventricle in an idealized geometry given by a prolate ellipsoid, with a pointwise set of coordinates representing the muscle fiber orientation $(\mathbf{f}_0, \mathbf{s}_0, \mathbf{t}_0)$ obtained through rule based methods [BBPT12]. We consider a Guccione hyperelastic potential [GMW91], given by

$$\begin{aligned}\Psi(\mathbf{F}) &= \frac{C}{2} \left(e^{Q(\mathbf{F})} - 1 \right) + \frac{B}{2} (J - 1) \log J, \\ Q(\mathbf{F}) &= b_{ff} E_{ff}^2 + b_{ss} E_{ss}^2 + b_{nn} E_{nn}^2 + b_{fs} (E_{fs}^2 + E_{sf}^2) + b_{fn} (E_{fn}^2 + E_{nf}^2) + b_{sn} (E_{sn}^2 + E_{ns}^2),\end{aligned}$$

where $\mathbf{E} = \frac{1}{2} (\mathbf{F}^T \mathbf{F} - \mathbf{I})$ and $E_{ab} = \mathbf{E} \mathbf{a} \cdot \mathbf{b}$ for $a, b \in \{f, s, n\}$, which are the components of \mathbf{E} in the fiber-induced frame of reference $\{\mathbf{f}, \mathbf{s}, \mathbf{n}\}$. See [GMW91] for reference values of the related parameters. The Piola stress tensor is enriched with a fiber-wise force known as active stress that models the contraction of the muscle cells (cardiomyocytes):

$$\mathbf{P}(\mathbf{F}) = \frac{\partial \Psi}{\partial \mathbf{F}}(\mathbf{F}) + \mathbf{P}_{\text{act}}(\mathbf{F}), \quad \mathbf{P}_{\text{act}}(\mathbf{F}) = \gamma(t) \frac{(\mathbf{F} \mathbf{f}) \otimes \mathbf{f}}{|\mathbf{F} \mathbf{f}|},$$

where γ is a given function known as the activation function. For simplicity we consider an analytical activation given by

$$\gamma(t) = C_{\text{PA}} \max\{\sin(2\pi t/T), 0\},$$

where the peak activation constant is $C_{\text{PA}} = 10^4$ and the period is $T = 0.8$. More details on the generation of the fibers and the activation function can be found in [BBPT12] and [RSA⁺20] respectively. Boundary conditions are given by homogeneous Neumann on the endocardium $\partial\Omega_{\text{endo}}$ and the base $\partial\Omega_{\text{base}}$, whereas the epicardium $\partial\Omega_{\text{epi}}$ considers a Robin condition that models the friction of the epicardium with the pericardium [ULM02] and is given by

$$\mathbf{P}(\mathbf{F}) \mathbf{n} =: -(\mathbf{n} \otimes \mathbf{n}) \left(K_{\perp}^{\text{epi}} \mathbf{d} + C_{\perp}^{\text{epi}} \dot{\mathbf{d}} \right) - (\mathbf{I} - \mathbf{n} \otimes \mathbf{n}) \left(K_{\parallel}^{\text{epi}} \mathbf{d} + C_{\parallel}^{\text{epi}} \dot{\mathbf{d}} \right) = \mathbf{0} \quad \text{on } \partial\Omega_{\text{epi}}. \quad (4)$$

More information of the physical meaning of these boundaries and the motivation for these choices can be found in [QLRRB17]. The Krylov solver in this case is identical to the one used in the Cook test. We show the solution at $t = 0.2$ in Figure 2 (c).

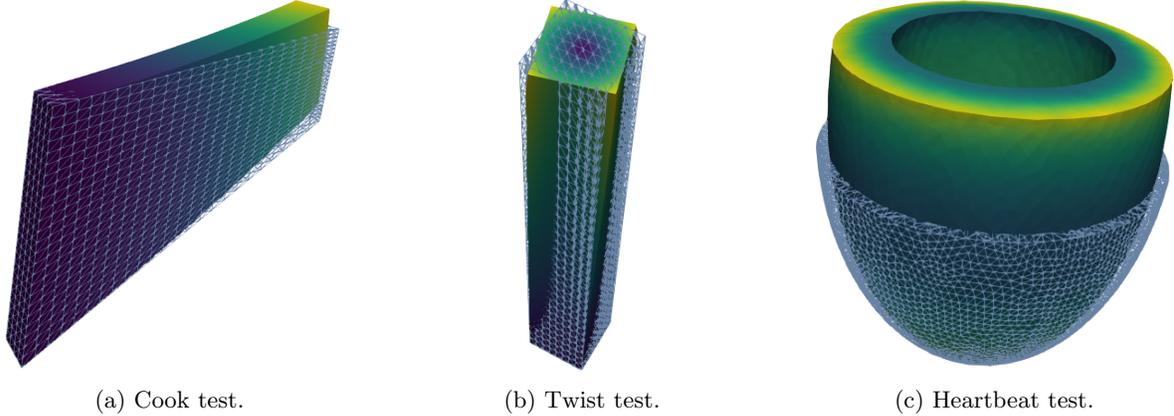


Figure 2: Solutions of the benchmark tests used. Color is scaled according to displacement norm $|\mathbf{d}|$, and the reference configuration is depicted with its mesh. In (c), the mesh is cut roughly at half the height of the geometry.

5 Theoretical results and solvers expected behavior

The purpose of this section is to show that the hypotheses from the convergence analysis, despite not being directly applicable to this problem, provide fundamental insight to guide the choice of the nonlinear solver. To state the results we consider an arbitrary scalar function $f : \mathbb{R}^m \rightarrow \mathbb{R}$ to be minimized and an initial point \mathbf{x}^0 . The results and their proofs can be found in [WN99]. First we state the convergence result of Newton’s method, where we remark that the inexact variant requires no additional hypotheses.

Theorem 1 (Newton methods) *Assume that there exists a point \mathbf{x}^* such that $\nabla f(\mathbf{x}^*) = 0$, $\nabla^2 f$ is Lipschitz continuous in a neighborhood of \mathbf{x}^* and $\nabla^2 f(\mathbf{x}^*)$ is positive definite. Then, \mathbf{x}^* is a local minimizer and, if the initial guess \mathbf{x}^0 is sufficiently close to \mathbf{x}^* , the iterates obtained by Newton’s method converge quadratically to \mathbf{x}^* .*

Now we state the general convergence theorem for the BFGS method.

Theorem 2 (Quasi-Newton methods) *Assume that f is twice continuously differentiable and that $\nabla^2 f$ is Lipschitz continuous at the minimizer \mathbf{x}^* . Assume also that the level*

$$\mathcal{L} = \{\mathbf{x} \in \mathbb{R}^m : f(\mathbf{x}) \leq f(\mathbf{x}^0)\},$$

is convex, that $\nabla^2 f$ is positive definite on \mathcal{L} and that the initial matrix \mathbf{B}^0 is positive definite. Then, the sequence of iterates obtained by the BFGS method converges superlinearly to \mathbf{x}^ .*

We make the following observations, which are fundamental to understand the performance of these methods.

- Both methods require the initial point to be sufficiently close. This can be seen in the BFGS method through the convexity of \mathcal{L} .
- The BFGS method is better suited for convex problems, given that the Hessian is required to be definite positive not only at the minimizer as Newton, but also in the entire level set \mathcal{L} .
- Inexact Newton-Krylov convergence holds under the same hypotheses of Newton’s method, plus some hypotheses on the relative tolerances η_k that guarantee the superlinear convergence. We consider tolerances that satisfy such hypotheses [EW94], and note that they can sometimes yield global convergence, unlike Newton’s method. This makes inexact Newton-Krylov methods attractive as they can be potentially more robust than a classic Newton method.

- The initial matrix approximation \mathbf{B}^0 is variable in the inexact BFGS method, meaning that its performance might not necessarily be better, nor more robust, than the case in which only the action of the preconditioner is considered. This is due to this case not being covered by the theory, so that in principle there could be additional hypotheses on the action of a variable \mathbf{B}^0 .
- Notes for different types of problems:
 - **Cook test:** In this case, the problem being solved is polyconvex, but if a sufficiently close initial guess is considered, we can possibly fall in an attraction basin. This means that both Newton and BFGS methods should converge, but both BFGS might possibly be less robust due to the convexity requirement of \mathcal{L} .
 - **Twist test:** This is a saddle point problem, so BFGS is not guaranteed to converge even if a sufficiently good initial approximation is considered. This can be seen from the positive definiteness hypothesis of the Hessian on \mathcal{L} .
 - **Heartbeat test:** The inertia term in elastodynamics results in a convex contribution to the variational principle associated to the problem, so we expect this case to be easier to solve in spite of the more complex nonlinearities involved. Instead, the nonlinearities should impact the number of nonlinear iterations.

6 Numerical results

In this section, we compare the performance of the solvers introduced in Sec. 3, for all tests under consideration. The scope of these tests is to assess the sensitivity of the methods with respect to the problem size, some of the problem parameters, and the number of computational cores. For the heartbeat problem, we consider four different meshes with the degrees of freedom shown in Table 1. We note that meshes 3 and 4 are refinements of meshes 1 and 2 respectively. Note as well that since BFGS-preonly does not solve any linear system, we do not report the linear iterations for this method. All models are implemented using the FEniCS library [ABH⁺15], and all interfaces with the underlying PETSc library are done by using the `petsc4py` interface [BAA⁺21]. The computations were performed on the EOS, INDACO and Galileo100 supercomputers.

Mesh name	DoFs \mathbb{P}_1	DoFs \mathbb{P}_2
Mesh 1	9375	60081
Mesh 2	20709	149511
Mesh 3	60081	421191
Mesh 4	149511	1123515

Table 1: Heartbeat test: Degrees of freedom yielded by each of the meshes under consideration.

6.1 Sensitivity with respect to problem size

To compare the performance of all methods, we report the total nonlinear iterations counts, the average Krylov iterations per nonlinear step and the solution time as we increase the number of degrees of freedom.

Cook test. We present the nonlinear iterations, linear iterations and CPU time in Table 2, varying the dimension of the problem from about 10 thousand to 1.5 million degrees of freedom. We first observe that notably all methods exhibit a robust behavior with respect to the degrees of freedom, except for the BFGS-preonly with second order finite elements. Indeed, the total and average linear iterations reported are roughly constant, but in general the inexact-BFGS yields lower average linear iterations. All methods exhibit an approximately linear increase of CPU time with respect to the degrees of

\mathbb{P}_1 discretization											
DoFs	NK			iNK			iB			B	
	nit	lit	T_{sol}	nit	lit	T_{sol}	nit	lit	T_{sol}	nit	T_{sol}
9375	4	11.8	0.7	5	9.4	0.8	6	8.7	0.4	112	4.0
64827	5	11.6	5.4	5	11.6	5.5	6	9.7	2.7	178	46.5
207831	6	10.3	23.2	5	10.6	20.3	7	10.0	12.7	59	17.2
479859	5	13.0	51.0	5	12.6	52.0	7	9.6	30.5	46	31.8
922383	5	13.4	105.0	5	13.0	103.2	7	11.1	69.5	43	52.7
1576875	5	13.8	178.6	5	13.4	183.2	7	10.1	115.7	48	112.7

\mathbb{P}_2 discretization											
DoFs	NK			iNK			iB			B	
	nit	lit	T_{sol}	nit	lit	T_{sol}	nit	lit	T_{sol}	nit	T_{sol}
9375	5	15.8	1.6	5	18.6	1.6	7	14.6	0.8	89	1.7
64827	4	20.2	11.5	5	17.4	13.6	7	14.1	7.8	203	65.3
207831	5	16.8	50.5	5	17.2	50.7	7	14.0	31.8	65	31.1
479859	4	20.0	102.2	5	18.2	118.0	7	14.1	79.0	326	1421.0
922383	4	20.8	198.7	5	17.6	226.0	8	13.3	168.8	332	2692.2
1576875	4	20.5	358.3	5	17.0	406.6	7	13.3	268.8	64	240.3

Table 2: Sensitivity with respect to problem size, Cook test. nit := nonlinear iterations, lit := average linear iterations per nonlinear iteration, T_{sol} := CPU time in seconds.

DoFs	NK			iNK		
	nit	lit	T_{sol}	nit	lit	T_{sol}
131163	5	551.8	836.7	15	162.8	980.2
223347	5	569.2	2045.5	15	170.1	1989.2
350955	6	699.7	5005.9	15	215.1	4100.7
519747	6	834.7	9332.2	16	233.7	7347.4

Table 3: Sensitivity with respect to problem size, Twist test. nit := nonlinear iterations, lit := average linear iterations per nonlinear iteration, T_{sol} := CPU time in seconds.

\mathbb{P}_1 discretization											
DoFs	NK			iNK			iB			B	
	nit	lit	T_{sol}	nit	lit	T_{sol}	nit	lit	T_{sol}	nit	T_{sol}
9375	3.6	7.6	0.4	4.8	5.6	0.5	7.9	3.0	0.2	24.6	0.3
20709	3.6	7.1	1.1	4.6	5.0	1.3	8.1	2.8	0.6	22.1	0.8
60081	3.6	8.3	3.0	5.0	5.8	3.7	8.3	3.0	1.8	25.9	2.3
149511	3.7	8.0	9.9	4.9	5.7	11.9	8.1	3.2	6.2	24.4	7.2
\mathbb{P}_2 discretization											
DoFs	NK			iNK			iB			B	
	nit	lit	T_{sol}	nit	lit	T_{sol}	nit	lit	T_{sol}	nit	T_{sol}
60081	3.7	10.4	4.4	5.0	7.0	5.1	9.6	3.9	2.6	32.7	2.5
149511	3.7	8.9	13.2	5.0	6.0	15.4	9.4	3.3	8.2	26.5	6.9
421191	3.8	10.8	41.4	5.2	7.1	45.7		F		33.0	23.5
1123515	3.8	11.0	125.5	5.1	7.2	137.0		F		32.2	70.9

Table 4: Sensitivity with respect to problem size, Heartbeat test. nit := nonlinear iterations, lit := average linear iterations per nonlinear iteration, T_{sol} := CPU time in seconds.

freedom, with the exception of BFGS-preonly with second order elements. In general, BFGS methods outperform Newton methods in terms of the execution time, but attention must be given to the second order scenario. The ratio between the CPU times of the best method (BFGS-preonly) and Newton’s method in the case with the largest amount of DoFs is 0.63 for \mathbb{P}_1 and 0.67 for \mathbb{P}_2 elements.

Twist test. We report the nonlinear iterations, average linear iterations per nonlinear step, and CPU times in Table 3. We observe that BFGS methods do not converge for this problem, so we can not report results for them. This holds not only for both methods presented (BFGS, inexact-BFGS), as we also tried the BFGS-exact variant described in Section 3, which was ineffective as well. Both Newton-Krylov and inexact Newton-Krylov methods present a very mild increase in the nonlinear iterations, but the inexact variant presents a plateau of average linear iterations, whereas the Newton-Krylov method presents a monotonic increase in the average linear iterations. This results in the inexact method being a faster method, and the difference in CPU times increases with the problem size. The ratio between the CPU times of the inexact Newton-Krylov method and Newton-Krylov method in the case with the largest amount of DoFs is 0.79.

Heartbeat test. We report the nonlinear iterations, average linear iterations per nonlinear iteration, and CPU time for first and second order finite elements averaged during the first timestep in Table 4 and their time evolution in Figure 3. We note that all methods are robust with respect to the degrees of freedom, except the inexact-BFGS, as it does not converge for fine meshes when using second order elements. Both Newton methods present similar nonlinear iteration counts, and all methods have a similar average linear iteration count. This results in the inexact-BFGS having a lower average linear iteration count overall as it requires more nonlinear iterations. This has an impact in the overall solution time, where BFGS methods perform much better than Newton methods. The ratio between the CPU times of the best method (inexact-BFGS in \mathbb{P}_1 , BFGS-preonly in \mathbb{P}_2) and the Newton-Krylov method is 0.63 for \mathbb{P}_1 and 0.57 for \mathbb{P}_2 elements.

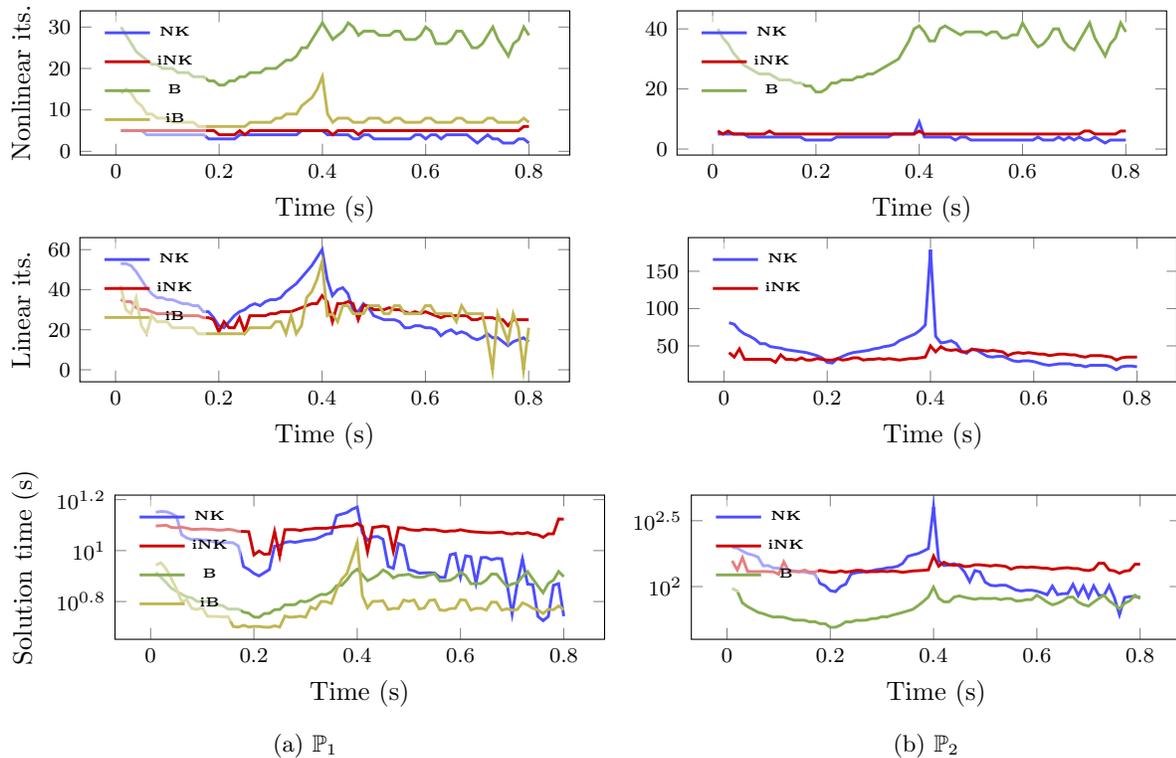


Figure 3: Sensitivity with respect to problem size, Heartbeat test. Time evolution of nonlinear iterations (nit), average linear iterations per nonlinear iteration (lit) and CPU time in seconds for \mathbb{P}_1 and \mathbb{P}_2 discretizations in mesh 4.

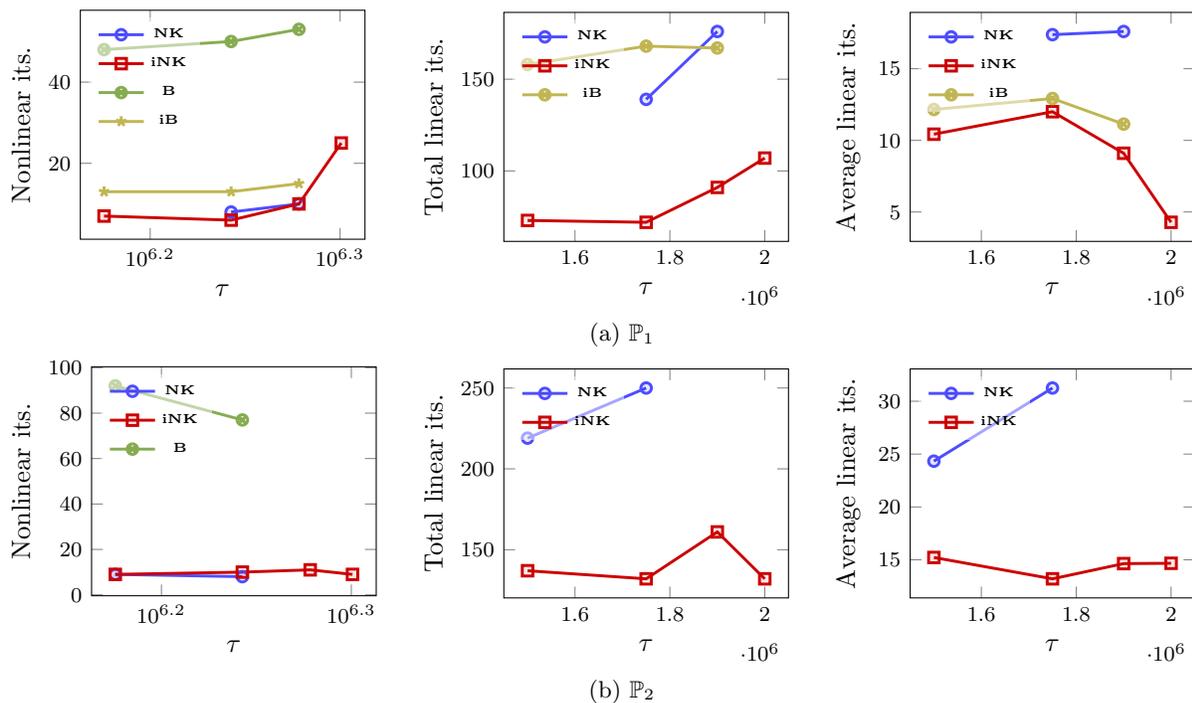


Figure 4: Robustness with respect to the data, Cook test. Comparison of the nonlinear, total linear and average linear iterations.

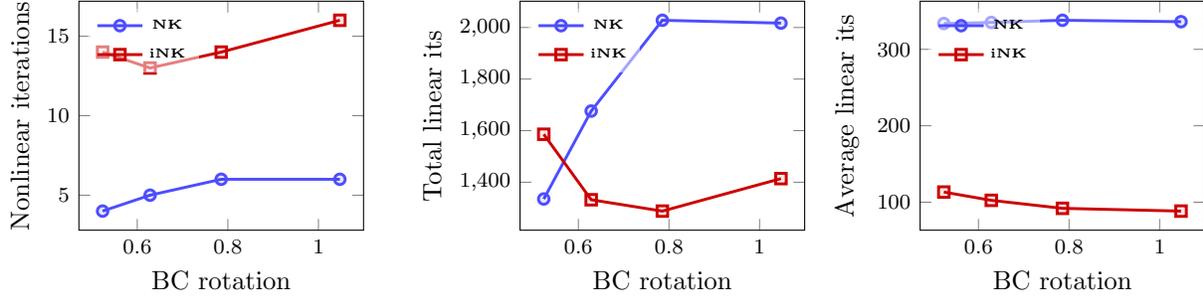


Figure 5: Robustness with respect to the data, Twist test. Nonlinear, total linear and average linear iterations with respect to the rotation angle.

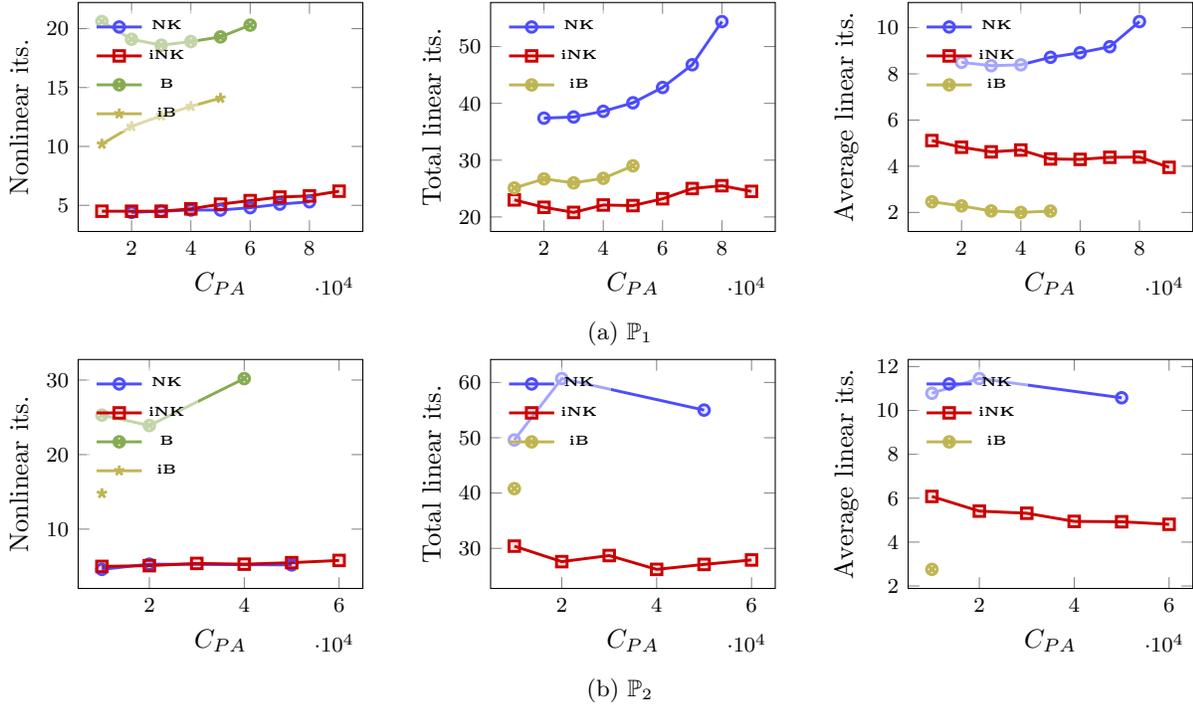


Figure 6: Robustness with respect to the data, Heartbeat test. Comparison of the nonlinear, total linear and average linear iterations, averaged over the first 10 timesteps.

6.2 Robustness with respect to the data

To study the robustness, we varied in each test case the parameter which resulted in a larger deformation. In the Cook test this refers to the vertical load τ , in the Twist case this refers to the angle of rotation in the boundary conditions, and on the Heartbeat case this refers to the peak activation constant C_{AP} .

Cook test. We vary the magnitude of the distributed traction term τ on the right from $1.5 \cdot 10^6$ Pa to $2 \cdot 10^6$ Pa, where values over the upper bound considered make all methods diverge. The nonlinear iterations, total linear iterations and average linear iterations per nonlinear step are shown in Figure 4. We have computed the solution for four values of τ , and plotted the result only when the solver converged. Most notably, only the inexact Newton-Krylov method is able to solve the highest load

value, confirming that it can show more robust behavior than the Newton-Krylov method. Moreover, the inexact-BFGS does not converge for some values in case of second order elements, making it the least robust method in this scenario. In general, the difference in robustness is much clearer with second order elements, as also BFGS-preonly and Newton-Krylov converge only for two of the four values considered for τ . When looking at the linear iteration counts, we see that the methods that converge tend to maintain a roughly constant number of average linear iterations, so that the increase in total linear iterations is given by the additional nonlinear iterations incurred. The only exception is the Newton-Krylov method for second order elements. Interestingly, inexact-BFGS is more robust than BFGS-preonly for first order elements, but instead BFGS-preonly is slightly more robust than inexact-BFGS for second order elements. There is no advantage in considering a Newton-Krylov method, but instead its inexact variant is much more robust with respect to τ .

Twist test. We vary the angle in the boundary conditions from $\pi/6$ to $\pi/2$, and show the results in Figure 5. We note that both Newton methods converge for all values under consideration, and diverge when higher ones are considered, so there is no significant difference in terms of robustness between them. Still, the Newton-Krylov method presents an increase in the total linear iterations, whereas the inexact variant shows a more robust iteration count.

Heartbeat test. We vary the peak activation C_{PA} from 10^4 to 10^5 , and plot only the results when the methods converged. In this case, we note that Newton methods provide an improved robustness for both first and second order finite elements, and as in the other tests, the inexact variant is more robust. This effect is inverted in the BFGS methods, which are less robust and instead the BFGS-preonly shows an improved robustness with respect to the activation force. We again observe that the total linear iterations of the Newton-Krylov method increase with the problem difficulty, whereas the inexact Newton-Krylov method maintains a roughly constant number of linear iterations despite the increase of the problem difficulty.

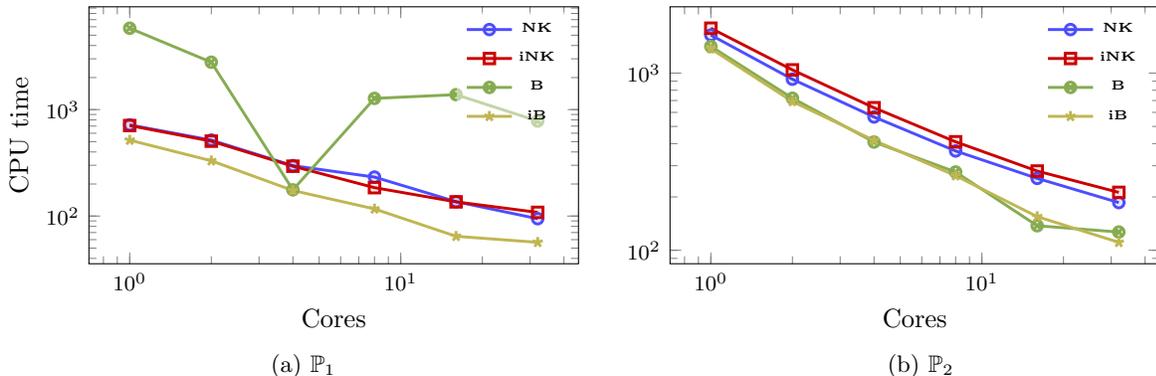


Figure 7: Scalability, Cook test. CPU times with respect to the number of cores.

\mathbb{P}_1 discretization, DoFs = 1576875											
cores	NK			iNK			iB			B	
	nit	lit	T_{sol}	nit	lit	T_{sol}	nit	lit	T_{sol}	nit	T_{sol}
1	5	12.4	718.2	5	12.0	708.1	7	10.4	516.4	184	5815.2
2	5	13.0	616.5	5	12.0	505.2	7	10.7	330.6	132	2786.8
4	5	13.2	297.5	5	12.6	293.9	7	10.1	173.8	43	176.2
8	6	11.0	232.4	5	13.0	184.9	7	12.1	116.7	146	1274.1
16	5	13.6	136.5	5	12.4	135.9	7	9.4	64.5	203	1382.1
32	4	16.3	94.4	5	13.8	108.3	7	11.4	56.6	146	780.8

\mathbb{P}_2 discretization, DoFs = 1576875											
cores	NK			iNK			iB			B	
	nit	lit	T_{sol}	nit	lit	T_{sol}	nit	lit	T_{sol}	nit	T_{sol}
1	4	20.3	1639.2	5	15.8	1791.4	7	13.4	1372.3	62	1414.9
2	4	20.0	923.4	5	16.0	1042.3	7	13.4	691.0	62	721.5
4	4	20.5	565.3	5	16.4	637.3	7	13.8	417.0	62	408.2
8	4	20.8	363.5	5	16.8	409.4	7	14.6	263.5	61	276.6
16	4	20.8	254.5	5	16.8	279.9	7	15.1	154.9	62	137.6
32	4	21.5	185.9	5	17.2	212.2	7	14.0	110.9	63	126.7

Table 5: Scalability, Cook test. nit := nonlinear iterations, lit := average linear iterations per nonlinear iteration, T_{sol} := CPU time in seconds.

DoFs = 519747						
cores	NK			iNK		
	nit	lit	T_{sol}	nit	lit	T_{sol}
1	6	800.0	$3.68 \cdot 10^4$	16	240.4	$2.61 \cdot 10^4$
2	6	806.2	$2.13 \cdot 10^4$	16	225.3	$1.62 \cdot 10^4$
4	6	806.0	$1.37 \cdot 10^4$	16	222.4	$1.06 \cdot 10^4$
8	6	826.3	$8.02 \cdot 10^3$	16	237.3	$6.56 \cdot 10^3$
16	6	835.5	$4.62 \cdot 10^3$	16	240.3	$3.94 \cdot 10^3$
32	6	828.5	$3.44 \cdot 10^3$	16	244.8	$3.10 \cdot 10^3$

Table 6: Scalability, Twist test. nit := nonlinear iterations, lit := average linear iterations per nonlinear iteration, T_{sol} := CPU time in seconds.

\mathbb{P}_1 discretization, DoFs = 149511											
cores	NK			iNK			iB			B	
	nit	lit	T_{sol}	nit	lit	T_{sol}	nit	lit	T_{sol}	nit	T_{sol}
1	4.5	9.6	111.0	5.0	6.0	113.6	10.9	2.4	69.4	23.2	87.7
2	4.5	9.6	60.9	5.0	6.0	62.3	11.0	2.5	30.1	23.2	45.5
4	4.5	9.6	26.5	5.0	6.0	33.2	10.8	2.7	18.8	23.3	23.3
8	4.5	9.8	18.2	5.0	6.2	18.9	11.0	2.6	10.3	23.9	12.6
16	4.5	9.7	11.1	5.0	6.1	11.5	10.8	2.6	4.8	23.5	6.8
32	4.5	9.7	7.6	5.0	6.2	7.9	10.9	2.8	3.1	23.6	4.0

\mathbb{P}_2 discretization, DoFs = 1123515											
cores	NK			iNK			iB			B	
	nit	lit	T_{sol}	nit	lit	T_{sol}	nit	lit	T_{sol}	nit	T_{sol}
1	4.6	13.4	1061.0	5.3	6.5	965.4	14.1	3.5	637.4	30.2	515.1
2	4.6	13.5	512.8	5.3	6.4	569.0	13.4	3.6	287.4	30.4	287.7
4	4.6	13.5	324.5	5.0	6.5	294.3	15.3	3.4	184.9	30.5	149.8
8	4.6	13.4	184.1	5.3	6.6	176.0	14.5	3.6	100.7	30.1	81.7
16	4.6	13.4	113.2	5.2	6.6	107.8	14.5	3.3	59.1	30.5	40.5
32	4.6	13.4	82.0	5.1	6.5	74.5	14.2	3.5	43.3	30.3	33.6

Table 7: Scalability, Heartbeat test. nit := nonlinear iterations, lit := average linear iterations per nonlinear iteration, T_{sol} := CPU time in seconds.

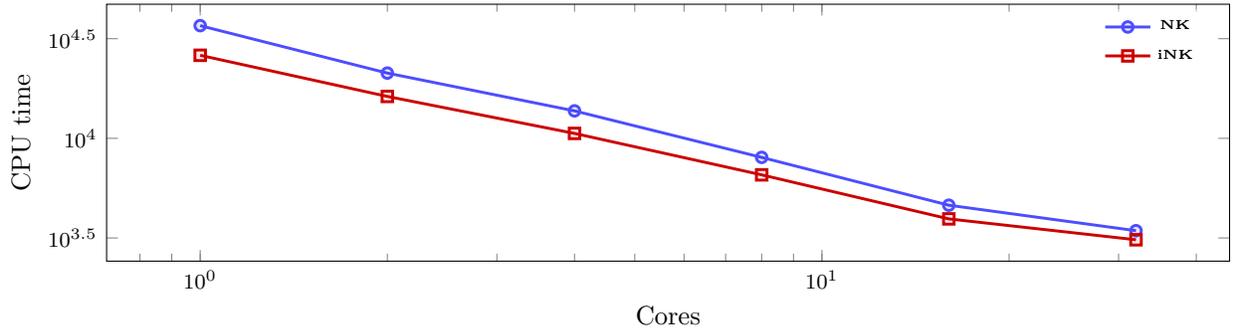


Figure 8: Scalability, Twist test. CPU time with respect to the number of cores.

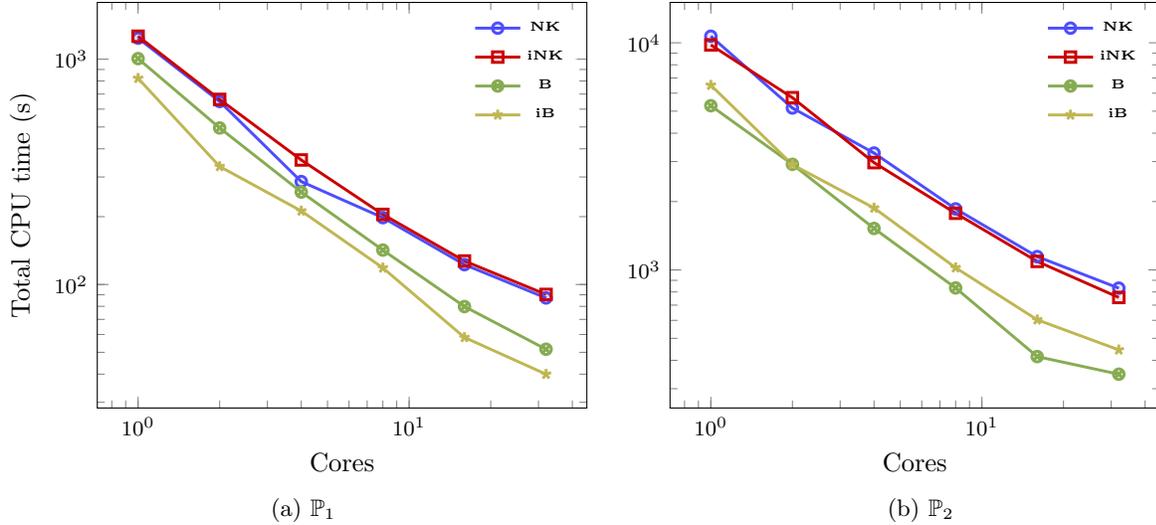


Figure 9: Scalability, Heartbeat test. CPU times with respect to the number of cores.

6.3 Scalability

We test the strong scalability with respect to the finest mesh used in the performance tests. In the heartbeat test, we consider only the first 10 timesteps.

Cook test. We report the nonlinear and average linear iterations and the CPU times in Table 5. The CPU times are also displayed in Figure 7. We note that all methods are robust with respect to the number of cores, as could have been expected in virtue of the robustness of the preconditioner. Still, it is very interesting to see that there is a point in which BFGS-preonly converges with a much lower iteration number in the \mathbb{P}_1 case, namely when using 8 processors. This is the only case in which the CPU time is comparable to that of the inexact-BFGS method. For any other number of cores, the inexact-BFGS outperforms all other methods for both \mathbb{P}_1 and \mathbb{P}_2 elements. As already observed in the performance test, the BFGS-preonly method is less robust than the other methods.

Twist test. We report the nonlinear and average linear iteration counts, together with the CPU times in Table 6. The CPU times are also displayed in Figure 8. We note that this test in particular is in complete agreement with the previous ones, meaning that both Newton-Krylov methods perform similarly and exhibit an overall robust behavior. Both methods present adequate strong scaling, with the inexact method showing better solution times.

Heartbeat test. For this test, we report the nonlinear and average linear iterations, and CPU times averaged for the first 10 timesteps in Table 7. The total CPU times as function of the number of cores are displayed in Figure 9. We can appreciate that all methods are robust with respect to the number of cores, as these yield no significant variations. From the scalability curves shown in Figure 9, we note that all methods show adequate scaling, with the BFGS methods presenting an overall better performance. More specifically, inexact-BFGS is faster for \mathbb{P}_1 , and BFGS-preonly is faster for \mathbb{P}_2 .

7 Conclusions

This work presents a detailed numerical study of some of the main numerical solvers used in large scale simulations of nonlinear mechanics, where all of the fundamental components of the methods are included: the nonlinear solver, the linear solver and the preconditioner. The main finding from this study is that superlinear solvers such as inexact Newton-Krylov and BFGS methods present an overall robustness similar

to the gold standard represented by the Newton-Krylov method in nonlinear elasticity, but with improved solution times.

In difficult problems such as static and incompressible mechanics, the choice of the solver can be defined a-priori according to the convergence hypotheses required for each method, where in fact we have seen the BFGS methods fail for incompressible mechanics, as they usually require some form of convexity to guarantee convergence. Indeed, BFGS methods showed the best solution times in static mechanics, while inexact Newton-Krylov showed the best performance for the incompressible tests. The time-dependent heartbeat test showed us that the convexity contributed by the inertia term can further improve the performance of the BFGS method, which showed much more robust iteration counts in this scenario than those shown in the static mechanics case. This comes as a surprise, given that the constitutive model used in the heartbeat test is much more complex than the models of the other tests. These nonlinear solvers present no drawbacks in terms of both performance and strong scalability, suggesting that in the context of nonlinear elasticity inexact Newton-Krylov and quasi-Newton methods should be preferred to standard Newton-Krylov methods.

We conclude by mentioning two possible extensions of this work. The first one is the inclusion in this study of Jacobian lagging techniques [BB13], which consist in avoiding the reassembly of the Jacobian for some Newton iterations. These lagging techniques could potentially further improve the performance of the Newton methods considered. Instead, inspired by our inexactness taxonomy, we note that a Jacobian reassembly could be considered for the BFGS methods, where the Jacobian gets sometimes reassembled to improve the quality of the initial Jacobian. Both strategies are non trivial to implement, and require a detailed problem-specific study to obtain the best performance. The second interesting extension is the inclusion in our nonlinear mechanics study of a polyconvex potential [Bal76], even if solvers for this specific type of potential are not yet available, to the best of our knowledge. One interesting work in this direction is [BGO15], where a Hu-Washizu [Was68] formulation is used to exploit the convexity of the auxiliary function arising from the polyconvex potential. Still, this formulation results in a problem with many auxiliary variables, and its efficient numerical approximation remains a research topic.

A PETSc options

The commands used to invoke each of the methods described in Section 3 is detailed in what follows.

- **Newton-Krylov (NK)**

```
-snes_type newtonls
-ksp_type gmres # minres should work as well
-pc_type hypre
-ksp_atol 1e-10
-ksp_rtol 1e-6
-snes_linesearch_type basic
```

Listing 1: PETSc commands to use Newton-Krylov.

- **Inexact Newton-Krylov (iNK)**

```
-snes_type newtonls
-ksp_type gmres
-pc_type hypre
-ksp_atol 1e-14
-ksp_rtol 1e-2 # Should not be used, set for safety
-snes_ksp_ew
-snes_ksp_ew_rtol0 1e-1
-snes_ksp_ew_rtolmax 0.1
-snes_linesearch_type basic
```

Listing 2: PETSc commands to use inexact Newton-Krylov.

- **Inexact-BFGS (iB)**

```

-snes_type qn
-ksp_type gmres
-pc_type hypre
-ksp_atol 1e-14
-ksp_rtol 1e-2
-snes_qn_type lbfgs
-snes_qn_m 50
-snes_qn_scale_type jacobian
-snes_linesearch_type basic

```

Listing 3: PETSc commands to use inexact-BFGS.

- **BFGS-preonly (B)**

```

-snes_type qn
-ksp_type preonly
-pc_type hypre
-snes_qn_type lbfgs
-snes_qn_m 50
-snes_qn_scale_type jacobian
-snes_linesearch_type basic

```

Listing 4: PETSc commands to use BFGS.

B Schur complement preconditioners

A Schur complement preconditioner is one that arises from a block LU factorization according to two (arbitrary) index sets of a matrix. If we consider a block matrix \mathbf{M} given by the general structure

$$\mathbf{M} = \begin{bmatrix} \mathbf{A} & \mathbf{B}_1 \\ \mathbf{B}_2 & \mathbf{C} \end{bmatrix},$$

with \mathbf{A} invertible, a Gaussian elimination procedure yields

$$\mathbf{M} = \begin{bmatrix} I & \mathbf{0} \\ \mathbf{B}_2\mathbf{A}^{-1} & \mathbf{0} \end{bmatrix} \begin{bmatrix} \mathbf{A} & \mathbf{0} \\ \mathbf{0} & \mathbf{C} - \mathbf{B}_2\mathbf{A}^{-1}\mathbf{B}_1 \end{bmatrix} \begin{bmatrix} I & \mathbf{A}^{-1}\mathbf{B}_1 \\ \mathbf{0} & I \end{bmatrix}. \tag{5}$$

We note that if \mathbf{C} is invertible, the same procedure can be applied with respect to it. Schur complement based preconditioners enjoy excellent theoretical properties, as the preconditioned system possesses at most 3 distinct eigenvalues [MGW00], implying that it converges in at most 3 iterations of a Krylov subspace method. This is true whenever the Schur complement $\mathbf{S} = \mathbf{C} - \mathbf{B}_2\mathbf{A}^{-1}\mathbf{B}_1$ is evaluated exactly, which is usually computationally intractable. One simple approximation, which is the one we use for the preconditioning the incompressible mechanics problem, is to consider the approximation $\mathbf{A}^{-1} \approx \text{diag}^{-1}(\mathbf{A})$. Another important point is that in general using all three blocks arising from the factorization in (5) is not necessary, and instead it is sufficient to consider a lower factorization (first two blocks) or an upper factorization (last two blocks). Whenever $\mathbf{C} = \mathbf{0}$, it is also possible to use a diagonal factorization (middle block only).

Acknowledgments

N. Barnafi and L. F. Pavarino have been supported by grants of MIUR (PRIN 2017AXL54F_002) and INdAM-GNCS. N. Barnafi and S. Scacchi have been supported by grants of MIUR (PRIN 2017AXL54F_003) and INdAM-GNCS. The Authors are also grateful to the University of Pavia, the University of Milan, and the CINECA laboratory for the usage of the EOS, INDACO and Galileo100 clusters, respectively.

References

- [ABH⁺15] M Alnæs, J Blechta, J Hake, A Johansson, B Kehlet, A Logg, C Richardson, J Ring, ME Rognes, and GN Wells. The FEniCS project version 1.5. *Archive of Numerical Software*, 3(100), 2015.
- [ADVL⁺13] F Auricchio, LB Da Veiga, C Lovadina, A Reali, RL Taylor, and P Wriggers. Approximation of incompressible large deformation elastic problems: some unresolved issues. *Computational Mechanics*, 52(5):1153–1167, 2013.
- [BAA⁺21] S Balay, S Abhyankar, MF Adams, J Brown, P Brune, K Buschelman, L Dalcin, A Dener, V Eijkhout, W Gropp, D Karpeyev, D Kaushik, M Knepley, D May, L Curfman McInnes, R Mills, T Munson, K Rupp, P Sanan, B Smith, S Zampini, H Zhang, and H Zhang. PETSc users manual. Technical Report ANL-95/11 - Revision 3.13, Argonne National Laboratory, 2021.
- [Bal76] JM Ball. Convexity conditions and existence theorems in nonlinear elasticity. *Archive for rational mechanics and Analysis*, 63(4):337–403, 1976.
- [BB13] J Brown and P Brune. Low-rank quasi-Newton updates for robust jacobian lagging in Newton methods. In *Proceedings of the 2013 International Conference on Mathematics and Computational Methods Applied to Nuclear Science and Engineering*, volume 201, pages 1–2, 2013.
- [BBPT12] JD Bayer, RC Blake, G Plank, and NA Trayanova. A novel rule-based algorithm for assigning myocardial fiber orientation to computational heart models. *Annals of Biomedical Engineering*, 40(10):2243–2254, 2012.
- [BGO15] J Bonet, AJ Gil, and R Ortigosa. A computational framework for polyconvex large strain elasticity. *Computer Methods in Applied Mechanics and Engineering*, 283:1061–1094, 2015.
- [BKM⁺12] J Brown, MG Knepley, DA May, LC McInnes, and B Smith. Composable linear solvers for multiphysics. In *2012 11th International Symposium on Parallel and Distributed Computing*, pages 55–62. IEEE, 2012.
- [BRKN17] M Borregales, FA Radu, K Kumar, and JM Nordbotten. Robust iterative schemes for non-linear poromechanics. *arXiv preprint arXiv:1702.00328*, 2017.
- [CD04] C Carstensen and G Dolzmann. An a priori error estimate for finite element discretizations in nonlinear elasticity for polyconvex materials under small loads. *Numerische Mathematik*, 97(1):67–80, 2004.
- [CFPS15] P Colli Franzone, LF Pavarino, and S Scacchi. Parallel multilevel solvers for the cardiac electro-mechanical coupling. *Appl. Numer. Math.*, 95:140–153, 2015.
- [CGK⁺98] X-C Cai, WD Gropp, DE Keyes, RG Melvin, and DP Young. Parallel Newton–Krylov–Schwarz algorithms for the transonic full potential equation. *SIAM J. Sci. Comput.*, 19(1):246–265, 1998.
- [CGKT94] X-C Cai, WD Gropp, DE Keyes, and MD Tidriri. Newton-Krylov-Schwarz methods in CFD. In *Numerical methods for the Navier-Stokes equations*, pages 17–30. 1994.
- [Cia21] PG Ciarlet. *Mathematical elasticity: Three-dimensional elasticity*. SIAM, 2021.
- [DES82] RS Dembo, SC Eisenstat, and T Steihaug. Inexact Newton methods. *SIAM Journal on Numerical analysis*, 19(2):400–408, 1982.
- [DS96] JE Dennis and RB Schnabel. *Numerical methods for unconstrained optimization and nonlinear equations*. SIAM, 1996.

- [EG13] A Ern and JL Guermond. *Theory and practice of finite elements*, volume 159. Springer Science & Business Media, 2013.
- [EHS⁺06] H Elman, VE Howle, J Shadid, R Shuttleworth, and R Tuminaro. Block preconditioners based on approximate commutators. *SIAM Journal on Scientific Computing*, 27(5):1651–1668, 2006.
- [EW94] SC Eisenstat and HF Walker. Globally convergent inexact Newton methods. *SIAM Journal on Optimization*, 4(2):393–422, 1994.
- [EW96] SC Eisenstat and HF Walker. Choosing the forcing terms in an inexact Newton method. *SIAM Journal on Scientific Computing*, 17(1):16–32, 1996.
- [FY02] RD Falgout and UM Yang. hypre: A library of high performance preconditioners. In *International Conference on Computational Science*, pages 632–641. Springer, 2002.
- [GMW91] JM Guccione, AD McCulloch, and LK Waldman. Passive material properties of intact ventricular myocardium determined from a cylindrical model. *Journal of biomechanical engineering*, 113(1):42–55, 1991.
- [GP88] JC Gelin and P Picart. Use of quasi-Newton methods for large strain elastic-plastic finite element computations. *Communications in applied numerical methods*, 4(4):457–469, 1988.
- [GVL96] GH Golub and CF Van Loan. *Matrix computations*, 1996.
- [Hac13] W Hackbusch. *Multi-grid methods and applications*, volume 4. Springer Science & Business Media, 2013.
- [Hol02] GA Holzapfel. Nonlinear solid mechanics: A continuum approach for engineering science. *Mechanica*, 37(4):489–490, 2002.
- [JHN11] N Jansson, J Hoffman, and M Nazarov. Adaptive simulation of turbulent flow past a full car model. In *SC’11: Proceedings of 2011 International Conference for High Performance Computing, Networking, Storage and Analysis*, pages 1–8. IEEE, 2011.
- [LBK17] T Liu, S Bouaziz, and L Kavan. Quasi-Newton methods for real-time simulation of hyperelastic materials. *ACM Transactions on Graphics (TOG)*, 36(3):1–16, 2017.
- [LLS05] S Linge, G Lines, and J Sundnes. Solving the heart mechanics equations with Newton and quasi Newton methods—a comparison. *Computer methods in biomechanics and biomedical engineering*, 8(1):31–38, 2005.
- [Man90] J Mandel. On block diagonal and schur complement preconditioning. *Numerische Mathematik*, 58(1):79–93, 1990.
- [MGW00] MF Murphy, GH Golub, and AJ Wathen. A note on preconditioning for indefinite linear systems. *SIAM Journal on Scientific Computing*, 21(6):1969–1972, 2000.
- [Noc80] J Nocedal. Updating quasi-Newton matrices with limited storage. *Mathematics of computation*, 35(151):773–782, 1980.
- [PM16] J-P Pelteret and A McBride. *The deal.ii code gallery: Quasi-static finite-strain compressible elasticity*, 2016.
- [QLRRB17] A Quarteroni, T Lassila, S Rossi, and R Ruiz-Baier. Integrated Heart—Coupling multiscale and multiphysics models for the simulation of the cardiac function. *Computer Methods in Applied Mechanics and Engineering*, 314:345–407, 2017.

- [RSA⁺20] F Regazzoni, M Salvador, PC Africa, M Fedele, L Dede, and A Quarteroni. A cardiac electromechanics model coupled with a lumped parameters model for closed-loop blood circulation. part I: model derivation. *arXiv e-prints*, 2020.
- [Rud16] S Ruder. An overview of gradient descent optimization algorithms. *arXiv preprint arXiv:1609.04747*, 2016.
- [Saa03] Y Saad. *Iterative methods for sparse linear systems*. SIAM, 2003.
- [SHOL06] T Smith, R Hooper, C Ober, and A Lorber. Intelligent nonlinear solvers for computational fluid dynamics. In *44th AIAA Aerospace Sciences Meeting and Exhibit*, page 1483, 2006.
- [ULM02] TP Usyk, IJ LeGrice, and AD McCulloch. Computational model of three-dimensional cardiac electromechanics. *Computing and Visualization in Science*, 4(4):249–257, Jul 2002.
- [Was68] K Washizu. Variational methos in elasticity and plasticity. *International Series of Monographs in Aeronautics and Astronautics*, 1968.
- [WDE07] M Weiser, P Deuffhard, and B Erdmann. Affine conjugate adaptive Newton methods for non-linear elastomechanics. *Optimisation Methods and Software*, 22(3):413–431, 2007.
- [WN99] S Wright and J Nocedal. Numerical optimization. *Springer Science*, 35(67-68):7, 1999.
- [ZTNZ77] OC Zienkiewicz, RL Taylor, P Nithiarasu, and JZ Zhu. *The finite element method*, volume 3. McGraw-hill London, 1977.