BENCHMARKING NUMERICAL ALGORITHMS FOR HARMONIC MAPS INTO THE SPHERE

SÖREN BARTELS[1], KLAUS BÖHNLEIN[2], CHRISTIAN PALUS[1], OLIVER SANDER[2]

[1] *Department of Applied Mathematics, University of Freiburg*
[2] *Institute of Numerical Mathematics, Technische Universität Dresden*

ABSTRACT. We numerically benchmark methods for computing harmonic maps into the unit sphere, with particular focus on harmonic maps with singularities. For the discretization we compare two different approaches, both based on Lagrange finite elements. While the first method enforces the unit-length constraint only at the Lagrange nodes, the other one adds a pointwise projection to fulfill the constraint everywhere. For the solution of the resulting algebraic problems we compare a nonconforming gradient flow with a Riemannian trust-region method. Both are energy-decreasing and can be shown to converge globally to stationary points of the discretized Dirichlet energy. We observe that while the nonconforming and the conforming discretizations both show similar behavior for smooth problems, the nonconforming discretization handles singularities better. On the solver side, the second-order trust-region method converges after few steps, whereas the number of gradient-flow steps increases proportionally to the inverse grid element diameter.

1. INTRODUCTION

Harmonic maps are stationary configurations of the Dirichlet energy [24, 33]. In this work we focus on maps into the unit sphere $S^{m-1} := \{x \in \mathbb{R}^m \ : \ |x| = 1\}$. More formally, given a Lipschitz domain $\Omega \subset \mathbb{R}^n$ with $n \in \{2, 3\}$, we seek stationary configurations $\mathbf{u} \colon \Omega \to S^{m-1}$ with $m \in \{2, 3\}$, $m \geq n$, of the Dirichlet energy

$$E[\mathbf{u}] := \frac{1}{2} \int_\Omega |\nabla \mathbf{u}|^2 \, \mathrm{d}x.$$

Equivalently, we will sometimes regard the problem as looking for vector fields $\mathbf{u} : \Omega \to \mathbb{R}^m$ that make $E[\cdot]$ stationary while fulfilling the constraint

$$|\mathbf{u}|^2 = 1 \text{ almost everywhere in } \Omega.$$

For well-posedness we require Dirichlet boundary conditions

$$\mathbf{u} = \mathbf{u}_\mathrm{D} \qquad \text{on } \partial\Omega,$$

for a function $\mathbf{u}_\mathrm{D} : \Omega \to S^{m-1}$ of suitable smoothness [33, Chap. 3.4].

Harmonic maps into spheres have numerous practical applications. They appear in models of liquid crystal materials, which, on a microscopic level, consist of rod-shaped molecules that exhibit a natural desire for mutual alignment. Popular macroscopic liquid crystal models are based on the Oseen–Frank energy, which is the Dirichlet energy of the field of molecule orientations.

The emergence of new technologies in the manufacturing of liquid crystal compound materials as well as new ideas for technical applications [21, 35, 53, 55] have led to an ongoing interest also in corresponding simulation schemes [14, 42, 52].

Harmonic maps into spheres also play a role in micromagnetics, which models orientation fields of magnetic dipoles. The Dirichlet energy is the simple-most representative of a family of different energy-based models [22]. It also serves as a prototype energy that already captures a range of interesting effects. Furthermore, harmonic maps serve as the basis for more complicated energies such as the Ginzburg–Landau model or the chiral skyrmions discussed in [39]. In this context singular harmonic maps are particularly interesting, since they arise as limits of minimizers of the Ginzburg–Landau energy [40].

Finally, from a mathematical point of view harmonic maps are interesting in their own right, and a considerable body of literature exists with investigations of mathematical properties of harmonic maps. Overviews and further literature can be found in [18, 33, 49].

The construction of finite element (FE) methods for the approximation of harmonic maps requires special care. The central challenge is the handling of the non-Euclidean image space S^{m-1}, or, equivalently, of the nonlinear, nonconvex constraint $|\mathbf{u}(x)|^2 = 1$ for (almost) all x in the domain Ω. This affects both the discretization of the problem as well as solution strategies for the resulting algebraic systems. In the last decades various approaches have been proposed that can be employed in the numerical approximation of harmonic maps. For example, several authors have proposed finite-difference approximations of the Dirichlet energy [19, 54], as well as point-relaxation methods [37] and gradient-type methods [5, 8]. For the constraint, approaches like parametrizations [51], Lagrange multipliers [19, 34] and penalization [23, 38] have been proposed. Several works also treat numerical methods for more general problems such as p-harmonic maps [51] or fractional harmonic maps [6].

While all these methods warrant thorough benchmarking, for reasons of space we limit this paper to two particular discretizations and two particular solver algorithms. In [9, 13], Bartels and coworkers proposed a nonconforming discretization method that consists of $\mathbb{R}^m$-valued Lagrange finite elements that are constrained to fulfill the unit-length constraint only at the Lagrange points. Stationary points of the discretized energy (1) are approximated via an implicit discrete gradient flow employing a linearization of the unit-length constraint (2). This solver approach is nonconforming in the sense that the iterates slowly accumulate a violation of the constraints. Bartels et al. showed, however, that this constraint violation remains bounded in terms of the time step size τ, and that the discrete solutions weakly converge to stationary points of the Dirichlet energy as the grid element size h and the time step size τ go to zero. The energy monotonicity property of the discrete gradient flow implies a convergence rate of $\mathcal{O}\big((\tau k)^{-\frac{1}{2}}\big)$ for the norms of the corrections, where k denotes the number of solver steps.

The second approach has been proposed by Sander et al. [43, 44], who aimed at a completely conforming method. The authors construct so-called geometric finite elements, which are generalizations of piecewise polynomial functions (of arbitrary order) that map into S^{m-1} at any point in the domain. For problems with sufficiently smooth solutions, Hardering et al. showed optimal convergence rates for the H^1- and L^2-discretization errors for these elements, for any polynomial order [29, 32]. Remarkably, such finite elements can also represent singular maps to a certain extent, which suggests to employ them for simulating harmonic maps with singularities. To solve the discrete problems Sander et al. interpreted them as optimization problems on a product manifold, and solved them using a Riemannian trust-region method [1, 43]. The convergence of this algebraic solver follows from general results for optimization methods on manifolds [1].

The goal of this paper is to compare the practical properties of the two different discretizations and solver algorithms for the numerical approximation of harmonic maps into the unit sphere.

The overall outcome is not a priori clear: While the conforming methods preserve more of the mathematical structure of the problem, the nonconforming methods are simpler to implement. Also, the behavior in the presence of singularities is hard to predict; indeed, previous works already observe that computing the exact placement of singularities is difficult [34]. As both solver methods are based on energy minimization they are unlikely to find anything but locally minimizing harmonic maps.

To test the relative merits of the methods we define a set of benchmark problems. This set includes smooth harmonic maps, but also maps with different types of singularities. In our measurements we focus on the discretization error, the constraint violation of the nonconforming methods, and the solver speed. We find that the approximation power of the two discretizations is the same for smooth problems, but the conforming discretization shows slightly better convergence orders than the nonconforming one in the presence of singularities. For problems with smooth solutions, discretization error convergence orders are clearly observed, and they match the theoretical predictions. For problems with singular solutions, it is much more difficult to determine an order, and the precise behavior depends on the exact position of the singularity. Both solvers converge, but for three-dimensional problems with a singularity, we observe that the placement of the final singularity can depend on the initial iterate. This effect, which was already noted in [34] for two-dimensional problems, is more pronounced for the conforming discretization than for the nonconforming one, and remains to be addressed in future work.

As predicted by theory, the constraint violation produced by the gradient-flow solver remains bounded as long as the step size is chosen to be proportional to the element diameter. The values we measured are small enough to be unproblematic for most practical applications. By construction, the trust-region solver does not introduce any constraint violation at all.

Considerable differences show when comparing the solver speeds. The iteration numbers of the trust-region method are low, and in most situations they seem to remain bounded independently of the grid resolution. The gradient-flow solver, on the other hand, has to tie its step size to the grid element diameter. That way, the total number of iterations increases as the grid is refined. As the iterations of both methods have comparable cost, this leads to a large speed advantage for the trust-region solver. Experiments that combine the gradient-flow method with a Newton solver can be found in [10].

The paper is structured as follows: Chapter 2 briefly reviews the different notions of harmonic maps used in this text. Chapter 3 then introduces the two discretization methods, and Chapter 4 does the same for the algebraic solvers. The different benchmark problems are presented in Chapter 5. Finally, the last two chapters contain the actual numerical results, with Chapter 6 investigating the discretizations, and Chapter 7 the solvers.

2. Harmonic maps into the sphere

Let Ω be an open, bounded domain in $\mathbb{R}^n$. We use the standard notation $H^k(\Omega; \mathbb{R}^m)$ for vector-valued Sobolev spaces defined on Ω with $k \geq 1$, and we denote the k-th order Sobolev vector fields with vanishing boundary trace by $H_0^k(\Omega; \mathbb{R}^m)$. For sphere-valued problems, we introduce the subspace

$$H^k(\Omega; S^{m-1}) := \left\{ \mathbf{v} \in H^k(\Omega; \mathbb{R}^m) : \mathbf{v}(x) \in S^{m-1} \text{ a.e.} \right\},$$

and its subspace of functions that satisfy the boundary conditions (3) in a trace sense

$$H_D^k(\Omega; S^{m-1}) := \left\{ \mathbf{u} \in H^k(\Omega; S^{m-1}) \ : \ \mathbf{u}|_{\partial\Omega} = \mathbf{u}_{\mathrm{D}}|_{\partial\Omega} \right\}.$$

There are various related definitions of harmonic maps. This text uses three of them, which we review here briefly. More details can be found in [18, 24, 33]. In the following, $(\cdot, \cdot)$ denotes the standard L^2 scalar product.

Definition 1 (Stationary harmonic map). A function $\mathbf{u} \in H^1(\Omega; S^{m-1})$ is called *stationary harmonic map* if it satisfies the weak Euler–Lagrange equation

$$(4) \qquad\qquad (\nabla \mathbf{u}, \nabla \boldsymbol{\varphi}) = 0$$

for all $\boldsymbol{\varphi} \in H_0^1(\Omega; \mathbb{R}^m)$ with $\mathbf{u} \cdot \boldsymbol{\varphi} = 0$ a.e. in Ω.

The tangentiality condition $\mathbf{u} \cdot \boldsymbol{\varphi} = 0$ can be avoided by the equivalent formulation

$$(\nabla \mathbf{u}, \nabla \boldsymbol{\varphi}) = (|\nabla \mathbf{u}|^2 \mathbf{u}, \boldsymbol{\varphi})$$

where we now test with all $\boldsymbol{\varphi} \in H_0^1(\Omega; \mathbb{R}^m) \cap L^\infty(\Omega; \mathbb{R}^m)$, see [12].

Definition 2 ((Locally) minimizing harmonic map). A map $\mathbf{u} \in H^1(\Omega; S^{m-1})$ is called *(locally) minimizing harmonic map* if it is a (local) minimizer of the Dirichlet energy (1).

In general, a minimizing harmonic map of sufficient smoothness is also a stationary harmonic map, but the converse is not always the case.

The characterization (4) can be generalized via partial integration to cover cases of lesser regularity.

Definition 3 (Distributional harmonic map). A map $\mathbf{u} \in L^\infty(\Omega; \mathbb{R}^m)$ with $|\mathbf{u}| = 1$ a.e. in Ω is called *distributional harmonic map* if

$$(\mathbf{u}, \Delta \boldsymbol{\varphi}) = 0$$

for all $\boldsymbol{\varphi} \in C_0^\infty(\Omega; \mathbb{R}^m)$ with $\mathbf{u} \cdot \boldsymbol{\varphi} = 0$ almost everywhere in Ω.

It is a particularity of sphere-valued problems that the set of continuous maps $\Omega \to S^{m-1}$ is not connected. In other words, given two continuous functions $\mathbf{v}_1, \mathbf{v}_2 : \Omega \to S^{m-1}$ with identical Dirichlet boundary values it may not be possible to continuously deform one into the other. The underlying reason is the fact that the homotopy groups of the sphere are not all trivial. Indeed, maps from a simply connected n-dimensional domain $\Omega \subset \mathbb{R}^n$ into S^{m-1} are closely related to the n-th homotopy group of S^{m-1}. This group is trivial in the case $n = m = 2$, but isomorphic to $\mathbb{Z}$ in the cases $n = m = 3$ and $n = 2, m = 3$ considered here. By this isomorphism the connected components of the function spaces can be labeled by an integer (which is is sometimes called topological degree or quantum number [50]). The disconnected nature of the function space is a structure that may or may not be preserved by discretizations.

The issue of connectedness is more subtle for the Sobolev spaces $H^1(\Omega; S^{m-1})$ that underly finite element theory. In particular, the homotopy classes of $W^{1,p}$ may not all be weakly closed, and therefore some of them may not contain minimizing harmonic maps [33, Chap. 5.3]. A discussion about the connected components of $W^{1,p}(M, N)$ for general manifolds M, N can be found in [33, Chap. 3.3].

3. Discretizations

We now present the two discretization approaches of [9, 13] and [43, 44] for maps from the flat domain Ω into $S^{m-1} \subset \mathbb{R}^m$. The central challenge is the fact that it is impossible to satisfy the unit-length constraint (2) everywhere in Ω if piecewise polynomials are used for the approximation. In the following we assume Ω to be polyhedral, and we let $\mathcal{T}_h$ be a shape-regular triangulation of Ω with triangles or tetrahedra of diameter no larger than $h > 0$. The set of

polynomials with degree at most $p \geq 0$ on a simplex $T \in \mathcal{T}_h$ is denoted with $P_p(T)$, and the space of continuous, piecewise polynomial m-vector fields is

$$\mathcal{S}^p(\mathcal{T}_h; \mathbb{R}^m) := \left\{ \mathbf{v}_h \in C(\Omega; \mathbb{R}^m) \ : \ \mathbf{v}_h|_T \in P_p(T)^m \text{ for all } T \in \mathcal{T}_h \right\}.$$

We use the notation $\mathcal{I}_h^p : H^2(\Omega; \mathbb{R}^m) \to \mathcal{S}^p(\mathcal{T}_h; \mathbb{R}^m)$ for the p-th order Lagrange interpolation operator associated to a set of Lagrange points $\mathcal{L}_h$, of size $N := |\mathcal{L}_h|$.

3.1. **Geometrically nonconforming Lagrange finite elements.** A natural approach to cope with the nonlinear nature of the image space S^{m-1} is to approximate sphere-valued maps with functions from $\mathcal{S}^p(\mathcal{T}_h; \mathbb{R}^m)$ and require the unit-length constraint (2) only at the Lagrange points. For first-order finite elements this has been proposed and analyzed in [9]. We obtain the following admissible set, which also enforces the boundary conditions:

$$\mathcal{A}_h^{\mathrm{nc}} := \left\{ \mathbf{v}_h \in \mathcal{S}^p(\mathcal{T}_h; \mathbb{R}^m) : |\mathbf{v}_h(z)| = 1 \text{ for all } z \in \mathcal{L}_h \text{ and } \mathbf{v}_h(z) = \mathbf{u}_{\mathrm{D}}(z) \text{ for all } z \in \mathcal{L}_h \cap \partial\Omega \right\}.$$

This is a smooth nonlinear submanifold of the vector space $\mathcal{S}^p(\mathcal{T}_h; \mathbb{R}^m)$, but observe that $\mathcal{A}_h^{\mathrm{nc}}$ is a connected set. It therefore waives one of the central features of the actual solution space, which consists of disconnected homotopy classes.

Discrete approximations of minimizing harmonic maps can now be defined as minimizers of the Dirichlet energy (1) in the discrete admissible set $\mathcal{A}_h^{\mathrm{nc}}$. This works because while $\mathcal{A}_h^{\mathrm{nc}}$ is not a subset of the actual solution space $H^1(\Omega; S^{m-1})$, the Dirichlet functional (1) extends naturally to $H^1(\Omega; \mathbb{R}^m) \supset \mathcal{A}_h^{\mathrm{nc}}$.

The Definition 1 of weak stationary harmonic maps involves test functions. As the space of admissible functions is nonlinear, the corresponding test function spaces differ from point to point. Formally, the test functions for a map $\mathbf{v} \in H^1(\Omega; S^{m-1})$ are elements of the tangent space $T_\mathbf{v} H^1(\Omega; S^{m-1})$, and likewise for discrete maps $\mathbf{v}_h \in \mathcal{A}_h^{\mathrm{nc}}$. More practically, test functions are constructed as variations of admissible functions [45]. For a given $\mathbf{v}_h \in \mathcal{A}_h^{\mathrm{nc}}$ let therefore $\gamma : [-\epsilon, \epsilon] \to \mathcal{A}_h^{\mathrm{nc}}$ be a differentiable path with $\gamma(0) = \mathbf{v}_h$. The set of test functions for $\mathcal{A}_h^{\mathrm{nc}}$ at $\mathbf{v}_h$ is the set of all functions $\boldsymbol{\varphi}_h : \Omega \to \mathbb{R}^m$ that can be represented as

$$\boldsymbol{\varphi}_h = \left. \frac{d\gamma(t)}{dt} \right|_{t=0}$$

with such a path γ. For the geometrically nonconforming discretization, this construction results in a space of piecewise polynomial vector fields that are orthogonal to $\mathbf{v}_h$ at the Lagrange points, and zero at the boundary

$$T_h^{\mathrm{nc}}(\mathbf{v}_h) := \left\{ \boldsymbol{\varphi}_h \in \mathcal{S}^p(\mathcal{T}_h; \mathbb{R}^m) : \mathbf{v}_h(z) \cdot \boldsymbol{\varphi}_h(z) = 0 \text{ for all } z \in \mathcal{L}_h, \ \boldsymbol{\varphi}_h(z) = 0 \text{ for all } z \in \mathcal{L}_h \cap \partial\Omega \right\}.$$

Finding discrete stationary harmonic maps then means finding functions $\mathbf{u}_h \in \mathcal{A}_h^{\mathrm{nc}}$ such that

$$(\nabla \mathbf{u}_h, \nabla \boldsymbol{\varphi}_h) = 0$$

for all $\boldsymbol{\varphi}_h \in T_h^{\mathrm{nc}}(\mathbf{u}_h)$. This corresponds to the continuous problem (4).

The nonconforming discretization of harmonic maps allows for different convergence theories. All results in the literature concern the first-order case $p = 1$ only, but the convergence theories from [13] apply verbatim to higher-order discretizations. Most generally, one can establish the Γ-convergence of the discretizations of the Dirichlet energy E from (1) defined by

$$E_h[\mathbf{u}_h] := \begin{cases} E[\mathbf{u}_h] & \text{if } \mathbf{u}_h \in \mathcal{A}_h^{\mathrm{nc}}, \\ \infty & \text{else,} \end{cases}$$

to the continuous minimization problem. A discrete compactness result then implies the accumulation of approximations at solutions of the continuous problem. A proof for the case $m = 3$ is given in [12].

Theorem 4 (Discrete compactness, [12, Theorem 7.6])**.** *Let $p = 1$ and let $(\mathbf{u}_h)_{h>0} \subset H^1(\Omega; \mathbb{R}^3)$ be a bounded sequence such that each $\mathbf{u}_h \in \mathcal{A}_h^{\mathrm{nc}}$ and*

$$(\nabla \mathbf{u}_h, \nabla \boldsymbol{\varphi}_h) = \mathcal{R}_h(\boldsymbol{\varphi}_h)$$

for all $\boldsymbol{\varphi}_h \in T_h^{\mathrm{nc}}(\mathbf{u}_h)$ for functionals $\mathcal{R}_h \in H_0^1(\Omega; \mathbb{R}^3)'$ with $\mathcal{R}_h \to 0$ in $H_0^1(\Omega; \mathbb{R}^3)'$ as $h \to 0$. Then every weak accumulation point of $(\mathbf{u}_h)_{h>0}$ is a stationary harmonic map.

Sketch of the proof. The satisfaction of the unit-length constraint by the accumulation points follows from standard interpolation and inverse estimates. The fact that weak limits are harmonic maps follows in the case $m = 3$ from the equivalent characterization

$$(\nabla \mathbf{u}, \nabla[\mathbf{u} \times \boldsymbol{\phi}]) = 0$$

for all $\boldsymbol{\phi} \in H_0^1(\Omega; \mathbb{R}^3) \cap L^\infty(\Omega; \mathbb{R}^3)$, by choosing the interpolant of $\mathbf{u}_h \times \boldsymbol{\phi}$ for given $\boldsymbol{\phi} \in C^\infty(\overline{\Omega}; \mathbb{R}^3)$ as a test function and then passing to the limit. $\qquad\square$

This result can be extended to manifolds that are more general than the unit-sphere [11].

Remark 5. In the proof of Theorem 4 one establishes the fact that for a sequence of discrete vector fields whose nodal values belong to S^{m-1}, any weak accumulation point in $H^1(\Omega)$ has values in S^{m-1} almost everywhere. The assumption of the nodal values belonging to S^{m-1} can be further weakened: Given a sequence of discrete vector fields whose nodal values are not necessarily on S^{m-1} but approach unit length as $h \to 0$, it still follows that any accumulation point in $H^1(\Omega; \mathbb{R}^m)$ satisfies the constraint exactly; cf. [13]. This allows to prove convergence of solutions obtained with the nonconforming discrete gradient flow discussed in Section 4.1, which does not preserve the constraint exactly.

Recently, discretization error bounds have also been proved. For piecewise linear Lagrange elements and $n = m = 2$, optimal error estimates in the energy norm have been derived for a corresponding saddle-point formulation in [34]. This result is extended in [15] to the cases $m = 3$ and $n > 2$.

3.2. **Geometrically conforming projection-based finite elements.** The second discretization constructs finite elements that map into the sphere S^{m-1} everywhere, not just at the Lagrange points. It does so by adapting the notion of a polynomial. The construction has been proposed in [29] and [27] under the name of *projection-based finite elements*, and is a member of the larger family of *geometric finite elements* [32].

In essence, projection-based finite elements are defined by projecting Lagrange finite elements with nodal values in S^{m-1} pointwise onto S^{m-1}. More formally, consider the closest-point projection from $\mathbb{R}^m$ onto S^{m-1} defined by

$$P : \mathbb{R}^m \setminus \{0\} \to S^{m-1}, \qquad P(\xi) := \frac{\xi}{|\xi|}.$$

This projection induces a superposition operator [7]

$$\mathcal{P} : C(\Omega; \mathbb{R}^m) \to W^{k,q}(\Omega; S^{m-1}) \qquad \mathcal{P}\mathbf{v}(x) := P(\mathbf{v}(x)), \qquad \forall x \in \Omega,$$

for suitable k and q. With its help we define the space of p-th order projection-based finite elements as

$$\mathcal{S}^{p,\mathrm{proj}}(\mathcal{T}_h; S^{m-1}) := \left\{ \mathcal{P}\mathbf{v}_h \ : \ \mathbf{v}_h \in \mathcal{S}^p(\mathcal{T}_h; \mathbb{R}^m) \text{ and } \mathbf{v}_h(z) \in S^{m-1} \text{ for all } z \in \mathcal{L}_h \right\}.$$
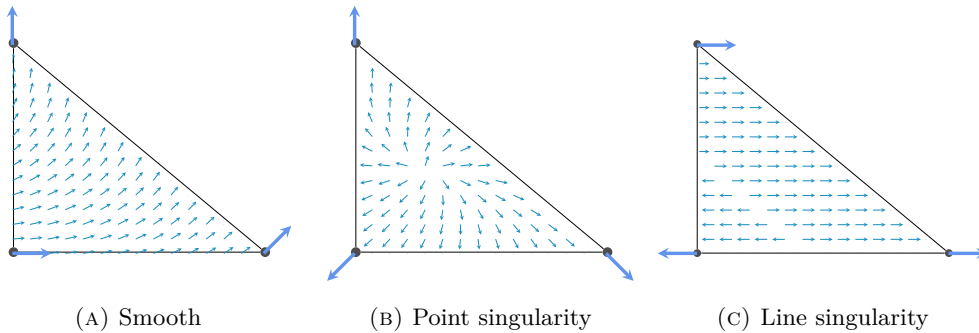
(A) Smooth   (B) Point singularity   (C) Line singularity

FIGURE 1. Three first-order projection-based finite element functions on a triangle, with values in S^1. The prescribed values at the Lagrange points are enlarged for better visibility.

A canonical interpolation operator into $\mathcal{S}^{p,\mathrm{proj}}(\Omega; S^{m-1})$ is given by

$$\mathcal{I}_h^{p,\mathrm{proj}} : C(\Omega; \mathbb{R}^m) \to \mathcal{S}^{p,\mathrm{proj}}(\mathcal{T}_h; S^{m-1})$$

$$\mathbf{v} \mapsto \mathcal{P}(\mathcal{I}_h^p \mathbf{v}).$$

Similar to the nonconforming discretization we then introduce the subspace of functions that comply with the boundary condition at the boundary Lagrange points

$$\mathcal{A}_h^{\mathrm{proj}} := \left\{ \mathbf{v}_h \in \mathcal{S}^{p,\mathrm{proj}}(\Omega; S^{m-1}) \; : \; \mathbf{v}_h(z) = \mathbf{u}_{\mathrm{D}}(z) \text{ for all } z \in \mathcal{L}_h \cap \partial\Omega \right\}.$$

Note that the correspondence between functions in $\mathcal{S}^{p,\mathrm{proj}}$ and their sets of values $c \in (S^{m-1})^N$ (with $N := |\mathcal{L}_h|$) at the Lagrange points is one-to-one. This allows to treat projection-based finite elements algorithmically as such sets of values, just as in the case of standard finite elements.

It was shown in [29] that elements of $\mathcal{S}^{p,\mathrm{proj}}$ are in $H^1(\Omega; S^{m-1})$ as long as configurations are avoided where the projection $P(\xi) := \xi/|\xi|$ becomes undefined. By this embedding property, the subset of continuous projection-based finite elements (with fixed boundary conditions) decomposes into disconnected homotopy classes just like $C(\Omega; S^{m-1})$ does.

However, as a particularity, the space $\mathcal{S}^{p,\mathrm{proj}}(\mathcal{T}_h; S^{m-1})$ also contains maps with singularities. These appear whenever the values at the Lagrange nodes on an element T are such that T contains a zero of the Lagrange interpolation in $\mathbb{R}^m$, because there, $P : \mathbb{R}^m \to S^{m-1}$ is not defined. Figure 1 shows two such configurations for the case $m = n = 2$. It is easily seen that finite element maps with singularities can be classified according to the dimension of the support of the singularity. For example, on a triangle, the Lagrange interpolation can be zero in a single point or on a line intersecting the triangle. Consequently, the projection-based finite element functions can have point or line singularities. However, as the values at the Lagrange points are constrained to have unit length, the singularities cannot appear everywhere within T. In particular, singularities at Lagrange points are not possible.

The precise regularity of these singular finite element functions remains an open question. The line-singular function on the right of Figure 1 is discontinuous on a line, and therefore clearly not an element of $W^{1,q}$ (because it is not absolutely continuous on almost every line parallel to the coordinate directions [25]). However, the precise Sobolev regularity of point-singular finite element functions has not been worked out yet.

Depending on the viewpoint, the singular projection-based finite element functions are a nuisance or a feature. If the aim is to approximate smooth functions, then the appearance of singular finite element functions can be problematic, because one is leaving the realm of established error

theory, and implementations have to guard against division by zero. Luckily, this happens only rarely in practice, usually when the grid is too coarse.

When trying to approximate singular functions, on the other hand, having singular approximation functions at hand may be considered an asset. Indeed, singular functions are difficult to approximate in standard finite element spaces, and introducing additional bespoke singular finite element functions (as, e.g., in XFEM [26]) is a standard way to increase the approximation power. As we will see below, for harmonic maps we do not observe a clear advantage, and the topic needs further investigation. An additional technical problem is that special quadrature rules are required to integrate the singular functions. We have experimented with Gauß–Legendre und low-order composite rules, and did not find one to be clearly superior to the other. Ideally, some sort of adaptive rule would be used.

As for the nonconforming discretization, test functions are constructed as variations of finite element functions. As functions in $\mathcal{S}^{p,\mathrm{proj}}$ map into S^{m-1} everywhere, such variations are vector fields $\boldsymbol{\varphi}_h : \Omega \to \mathbb{R}^m$ for which $\mathbf{u}_h(x) \cdot \boldsymbol{\varphi}_h(x) = 0$ for all $x \in \Omega$.[1] Using a construction similar to the one given in [45], one can show that the test function space at a function $\mathbf{v}_h \in \mathcal{S}^{p,\mathrm{proj}}$ is isomorphic to $\prod_{z \in \mathcal{L}_h} T_{\mathbf{v}_h(z)} S^{m-1}$. To construct this isomorphism, let I^{proj} be the operator that maps coefficient sets in $(S^{m-1})^N$ to functions in $\mathcal{S}^{p,\mathrm{proj}}$ and let $\mathbf{v}_h \in \mathcal{S}^{p,\mathrm{proj}}$ with coefficients $v_1, \ldots, v_N$. Then, for all $b_i \in T_{v_i} S^{m-1}$, $i = 1, \ldots, N$ and $x \in \Omega$, the corresponding test function $\boldsymbol{\varphi}_h$ can be evaluated by

$$\boldsymbol{\varphi}_h[b_1, \ldots, b_N](x) = \sum_{i=1}^{N} \frac{\partial I^{\mathrm{proj}}(v_1, \ldots, v_N; x)}{\partial v_i} \cdot b_i.$$

In line with Section 3.1 we define $T_h^{\mathrm{proj}}(\mathbf{v}_h)$ as the set of test functions at $\mathbf{v}_h \in \mathcal{A}_h^{\mathrm{proj}}$ that vanish on the boundary nodes.

As the admissible set $\mathcal{A}_h^{\mathrm{proj}}$ is not a vector space, the standard approximation error theory of finite elements does not apply. However, for problems with smooth solutions there are rigorous optimal a priori discretization error estimates both for the L^2 norm and the H^1 norm [29,32]. The proofs generalize results like the Bramble–Hilbert lemma in order to prove optimal interpolation error estimates. Then, the H^1 discretization error is estimated with a nonlinear version of the Céa Lemma [28]. To get bounds on the L^2 error, a generalized Aubin–Nitsche lemma for predominantly quadratic energies such as the Dirichlet energy is proved [30,31,32]. Here as well, the theoretical results are not restricted to the unit sphere, but hold for more general Riemannian manifolds [32].

4. Solvers for the discrete problems

We now discuss two solver algorithms for the algebraic problems that arise from the discretizations in Section 3. Both solvers presented here are based on energy minimization, and they are therefore unlikely to find stationary harmonic maps that are not locally minimizing. One of them is a conforming algorithm, i.e., each iterate is guaranteed to be an element of the discrete admissible set. For the other one, this only holds in the limit of vanishing step size. We use the symbol $\mathcal{A}_h$ to denote either the admissible set $\mathcal{A}_h^{\mathrm{nc}}$ or $\mathcal{A}_h^{\mathrm{proj}}$, and $T_h(\mathbf{v}_h)$ to denote the corresponding test function space at $\mathbf{v}_h \in \mathcal{A}_h$, when the difference is irrelevant.

4.1. **Nonconforming discrete gradient flow.** The first solver interprets the minimization problem for the energy (1) over $\mathcal{A}_h^{\mathrm{nc}}$ as being posed on $(\mathbb{R}^m)^N$, $N := |\mathcal{L}_h|$, subject to the constraint $|\mathbf{u}(z)| = 1$ for each Lagrange point $z \in \mathcal{L}_h$. Its approach to deal with the constraint

---

[1]Here we tacitly interpret tangent vectors of S^{m-1} as elements of the surrounding space $\mathbb{R}^m$.
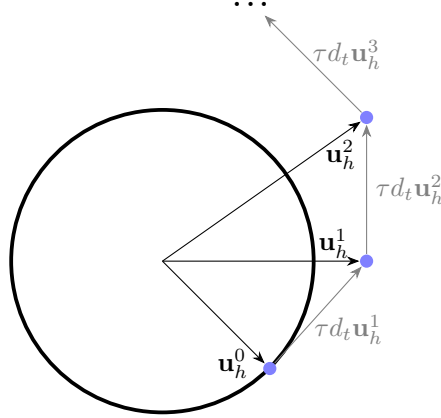
FIGURE 2. Steps of the discrete gradient flow

is to linearize it and then, starting from some initial configuration $\mathbf{u}_h^0 \in \mathcal{A}_h$, follow a discrete gradient flow for the energy (1), imposing the linearized constraint in the Lagrange points of the finite element space. In order to define the discrete gradient flow, we use the L^2 scalar product of the gradients $(\nabla \cdot, \nabla \cdot)$. Writing $P_{\mathbf{v}_h}$ for the H^1-projection into the linear space $T_h(\mathbf{v}_h)$ at a $\mathbf{v}_h \in \mathcal{S}^p(\mathcal{T}_h; \mathbb{R}^m)$, and introducing a pseudo-time variable t, the gradient flow is

$$\frac{\partial \mathbf{u}_h}{\partial t} = -P_{\mathbf{u}_h}(\nabla_{H^1} E[\mathbf{u}_h]),$$

with weak form

$$\left(\nabla \frac{\partial \mathbf{u}_h}{\partial t}, \nabla \boldsymbol{\varphi}_h\right) = -(\nabla \mathbf{u}_h, \nabla \boldsymbol{\varphi}_h).$$

The proposed algorithm is then an implicit Euler method for this flow.

Algorithm 1: Discrete gradient flow

Input: initial iterate $\mathbf{u}_h^0 \in \mathcal{A}_h^{\mathrm{nc}}$, step size $\tau > 0$, stopping threshold $\epsilon_{\mathrm{stop}} > 0$

1 for $k = 0, 1, 2, \ldots$ **do**

2   determine $d_t\mathbf{u}_h^{k+1} \in T_h(\mathbf{u}_h^k)$ such that

$$(\nabla d_t\mathbf{u}_h^{k+1}, \nabla \boldsymbol{\varphi}_h) = -(\nabla \mathbf{u}_h^k, \nabla \boldsymbol{\varphi}_h) - \tau(\nabla d_t\mathbf{u}_h^{k+1}, \nabla \boldsymbol{\varphi}_h) \text{ for all } \boldsymbol{\varphi}_h \in T_h(\mathbf{u}_h^k)$$

3   set $\mathbf{u}_h^{k+1} = \mathbf{u}_h^k + \tau d_t\mathbf{u}_h^{k+1}$

4   **if** $|d_t\mathbf{u}_h^{k+1}|_{H^1} \leq \epsilon_{\mathrm{stop}}$ **then**

5   **return** $\mathbf{u}_h^{k+1}$

6   **end**

7 end

Determining the tangential correction in Line 2 is a linear elliptic equation on the vector space $T_h(\mathbf{u}_h^k)$. There are several ways how to enforce the restriction to the tangent space. Following own previous work, our implementation uses Lagrange multipliers. At each step k, we solve the linear system

$$\begin{pmatrix} S & (A^k)^\top \\ A^k & 0 \end{pmatrix} \begin{pmatrix} d^{k+1} \\ \lambda^{k+1} \end{pmatrix} = \begin{pmatrix} b \\ 0 \end{pmatrix},$$

where $d^{k+1} \in (\mathbb{R}^m)^N$ is a representation of the discrete function $d_t\mathbf{u}_h^{k+1}$ in the nodal basis of $\mathcal{S}^p(\mathcal{T}_h; \mathbb{R}^m)$, the matrix S is the algebraic representation of the scalar product $(\nabla\cdot, \nabla\cdot)_{L^2}$ in the finite element space, $b \in \mathbb{R}^{mN}$ encodes the explicit terms on the right hand side of Line 2 of Algorithm 1, and A^k is a matrix encoding the linearized nodal constraints $d_t\mathbf{u}_h^{k+1}(z) \cdot \mathbf{u}_h^k(z) = 0$, $z \in \mathcal{L}_h$, with the corresponding Lagrange multiplier $\lambda^{k+1} \in \mathbb{R}^N$. The resulting linear systems are symmetric but indefinite, and are solved with a direct solver if $\dim\Omega = 2$, and with a GMRES solver if $\dim\Omega = 3$.

Rigorous convergence analysis for this algorithm for the case of the first-order nonconforming discretization of Chapter 3.1 appears in [9, 13]. To obtain convergence, the time step size τ has to tend to zero with the mesh size h. Monotonicity of the method implies that the norms of the corrections $d_t u_h^k$ decay like $\mathcal{O}\big((\tau k)^{-\frac{1}{2}}\big)$. In this case the successive violations of the constraint, shown in Figure 2, are controlled by the energy of the initial iterate and the step size,

$$(5) \qquad \max_{k=0,1,2,\dots} \int_\Omega \mathcal{I}_h^1\big(\big||\mathbf{u}_h^k|^2 - 1\big|\big)\,\mathrm{d}x \le c\tau E[\mathbf{u}_h^0].$$

In particular, the maximum violation is independent of the number of iterations, and the bound does not depend on structural properties of the underlying triangulation. This observation holds for all target manifolds that are given as level sets of suitably regular functions.

In order to maintain the optimal convergence of the discretization, a violation of order $\mathcal{O}(h)$ in the constraint has to be guaranteed, which requires choosing $\tau = \mathcal{O}(h)$. This however reduces the speed of the convergence of the iteration which is $\mathcal{O}((\tau N)^{-\frac{1}{2}})$.

On the other hand, note that the linear convergence of the constraint violation potentially spoils the approximation error when a higher-order finite element space is employed. As a remedy a higher-order approximation in time may be used to define the discrete gradient flow. Recently, it has been shown that using a second-order backward differentiation formula leads to quadratic constraint consistency [3].

An earlier version of Algorithm [5] projected each tangential correction back onto the sphere at each Lagrange point. That way, no algebraic constraint violation was accumulated, and the algorithm remained conforming. However, such projections may lead to energy increase. Energy decrease can be guaranteed for certain types of triangulations, but the corresponding restrictions limit the applicability of the algorithm in three-dimensional situations [9]. We do not consider this variant of the algorithm any further in this text.

4.2. **Riemannian trust-region method.** For a method that does not violate the algebraic sphere constraints, we turn to the field of optimization on manifolds [1]. Here, the admissible set $\mathcal{A}_h \in \{\mathcal{A}_h^{\mathrm{nc}}, \mathcal{A}_h^{\mathrm{proj}}\}$ is viewed as the product manifold $(S^{m-1})^N$. The discrete problem for finding minimizing harmonic maps then has the form of a minimization problem for the Dirichlet energy (1) on that manifold.

In order to solve the minimization problem on $\mathcal{A}_h$ we employ a Riemannian trust-region method. Such a method generalizes standard trust-region methods to objective functionals defined on a Riemannian manifold [1]. In the spirit of Newton's method, the general idea is to consider a quadratic model of the objective functional around the current iterate. The algorithm then computes a correction φ_h^k in the tangent space of the current iterate $\mathbf{u}_h^k \in \mathcal{A}_h$. However, as the model is assumed to be accurate only in a neighborhood of 0 on $T_h(\mathbf{u}_h^k)$, the tangential correction is restricted to remain in a ball around 0 (the name-giving trust region), the radius of which is controlled adaptively. To be more precise, in each step k the quadratic trust-region

---

Algorithm 2: Riemannian trust-region method

Input: Initial iterate $\mathbf{u}_h^0 \in \mathcal{A}_h$, initial trust-region radius $\Delta_0 > 0$, tolerances $\beta_1 > \beta_2 > 0$ and stopping threshold $\epsilon_{\text{stop}} > 0$

1 for $k = 0, 1, 2, \ldots$ **do**

2 $\quad$ Solve (6) for $\boldsymbol{\varphi}_h \in T_h(\mathbf{u}_h^k)$ $\qquad\qquad\qquad$ // Compute tangential correction

3 $\quad$ **if** $|\boldsymbol{\varphi}_h|_{H^1} < \epsilon_{\text{stop}}$ **then**

4 $\quad\quad$ **return** $\mathbf{u}_h^k$

5 $\quad$ **else**

6 $\quad\quad$ Evaluate ρ_k from (7) $\qquad\qquad\qquad\qquad$ // Estimate model quality

7 $\quad\quad$ **if** $\rho_k > \beta_1$ **then**

8 $\quad\quad\quad$ $\mathbf{u}_h^{k+1} = \text{Exp}_{\mathbf{u}_h^k}(\boldsymbol{\varphi}_h)$ and $\Delta_{k+1} = 2\Delta_k$ $\qquad$ // 'Very successful' step

9 $\quad\quad$ **else if** $\rho_k > \beta_2$ **then**

10 $\quad\quad\quad$ $\mathbf{u}_h^{k+1} = \text{Exp}_{\mathbf{u}_h^k}(\boldsymbol{\varphi}_h)$ and $\Delta_{k+1} = \Delta_k$ $\qquad$ // 'Successful' step

11 $\quad\quad$ **else**

12 $\quad\quad\quad$ $\mathbf{u}_h^{k+1} = \mathbf{u}_h^k$ and $\Delta_{k+1} = \frac{1}{2}\Delta_k$ $\qquad\qquad$ // 'Unsuccessful' step

13 $\quad\quad$ **end**

14 $\quad$ **end**

15 end

---

subproblem at an iterate $\mathbf{u}_h^k$ is

$$(6) \qquad \min_{\boldsymbol{\varphi}_h \in T_h(\mathbf{u}_h^k)} m_k(\boldsymbol{\varphi}_h) := E[\mathbf{u}_h^k] + \left\langle \text{Grad}\, E[\mathbf{u}_h^k], \boldsymbol{\varphi}_h \right\rangle + \frac{1}{2} \left\langle \text{Hess}\, E[\mathbf{u}_h^k](\boldsymbol{\varphi}_h), \boldsymbol{\varphi}_h \right\rangle$$

subject to

$$\|\boldsymbol{\varphi}_h\|_{\mathbf{u}_h^k} \le \Delta_k,$$

where $\|\cdot\|_{\mathbf{u}_h^k}$ is a suitable norm, and Δ_k is the current trust-region radius. The terms $\text{Grad}\, E$ and $\text{Hess}\, E$ denote the Riemannian gradient and Hessian, respectively. These can be computed from simple modifications of the corresponding Euclidean quantities of the functional extended into the surrounding Euclidean space [2], which take a particularly simple form for the unit sphere as a target manifold. In a finite element context, one can conveniently solve the quadratic minimization problems (6) by choosing the infinity norm $\|\cdot\|_{\mathbf{v}_h^k} = \|\cdot\|_\infty$ with respect to some coordinate system to define the trust-region, and then use a monotone multigrid (MMG) method for the minimization. Detailed descriptions are given in [36, 43, 47].

Given an appropriate tangential update $\boldsymbol{\varphi}_h^k$, the next iterate is then $\mathbf{u}_h^{k+1} = \exp_{\mathbf{u}_h^k}(\boldsymbol{\varphi}_h^k) \in \mathcal{A}_h$, where $\exp_{\mathbf{u}_h^k}$ is the Riemannian exponential map of $\mathcal{A}_h$ at $\mathbf{u}_h^k$. As $\mathcal{A}_h$ is a product manifold, the exponential map acts on each factor space S^{m-1} separately. Conveniently, the exponential map of the sphere has the closed-form expression

$$\exp_\xi v = \cos|v| \cdot \xi + \frac{\sin|v|}{|v|} \cdot v \qquad \text{where } \xi \in S^{m-1} \text{ and } v \in T_\xi S^{m-1}.$$

Like the gradient-flow method of the previous section, the trust-region method computes an update in a tangent space. However, while the gradient-flow algorithm uses Lagrange multipliers to enforce tangentiality, the trust-region method introduces a basis for $T_h(\mathbf{u}_h^k)$ and solves the quadratic minimization problem (6) with respect to this basis. As a consequence, the update problem involves fewer variables, and it is elliptic instead of a saddle-point problem. For coefficients in S^2, the basis is constructed as the push-forward of the canonical basis of $\mathbb{R}^2$ under the

inverse stereographic projection, which is conformal. While the tangent problem of the gradient-flow method could be easily formulated in terms of a basis for $T_h(\mathbf{u}_h^k)$, the trust-region problem involves an additional inequality constraint. Solving minimization problems with inequality and equality constraints is possible, but much more complicated.

Whether the new iterate $\mathbf{u}_h^{k+1} = \mathrm{Exp}_{\mathbf{u}_h^k}(\boldsymbol{\varphi}_h)$ is actually accepted depends on whether it realizes sufficient energy decrease. This is measured by the ratio of actual and expected energy decrease

$$(7) \qquad \rho_k := \frac{E\big[\mathrm{Exp}_{\mathbf{u}_h^k}(0)\big] - E\big[\mathrm{Exp}_{\mathbf{u}_h^k}(\boldsymbol{\varphi}_h)\big]}{m_k(0) - m_k(\boldsymbol{\varphi}_h)}.$$

This quantity also controls the evolution of the trust-region radius [20].

The global convergence of this method to first-order stationary points has been proved by Absil et al. [1]. Moreover, depending on the accuracy of the inner solver, the local convergence is superlinear or even quadratic. We refer to [1] for more details. More applications of the method to manifold-valued finite element problems appear in [41, 47].

5. Benchmarks problems

In this section we present the benchmark problems that will serve to test the discretizations and solver algorithms. All problems are posed on the open n-dimensional square $\Omega = (-\frac{1}{2}, \frac{1}{2})^n$.

5.1. **Inverse stereographic projection.** The first test problem constructs a harmonic map that is C^∞ in the domain.

Problem 1 (Inverse stereographic projection). Let $\Omega = (-\frac{1}{2}, \frac{1}{2})^2$ and set $m = 3$, i.e., the image space is $S^2 \subset \mathbb{R}^3$. Denote the stereographic projection (from the north pole) by

$$\pi_{\mathrm{st}} : S^2 \setminus \big\{(0,0,1)^T\big\} \to \mathbb{R}^2 \qquad \pi_{\mathrm{st}}(\xi_1, \xi_2, \xi_3) := \left(\frac{\xi_1}{1-\xi_3}, \frac{\xi_2}{1-\xi_3}\right)^T,$$

and note that it is invertible with

$$\pi_{\mathrm{st}}^{-1}(x_1, x_2) = \left(x_1^2 + x_2^2 + 1\right)^{-1} \begin{pmatrix} 2x_1 \\ 2x_2 \\ x_1^2 + x_2^2 - 1 \end{pmatrix}.$$

Find a harmonic map $\mathbf{u} : \Omega \to S^2$ such that $\mathbf{u} = \pi_{\mathrm{st}}^{-1}$ on the domain boundary $\partial\Omega$.

Direct calculations show that π_{st}^{-1} itself is a stationary map, and it is in C^∞. Its minimization properties are discussed in [17].

5.2. **Radial projection.** The next two problems compute harmonic maps with a single singularity. Both use the radial projection map

$$\mathbf{u}_\odot : \Omega \to S^{\dim \Omega}, \qquad \mathbf{u}_\odot(x) := \frac{x}{|x|}$$

as the boundary condition function.

Problem 2 (Radial projection).
  (a) Find a harmonic map on $\Omega := \left(-\frac{1}{2}, \frac{1}{2}\right)^3$ with image in S^2 such that $\mathbf{u} = \mathbf{u}_\odot$ on $\partial\Omega$.
  (b) Find a harmonic map on $\Omega := \left(-\frac{1}{2}, \frac{1}{2}\right)^2$ with image in S^1 such that $\mathbf{u} = \mathbf{u}_\odot$ on $\partial\Omega$.

(a) $\mathbb{R}^2 \supset \Omega \to S^2$ \qquad (b) $\mathbb{R}^3 \supset \Omega \to S^2$ \qquad (c) $\mathbb{R}^2 \supset \Omega \to S^1$
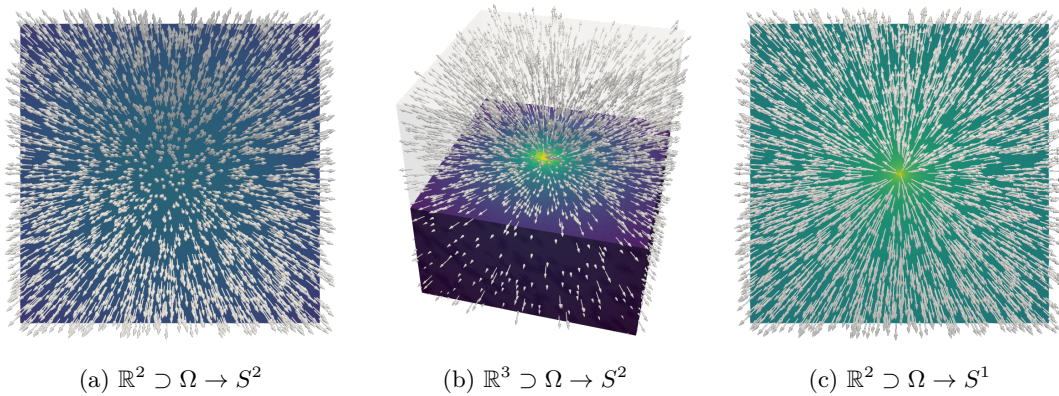
FIGURE 3. Discrete solutions $\mathbf{u}_h$ for Problems 1, 2(a) and 2(b)

Both problems have $\mathbf{u}_\odot$ as a solution. A direct computations in polar coordinates shows that $\mathbf{u}_\odot$ is an element of $W^{1,q}(\Omega; S^{m-1})$ for $q < n$ (see also [33, Chap. 3.2]) and that it also belongs to $W^{2,q}(\Omega; S^{m-1})$ for $q < \frac{n}{2}$. Further properties depend on the space dimensions:

(a) If $n = m = 3$, the radial projection $\mathbf{u}_\odot$ is the unique minimizer of the Dirichlet energy in $H^1(\Omega, S^2)$. This was proven by Brezis, Coron and Lieb [18, Corollary 7.9], based on the work of Schoen and Uhlenbeck [48]. The corresponding minimal energy on Ω is

$$E[\mathbf{u}_\odot] = 6 \int\limits_{-\pi/4}^{\pi/2} \frac{\pi - 2\arctan\left(\frac{1}{\sin(\vartheta)}\right)}{\sin(\vartheta)} \, \mathrm{d}\vartheta \approx 7.674124,$$

which can be verified with a computer algebra system. As $2 \not< \frac{n}{2}$ if $n = 3$, we have $\mathbf{u}_\odot \notin H^2(\Omega; S^2)$ in this case.

(b) If $n = m = 2$, the radial projection $\mathbf{u}_\odot$ is not even in $H^1(\Omega; S^1)$ (see also [49, p. 71]), and therefore $E[\mathbf{u}_\odot]$ is not well-defined. However, it is still an element of $L^2(\Omega; S^1)$.

The radial projection $\left(-\frac{1}{2}, \frac{1}{2}\right) \to S^1$ is still harmonic in the distributional sense, even though its Dirichlet energy is undefined. We give a short proof of this result, which we have not found elsewhere.

Lemma 6. *If $n = m = 2$ the radial projection $\mathbf{u}_\odot : x \mapsto \frac{x}{|x|}$ is harmonic in the distributional sense, i.e., $(\mathbf{u}_\odot, \Delta \varphi) = 0$ for any $\varphi \in C_0^\infty(\Omega; \mathbb{R}^m)$ with $\mathbf{u}_\odot \cdot \varphi = 0$ a.e.*

Proof. One directly checks that $\partial_i \mathbf{u}_\odot$, $i = 1, 2$, and $-\Delta \mathbf{u}_\odot$ are, respectively, perpendicular and parallel to $\mathbf{u}_\odot$, which itself coincides with the unit normal $\mathbf{n}$ to the ball $B_\varepsilon(0)$ of radius ε around the origin. With this, a splitting of the integral and two integrations by parts lead to

$$\int_\Omega \mathbf{u}_\odot \cdot \Delta \varphi \, \mathrm{d}x = \int_{B_\varepsilon(0)} \mathbf{u}_\odot \cdot \Delta \varphi \, \mathrm{d}x + \int_{\partial(\Omega \setminus B_\varepsilon(0))} \mathbf{u}_\odot \cdot (\nabla \varphi \mathbf{n}) \, \mathrm{d}s$$

$$- \int_{\partial(\Omega \setminus B_\varepsilon(0))} (\nabla \mathbf{u}_\odot \mathbf{n}) \cdot \varphi \, \mathrm{d}s + \int_{\Omega \setminus B_\varepsilon(0)} \Delta \mathbf{u}_\odot \cdot \varphi \, \mathrm{d}x,$$

for any $\varphi \in C_0^\infty(\Omega; \mathbb{R}^m)$. The first integral on the right-hand side is bounded in terms of ε^2, the second integral is bounded in terms of ε, the third integral vanishes since $\nabla \mathbf{u}_\odot \mathbf{n} = 0$, and the fourth integral vanishes because $\mathbf{u}_\odot$ is harmonic outside of $B_\varepsilon(0)$. $\qquad \square$
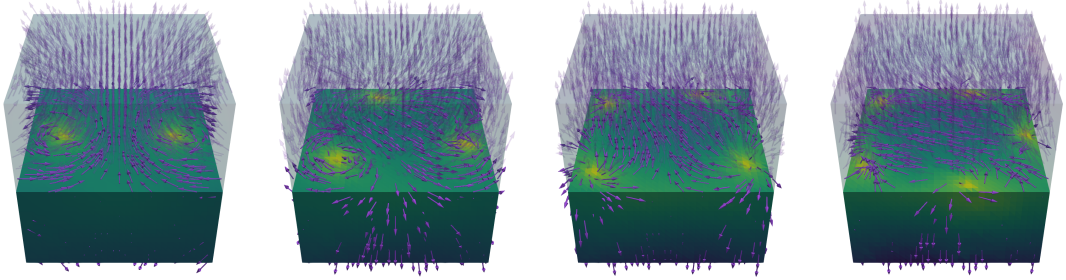
FIGURE 4. Stationary configurations of harmonic maps from $\left(-\frac{1}{2}, \frac{1}{2}\right)^3$ to S^2 with multiple singularities, obtained with the nonconforming discretization and Algorithm 1 with initial singularity degrees $\kappa = 2, 3, 4, 5$ (left to right). Singularities are located on the horizontal mid-surface, which is colored by the Frobenius norm of the discrete solution gradients.

5.3. **Multiple singularities.** The final benchmark investigates harmonic maps with multiple singularities. It also appears in [5].

Recall that if a map $\mathbf{v} \colon \Omega \to S^2$ is continuous in a ball $B_\varepsilon(x)$ around a point $x \in \Omega$ except at x itself, the degree of the singularity at x is defined as the topological winding number of $\mathbf{v}|_{\partial B_\varepsilon(x)}$ with respect to x, see [18].

Problem 3 (Multiple singularities in $\mathbb{R}^3$). Interpret the stereographic projection π_{st} as a map into the complex plane, and define

$$(8) \qquad \mathbf{u}_D^\kappa(x) := \pi_{\mathrm{st}}^{-1} \circ g_\kappa \circ \pi_{\mathrm{st}} \circ \mathbf{u}_\odot,$$

where

$$g_\kappa : \mathbb{C} \to \mathbb{C}, \qquad g_\kappa(z) := z^\kappa, \qquad \forall \kappa \in \mathbb{N}.$$

For given $\kappa \in \mathbb{N}$, find a harmonic map $\mathbf{u} : \Omega = \left(-\frac{1}{2}, \frac{1}{2}\right)^3 \to S^2$ such that $\mathbf{u} = \mathbf{u}_D^\kappa$ on $\partial\Omega$.

The map $\mathbf{u}_D^\kappa$ has a degree-κ singularity at the origin. For $\kappa = 1$ this is identical to Problem 2(a). However, according to Brezis, Coron, and Lieb [18] (locally) minimizing harmonic maps cannot have singularities of (absolute) degree greater than one. Hence, the maps $\mathbf{u}_D^\kappa$ are unstable configurations for $\kappa > 1$, and will typically not be observed in simulations using energy-descent-based solver algorithms. Figure 4 shows configurations for the cases $\kappa = 2, \ldots, 5$ obtained by numerical simulations using the nonconforming discretization and the discrete gradient-flow solver. One can see that the initial degree κ singularity of $\mathbf{u}_D^\kappa$ splits up, and the system runs into a stable state with κ isolated singularities of unit degree each.

6. BENCHMARKING THE DISCRETIZATIONS

We now numerically compare the two discretizations of Section 3 on the benchmark Problems 1, 2(a) and 2(b) defined in Section 5. We omit Problem 3, which involves harmonic maps with multiple singularities. The configurations of Problem 3 are essentially multiple copies of the single-singularity case of Problem 2, and we therefore do not expect to see any new effects. Also, this problem has no closed-form solution, but solving it numerically is also problematic (see Chapter 7.3 below). As a consequence, we do not have a good reference to test the discretization errors with.

6.1. **Benchmarking procedure.** For the numerical experiments we employ uniform triangulations $\mathcal{T}_h$ of $\Omega \subset \mathbb{R}^n$ into simplices. Starting from triangulations with only 2 triangles (if Ω is two-dimensional) or 6 tetrahedra (if Ω is three-dimensional) we obtain sequences of test grids by uniform refinement. We label these grids by their refinement level $r \in \mathbb{N}$. A grid for $\Omega = \left(-\frac{1}{2}, \frac{1}{2}\right)^n$ with refinement level r has an element diameter of $h = 2^{-r}\sqrt{n}$.

We solve the algebraic systems resulting from the different discretizations to machine precision with the Riemannian trust-region solver of Section 3.2. That way, no algebraic constraint violation occurs, and the algebraic error introduced by the iterative nature of the solver remains negligible. We use the parameters $\Delta_0 = \frac{1}{2}$, $\delta_1 = 0.9$ and $\delta_2 = 10^{-2}$ for the outer solver. The inner multigrid solver is set to iterate until the H^1 seminorm of the correction drops below 10^{-10}. The initial iterates will be given below. The implementation is a hand-written C++ code based on the DUNE libraries [16, 46]. The `dune-alugrid` [4] extension module is used as the grid data structure, and the `dune-gfe` module[2] provides the implementation of the projection-based finite elements.

Let $\mathbf{u}$ denote the known exact solution of a problem, and let $\mathbf{u}_h$ be a finite element approximation on a grid of maximal edge length h. We can directly compute the errors in the L^2-norm $\|\mathbf{u}_h - \mathbf{u}\|_{L^2}$ and the H^1-seminorm $|\mathbf{u}_h - \mathbf{u}|_{H^1}$. On our hierarchy of grids obtained by uniform refinement, we can estimate the corresponding orders of convergence by

$$\mathrm{EOC}_h^{L^2} := \log_2\left(\frac{\|\mathbf{u}_{2h} - \mathbf{u}\|_{L^2}}{\|\mathbf{u}_h - \mathbf{u}\|_{L^2}}\right) \qquad \text{and} \qquad \mathrm{EOC}_h^{H^1} := \log_2\left(\frac{|\mathbf{u}_{2h} - \mathbf{u}|_{H^1}}{|\mathbf{u}_h - \mathbf{u}|_{H^1}}\right).$$

Note that for the nonconforming discretization, the approximate solutions $\mathbf{u}_h$ are piecewise polynomial functions, and therefore the harmonic energy can be integrated exactly. However, the reference solutions $\mathbf{u}$ are not polynomial, and therefore computing the errors and convergence orders always involves a quadrature error. The projection-based finite elements, however, are not piecewise polynomials themselves, and no exact quadrature formula is known even for the harmonic energy. Here, we use Gauss–Legendre quadrature of second and sixth order for the harmonic energy and the EOCs, respectively. Finding more appropriate quadrature formulas for geometric finite elements remains an interesting research subject of its own.

6.2. **Inverse stereographic projection.** We begin with Problem 1, i.e., the computation of a harmonic map from $\Omega = \left(-\frac{1}{2}, \frac{1}{2}\right)^2$ to S^2, with boundary data given by the restriction of the inverse stereographic projection π_{st}^{-1} to Ω. One solution to this is the inverse stereographic projection itself, which is smooth.

We measure the discretization errors for finite elements of orders $p = 1$ and $p = 2$, starting the solver from the interpolant of the inverse stereographic projection. The results are shown in Table 1. The rates agree with what would be expected for a linear problem, i.e., they are iin $\mathcal{O}(h^{p+1})$ for the L^2-error and in $\mathcal{O}(h^p)$ for the H^1-error. For the projection-based finite element discretization this has been proven in [29, 32]. For nonconforming discretizations with $p = 1$, optimal error estimates in the energy norm have been derived for a corresponding saddle-point formulation in [34] for the case $m = n = 2$, and for more general two- and three-dimensional settings in [15].

Table 1 also shows the discrete energies of the minimizers for the different grids. The values produced by the two different discretizations roughly agree—more so for the second-order finite elements. Observe how the minimal energy increases with increasing mesh refinement for the nonconforming discretization, whereas it decreases for the conforming discretization. The latter would be the expected behavior for nested finite element spaces. However, neither the conforming

---

[2]`https://gitlab.mn.tu-dresden.de/osander/dune-gfe/`

		nonconforming			conforming		
r	$|\mathcal{T}_h|$	$E[\mathbf{u}_h]$	$\mathrm{EOC}_h^{L^2}$	$\mathrm{EOC}_h^{H^1}$	$E[\mathbf{u}_h]$	$\mathrm{EOC}_h^{L^2}$	$\mathrm{EOC}_h^{H^1}$
		order $p = 1$					
1	8	2.667	-	-	3.192	-	-
2	32	2.920	1.930	0.949	3.062	2.152	1.044
3	128	2.986	1.984	0.988	3.023	2.005	1.002
4	512	3.003	1.996	0.997	3.013	1.997	1.000
5	2048	3.008	1.999	0.999	3.010	1.999	1.000
6	8192	3.009	2.000	1.000	3.009	2.000	1.000
7	32 768	3.009	2.000	1.000	3.009	2.000	1.000
8	131 072	3.009	2.000	1.000	3.009	2.000	1.000
		order $p = 2$					
1	8	3.004	-	-	3.030	-	-
2	32	3.009	2.962	1.938	3.011	2.924	1.927
3	128	3.009	2.990	1.987	3.009	3.027	1.984
4	512	3.009	2.998	1.997	3.009	3.013	1.996
5	2048	3.009	2.999	1.999	3.009	3.004	1.999
6	8192	3.009	3.000	2.000	3.009	3.001	2.000
7	32 768	3.009	3.000	2.000	3.009	3.000	2.000

TABLE 1. Experimental discretization error convergence orders for Problem 1

nor the nonconforming finite elements form nested approximation space hierarchies with respect to the refinement r.

6.3. **Radial projection.** We repeat the same experiment for Problems 2(a) and 2(b), which both lead to harmonic maps with a singularity.

6.3.1. *The case $n = m = 3$.* We first test the discretizations on Problem 2(a), the solution of which is the radial projection $\mathbf{u}_\odot : x \mapsto \frac{x}{|x|}$ on the three-dimensional domain $\Omega = \left(-\frac{1}{2}, \frac{1}{2} \right)^3$. Recall that the radial projection map $\mathbf{u}_\odot$ is in $H^1(\Omega; S^2)$, but not in $H^2(\Omega; S^2)$ (more details in Chapter 5.2).

We cannot pick the solution as the initial iterate, and instead we choose

$$(9) \qquad \mathbf{u}_h^0 \in \mathcal{A}_h, \qquad \mathbf{u}_h^0(x) = \begin{cases} \mathbf{u}_\odot(x) & \text{if } x \in \mathcal{L}_h \text{ and } x \neq (0,0,0)^T, \\ (0,0,1)^T & \text{if } x \in \mathcal{L}_h \text{ and } x = (0,0,0)^T. \end{cases}$$

The special treatment for the value at $(0,0,0)^T \in \Omega$ is necessary because it is a Lagrange point, but $\mathbf{u}_\odot$ is not defined there. (Even though projection-based finite elements can represent singular functions, they cannot represent functions that are singular at Lagrange points.) The choice of $(0,0,1)^T$ for the value at $x = (0,0,0)^T$ is arbitrary; note that it breaks some of the problem's inherent symmetry.

The experimental results are listed in Table 2. The measured orders of convergence are much lower now that the solution is not in $H^2(\Omega; S^2)$, and they have a larger variance. For both discretizations we obtain orders that are a little under 1 for the L^2-error and around 0.4 for the H^1-error. Given the regularity of the solution $\mathbf{u}_\odot$, these convergence orders are plausible. However, no rigorous a priori discretization error bounds exist currently for either discretization. The conforming discretization seems to perform a bit better than the nonconforming one, but this may be coincidental. Its L^2-error is much closer to 1, but it has an outlier at ≈ 0.5 for the grid with $r = 3$. The reason for this is unclear.

		nonconforming			conforming		
r	$\|\mathcal{T}_h\|$	$E[\mathbf{u}_h]$	$\mathrm{EOC}_h^{L^2}$	$\mathrm{EOC}_h^{H^1}$	$E[\mathbf{u}_h]$	$\mathrm{EOC}_h^{L^2}$	$\mathrm{EOC}_h^{H^1}$
				order $p=1$			
1	48	5.308	-	-	7.858	-	-
2	384	6.610	1.052	0.312	8.001	1.023	0.435
3	3072	7.164	0.776	0.364	7.937	0.555	0.479
4	24 576	7.423	0.720	0.376	7.820	0.863	0.409
5	196 608	7.549	0.805	0.406	7.751	0.928	0.441
6	1 572 864	7.612	0.888	0.438	7.714	0.962	0.464
				order $p=2$			
1	48	6.642	-	-	7.920	-	-
2	384	7.164	0.861	0.363	7.833	1.258	0.480
3	3072	7.409	0.894	0.414	7.753	1.205	0.495
4	24 576	7.538	0.925	0.438	7.713	1.138	0.496
5	196 608	7.605	0.954	0.459	7.694	1.080	0.497

TABLE 2. Experimental discretization error convergence orders for Problem 2(a), with the singularity at $(0,0,0)^T$ and using the initial iterate defined in (9)

Again, Table 2 also shows the energies of the minimizers for the two discretizations. The energies increase as the grid is refined for the nonconforming discretization (with one exception), whereas it decreases for the nonconforming discretization. This is the same behavior as in the smooth case.

6.3.2. *Moving the singularity off the Lagrange point.* The situation of the previous section is somewhat special: The singularity of the exact solution is right on a Lagrange point for all grid refinements, and it therefore cannot be represented even with a projection-based finite element function.

To study whether this has harmful consequences, we redo the previous experiment with a singularity that is slightly shifted. More precisely, we keep the domain $\Omega = \left(-\frac{1}{2}, \frac{1}{2}\right)^3$, but we set the Dirichlet boundary values such that

$$\mathbf{u}_{\odot,\delta} := \frac{x-\delta}{|x-\delta|}, \qquad \delta := 2^{-6}\left(\frac{1}{3}, \frac{1}{3}, \frac{1}{3}\right)^T$$

becomes a solution. With this choice of δ, the singularity is not on a Lagrange point for all grids that we test with, and we can therefore use it as initial iterate without modification.

Table 3 shows the results. One can see that the new position of the singularity does have an influence on the convergence behavior. For both types of discretizations, the measured orders are a little higher. On the other hand, the orders are still highly unstable from level to level, and it is not possible to single out a fixed value as the definite order. There is even one outlier where the order becomes negative. Fully understanding the behavior of the discretizations in such a nonregular scenario certainly requires further work.

6.3.3. *The case $n = m = 2$.* In the situation of Problem 2(b) the solution is even less regular. The radial projection $\mathbf{u}_{\odot} : \left(-\frac{1}{2}, \frac{1}{2}\right) \to S^1$ is still a harmonic map, but it is so only in the distributional sense (Lemma 6). As a consequence, the Dirichlet energy of $\mathbf{u}_{\odot}$ is not a finite number. The problem is still well-defined in finite element spaces when using numerical quadrature, but we cannot hope for convergence in the H^1 sense.

To compute finite element solutions we use the same approach as for Problem 2(a) with $n = m = 3$. The singularity is on a Lagrange point again, and we set the initial iterate to the

		nonconforming			conforming		
r	$|\mathcal{T}_h|$	$E[\mathbf{u}_h]$	$\mathrm{EOC}_h^{L^2}$	$\mathrm{EOC}_h^{H^1}$	$E[\mathbf{u}_h]$	$\mathrm{EOC}_h^{L^2}$	$\mathrm{EOC}_h^{H^1}$
		order $p=1$					
1	48	5.272	-	-	7.814	-	-
2	384	6.590	1.073	0.321	7.986	1.048	0.436
3	3072	7.151	0.820	0.385	7.978	0.349	0.354
4	24 576	7.416	0.815	0.391	7.834	0.962	0.403
5	196 608	7.545	1.027	0.532	7.747	1.462	0.778
6	1 572 864	7.609	1.472	0.407	7.712	1.698	0.549
		order $p=2$					
1	48	6.638	-	-	7.911	-	-
2	384	7.162	0.849	0.442	7.828	1.366	0.524
3	3072	7.403	0.918	0.466	7.751	1.530	0.592
4	24 576	7.533	1.171	0.512	7.712	1.790	0.546
5	196 608	7.602	1.783	0.806	7.693	-0.255	0.967

TABLE 3. Experimental discretization error convergence orders for Problem 2(a), with singularity at $\delta = 2^{-6}(\frac{1}{3}, \frac{1}{3}, \frac{1}{3})^T$

		nonconforming			conforming		
r	$|\mathcal{T}_h|$	$E[\mathbf{u}_h]$	$\mathrm{EOC}_h^{L^2}$	$|\mathbf{u}_h - \mathbf{u}|_{H^1}$	$E[\mathbf{u}_h]$	$\mathrm{EOC}_h^{L^2}$	$|\mathbf{u}_h - \mathbf{u}|_{H^1}$
		order $p=1$					
1	8	5.172	-	4.301	5.528	-	5.498
2	32	6.866	0.652	4.373	6.794	0.676	6.695
3	128	8.834	0.621	4.483	8.791	0.824	6.842
4	512	$1.096 \cdot 10^1$	0.761	4.529	$1.094 \cdot 10^1$	0.878	6.853
5	2048	$1.312 \cdot 10^1$	0.832	4.544	$1.311 \cdot 10^1$	0.897	6.855
6	8192	$1.529 \cdot 10^1$	0.868	4.548	$1.529 \cdot 10^1$	0.910	6.856
7	32 768	$1.747 \cdot 10^1$	0.890	4.549	$1.747 \cdot 10^1$	0.920	6.856
8	131 072	$1.965 \cdot 10^1$	0.905	4.550	$1.965 \cdot 10^1$	0.928	6.856
		order $p=2$					
1	8	7.430	-	3.636	$1.019 \cdot 10^1$	-	3.850
2	32	9.452	0.643	3.763	$1.230 \cdot 10^1$	0.933	3.883
3	128	$1.159 \cdot 10^1$	0.787	3.805	$1.447 \cdot 10^1$	0.921	3.885
4	512	$1.376 \cdot 10^1$	0.845	3.818	$1.664 \cdot 10^1$	0.923	3.886
5	2048	$1.593 \cdot 10^1$	0.876	3.822	$1.882 \cdot 10^1$	0.929	3.887
6	8192	$1.811 \cdot 10^1$	0.895	3.823	$2.100 \cdot 10^1$	0.935	3.887
7	32 768	$2.028 \cdot 10^1$	0.909	3.823	$2.318 \cdot 10^1$	0.941	3.887

TABLE 4. Experimental discretization error convergence orders for Problem 2(b)

vector $(1, 0)^T$ there. The results are shown in Table 4. For both discretizations we still observe convergence in the L^2-sense to $\mathbf{u}_\odot$. The orders are even a bit less oscillatory than in the three-dimensional situation. Indeed, the orders of the conforming discretization are slightly better than the orders of the nonconforming one, and both seem to be increasing slowly with increasing mesh resolution. All orders stay below 1, and it is unclear whether the orders would approach a limit on even finer grids. In similar experiments in [34], Hu, Tai, and Winther observed the same linear convergence.

However, in the H^1-sense there is no convergence at all. Table 4 shows the errors, and one can observe that they converge to a fixed value not equal to zero.

(a) $\mathbb{R}^2 \supset \Omega \to S^2$ (b) $\mathbb{R}^3 \supset \Omega \to S^2$ (c) $\mathbb{R}^2 \supset \Omega \to S^1$
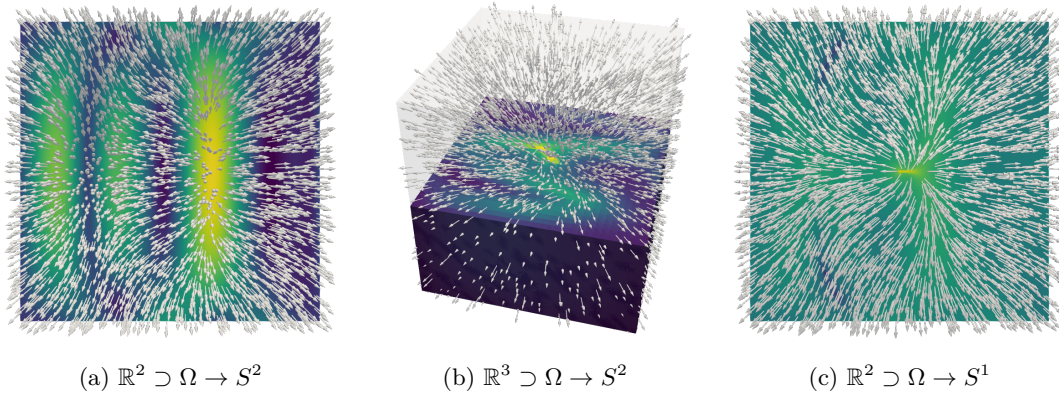
FIGURE 5. Initial configurations $\mathbf{u}_h^0$ for Problems 1, 2(a) and 2(b). These are perturbations of the problem solutions.

7. BENCHMARKING THE SOLVERS

In this second chapter of numerical tests we benchmark the solver algorithms of Chapter 4. The intention is to test both solvers of Chapter 4 with both discretizations of Chapter 3. However, the nonconforming gradient-flow solver cannot be combined with the conforming discretization, which cannot, without modifications, handle values not on the sphere. We therefore end up with three combinations:

(1) The nonconforming gradient-flow solver of Chapter 4.1 for problems discretized with the nonconforming discretization (henceforth abbreviated as Solve(GF)/Discr(N)),
(2) the Riemannian trust-region solver of Chapter 4.2 with the nonconforming discretization (Solve(TR)/Discr(N)),
(3) the Riemannian trust-region solver with the conforming discretization (Solve(TR)/Discr(C)).

We are interested in iteration numbers, wall times, and how the nonconformity of the gradient-flow solver influences the final result. In order to measure this nonconformity we use the quantity

$$\delta_1[\mathbf{v}_h] := \int_\Omega \mathcal{I}_h^1\big(\big||\mathbf{v}_h|^2 - 1\big|\big) \, \mathrm{d}x,$$

which is the quantity that appears in the theoretical constraint violation bound (5). As δ_1 is zero for all iterates provided by the Riemannian trust-region method, we show it only for the gradient-flow solver.

We again use the C++/DUNE implementation of the previous chapter. For the wall-time measurements we used a standard laptop computer with an AMD Ryzen 7 processor and 16 GB of DDR4 RAM. We set both solvers to iterate until the H^1-seminorm of the correction drops below $\epsilon_{\text{stop}} = 10^{-3}$. Other than that, we used the same settings for the Riemannian trust-region solver as before, i.e., initial trust-region radius $\Delta_0 = \frac{1}{2}$ and step acceptance parameters $\beta_1 = 0.9$ and $\beta_2 = 10^{-2}$. For the gradient-flow solver we set the pseudo time step size τ to four times the maximum element diameter h of the grid.

7.1. **Inverse stereographic projection.** For the first test we again consider Problem 1, where both domain $\Omega = (-\frac{1}{2}, \frac{1}{2})^2$ and image space S^2 are two-dimensional. This situation has a smooth solution in form of the inverse stereographic projection π_{st}^{-1}.

7.1.1. *Starting from close to the solution.* As a first test, we start the solvers from the nodal interpolation of the function π_{st}^{-1} itself. This is the approach used for the discretization error

| | | Solve(GF)/Discr(N) | | | | Solve(TR)/Discr(N) | | Solve(TR)/Discr(C) | |
| | | $(p=1)$ | | $(p=2)$ | | $(p=1)$ | $(p=2)$ | $(p=1)$ | $(p=2)$ |
r	$|\mathcal{T}_h|$	#Iter	$\delta_1[\mathbf{u}_h]$	#Iter	$\delta_1[\mathbf{u}_h]$	#Iter	#Iter	#Iter	#Iter
1	8	1	$2.776 \cdot 10^{-17}$	3	$2.776 \cdot 10^{-17}$	1	2	1	2
2	32	5	$1.196 \cdot 10^{-6}$	2	$7.235 \cdot 10^{-9}$	2	2	2	2
3	128	4	$4.370 \cdot 10^{-8}$	1	$1.551 \cdot 10^{-11}$	2	1	2	1
4	512	1	$5.197 \cdot 10^{-10}$	1	$2.508 \cdot 10^{-14}$	1	1	1	1
5	2 048	1	$1.002 \cdot 10^{-11}$	1	$1.566 \cdot 10^{-16}$	1	1	1	1
6	8 192	1	$1.756 \cdot 10^{-13}$	1	$1.256 \cdot 10^{-16}$	1	1	1	1
7	32 768	1	$2.986 \cdot 10^{-15}$	1	$1.226 \cdot 10^{-16}$	1	1	1	1
8	131 072	1	$1.510 \cdot 10^{-16}$			1		1	

TABLE 5. Iteration numbers and unit-length constraint violation δ_1 for Problem 1, starting close to the solution

measurements of the previous chapter, only the termination criterion is now less strict. Table 5 shows the iteration numbers. Not surprisingly, convergence is very fast. Indeed, both solvers rarely need more than a single iteration to reach a situation where the termination criterion holds. Also, the constraint violation introduced by the gradient-flow solver remains very small.

As the iteration histories are so short we omit the wall-time measurements.

7.1.2. *Starting from further away.* To challenge the solvers a bit more, we now construct an initial iterate that is further away from the solution. For this, define the scalar perturbation function

$$\widetilde{p}_n(x) := \cos(3\pi x_1) \cdot 4^n \prod_{i=1}^{n} \left(x_i^2 - \frac{1}{4} \right), \tag{10}$$

and use it to define the initial iterate

$$\mathbf{u}_h^0 := \mathcal{I}_h \left(\frac{\left(\pi_{\mathrm{st}}^{-1} + \widetilde{p}_2 \cdot \begin{pmatrix} 1 \\ 0 \\ 0 \end{pmatrix} \right)}{\left| \pi_{\mathrm{st}}^{-1} + \widetilde{p}_2 \cdot \begin{pmatrix} 1 \\ 0 \\ 0 \end{pmatrix} \right|} \right). \tag{11}$$

See Figure 5a for how this looks like. The perturbation function (10) vanishes on the domain boundary $\partial\Omega$, and therefore the initial iterate still satisfies the boundary condition $\mathbf{u}_h^0 = \pi_{\mathrm{st}}^{-1}$ on $\partial\Omega$. Also, the new initial iterate is still in the same homotopy class as π_{st}^{-1}.

Table 6 shows the solver performance results. One can see that the trust-region method still only needs a low number of iterations (4–5) to reach the required accuracy, independent of the grid refinement and the approximation order. The gradient-flow solver, on the other hand, needs much larger iteration numbers to reach the same accuracy. This is caused by the step size choice $\tau = 4h$, which leads to very small steps τ for finer grids. Indeed, the iteration numbers seem to double from one grid to the next, reflecting the linear dependence of τ on h. The coupling of τ and h, however, is necessary in order to keep a bound on the constraint violation δ_1 according to the estimate (5).

Table 6 shows that bounding the constraint violation in this way does work in practice. The constraint violations of the minimizers computed by the gradient-flow method remain in the range of 10^{-4} to 10^{-3}, which will be negligible for many practical purposes. In fact, as predicted

		Solve(GF)/Discr(N)				Solve(TR)/Discr(N)		Solve(TR)/Discr(C)	
		$(p=1)$		$(p=2)$		$(p=1)$	$(p=2)$	$(p=1)$	$(p=2)$
r	$\|\mathcal{T}_h\|$	#Iter	$\delta_1[\mathbf{u}_h]$	#Iter	$\delta_1[\mathbf{u}_h]$	#Iter	#Iter	#Iter	#Iter
1	8	14	$3.630 \cdot 10^{-2}$	10	$7.331 \cdot 10^{-2}$	4	4	4	5
2	32	15	$4.485 \cdot 10^{-2}$	15	$4.421 \cdot 10^{-2}$	4	4	4	5
3	128	24	$1.837 \cdot 10^{-2}$	25	$1.798 \cdot 10^{-2}$	4	4	4	4
4	512	43	$1.032 \cdot 10^{-2}$	43	$1.026 \cdot 10^{-2}$	4	4	4	4
5	2048	79	$5.480 \cdot 10^{-3}$	79	$5.472 \cdot 10^{-3}$	4	4	4	4
6	8192	151	$2.832 \cdot 10^{-3}$	151	$2.831 \cdot 10^{-3}$	4	4	4	4
7	32 768	296	$1.441 \cdot 10^{-3}$	296	$1.440 \cdot 10^{-3}$	4	4	4	4
8	131 072	585	$7.267 \cdot 10^{-4}$			4		4	

TABLE 6. Iteration numbers and unit-length constraint violation δ_1 for Problem 1 with initial iterate $\mathbf{u}_0^h$ given by (11)

		Solve(GF)/Discr(N)		Solve(TR)/Discr(N)		Solve(TR)/Discr(C)	
		$(p=1)$	$(p=2)$	$(p=1)$	$(p=2)$	$(p=1)$	$(p=2)$
r	$\|\mathcal{T}_h\|$	time [s]	time [s]	time [s]	time [s]	time [s]	time [s]
4	512	<1	1	0.444	0.566	0.480	0.621
5	2 048	2	8	0.685	1.217	0.819	1.389
6	8 192	13	99	1.647	3.409	1.992	4.185
7	32 768	152	1269	5.385	12.494	6.534	15.726
8	131 072	1880		20.466		25.159	

TABLE 7. Wall-times for Problem 1 with perturbation, with initial iterate $\mathbf{u}_0^h$ given by (11)

by (5), the constraint violation is proportional to the time-step size: As we have coupled τ to be proportional to h, the violation is roughly reduced by a factor of 2 for each grid refinement.

High iteration numbers would be unproblematic if the gradient-flow iterations were cheap. However, both solvers have roughly the same cost per iteration, dominated by having to solve a linear system of equations for each iteration. The matrix of the trust-region method is symmetric and positive definite, but it is iteration-dependent, and therefore has to be reassembled at each iteration. The linear system of the gradient-flow method is a saddle-point problem, on the other hand, because (in our implementation) the tangentiality of the correction is enforced via Lagrange multipliers. The benefit of this is that the two diagonal blocks of the matrix are independent of the iteration, and have to be assembled only once. On the downside, the problem has more unknowns than the formulation in local coordinates of the tangent space. An alternative implementation could formulate the tangent problem of the gradient-flow method in local coordinates of the tangent space, which would get rid of the Lagrange multipliers. Then, however, the entire matrix would have to be reassembled at each step as well.

The net effect of iteration numbers and time-per iteration can be seen in Table 7. The wall-time of the trust-region solver scales roughly linearly with the number of degrees of freedom. This is the combination of the fact that the outer trust-region solver needs a resolution-independent number of iterations, and that the inner solver is a multigrid solver with optimal complexity. The gradient-flow solver does not have these features, and therefore the wall-time it requires increases much faster as the grid gets finer. The difference between the two is in the range of two orders of magnitude on the finer grids.

| | | Solve(GF)/Discr(N) | | | | Solve(TR)/Discr(N) | | Solve(TR)/Discr(C) | |
| | | $(p=1)$ | | $(p=2)$ | | $(p=1)$ | $(p=2)$ | $(p=1)$ | $(p=2)$ |
r	$|\mathcal{T}_h|$	#Iter	$\delta_1[\mathbf{u}_h]$	#Iter	$\delta_1[\mathbf{u}_h]$	#Iter	#Iter	#Iter	#Iter
1	48	1	$4.163 \cdot 10^{-17}$	58	$2.015 \cdot 10^{-3}$	1	4	6	5
2	384	353	$1.171 \cdot 10^{-3}$	90	$4.888 \cdot 10^{-4}$	5	8	6	5
3	3072	244	$2.832 \cdot 10^{-4}$	159	$6.022 \cdot 10^{-5}$	5	8	7	5
4	24576	387	$4.222 \cdot 10^{-5}$	299	$6.850 \cdot 10^{-6}$	6	7	7	5
5	196608	750	$5.323 \cdot 10^{-6}$			6	7	7	5

TABLE 8. Iteration numbers and unit-length constraint violation for Problem 2(a), starting from close to the solution

| | | Solve(GF)/Discr(N) | | Solve(TR)/Discr(N) | | /Solve(TR)/Discr(C) | |
| | | $(p=1)$ | $(p=2)$ | $(p=1)$ | $(p=2)$ | $(p=1)$ | $(p=2)$ |
r	$|\mathcal{T}_h|$	time [s]	time [s]	time [s]	time [s]	time [s]	time [s]
3	3072	10.00	166.0	1.113	7.294	1.409	5.487
4	24576	177.0	3358	6.658	69.904	9.490	53.473
5	196608	3509		60.615	776.859	83.515	649.919

TABLE 9. Wall-times for Problem 2(a), starting from close to the solution

7.2. **Radial projection.** In the next sequence of tests we consider the benchmark problems 2(a) and 2(b), which involve harmonic maps with one singularity. We will see that this singularity influences the solver behavior.

7.2.1. *The case $n = m = 3$.* We start with Problem 2(a), which asks for a harmonic map on $\Omega = (-\frac{1}{2}, \frac{1}{2})^3$ with image in S^2, with boundary data given by the radial projection $\mathbf{u}_{\odot} : x \mapsto \frac{x}{|x|}$. This projection is also a solution of the problem. It is singular at the origin, but nevertheless an element of $H^1(\Omega; S^2)$.

As the first test we start the solvers directly from (9), which is essentially the radial projection itself. Table 8 shows the number of iterations. One can see that this problem is more difficult than the previous one: The number of required trust-region iterations is still bounded independent from the grid resolution, but 5 to 8 iterations are now needed to reach the termination criterion even for this good initial iterate. The gradient-flow solver, on the other hand, quickly requires three-digit iteration numbers, and the inverse proportional dependence on the grid element diameter h can be observed again.

Table 8 also shows the constraint violation δ_1 of the gradient-flow solver. As in the previous example it decreases with each refinement step. Unlike previously, the reduction is even better than what is expected from the bound (5).

The wall-time measurements in Table 9 reflect the different iteration numbers. Even considering that each refinement now multiplies the number of degrees of freedom by 8, the time complexity of the trust-region is not quite optimal anymore. Still, it is between one and two orders of magnitude faster than the gradient-flow solver.

7.2.2. *Starting farther away from the solution.* As for the previous problem we now start the solver at initial data that is further away from the discrete solution. To construct an initial

(a) no perturbation, $r = 4$ (b) no perturbation, $r = 5$ (c) no perturbation, $r = 6$

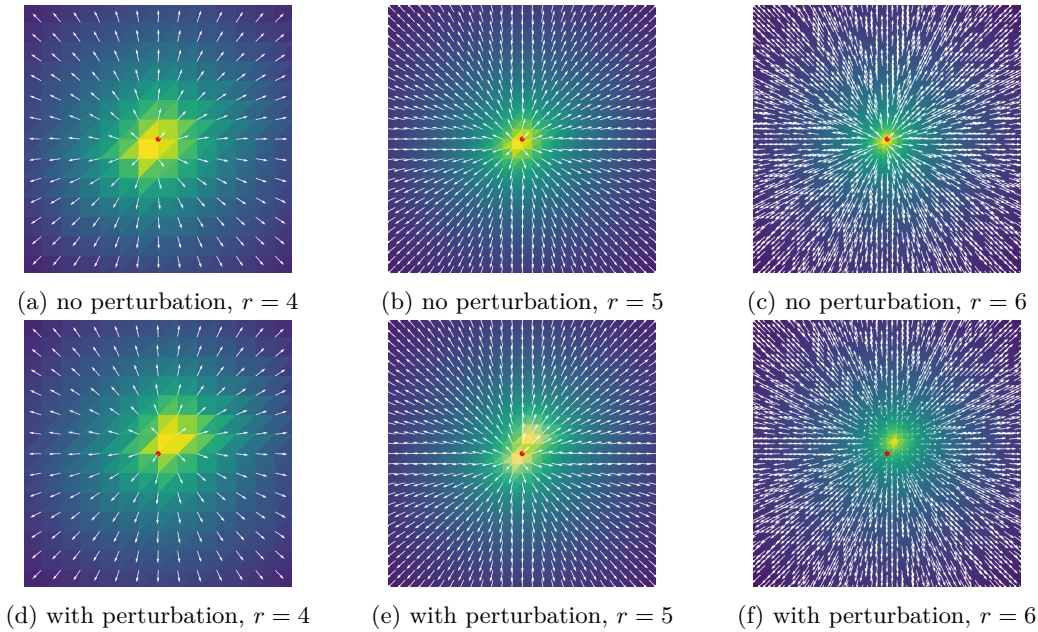(d) with perturbation, $r = 4$ (e) with perturbation, $r = 5$ (f) with perturbation, $r = 6$

FIGURE 6. Nonconforming discretization: Limit configurations $\mathbf{u}_h$ for Problem 2(a) for initial iterate (9) close to the solution (top row), and with initial iterate (12) (bottom row)

iterate, we reuse the scalar perturbation function $\widetilde{p}_n$ defined in (10), and define

$$(12) \qquad \mathbf{u}_h^0(x) := \mathcal{I}_h\left(\frac{x + \widetilde{p}_3(x)\left(\begin{smallmatrix} 1 \\ 1 \\ 0 \end{smallmatrix}\right)}{\left| x + \widetilde{p}_3(x)\left(\begin{smallmatrix} 1 \\ 1 \\ 0 \end{smallmatrix}\right) \right|} \right).$$

This new initial iterate is visualized in Figure 5b. The singularity has moved to about $(0.148, 0.148, 0)^T$, but by the construction of $\widetilde{p}_n$ in (10), the map $\mathbf{u}_h^0$ still satisfies the boundary conditions $\mathbf{u}_h^0 = \mathbf{u}_\odot$ on $\partial\Omega$.

Table 10 shows the iteration numbers and nonconformity, and Table 11 shows the wall times. The iteration numbers and wall times have increased further, but qualitatively the behavior is still the same as for the previous examples. However, looking at the computed configurations reveals a problem: Figures 6 and 7 show the limit configurations on three different grids for the nonconforming and the conforming discretization, respectively. In all pictures, a red dot marks the origin, i.e., the place where the singularity should be. One can see that its final computed position is not at the origin when starting from (12). Rather, it seems to be stuck at where it was in $\mathbf{u}_h^0$. This is not caused by the relatively weak termination criterion—the results do not change when setting ϵ_{stop} to 10^{-15}. The effect is markedly stronger for the conforming discretization, but it is also visible for the nonconforming one. There seems to be an effect that obstructs the movement of singularities. One conjecture is that the fact that singularities are restricted to certain parts of an element in a projection-based finite element computation introduces these obstructions, but we have no direct justification for this yet.

Our results are in line with [34], where, in a two-dimensional example, the authors also observed different limit configurations depending on the initial iterate.

(a) no perturbation, $r = 4$ (b) no perturbation, $r = 5$ (c) no perturbation, $r = 6$

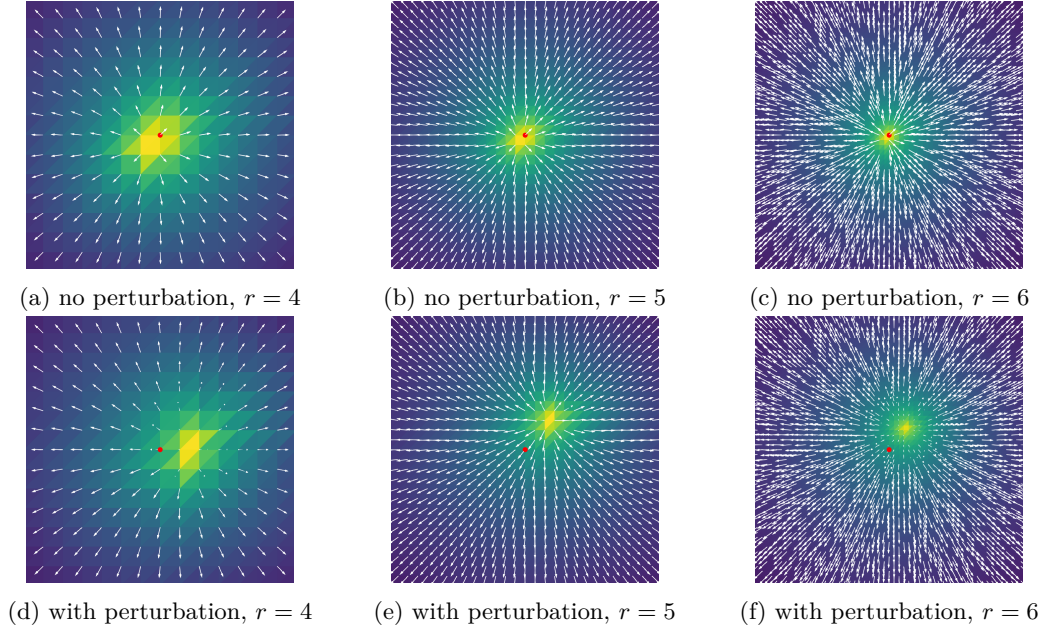(d) with perturbation, $r = 4$ (e) with perturbation, $r = 5$ (f) with perturbation, $r = 6$

FIGURE 7. Conforming discretization: Limit configurations $\mathbf{u}_h$ for Problem 2(a) for initial iterate (9) close to the solution (top row), and with initial iterate (12) (bottom row)

| | | Solve(GF)/Discr(N) | | | | Solve(TR)/Discr(N) | | Solve(TR)/Discr(C) | |
| | | $(p = 1)$ | | $(p = 2)$ | | $(p = 1)$ | $(p = 2)$ | $(p = 1)$ | $(p = 2)$ |
r	$|\mathcal{T}_h|$	#Iter	$\delta_1[\mathbf{u}_h]$	#Iter	$\delta_1[\mathbf{u}_h]$	#Iter	#Iter	#Iter	#Iter
1	48	1	$4.626 \cdot 10^{-17}$	115	$2.025 \cdot 10^{-2}$	1	9	6	6
2	384	237	$1.625 \cdot 10^{-2}$	79	$2.287 \cdot 10^{-2}$	10	8	10	14
3	3072	125	$7.342 \cdot 10^{-3}$	361	$7.507 \cdot 10^{-3}$	8	10	10	13
4	24576	920	$3.894 \cdot 10^{-3}$	1345	$3.897 \cdot 10^{-3}$	11	13	9	8
5	196608	5077	$2.009 \cdot 10^{-3}$			10	20	11	10

TABLE 10. Iteration numbers and unit-length violation for Problem 2(a) with initial iterate (12)

| | | Solve(GF)/Discr(N) | | Solve(TR)/Discr(N) | | Solve(TR)/Discr(C) | |
| | | $(p = 1)$ | $(p = 2)$ | $(p = 1)$ | $(p = 2)$ | $(p = 1)$ | $(p = 2)$ |
r	$|\mathcal{T}_h|$	time [s]	time [s]	time [s]	time [s]	time [s]	time [s]
1	48	<1	3	0.780	1.454	0.213	0.455
2	384	6	7	1.065	1.986	0.587	3.277
3	3072	6	364	2.553	11.721	2.522	14.532
4	24576	413	15040	14.691	144.848	17.072	101.011
5	196608	23700		105.062	2201.44	141.482	1326.09

TABLE 11. Wall-times for Problem 2(a) with initial iterate (12)

		Solve(GF)/Discr(N)				Solve(TR)/Discr(N)		Solve(TR)/Discr(C)	
		$(p=1)$		$(p=2)$		$(p=1)$	$(p=2)$	$(p=1)$	$(p=2)$
r	$|\mathcal{T}_h|$	#Iter	$\delta_1[\mathbf{u}_h]$	#Iter	$\delta_1[\mathbf{u}_h]$	#Iter	#Iter	#Iter	#Iter
1	8	1	$5.551 \cdot 10^{-17}$	12	$5.551 \cdot 10^{-17}$	1	5	4	5
2	32	18	$4.752 \cdot 10^{-3}$	21	$4.611 \cdot 10^{-4}$	5	6	7	7
3	128	33	$1.507 \cdot 10^{-3}$	36	$1.429 \cdot 10^{-4}$	5	5	9	7
4	512	61	$3.042 \cdot 10^{-4}$	64	$3.003 \cdot 10^{-5}$	6	5	9	7
5	2048	117	$5.116 \cdot 10^{-5}$	118	$5.238 \cdot 10^{-6}$	6	5	9	7
6	8192	226	$7.817 \cdot 10^{-6}$	228	$8.241 \cdot 10^{-7}$	6	5	9	7
7	32768	445	$1.130 \cdot 10^{-6}$	446	$1.220 \cdot 10^{-7}$	6	6	9	7
8	131072	883	$1.581 \cdot 10^{-7}$			6		9	

TABLE 12. Iteration numbers and unit-length violation for Problem 2(b), starting close to the solution

		Solve(GF)/Discr(N)		Solve(TR)/Discr(N)		Solve(TR)/Discr(C)	
		$(p=1)$	$(p=2)$	$(p=1)$	$(p=2)$	$(p=1)$	$(p=2)$
r	$|\mathcal{T}_h|$	time [s]	time [s]	time [s]	time [s]	time [s]	time [s]
4	512	<1	1	0.725	1.400	1.371	2.483
5	2048	1	6	1.502	4.111	2.703	6.492
6	8192	9	76	4.476	13.263	7.368	18.874
7	32768	95	858	15.06	51.086	22.476	60.835
8	131072	1244		48.234		71.807	

TABLE 13. Wall-times for Problem 2(b), starting close to the solution

7.2.3. *The case $n = m = 2$.* Next, we investigate Problem 2(b), the approximation of a harmonic map from $(-\frac{1}{2}, \frac{1}{2})^2$ to S^1 with boundary data given by the radial projection $\mathbf{u}_\odot$ now on this two-dimensional domain. Since the solution of this, i.e., $\mathbf{u}_\odot$ itself, is now even less regular than in the three-dimensional scenario, we expect additional numerical difficulties. However, the numerical results suggest that these fears are unfounded.

We again start from the nodal interpolation of the solution $x \mapsto \frac{x}{|x|}$, with the modification described in Chapter 6.3.3. Table 12 shows the total number of iterations as well as the unit-length constraint violation δ_1. The number of iterations of both methods is comparable to the three-dimensional situation of Chapter 7.2.1, which had an H^1-solution (Table 8). However, the nonconforming discretization seems to slightly reduce the number of total iteration steps of the Riemannian trust-region method.

The wall times, shown in Table 13, match the results of the earlier test problems: While the discrete gradient flow solver is faster for smaller problems, the trust-region solver outperforms it by orders of magnitude once the grid reaches a certain size.

7.2.4. $n = m = 2$: *Starting farther away from the solution.* Similar to the previous examples we use a perturbation of $\mathbf{u}_\odot$ as a starting configuration (see Figure 3c)

$$\mathbf{u}_h^0(x) := \mathcal{I}_h \left(\frac{x + \widetilde{p_2}(x) \binom{1}{0}}{|x + \widetilde{p_2}(x) \binom{1}{0}|} \right).$$

Very little changes for the trust-region solver: The iteration numbers increase only a bit, and the wall time is increased by about 30 %. For the gradient-flow solver, the increase of iteration numbers is above 500 with a corresponding increase of wall time. Finally, one can see from

| | | Solve(GF)/Discr(N) | | | | Solve(TR)/Discr(N) | | Solve(TR)/Discr(C) | |
| | | $(p=1)$ | | $(p=2)$ | | $(p=1)$ | $(p=2)$ | $(p=1)$ | $(p=2)$ |
r	$|\mathcal{T}_h|$	#Iter	$\delta_1[\mathbf{u}_h]$	#Iter	$\delta_1[\mathbf{u}_h]$	#Iter	#Iter	#Iter	#Iter
1	8	1	$5.551 \cdot 10^{-17}$	32	$1.284 \cdot 10^{-2}$	1	5	4	7
2	32	66	$2.204 \cdot 10^{-2}$	42	$2.346 \cdot 10^{-2}$	5	6	6	10
3	128	68	$9.836 \cdot 10^{-3}$	95	$9.501 \cdot 10^{-3}$	7	7	5	6
4	512	157	$5.705 \cdot 10^{-3}$	215	$5.763 \cdot 10^{-3}$	7	7	9	11
5	2048	378	$3.105 \cdot 10^{-3}$	492	$3.199 \cdot 10^{-3}$	7	9	8	9
6	8192	903	$1.661 \cdot 10^{-3}$	1130	$1.704 \cdot 10^{-3}$	10	9	7	9
7	32 768	2132	$8.669 \cdot 10^{-4}$	2591	$8.817 \cdot 10^{-4}$	9	9	1	9
8	131 072	4965	$4.443 \cdot 10^{-4}$			8		9	

TABLE 14. Iteration numbers and unit-length violation for Problem 2(b) starting away from the solution

| | | Solve(GF)/Discr(N) | | Solve(TR)/Discr(N) | | Solve(TR)/Discr(C) | |
| | | $(p=1)$ | $(p=2)$ | $(p=1)$ | $(p=2)$ | $(p=1)$ | $(p=2)$ |
r	$|\mathcal{T}_h|$	time [s]	time [s]	time [s]	time [s]	time [s]	time [s]
4	512	<1	2	0.975	2.166	1.334	3.215
5	2048	3	25	2.146	7.623	2.405	7.712
6	8192	38	368	7.262	23.605	6.063	24.152
7	32 768	450	4816	22.632	79.261	27.029	81.243
8	131 072	6941		67.557		77.2647	

TABLE 15. Wall-times for Problem 2(b), starting away from the solution

Table 14 that the constraint violation accumulated by the gradient-flow solver is now a factor 1000 larger than when starting from the radial projection. In absolute numbers it is still small, though.

Curiously, in this two-dimensional setting we could not reproduce the effect shown in Figure 7, where the singularity would get stuck away from its optimal position, regardless of the grid resolution and solver precision. In the two-dimensional experiments, the limit position of the singularity produced by both solvers and discretizations would always be right at the origin, where it should be.

7.3. **Harmonic maps with multiple singularities.** As the final test we measure iteration numbers and constraint violation for a harmonic map with several singularities (Problem 3). Here, the start iterate $\mathbf{u}_h^0$ is the function $\mathbf{u}_D^\kappa$ defined in (8), which has a singularity of degree κ. As such a configuration cannot stable for $\kappa \geq 2$ [18], we expect the solvers to converge to configurations with κ isolated singularities of degree 1. Figure 8 shows that this does actually happen: The final configuration consists of κ singularities placed at roughly equal distances on a circle. Unfortunately, one notices that the final configurations differ considerably depending on which discretization is used. Recalling from Chapter 7.2.2 that the methods sometimes converge to configurations where the singularity was not at the correct spot (in particular when using the conforming discretization), the configurations shown in Figure 8 should not receive a lot of trust. A deeper understanding of the behavior of discrete harmonic map models is needed before this issue can be fully resolved.

Regarding the solver performance, the usual metrics are shown in Table 16 for first-order discretizations. There, one sees that numerically, this scenario does seem to be more difficult than the others. In particular, for the trust-region methods the iteration numbers are not bounded
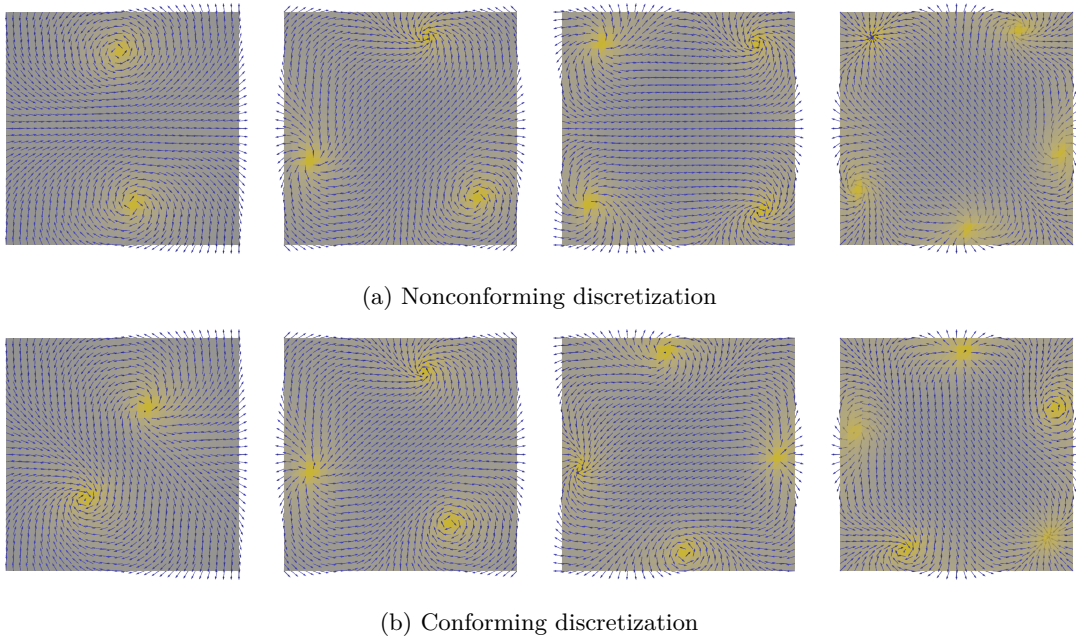
(a) Nonconforming discretization

(b) Conforming discretization

FIGURE 8. Cut through midsurface $x_3 = 0$ for $\kappa = 2, \ldots, 5$ in Problem 3, using the trust-region solver

anymore, but increase slowly with increasing mesh size. Unlike in earlier situations, more iterations are now needed for the conforming discretization than for the nonconforming one. The iteration numbers for the gradient-flow method are again much higher, but do not show a clear pattern as in earlier case. The constraint violation remains in a reasonable range, and gets smaller with decreasing $\tau = 4h$.

As the number of required trust-region iterations has increased, the total wall-time difference between the two solver algorithms is not as big anymore, at least for grid sizes in the range that we could measure. However, the difference is still considerable.

REFERENCES

[1] P.-A. ABSIL, R. MAHONY, AND R. SEPULCHRE, *Optimization algorithms on matrix manifolds*, Princeton University Press, 2009.

[2] P.-A. ABSIL, R. MAHONY, AND J. TRUMPF, *An extrinsic look at the Riemannian Hessian*, in International Conference on Geometric Science of Information, Springer, 2013, pp. 361–368.

[3] G. AKRIVIS, S. BARTELS, AND C. PALUS, *Quadratic constraint consistency in the projection-free approximation of harmonic maps and bending isometries*, Mathematics of Computation, online (2024).

[4] M. ALKÄMPER, A. DEDNER, R. KLÖFKORN, AND M. NOLTE, *The dune-alugrid module*, Archive of Numerical Software, 4 (2016), pp. 1–28.

[5] F. ALOUGES, *A new algorithm for computing liquid crystal stable configurations: The harmonic mapping case*, SIAM J. Numer. Anal., 34 (1997), pp. 1708–1726.

[6] H. ANTIL, S. BARTELS, AND A. SCHIKORRA, *Approximation of fractional harmonic maps*, IMA Journal of Numerical Analysis, 43 (2023), pp. 1291–1323.

		Solve(GF)/Discr(N)			Solve(TR)/Discr(N)		Solve(TR)/Discr(C)	
r	$|\mathcal{T}_h|$	#Iter	$\delta_1[\mathbf{u}_h]$	time [s]	#Iter	time [s]	#Iter	time [s]
				degree $\kappa = 2$				
1	48	1	$2.776 \cdot 10^{-17}$	<1	1	0.031	6	0.113
2	384	238	$6.920 \cdot 10^{-3}$	6	8	0.623	11	0.664
3	3072	662	$2.970 \cdot 10^{-3}$	27	9	2.082	20	4.622
4	24576	801	$8.260 \cdot 10^{-4}$	296	21	25.459	17	23.213
5	196608	3004	$2.068 \cdot 10^{-4}$	13960	19	164.47	33	368.75
				degree $\kappa = 3$				
1	48	1	$7.864 \cdot 10^{-17}$	<1	1	0.709	3	0.075
2	384	867	$1.730 \cdot 10^{-2}$	22	12	4.709	19	1.566
3	3072	619	$7.047 \cdot 10^{-3}$	26	11	4.513	16	3.853
4	24576	1310	$2.040 \cdot 10^{-3}$	491	14	15.818	29	39.086
5	196608	3058	$5.193 \cdot 10^{-4}$	14930	23	193.976	38	453.057
				degree $\kappa = 4$				
1	48	23	$4.919 \cdot 10^{-2}$	<1	4	0.082	7	0.135
2	384	194	$2.267 \cdot 10^{-2}$	5	9	1.378	15	1.915
3	3072	1198	$1.129 \cdot 10^{-2}$	48	11	4.025	24	5.653
4	24576	1238	$3.451 \cdot 10^{-3}$	483	27	33.696	29	42.064
5	196608	2908	$8.699 \cdot 10^{-4}$	15000	38	285.293	44	530.704
				degree $\kappa = 5$				
1	48	1	$4.857 \cdot 10^{-17}$	<1	1	0.712	3	0.076
2	384	303	$2.521 \cdot 10^{-2}$	7	13	2.170	13	0.917
3	3072	828	$1.535 \cdot 10^{-2}$	34	11	6.159	25	6.809
4	24576	1197	$4.775 \cdot 10^{-3}$	468	16	23.111	33	43.066
5	196608	3005	$1.196 \cdot 10^{-3}$	19660	22	178.889	56	655.635

TABLE 16. Iteration numbers, unit-length violation and wall-times for problem 3

[7] J. APPELL AND P. ZABREJKO, *Nonlinear superposition operators*, Cambridge University Press, 1990.

[8] S. BARTELS, *Robust a priori error analysis for the approximation of degree-one Ginzburg–Landau vortices*, ESAIM: Mathematical Modelling and Numerical Analysis, 39 (2005), pp. 863–882.

[9] ———, *Stability and convergence of finite-element approximation schemes for harmonic maps*, SIAM J. Numer. Anal., 43 (2005), pp. 220–238.

[10] ———, *Combination of global and local approximation schemes for harmonic maps into spheres*, J. Comput. Math., 27 (2009), pp. 170–183.

[11] ———, *Numerical analysis of a finite element scheme for the approximation of harmonic maps into surfaces*, Mathematics of Computation, 79 (2010), pp. 1263–1301.

[12] ———, *Numerical Methods for Nonlinear Partial Differential Equations*, Springer, 2015.

[13] ———, *Projection-free approximation of geometrically constrained partial differential equations*, Math. Comp., 85 (2016), pp. 1033–1049.

[14] S. BARTELS, M. GRIEHL, S. NEUKAMM, D. PADILLA-GARZA, AND C. PALUS, *A nonlinear bending theory for nematic LCE plates*, Mathematical Models and Methods in Applied Sciences, 33 (2023), pp. 1437–1516.

[15] S. BARTELS, C. PALUS, AND Z. WANG, *Quasi-optimal error estimates for the finite element approximation of stable harmonic maps with nodal constraints*, SIAM J. Numer. Anal., 61 (2023), pp. 1819–1834.

[16] P. BASTIAN, M. BLATT, A. DEDNER, N.-A. DREIER, C. ENGWER, R. FRITZE, C. GRÄSER, C. GRÜNINGER, D. KEMPF, R. KLÖFKORN, ET AL., *The Dune framework: Basic concepts and recent developments*, Computers & Mathematics with Applications, 81 (2021), pp. 75–112.

[17] H. BREZIS AND J.-M. CORON, *Large solutions for harmonic maps in two dimensions*, Communications in Mathematical Physics, 92 (1983), pp. 203–215.

[18] H. BREZIS, J.-M. CORON, AND E. H. LIEB, *Harmonic maps with defects*, Communications in Mathematical Physics, 107 (1986), pp. 649–705.

[19] R. Cohen, R. Hardt, D. Kinderlehrer, S.-Y. Lin, and M. Luskin, *Minimum energy configurations for liquid crystals: Computational results*, in Theory and Applications of Liquid Crystals, Springer, 1987, pp. 99–121.

[20] A. Conn, N. Gould, and P. Toint, *Trust-Region Methods*, SIAM, 2000.

[21] L. T. de Haan, V. Gimenez-Pinto, A. Konya, T.-S. Nguyen, J. M. Verjans, C. Sánchez-Somolinos, J. V. Selinger, R. L. Selinger, D. J. Broer, and A. P. Schenning, *Accordion-like actuators of multiple 3d patterned liquid crystal polymer films*, Advanced Functional Materials, 24 (2014), pp. 1251–1258.

[22] A. DeSimone, R. V. Kohn, S. Müller, and F. Otto, *Recent analytical developments in micromagnetics*, Preprint 80, Max-Planck-Institut für Mathematik in den Naturwissenschaften, 2004.

[23] Q. Du, M. D. Gunzburger, and J. S. Peterson, *Analysis and approximation of the Ginzburg–Landau model of superconductivity*, SIAM Review, 34 (1992), pp. 54–81.

[24] J. Eells and L. Lemaire, *Two reports on harmonic maps*, World Scientific, 1995.

[25] L. Evans, *Measure theory and fine properties of functions*, Routledge, 2018.

[26] T.-P. Fries and T. Belytschko, *The extended/generalized finite element method: An overview of the method and its applications*, Internat. J. Numer. Methods Engrg., 84 (2010), pp. 253–304.

[27] E. S. Gawlik and M. Leok, *Embedding-based interpolation on the special orthogonal group*, SIAM Journal on Scientific Computing, 40 (2018), pp. A721–A746.

[28] P. Grohs, H. Hardering, and O. Sander, *Optimal a priori discretization error bounds for geodesic finite elements*, Foundations of Computational Mathematics, 15 (2015), pp. 1357–1411.

[29] P. Grohs, H. Hardering, O. Sander, and M. Sprecher, *Projection-based finite elements for nonlinear function spaces*, SIAM J. Numer. Anal., 57 (2019), pp. 404–428.

[30] H. Hardering, *The Aubin–Nitsche trick for semilinear problems*, arXiv:1707.00963, (2017).

[31] ———, L_2-*discretization error bounds for maps into Riemannian manifolds*, Numerische Mathematik, 139 (2018), pp. 381–410.

[32] H. Hardering and O. Sander, *Geometric finite elements*, in Handbook of Variational Methods for Nonlinear Geometric Data, Springer, 2020, pp. 3–49.

[33] F. Hélein and J. C. Wood, *Harmonic maps*, Handbook of Global Analysis, 1213 (2008), pp. 417–491.

[34] Q. Hu, X.-C. Tai, and R. Winther, *A saddle point approach to the computation of harmonic maps*, SIAM J. Numer. Anal., 47 (2009), pp. 1500–1523.

[35] R. R. Kohlmeyer and J. Chen, *Wavelength-selective, IR light-driven hinges based on liquid crystalline elastomer composites*, Angewandte Chemie, 125 (2013), pp. 9404–9407.

[36] R. Kornhuber, *Adaptive monotone multigrid methods for nonlinear variational problems*, Vieweg+ Teubner Verlag, 1997.

[37] S.-Y. Lin, *Numerical analysis for liquid crystal problems*, PhD thesis, University of Minnesota, 1987.

[38] C. Liu and N. J. Walkington, *Approximation of liquid crystal flows*, SIAM J. Numer. Anal., 37 (2000), pp. 725–741.

[39] C. Melcher, *Chiral skyrmions in the plane*, Proc. of the Royal Society A, (2014), pp. 1–17.

[40] A. Monteil, R. Rodiac, and J. Van Schaftingen, *Ginzburg–Landau relaxation for harmonic maps on planar domains into a general compact vacuum manifold*, Archive for Rational Mechanics and Analysis, 242 (2021), pp. 875–935.

[41] L. J. Nebel, O. Sander, M. Bîrsan, and P. Neff, *A geometrically nonlinear Cosserat shell model for orientable and non-orientable surfaces: Discretization with geometric finite elements*, Comput. Methods Appl. Mech. Engrg., 416 (2023).

[42] R. H. Nochetto, S. W. Walker, and W. Zhang, *A finite element method for nematic liquid crystals with variable degree of orientation*, SIAM J. Numer. Anal., 55 (2017), pp. 1357–1386.

[43] O. Sander, *Geodesic finite elements on simplicial grids*, Internat. J. Numer. Methods Engrg., 92 (2012), pp. 999–1025.

[44] ———, *Geodesic finite elements of higher order*, IMA J. Numer. Anal., 36 (2016), pp. 238–266.

[45] ———, *Test function spaces for geometric finite elements*, arXiv preprint arXiv:1607.07479, (2016).

[46] ———, *DUNE—The Distributed and Unified Numerics Environment*, Springer, 2020.

[47] O. Sander, P. Neff, and M. Bîrsan, *Numerical treatment of a geometrically nonlinear planar Cosserat shell model*, Computational Mechanics, 57 (2016), pp. 817–841.

[48] R. Schoen and K. Uhlenbeck, *A regularity theory for harmonic maps*, Journal of Differential Geometry, 17 (1982), pp. 307–335.

[49] M. Struwe, *Variational methods*, Springer, 2000.

[50] D. J. Thouless, *Topological quantum numbers in nonrelativistic physics*, World Scientific, 1998.

[51] L. A. Vese and S. J. Osher, *Numerical methods for p-harmonic flows and applications to image processing*, SIAM J. Numer. Anal., 40 (2002), pp. 2085–2104.

[52] S. W. WALKER, *A finite element method for the generalized Ericksen model of nematic liquid crystals*, ESAIM Math. Model. Numer. Anal., 54 (2020), pp. 1181–1220.

[53] T. H. WARE, J. S. BIGGINS, A. F. SHICK, M. WARNER, AND T. J. WHITE, *Localized soft elasticity in liquid crystal elastomers*, Nature Communications, 7 (2016), pp. 1–7.

[54] A. WEINMANN, L. DEMARET, AND M. STORATH, *Total variation regularization for manifold-valued data*, SIAM Journal on Imaging Sciences, 7 (2014), pp. 2226–2257.

[55] T. J. WHITE AND D. J. BROER, *Programmable and adaptive mechanics with liquid crystal polymer networks and elastomers*, Nature Materials, 14 (2015), p. 1087.