Numer. Math. Theor. Meth. Appl.
doi: 10.4208/nmtma.2014.1231nm

Vol. 7, No. 1, pp. 58-106
February 2014

pde2path – A Matlab Package for Continuation and Bifurcation in 2D Elliptic Systems

Hannes Uecker[1,*], Daniel Wetzel[1] and Jens D. M. Rademacher[2]

[1] *Institut für Mathematik, Universität Oldenburg, 26111 Oldenburg, Germany.*
[2] *Universität Bremen, Fachbereich Mathematik, Postfach 33 04 40, 28359 Bremen, Germany.*

Received 26 September 2012; Accepted (in revised version) 2 May 2013

Available online 24 January 2014

Abstract. pde2path is a free and easy to use Matlab continuation/bifurcation package for elliptic systems of PDEs with arbitrary many components, on general two dimensional domains, and with rather general boundary conditions. The package is based on the FEM of the Matlab pdetoolbox, and is explained by a number of examples, including Bratu's problem, the Schnakenberg model, Rayleigh-Bénard convection, and von Karman plate equations. These serve as templates to study new problems, for which the user has to provide, via Matlab function files, a description of the geometry, the boundary conditions, the coefficients of the PDE, and a rough initial guess of a solution. The basic algorithm is a one parameter arclength-continuation with optional bifurcation detection and branch-switching. Stability calculations, error control and mesh-handling, and some elementary time-integration for the associated parabolic problem are also supported. The continuation, branch-switching, plotting etc are performed via Matlab command-line function calls guided by the AUTO style. The software can be downloaded from www.staff.uni-oldenburg.de/hannes.uecker/pde2path, where also an online documentation of the software is provided such that in this paper we focus more on the mathematics and the example systems.

AMS subject classifications: 35J47, 35J60, 35B22, 65N30

Key words: Elliptic systems, continuation and bifurcation, finite element method.

1. Introduction

For algebraic systems, ordinary differential equations (ODEs), and partial differential equations (PDEs) in one spatial dimension there is a variety of software tools for the numerical continuation of families of equilibria and detection and following of bifurcations.

*Corresponding author. *Email addresses:* hannes.uecker@uni-oldenburg.de (H. Uecker), daniel.wetzel@uni-oldenburg.de (D. Wetzel), rademach@math.uni-bremen.de (J. Rademacher)

 ©2014 Global-Science Press

These include, e.g., AUTO [11], XPPaut [10] (which relies on AUTO for the continuation part) and MatCont [14], see also www.enm.bris.ac.uk/staff/hinke/dss/ for a comprehensive though somewhat dated list. Another interesting approach is the "general continuation core" coco, [31].

However, for elliptic systems of PDEs with two spatial dimensions there appear to be few general continuation/bifurcation tools and hardly any that work out-of-the-box for non-expert users. PLTMG [2] treats scalar equations, and there are many case studies using ad hoc codes, often based on AUTO using suitable expansions for the second spatial direction; for 2D systems there also is ENTWIFE [38], which however appears to be no longer maintained since 2001. For experts we also mention Loca [29], which is designed for large scale problems, and oomph [16], another large package which also supports continuation/bifurcation, though this is not yet documented.

Our software pde2path is intended to fill this gap. Its main design goals and features are:

- **Flexibility and versatility.** The software is based on the Matlab pdetoolbox and treats PDE systems

$$G(u,\lambda) := -\nabla\cdot(c\otimes\nabla u) + au - b\otimes\nabla u - f = 0, \tag{1.1}$$

where $u = u(x) \in \mathbb{R}^N$, $x \in \Omega \subset \mathbb{R}^2$ some bounded domain, $\lambda \in \mathbb{R}$ is a parameter, $c \in \mathbb{R}^{N\times N\times 2\times 2}$, $b \in \mathbb{R}^{N\times N\times 2}$ (see (1.4a), (1.4b) below), $a \in \mathbb{R}^{N\times N}$ and $f \in \mathbb{R}^N$ can depend on $x, u, \nabla u$, and, of course, parameters. The standard assumption is that c, a, f, b depend on $u, \nabla u, \cdots$, locally, e.g., $f(x,u) = f(x,u(x))$; however, the dependence of c, a, f, b on arguments *can* in fact be quite general, for instance involving global coupling, see Section 3.5. The current version supports "generalized Neumann" boundary conditions (BC) of the form

$$\mathbf{n}\cdot(c\otimes\nabla u) + qu = g, \tag{1.2}$$

where $\mathbf{n}$ is the outer normal and again $q \in \mathbb{R}^{N\times N}$ and $g \in \mathbb{R}^N$ may depend on x, u, ∇u and parameters. These boundary conditions include zero flux BC, and a "stiff spring" approximation of Dirichlet BC via large prefactors in q and g, that we found to work well.

There are a number of predefined functions to specify domains Ω and boundary conditions, or these can be exported from matlab's pdetoolbox GUI, thus making it easy to deal with (almost) arbitrary geometry and boundary conditions. The software can also be used to time-integrate parabolic problems of the form

$$\partial_t u = -G(u,\lambda), \tag{1.3}$$

with G as in (1.1). This is mainly intended to easily find initial conditions for continuation. Finally, any number of eigenvalues of the Jacobian $G_u(u,\lambda)$ can be computed, thus allowing stability inspection for stationary solutions of (1.3).

- **Easy usage.** The user has to provide a description of the geometry, the boundary conditions, the coefficients of the PDE, and a rough initial guess of a solution. There are a number of templates for each of these steps which cover some standard cases and should be easy to adapt. The software provides a number of `Matlab` functions which are called from the command line to perform continuation runs with bifurcation detection, branch switching, time integration, etc.
- **Easy hackability and customization.** While `pde2path` works "out-of-the-box" for a significant number of examples, already for algebraic equations and 1D boundary value problems it is clear that there cannot be a general purpose "solve-it-all" tool for parametrized problems, see, e.g., [33, Chapter 3]. Thus, given a particular problem the user might want to customize `pde2path`. We tried to make the data structures and code as modular and transparent as possible. When dealing with a trade off between speed and readability we usually opted for the latter, and thus we believe that the software can be easily modified to add new features. In fact, we give some examples of "customization" below. Here, of course, having the powerful `Matlab` machinery at our disposal is a great advantage.

Remark 1.1. The i^{th} components of $\nabla\cdot(c\otimes\nabla u)$, au and $b\otimes\nabla u$ in (1.1) are given by

$$[\nabla\cdot(c\otimes\nabla u)]_i := \sum_{j=1}^N [\partial_x c_{ij11}\partial_x + \partial_x c_{ij12}\partial_y + \partial_y c_{ij21}\partial_x + \partial_y c_{ij22}\partial_y]u_j, \tag{1.4a}$$

$$[au]_i = \sum_{j=1}^N a_{ij}u_j, \qquad [b\otimes\nabla u]_i := \sum_{j=1}^N [b_{ij1}\partial_x + b_{ij2}\partial_y]u_j, \tag{1.4b}$$

and $f = (f_1,\cdots,f_N)$ should be seen as a column vector. If, for instance, we want to implement $-D\Delta u = -(d_1\Delta u_1,\cdots,d_N\Delta u_N) = -\nabla\cdot(D\nabla u)$ with D a constant diagonal diffusion matrix, as it often occurs in applications, then

$$c_{ii11} = c_{ii22} = d_i, \quad i = 1,\cdots,N, \quad \text{and all other } c_{ijkl} = 0, \tag{1.5}$$

and there are special ways to encode this (and other symmetric situations for c and a) in the `pdetoolbox`. See the templates below, and Subsection 3.1.3. For c, a, f see also the `pdetoolbox` documentation, for instance `assempde` in the `Matlab` help, while $-b\otimes\nabla u$ has been added by us. `pde2path` also provides a simplified encoding for isotropic systems without mixed derivatives, see Section 4.1. Finally, the i^{th} component of the boundary term $\mathbf{n}\cdot(c\otimes\nabla u)$ is given by

$$[\mathbf{n}\cdot(c\otimes\nabla u)]_i = \sum_{j=1}^N [n_1(c_{ij11}\partial_x + c_{ij12}\partial_y) + n_2(c_{ij21}\partial_x + c_{ij22}\partial_y)]u_j, \tag{1.6}$$

where $\mathbf{n} = (n_1, n_2)$.

Remark 1.2. Clearly, the splitting between a and f (or b and f) in (1.3) is not unique, e.g., for $G(u) = -\Delta u - \lambda u + u^3$ we could use ($a = -\lambda$, $f = -u^3$) or ($a = 0$, $f = \lambda u - u^3$). Similarly, for, e.g., $G(u) = -\Delta u - \partial_x u$ we can use $b = (1,0)$ and $f = 0$ or $b = (0,0)$ and $f = \partial_x u$. This flexibility of (1.1) has the advantage that in most cases the needed derivatives G_u, G_λ can be assembled efficiently from suitable coefficients c, a, b, and no numerical Jacobians are needed.

Also note that (1.1) allows to treat equations in non divergence form, too. For instance, we may write a scalar equation $-c(u)\Delta u - f(u) = 0$ as $-\nabla\cdot(c(u)\nabla u) + (c'(u)\nabla u)\cdot\nabla u - f(u) = 0$, and set $b_{111}(u) = -c'(u)\partial_x u$ and $b_{112}(u) = -c'(u)\partial_y u$, or add $-(c'(u)\nabla u)\cdot\nabla u$ to f.

Currently, the main drawbacks of `pde2path` are:

- `pde2path` requires `Matlab` including the `pdetoolbox`. Its usage explains the form (1.1). One of its drawbacks is a somewhat slow performance, compared to, e.g., some Fortran implementations of the FEM. Another drawback is a somewhat unhandy non-GUI description of geometry and boundary conditions, but for these we provide fixes. See also, e.g., [27]. On the other hand, in addition to the `Matlab`-environment, the `pdetoolbox` has a number of nice features: it also takes care of the geometry and mesh generation, it is well documented, it is fully based on sparse linear algebra techniques (which are vital for large scale problems), it exports (sparse) mass and stiffness matrices, and it provides a number of auxiliary functions such as adaptive mesh-refinement, or various plot options.
- Presently, only one parameter continuation is supported, and only bifurcations via simple eigenvalues are detected, located, and dealt with. In case symmetries cause multiple eigenvalues, artificial symmetry breaking sometimes is a viable ad hoc solution for the latter. We plan to add new features as examples require them, and invite every user to do so as well.

In the following we first very briefly recall some basics of continuation and bifurcation. Then we explain design and usage of our software by a number of examples, mainly a modified Bratu problem as a standard scalar elliptic equation, some Allen-Cahn type equations, some pattern forming Reaction-Diffusion systems, including some animal coats intended for illustration of how to set up problems with complicated geometries. We give a rather detailed bifurcation diagram for the Schnakenberg system, and we consider three rather classical problems from physics: Rayleigh-Bénard convection, some multi-component Bose-Einstein systems, and the von Kármán plate equations. Thus, besides some mathematical aspects of continuation and the example systems, here we explain the syntax and usage of the software in a rather concise way. More comprehensive documentation of the data structures and functions is included in the software, or online at [26].

Remark 1.3. Please report any bugs to `pde2path@mathematik.uni-oldenburg.de`, as well desired additional features. We will appreciate any feedback, and will be happy to provide help with interesting applications. See also [26] for an online documentation of the software, updates, FAQ, and general further information.

2. Some basics of continuation and bifurcation

2.1. Arclength continuation

A standard method for numerical calculation of solution branches of $G(u,\lambda)=0$, where $G: X\times\mathbb{R}\to X$ is at least C^1, X a Banach space, is (pseudo)arclength continuation. For convenience and reference here we recall the basic ideas. Standard textbooks on continuation and bifurcation are [1, 15, 20, 33], see also [7, 19], and the "matrix-free" approach [12]. Consider a branch $z(s) := (u(s),\lambda(s)) \in X\times\mathbb{R}$ parametrized by $s\in\mathbb{R}$ and the extended system

$$H(u,\lambda)=\begin{pmatrix}G(u,\lambda)\\ p(u,\lambda,s)\end{pmatrix}=0\in X\times\mathbb{R}, \tag{2.1}$$

where p is used to make s an approximation to arclength on the solution arc. Assuming that X is a Hilbert space with inner product $\langle\cdot,\cdot\rangle$, the standard choice is as follows: given s_0 and a point $(u_0,\lambda_0):=(u(s_0),\lambda(s_0))$, and additionally knowing a tangent vector

$$\tau_0 := (\dot{u}_0,\dot{\lambda}_0) := \frac{\mathrm{d}}{\mathrm{d}s}(u(s),\lambda(s))|_{s=s_0}$$

we use, for s near s_0,

$$p(u,\lambda,s) := \xi\,\langle \dot{u}_0, u(s)-u_0\rangle + (1-\xi)\dot{\lambda}_0(\lambda(s)-\lambda_0)-(s-s_0). \tag{2.2}$$

Here $0<\xi<1$ is a weight, and τ_0 is assumed to be normalized in the weighted norm

$$\|\tau\|_\xi := \sqrt{\langle\tau,\tau\rangle_\xi}, \quad \left\langle \begin{pmatrix}u\\ \lambda\end{pmatrix}, \begin{pmatrix}v\\ \mu\end{pmatrix}\right\rangle_\xi := \xi\,\langle u,v\rangle + (1-\xi)\lambda\mu.$$

For fixed s and $\|\tau_0\|_\xi = 1$, $p(u,\lambda,s)=0$ thus defines a hyperplane perpendicular (in the inner product $\langle\cdot,\cdot\rangle_\xi$) to τ_0 at distance $\mathrm{d}s := s-s_0$ from (u_0,λ_0). We may then use a predictor $(u^1,\lambda^1)=(u_0,\lambda_0)+\mathrm{d}s\tau_0$ for a solution (2.1) on that hyperplane, followed by a corrector using Newton's method in the form

$$\begin{pmatrix}u^{l+1}\\ \lambda^{l+1}\end{pmatrix} = \begin{pmatrix}u^{l}\\ \lambda^{l}\end{pmatrix} - A(u^l,\lambda^l)^{-1}H(u^l,\lambda^l), \quad \text{where } A=\begin{pmatrix}G_u & G_\lambda\\ \xi\dot{u}_0 & (1-\xi)\dot{\lambda}_0\end{pmatrix}. \tag{2.3}$$

Since $\partial_s p = -1$, on a smooth solution arc we have

$$A(s)\begin{pmatrix}\dot{u}(s)\\ \dot{\lambda}(s)\end{pmatrix} = -\begin{pmatrix}0\\ \partial_s p\end{pmatrix} = \begin{pmatrix}0\\ 1\end{pmatrix}. \tag{2.4}$$

Thus, after convergence of (2.3) yields a new point (u_1,λ_1) with Jacobian A^1, the tangent direction τ_1 at (u_1,λ_1) with conserved orientation, i.e., $\langle\tau_0,\tau_1\rangle = 1$, can be computed from

$$A^1\tau_1 = \begin{pmatrix}0\\ 1\end{pmatrix}, \text{ with normalization } \|\tau_1\|_\xi = 1. \tag{2.5}$$

Alternatively to (2.3) we may also use a chord method, where $A = A(u^1, \lambda^1)$ is kept fixed during iteration,

$$\begin{pmatrix} u^{l+1} \\ \lambda^{l+1} \end{pmatrix} = \begin{pmatrix} u^{l} \\ \lambda^{l} \end{pmatrix} - A(u^1, \lambda^1)^{-1} H(u^l, \lambda^l). \tag{2.6}$$

This avoids the costly evaluation of G_u at the price of a usually modest increase of required iterations.

The role of ξ is twofold. First, if $\mathtt{G}(\mathtt{u}, \lambda) = 0$ comes from the discretization of a PDE $G(u, \lambda)$ such as (1.1) over a domain Ω with n_p spatial points, then $\mathtt{u} \in \mathbb{R}^p$ with large p, say $p = Nn_p$. Here $\mathtt{u}$ stands for the FEM approximation of u, and $\mathtt{G(u,}\lambda\mathtt{)}$ for the FEM approximation of G. Below, the difference between the two will be clear from the context, and thus we will mostly drop the different notations again. Typically, we want to choose ξ such that $\xi \|\mathtt{u}\|^2_{\mathbb{R}^p}$ is an approximation of $|\Omega|^{-1}\|u\|^2_{L^2(\Omega)}$. If (as usual), $u \equiv 1$ corresponds to $\mathtt{u}_j = 1$ for $j = 1, \cdots, n_p$, then a rough estimate can be obtained by assuming that each component $u_i \equiv 1$, $i = 1, \cdots, N$. Then

$$\frac{1}{|\Omega|}\|u\|^2_{(L^2)^N} = N \stackrel{!}{=} \xi \|\mathtt{u}\|^2_{\mathbb{R}^p} = \xi N n_p, \quad \text{hence } \xi = 1/n_p. \tag{2.7}$$

This gives the basic formula for our choice of ξ. It is important that different ξ may give different continuations: in the Newton loop, small ξ favors changes in u, while larger ξ favors λ, see Fig. 1 for a sketch. Moreover, ξ is also related to the scaling of the problem: if, e.g., we replace λ by, say, $\tilde{\lambda} := 100\lambda$, then ξ should be adapted accordingly, i.e., $\tilde{\xi} = \xi/100$. In summary, ξ should be considered as a parameter that can be used to tune the continuation, and that may also be changed during runs if appropriate.

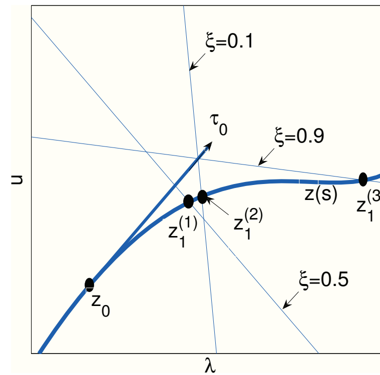

Figure 1: The role of ξ in a one-dimensional ($u \in \mathbb{R}$) sketch with $\tau_0 = (\dot{u}_0, \dot{\lambda}_0) = (1,1)$ (unnormalized). Depending on ξ we get different hyperplanes $\{(u,\lambda) \in \mathbb{R}^2 : \langle \tau_0, (u,\lambda)\rangle_\xi = \mathrm{ds}\}$ and consequently different "next points" $z_0^{(i)}$, $i = 1,2,3$, on the solution curve $z(s)$. Small ξ favors Newton search in u direction (and thus orthogonal to a "horizontal" branch), while large ξ favors the λ direction (parallel to a "horizontal" branch).

Given a weight ξ, a starting point (u_0, λ_0, τ_0), and an intended step size ds, the basic continuation algorithm thus reads as follows, already including some elementary stepsize control:

Algorithm 2.1: `cont`

1. **Predictor.** Set $(u^1, \lambda^1) = (u_0, \lambda_0) + \mathrm{ds}\tau_0$.
2. **Newton-corrector.** Iterate (2.3) (or (2.6)) until convergence; decrease ds if (2.3) fails to converge and return to 1; increase ds for the next step if (2.3) converges quickly.
3. **New tangent.** Calculate τ_1 from (2.5), set (u_0, λ_0, τ_0)=(u_1, λ_1, τ_1) and return to step 1.

Theoretically, this does not work at possible "bifurcation points" where A is singular (although generically continuation routines simply shoot past singular points). More specifically, we *define:*

B1. A *simple bifurcation point* is a point (u, λ) where det$A$ changes sign. The implicit assumption is that this happens due to a simple eigenvalue of A crossing zero.

This clearly excludes folds (also called turning points), where a simple eigenvalue of A reaches zero but det$A$ does not change sign, see [19]. However, folds are no problem for the algorithm and can easily be seen in the bifurcation diagram anyway. Therefore there is no special treatment of folds in the current 1-parameter version of `pde2path`. **B1** also excludes bifurcations via even numbers of eigenvalues crossing, which are more complicated to deal with.

Remark 2.1. Numerically, for **B1** we found it more robust to use $\xi = 1/2$ in the definition of A for bifurcation purposes. For algorithmic reasons, we only use the first part of **B1** for the detection of bifurcation points, i.e., the sign change of det$A$, which also occurs for an odd number of eigenvalues (counting multiplicities) crossing zero. Finally, by default we use the LU decomposition to calculate sign(det$A$), but there are other options, see Section 3.1.6.

After *detection* of a bifurcation between s_k and s_{k+1}, the bifurcation is *located* by a bisection method, with a secant, tangent, or quadratic predictor, see `p.biflocsw` below. Although this is a slow method for finding roots of continuous real functions [9, Chapter 2], in the setting of calculating sign changes of det$A$ via LU decomposition it seems difficult to improve. See also Section 3.1.6.

To switch branches we use "Method I" of [19] (page 379). Let (u_0, λ_0) be a simple bifurcation point, $G_u = G_u(u_0, \lambda_0)$, and $\tau_0 = (\dot{u}_0, \dot{\lambda}_0)$ be the tangent along the branch already computed. To obtain a tangent τ_1 along the other branch we proceed as follows:

2.2. Switching back and forth to the natural parametrization

If $\dot{\lambda} := \partial_s \lambda$ does not change sign, then we know that a branch also has the "natural parametrization" $(u(\lambda), \lambda)$, and, except at possible bifurcation points, $G_u(u, \lambda) u'(\lambda) =$

Algorithm 2.2: `swibra`

1. Calculate ϕ_1, ψ_1 with $G_u^0\phi_1 = 0$, $G_u^{0^T}\psi_1 = 0$, $\|\phi_1\| = 1$, $\langle\psi_1, \phi_1\rangle = 1$.
2. Let $\alpha_0 = \dot\lambda_0$, $\alpha_1 = \psi^T\dot u_0$, $\phi_0 = \alpha_0^{-1}(\dot u_0 - \alpha_1\phi_1)$.
3. Choose some small $\delta > 0$ and calculate the finite differences

$$a_1 = \frac{1}{\delta}\psi_1^T\left[G_u(u+\delta\phi_1, \lambda_0) - G_u^0\right]\phi_1,$$
$$b_1 = \frac{1}{\delta}\psi_1^T\left[\left[G_u(u_0+\delta\phi_1, \lambda_0) - G_u^0\right]\phi_0 + G_\lambda(u_0+\delta\phi_1, \lambda_0) - G_\lambda(u_0, \lambda_0)\right].$$

Assuming $\alpha_0 \neq 0$ (see [19] if this is not the case), set

$$\overline{\alpha}_1 = -\Big(\frac{a_1\alpha_1}{\alpha_0} + 2b_1\Big) \quad \text{and} \quad \tau_1 = \begin{pmatrix}\overline{\alpha}_1\phi_1 + a_1\phi_0 \\ a_1\end{pmatrix}.$$

Choose a weight ξ and a stepsize ds, set $\tau_0 = \tau_1/\|\tau_1\|_\xi$ and go to `cont`, Step 1. (If branch switching fails, i.e., if there is no convergence in `cont` or if the solution falls back onto the known branch, then it may help to change `ds` and/or ξ).

$-G_\lambda(u, \lambda)$ has the unique solution

$$u'(\lambda) = -G_u(u, \lambda)^{-1}G_\lambda(u, \lambda).$$

Thus, given (u_0, λ_0) we may use the predictor $(u^1, \lambda_1) = (u_0, \lambda_0) + \text{ds}\,(u'(\lambda_0), 1)$ and then correct with fixed $\lambda = \lambda_1$. Algorithmically, however, we choose to keep the predictor $(u^1, \lambda_1) = (u_0, \lambda_0) + \text{ds}\tau_0$, i.e., altogether,

$$(u^1, \lambda_1) = (u_0, \lambda_0) + \text{ds}\tau_0, \quad u^{l+1} = u^l - G_u(u^l, \lambda_1)^{-1}G(u^l, \lambda_1). \tag{2.8}$$

After convergence to (u_1, λ_1) we calculate the new tangent via (2.5), and **B1** can again be used as a check for bifurcation. Moreover, with $\text{tol}_{\dot\lambda} > 0$ a given tolerance, say $\text{tol}_{\dot\lambda} = 0.5$, this gives a criterion when to switch back and forth between the algorithms, namely:

$$\text{If } |\dot\lambda| > \text{tol}_{\dot\lambda}, \text{ then use (2.8), else use (2.3).} \tag{2.9}$$

Here again the weight ξ is important: for fixed $\xi = 1/2$ (say), $\dot\lambda \to 0$ as $n_p \to \infty$ unless the branch is strictly horizontal, i.e., $\dot u = 0$.

Remark 2.2. If applicable, (2.8) is usually slightly faster than (2.3), as expected. On the other hand, we found that even for "nearly horizontal" branches, locating the bifurcation point typically works better with arclength continuation (2.3).

3. Some scalar problems in `pde2path`

We now start the tutorial on `pde2path` by way of basic examples. The names in brackets refer to the sub-directory name of the directory `demos`, which contains the given example.

3.1. Bratu's problem (bratu)

Our first example is the scalar elliptic equation

$$-\Delta u - f(u,\lambda) = 0, \quad f(u,\lambda) = -10(u - \lambda e^u), \quad u = u(x) \in \mathbb{R}, \tag{3.1}$$

on the unit square with zero flux BC, i.e.,

$$x \in \Omega = \Big(-\frac{1}{2}, \frac{1}{2}\Big)^2, \quad \partial_n u|_{\partial\Omega} = 0. \tag{3.2}$$

This problem has the advantage that a number of results can immediately be obtained analytically, that there are some nontrivial numerical questions (see below), and that we can compare with previous results, see, e.g., [2,5,24].

There is a primary homogeneous branch $u \equiv u_h(s)$, $\lambda = \lambda(s)$, "starting" in $(0,0)$, on which $(u_h(s), \lambda(s))$ satisfies $f(u) = -10(u - \lambda e^u) = 0$. Bifurcation points (u_k, λ_k) are obtained from $G_u w = -\Delta w - f_u w \overset{!}{=} 0$ which yields $10(u_k - 1) = \mu_k$ where $\mu_k = (k_1^2 + k_2^2)\pi^2$, $k \in \mathbb{N}_0^2$, are the eigenvalues of $-\Delta$ on Ω, see Table 1. From Section 2.1, for arclength continuation the fold is nothing special and the two dimensional kernels $k = (k_1, k_2)$, $k_1 \neq k_2$, will go undetected using **B1**. The simple bifurcation points should be detected, and branch switching can be tried. On bifurcating branches, further bifurcations may be expected.

Table 1: Bifurcations from the homogeneous branch in (3.1).

k	(0,0)	(1,0),(0,1),	(1,1)	(2,1),(1,2)	(2,2)	$\cdots$
u_k	1	$1+\pi^2/10$	$1+\pi^2/5$	$1+\pi^2/2$	$1+4\pi^2/5$	$\cdots$
$\lambda_k = u_k e^{-u_k}$	$1/e \approx 0.3679$	≈ 0.2724	≈ 0.1520	≈ 0.0157	≈ 0.0012	$\cdots$
type	fold	double	simple	double	simple	$\cdots$

In pde2path, (3.1), (3.2) can be setup and run in a few steps explained next, to quickly obtain the (basic) bifurcation diagram and solution plots in Fig. 2 on page 71.

3.1.1. Installation and preparation

The basic pde2path installation consists of a root directory, called pde2path, with a subdirectory p2plib containing the actual software, a subdirectory demos with a further subdirectory for each problem, a subdirectory octcomp, providing some basic octave compatibility (see [26]), and one Matlab file setpde2path.m, which is a utility function to set the Matlab path. Each of the demos comes with a file *cmds.m, which contains the commands to run the example (and some comments), and which should be seen as a quarry for typical commands, and with a file *demo.m, which produces more verbose output. To start we recommend to run setpde2path (without arguments) in the root directory pde2path and then change into one of the demo-directories, e.g., type cd demos/bratu in Matlab. Then inspect the file bratucmds.m and copy paste the commands to the Matlab command line, or just execute bratucmds or bratudemo.

Table 2: Basic variables in structure p of a problem, and their initialization in bratuinit.

func. handles	meaning	example (in initbratu)
f	[c,a,f,b]=f(p,u,lam) PDE coefficients in (1.1)	p.f=@bratuf
jac	[cj,aj,glam,bj]=jac(p,u,lam) used to build G_u, G_λ from $f_u, f_\lambda, \cdots$, (see jsw below),	p.jac=@bratujac
bcf	bc=bcf(p,u,lam), boundary conditions function	p.bcf=@(p,u,lam) bc[a,c]
outfu	out=outfu(p,u,lam), defining output for bifurcation diagram, i.e., quantities saved on p.branch. Default setting out=stanbra(p,u,lam) gives out=[$\|u_1\|_\infty$; $\|u_1\|_{L^2}$];	p.outfu=@stanbra[b]
ufu	cstop=ufu(p,brout,ds), user function, called after each calculation of a point, e.g. for printing information and checking stopping criteria	p.ufu=@stanufu[b]
headfu	headfu(p) (no return arguments); print screen output headline	p.headfu=@stanheadfu[b]
blss, lss	(bordered) linear system solver, see Section 3.1.7	p.blss=@blss[b], p.lss=@lss[b].
various var.	**meaning**	**example (in initbratu)**
geo	geometry (and also BC in call to recnbc1)	[p.geo,bc]=recnbc1(0.5,0.5);[a,b,c]
points,edges,tria	mesh	p=stanmesh(p,lx,ly,nx)
bpoints,$\cdots$,btria	*b*ackground mesh (used for mesh-adaptation)	p=setbmesh(p)[b]
neq, np, nt	number of equations, mesh points, triangles	automatically from mesh
u,lam,tau,ds	essential continuation data	p.u=0.2*ones(p.np); $\cdots$;
branch, bifvals	lists generated via p.outfu, used to plot bif. diagrams	p.branch=[]; p.bifvals=[];
file names	**meaning**	**standard (in setfn)**
pre	name of subdir for files; by default automatically set to the struct name used in call to cont or *init	
pname,bpname	(base)filenames for output of points, bifurcation points; actual filenames augmented by counter	p.pname='p' for *p*oint, p.bpname='bp' for *bif.point*

[a] typical setup with bc independent of λ, u and hence defined in advance
[b] "standard" choices provided by p=stanparam(p)
[c] see bratuinit.m and documentation of pdetoolbox for explanation

3.1.2. General structure, initialization, and continuation runs

In pde2path, a continuation and bifurcation problem is described by a structure, henceforth called p (as in problem), which we now outline, see also Tables 2 and 4.* Essentially,

*In addition to the fields/variables listed, there are quite a few more within p. See stanparam.m in directory p2plib for these, and also for more comments on the ones which are listed. Some of the "control fields" are unlikely to be changed by the user, at least at the beginning, e.g., p.evopts.disp=0 (to suppress output during eigenvalue calculations), or p.pfig=1, p.brfig=2 (the figure numbers for plotting), and some of the additional fields/variables are only generated during computation, e.g., the residual res and the error

Table 3: Main "user" functions in `pde2path`.

main functions	purpose
p=cont(p(,msteps))	main continuation routine; aux. argument `msteps` overrules `p.nsteps`
p=swibra(pre,file,newpre)	branch-switching at bifurcation point from previous run, prefix set to newpre
p=meshadac(p,varargin)	adapt mesh in p, see Section 3.1.5 for varargin
p=findbif(p(,nbifpoints))	locate (nbifpoints) bifurcation point(s) based on G_u;
p=loadp(pre,file,npre)	load solution struct p from file and reset prefix to npre
plotbra(p,wnr,cmp(,aux)	plot component `cmp` of branch over λ, in figure `wnr`, various `aux` arguments
plotbraf(pre,file,wnr,cmp(,aux))	as plotbra but from file (saved previously, leave out the '.mat')
plotsol(p,wnr,cmp,pstyle)	plot solution, use plotsolf(pre,file,wnr,$\cdots$) to plot from file

`p` contains:

- function handles which describe the functions c, a, b, f and the BC (and possibly the Jacobian) in (1.1);
- fields which describe the geometry of the problem, including the FEM mesh;
- fields which hold the current solution, i.e., u, λ and the tangent τ;
- a number of variables, switches and further functions (i.e., function handles) controlling the behavior of the continuation and bifurcation algorithm, and filenames for file output.

Studying a continuation and bifurcation problem using `pde2path` thus consist of:

- Setting up a file defining the coefficient functions c, a, b and f (and usually a function for the Jacobians) in (1.1), e.g. `bratuf.m` (and `bratujac.m`). (Here we assume that the BC function p.bcf is defined inline as in `bratuinit.m` and many of the further examples.)
- Setting up an initialization function file, e.g., `bratuinit` filling `p`. This function is not called by `pde2path` but should be used as a convenience function by the user. Thus, it can have any sort of input/output arguments. As an example, in `bratuinit.m` we find it convenient to include the lengths l_x, l_y and the discretization n_x (from which n_y is calculated) as parameters, because, e.g., below we want to break the square symmetry of Ω to unfold the double eigenvalues associated to $k = (1,0), (0,1)$, $k = (2,1), (1,2)$ etc, see Section 3.1.6. The main initialization steps are (1) define p.f and p.jac, (2) define geometry and mesh, and usually the BC by an inline function, (3) set the parameters and provide a starting point. In many

estimate `err`, which are put into `p` for easy passing between subroutines and user access. Currently there are no global variables in `pde2path`, with the exceptions `pj,lamj` which are set for numerical differentiation in `resinj`. However, see also Section 3.1.7 and Section 3.5 for examples of using global variables to streamline calculations.

Table 4: Main switches and controls in a structure p used in cont, see stanparam.m for typical values and a number of additional switches with detailed comments. See also Section 4.3 for additional parameters controlling pmcont, the parallelmulti-continuation version.

Newton&Cont	meaning
imax,normsw	Newton: max number of steps and selection of norm ("norm-switch")
tol	stop-crit. for Newton, typically should be around 1e-10;
nsw	0 for Newton, 1 for chord
jsw	switch for derivatives (G_u, G_λ). 0: (G_u, G_λ) by (cj, aj, bj, g_λ), 1: G_u by cj, aj, bj, G_λ by FD 2: G_u by FD, G_λ by g_λ, 3: (G_u, G_λ) by FD.
dsmin,dsmax,dlammax	min/max stepsize in s, max stepsize in λ
lammin,lammax	min/max λ, preset to $\mp 10^6$, reset to use as stopping criteria
nsteps	number of steps to take
parasw,lamdtol	parametrization switch and tolerance: if parasw=0 resp. parasw=2 then always use (2.8) resp. (2.3). If parasw=1 then use (2.9) with tol_λ =lamdtol
amod, maxt, ngen	controls for mesh adaptation: adapt every amod-th step, aim at maxt triangles, in at most ngen refinement steps
errchecksw, errtol	switch and tol for a posteriori error estimate and handling, see Section 3.1.5.
Bif.,Plot&User-control	**meaning**
bifchecksw	0 for no checks, 1 for check via B1 with consistency with eigenvalues of G_u, 2 for B1 alone (default)
biflocsw	bisection method, 0 for secant, 1 for tangent, 2 for 2nd order (default)
neig	number of eigenvalues to calc. in spcalc for stability, default 50
neigdet	number of eigenvalues to calc. in bifdetec, default 0 (use LU decomp.)
eigsstart	0 to start eigs randomly, 1 to start with $(1, \cdots, 1)$ (default)
bisecmax	max number of bisections to locate bifurcation; turned off by bisecmax=0
spcalcsw	0/1 for eigenvalue calculation off/on (default 1)
pmod/pstyle	plot each pmod-th step in pstyle (1 mesh, 2 pcolor, 3 rendered 3D, $\cdots$)
pcmp, bpcmp	component to plot, component of branches to plot
smod	save every smod-th step
isw,vsw	interaction/verbosity switch: 0=none, 1=some, 2=much;
pfig,brfig,ifig	figure-numbers for u-plot, branch-plot and info-plot during cont.; in ifig we plot additional information, e.g., mesh after mesh adaptation, new tangent after swibra, or solution during time integration tint
timesw	if >0, print timing after cont. See stanparam.m for the timers in p.
nbp	number of user-components of branch to be printed on screen (p.ufu=@stanufu)
resfac, mst, pmimax	pmcont only, see Section 4.3

cases most parameters and switches can be set to "standard values". For this we provide the function p=stanparam(p), which should be called first, and afterwards individual parameters can be reset as needed. For the mesh generation we provide an elementary function p=stanmesh(p,hmax), where hmax is the maximal triangle side-length. This is based on initmesh from the pdetoolbox, i.e., a Delaunay algorithm. For rectangular domains the syntax p=stanmesh(p,nx,ny) is also allowed, which is based on poimesh from the pdetoolbox, with obvious meaning. If appli-

Table 5: The basic init-routine bratuinit.m, the definitions of PDE coefficients and Jacobian (see Section 3.1.3), and some selected commands (see bratucmds.m) to run pde2path. See also the files for detailed comments.

```
function p=bratuinit(p,lx,ly,nx) % init-routine, see bratuinit.m for more comments
p=stanparam(p); p.neq=1; p.f=@bratuf; p.jac=@bratujac; [p.geo,bc]=recnbc1(lx,ly);
p.bcf=@(p,u,lam) bc; % typical inline definition of the BC function
ny=round(nx*ly/lx);p=stanmesh(p,nx,ny); p=setbmesh(p); % mesh and "background" mesh
pre=sprintf('%s',inputname(1)); p=setfn(p,pre); % set filename (prefix)
p.xi=1/p.np; p.dlammax=0.02; p.lammin=0.02;
p.lam=0.2; p.u=0.1*ones(p.np,1); p.ds=0.05; % "trivial" branch
```

```
function [c,a,f,b]=bratuf(p,u,lam) % coeff for Bratu
u=pdeintrp(p.points,p.tria,u); c=1; a=0; f=-10*(u-lam*exp(u)); b=0;
```

```
function [cj,aj,glam,bj]=bratujac(p,u,lam) % Jacobian for Bratu
u=pdeintrp(p.points,p.tria,u); cj=1; aj=10*(1-lam*exp(u)); glam=-10*exp(u); bj=0;
```

```
% commands to run bratu in pde2path (selection, see also script file bratucmds.m)
p=[];p=bratuinit(p,0.5,0.5,20); p=cont(p);
q=swibra('p','bp1','q'); q.lammin=0.1;q.nsteps=20;q=cont(q);
plotbra(p,3,2,'ms',12,'lw',5,'fs',16,'cl','k');plotsolf('q','p20',4,1,1);
```

cable and used with care, this is faster and gives more regular meshes, i.e., nx/ny should correspond to L_x/L_y where L_x, L_y are the side lengths of the rectangle.

- Calling a number of pde2path functions. The basic call is p=cont(p); (a continuation run)*, which can be followed, e.g., by a repeated call to extend the branch. Or, in case a branch point has been found, a call to branch switching and subsequent continuation by q=swibra('p','bp1','q'); q=cont(q); where, e.g., ./p/bp1.mat is data written at a branch point during the previous run. In between runs, p can be modified from the command line, e.g., type p.imax=5 (say) to (re)set the maximal number of Newton-iterations, or call p=meshref(p) to refine the mesh before a subsequent run; afterwards, call p=cont(p) again. According to the settings, data is plotted and written to files in a sub-directory with name p.pre. There are also functions for further post processing, e.g., plotting of solutions and bifurcation diagrams, whose documentation is mainly provided within the corresponding matlab files, and by the calls in the example directories.

The fundamental user provided functions thus are the coefficient function p.f and the additionally recommended Jacobian function p.jac, see Section 3.1.3. To study a new problem, we recommend to edit copies of the files *init.m, *f.m and *jac.m of a suitable example (e.g., *=bratu for a scalar problem) in an empty directory, and start with calling p=[];p=newinit(p); p=cont(p). The bifurcation diagram in Fig. 2 and the solution plots are generated from the commands in Table 5, either given from the command line, or put into a Matlab script.

*also possible: p=cont(p,msteps), where msteps overrules (but does not replace) p.nsteps

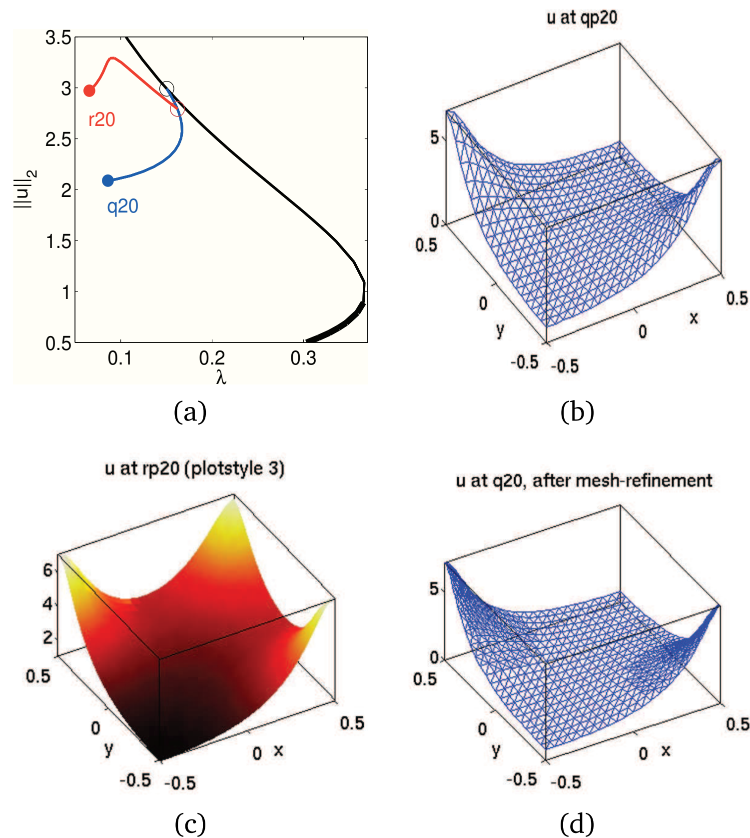

Figure 2: (a) Elementary bifurcation diagram for (3.1) generated by pde2path, over a uniform mesh with 800 triangles. Thick lines indicate stable (parts of) branches, thin lines unstable branches, ∘ bifurcation points. (b), (c) Some solution plots. (d) Preview of mesh refinement. See Section 3.1.5 for the quality of the mesh at the "ends" of branches q, r and the due mesh-refinement.

Remark 3.1. Here as in most of the examples we follow bifurcating branches only in one direction, since typically the other direction is related via some symmetry. In Fig. 2 both displayed bifurcations are pitchforks, and for instance the other direction of the q branch simply has maxima in $(x, y) = \pm(0.5, 0.5)$. The direction at bifurcation can conveniently be chosen by calling `swibra('p','bp1','q',ds)` with positive resp. negative `ds`.

3.1.3. The PDE-coefficients and Jacobians

The coefficient function `p.f` is of course mandatory, and the Jacobian function `p.jac` is recommended. The input argument `u` of both are the nodal values, $u_1(\cdot)$ =`u(1:p.np)`, $u_2(\cdot)$ =`u(p.np+1:2*p.np)`, $\cdots$, $u_N(\cdot)$ =`u((p.neq-1)*p.np+1:p.neq*p.np)`. For the outputs c, a, f, b (of `p.f`) we allow two forms, i.e., (arrays of) constants, or (arrays of) values *on the triangle midpoints* of the FEM-mesh, essentially as explained in [35]. There are two major ways to generate c, a, f, b from u. We first focus on f:

a) Use `u=pdeintrp(p.points,p.tria,u)` to first interpolate u to the triangle values (again called u), which yields a matrix

$$u = \begin{pmatrix} u_{11} & u_{12} & \cdots & u_{1,n_t} \\ \cdots & \cdots & \cdots & \cdots \\ u_{N1} & u_{N2} & \cdots & u_{Nn_t} \end{pmatrix},$$

where n_t is the number of triangles in the mesh. Then write, e.g., $f(u)$ in a standard `Matlab` way, i.e., `f=-10*(u-lam*exp(u));` see `bratuf` in Table 5.

b) First express, e.g., c, f as "Matlab text expression in x, y, u, ux, uy" (from [35]), with obvious meaning. For this, a parameter `lam` must be converted to a string. Afterwards, `pdetxpd` is called to evaluate the text expression on the triangle midpoints. See file `bratuft.m` for an example.

Option a) has the advantage that it is more "natural", more flexible, and, at least for simple expressions, slightly shorter. If, however, f depends on $x, u_x, \cdots$, then option b) might be shorter. To some extent it is a matter of taste which way to generate f is preferred (and similarly $c, \cdots$), therefore for (3.1) we provide both as templates. However, b) only allows local dependence $f(u,\cdots) = f(u(x),\cdots)$.

For `c`, which in principle is the $N \times N \times 2 \times 2 = 1{\times}1{\times}2{\times}2$ tensor

$$c_{11kl} = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}$$

we simply write `c=1` which corresponds to the simplest symmetric case.*

The tensor $b = b_{ijk}$, $i, j = 1, \cdots, N$, $k = 1, 2$ in (1.1) is not part of the `pdetoolbox`. Its coding and storage mimics that of `c`. In detail, `b` is an $2N^2 \times m$ array, where $m = 1$ (constant case) or $m =$`p.nt`, in the order

$$\begin{aligned} b = [&b_{111}; b_{112}; b_{211}; b_{212}; \cdots b_{N11}; b_{N12}; \\ &b_{121}; b_{122}; \cdots; b_{N21}; b_{N22}; \;\cdots\; b_{1N1}; \cdots; b_{NN2}], \end{aligned} \tag{3.3}$$

i.e., b_{ijk} is in row $2N(j-1) + 2i + k - 2$. Unlike the `pdetoolbox` setup for c and a there are currently no schemes to encode special situations such as, e.g., $b \otimes \nabla u = \alpha \partial_x u$ which corresponds to advection into direction x. Thus, in this case, for $N = 2$, `b` reads $b = [\alpha; 0; 0; 0;\; 0; 0; \alpha; 0]$, while, e.g., $\beta \partial_y u$ yields $b = [0; \beta; 0; 0;\; 0; 0; 0; \beta]$. Again we remark that advective terms can also be put into f such that `p.f` may simply return `b=0`. However, `b` is needed if there are advective terms and we want explicit Jacobians as explained next.

If `p.jsw`< 3 then the user must also set `p.jac` to a function handle with outputs `cj,aj,glam,bj` again defined on triangle midpoints. These are the coefficients for assembling `Gu` and `Glam`, where in fact `cj` is the same as `c` in `p.f`, compare Table 5.

*There are various special coding schemes for diagonal or symmetric cases, see [35]. For convenience here we just note that in the general case `c` is a $4N^2 \times$ `p.nt` matrix (or in case of constant coefficients a $4N^2$ column vector) with c_{ijkl} in row $4N(j-1) + 4i + 2l + k - 6$, $i, j = 1, \cdots, N$, $k, l = 1, 2$. In this scheme we have `c=[1;0;0;1]` to encode the Laplacian. Similarly, there are special storage schemes for symmetric a, but in general $a = a_{ij}$, $i, j = 1, \cdots, N$, is stored as a $N^2 \times$ `p.nt` matrix (resp. a N^2 column vector in case of constant coefficients) with a_{ij} in row $N(j-1) + i$. Note the somewhat non-lexicographical order.

Providing the same c=cj (and also often the same $\texttt{b}=\texttt{bj}$) twice (in p.f and p.jac) is redundant, and there are situations where only either cj,aj,bj or glam are needed, namely jsw=1 resp. jsw=2. However, we found the small overhead of recomputing cj,bj, resp. unnecessarily computing aj,bj resp. glam acceptable to have a clear code. On the other hand, splitting the calculation of PDE coefficients and Jacobian coefficients into two routines is reasonable since often at least for testing it is convenient to use jsw=3 where analytical Jacobians are never needed.

Remark 3.2. Given coefficients $c,a,f,b = c(u_0),a(u_0),f(u_0),b(u_0)$, the FEM transforms (1.1) into the algebraic system $\texttt{K}(\texttt{u}_0)\texttt{u}-\texttt{F}(\texttt{u}_0) \stackrel{!}{=} 0$ where K is called the stiffness matrix, assembled from c,a and b, and F is the FEM representation of f. Thus, u solves the FEM discretization of (1.1) if the residual $\texttt{r}=\texttt{resi}(\texttt{p},\texttt{u},\texttt{lam}) := \texttt{K}(\texttt{u})\texttt{u}-\texttt{F}(\texttt{u}) \stackrel{!}{=} 0$. The basic pdetoolbox routine to assemble K=$\texttt{K}_0$ (in case $b = 0$) and F is [K,F]=assempde(⋯,c,a,f), where "⋯" stands for boundary conditions and mesh-data. To assemble the advection matrix B we additionally provide B=assemadv(p,t,b). The full system-matrix then is $K = K_0 - B$.

Thus, if $a = b = 0$ and c does not depend on (u,λ), then with the local derivatives $\texttt{aj} = -f_u$ and $\texttt{glam} = -f_\lambda$ returned from p.jac, the Jacobian G_u and the derivative G_λ can be obtained from

$$[\texttt{Gu},\texttt{Glam}] = \texttt{assempde}(\cdots,\texttt{cj},\texttt{aj},\texttt{glam}), \tag{3.4}$$

and this is done for jsw=0. If $a \neq 0$ is independent of u (and still $b = 0$), then p.jac must return

$$\texttt{aj} = a - f_u.$$

If $b = 0$ but, e.g., a depends on u, then the formulas must be adapted accordingly, and if, e.g., $c = c(u)$ or $b = b(u) \neq 0$, then G_u can still be assembled using $\texttt{cj}(\texttt{u}) = \texttt{c}(\texttt{u})$ and suitable bj and aj, see Section 3.4 and Section 4.1 for examples. Similarly, glam must always be understood in a "generalized" sense, i.e., it must assemble to G_λ, even if this involves derivatives of u which originally were implemented via c.

For jsw=1, Gu is still assembled but Glam is calculated by finite differences. Since approximating G_λ by finite differences only takes one additional call of resi, speed is not an issue for choosing between jsw=0 and jsw=1. In the latter case, simply set flam=0 in p.jac. See Section 3.2.1 for an example. For $\texttt{jsw} = 2$ we use numjac to calculate Gu; to be efficient this requires a sparsity structure S, and here we assume that $\texttt{F}_i$ depends only on (all components of) u on the i-th node and all neighboring nodes, which corresponds to the sparsity structure obtained by [Gu,Glam]=assempde(⋯,0,a,0) with $a_{ij} = 1$ for $i,j = 1,\cdots,N$. Our experience is that numerical Jacobians are fast enough for moderate size problems, i.e., for up to a few thousand degrees of freedom. Of course this also depends on the structure of the problem: diagonal diffusion or not, weak or strong coupling of the different components of u. Still, assembling Jacobians is usually much faster. For $\texttt{jsw} = 3$, both Gu and Glam are approximated by finite differences. In any case, for both ($\texttt{jsw} \leq 1$ and $\texttt{jsw} \geq 2$) we assume local dependence of f on u. See, however, Section 3.5 for some modifications for the case of global coupling.

Remark 3.3. The boundary conditions, see Section 3.1.4, are updated from `bc=p.bcf(p,u,lam)` before assembling. In the (frequent) case that the BC do not change during continuation we set may `p.bcf=@(p,u,lam) bc` in the init-routine (after generating `bc`). See, however, Section 3.3 for examples with λ-dependent BC.

As mentioned, when applicable, assembling G_u via `cj,aj` and `bj` (`p.jsw=0,1`) gives a matrix $G_{u,a}$ and is faster (by orders of magnitude for large Nn_p) than numerical differentiation by `numjac` (`p.jsw=2,3`), which gives a matrix $G_{u,n}$ which is in general close to but not equal to $G_{u,a}$. Intuitively we might also expect $G_{u,a}$ to be "more accurate" than $G_{u,n}$. However, we need some caution: in fact, $G_{u,n}$ often gives better convergence of the Newton loop for the algebraic system $r(u) = K(u)u - F(u) \stackrel{!}{=} 0$. The reason is that $G_{u,a}$ involves interpolation of nodal values to triangle values in c, a, b and f_u, while for $G_{u,n}$ this is done on c, a, f, b, consistent with the definition of $r(u)$. This effect becomes prominent on poor (underresolved) meshes, where the relative error $e = \|G_{u,n} - G_{u,a}\|/\|G_{u,n}\|$ can be of order 0.05 or larger. However, $e \to 0$ for mesh spacing $h \to 0$. For convenience we provide the function `[Gua,Gun]=jaccheck(p)` which returns $G_{u,a}, G_{u,n}$, produces spy-plots of these matrices, and prints the timing and some diagnostics.

Thus, for small n_p it might appear that $G_{u,n}$ is favorable. On the other hand, $G_{u,a}$ is obtained in acceptable time on much finer meshes, where the FEM solution u_h should be much closer to a PDE solution u. In fact, using `jsw=2,3` can even be dangerous in the sense that it may mask the fact that a FEM solution u_h is not close to a PDE solution u. See Section 3.1.5.

3.1.4. The geometry and the boundary conditions

The domain Ω is typically described as a polygon. As the `pdetoolbox` syntax is somewhat unhandy, and the rectangular case is quite common we provide the function `geo=rec(lx,ly)` which yields $\Omega = [-l_x, l_x] \times [-l_y, l_y]$. An extension is the function `polygong`, see [27] for its syntax. Setting up "arbitrary complicated" geometries Ω is most convenient if there is a drawing `img.jpg` of Ω in the current directory. Type `im=imread('img.jpg'); figure(1);image(im); [x,y]=ginput;` which yields a crosshair. Click (counterclockwise) on $\partial\Omega$, stop with `return`. The obtained vectors `x,y` can be saved as a `*.mat` or `*.txt` file, and piped through `geo=polygong(x,y)`. Finally, geometries can also be exported from the `pdetoolbox` GUI.

The `pdetoolbox` syntax for the boundary conditions (1.2) is also somewhat unhandy. For scalar equations the most common BC are homogeneous Dirichlet or Neumann BC. The routine `[geo,bc]=recdbc1(lx,ly,qs)` approximates Dirichlet BC over rectangles via "stiff spring" Robin BC of the form $\mathbf{n}\cdot(c\otimes\nabla u)+q_s hu = 0$ with a large $q_s =$ `qs`. For `c` of order 1 typically $q_s = \mathcal{O}(10^2)$ or $q_s = \mathcal{O}(10^3)$ works well. For homogeneous Neumann BC we provide `[geo,bc]=recnbc1(lx,ly)`, with the extension `[geo,bc]=recnbc2(lx,ly)` to 2-component systems.

For the genuine systems case (or the case of non-rectangular domains) we provide the routines `bc=gnbc(pneq,varargin)` and `bc=gnbcs(pneq,varargin)`. For a system

with `neq` components and a domain with `nedges` edges, `bc=gnbc(neq,nedges,q,g)` creates "generalized Neumann BC" (1.2) that are given as numerical data. Different boundary conditions q_j, g_j at the edge with index j are generated by a call of the form `bc=gnbc(neq,q_1,g_1, ···, q_nedges,g_nedges)`. For g, q given in terms of an explicit formula, e.g., involving $x, u, \nabla u$, the function `gnbcs` accepts a string variable encoding of g and q or g_j, q_j and otherwise works in the same way. See the demos below.

3.1.5. Error estimates and mesh adaptation

As an ad hoc way to check whether a FEM solution u_h =`p.u` approximates a PDE solution u we provide `[q,ud]=meshcheck(p,cmp)`. This (adaptively) refines the FEM mesh in p to roughly the double number of triangles and calculates a new FEM solution $u_{h,\text{new}}$ from the old solution $u_{h,\text{old}}$. Then $u_{h,\text{old}}$ is interpolated to $\tilde{u}_{h,\text{new}}$ on the new mesh, $u_{\text{diff}} = u_{h,\text{new}} - \tilde{u}_{h,\text{new}}$ is formed, and $||u_{\text{diff}}||_\infty$ and the relative error $||u_{\text{diff}}||_\infty / ||u_{h,\text{new}}||_\infty$ are printed. The new solution structure q and the difference ud= u_{diff} are returned, and for `cmp`> 0 we additionally generate a plot of the `cmp`th component of u_{diff}. For instance, in Fig. 3(a) we check the mesh at point 20 on the q branch in Fig. 2 which indicates that the mesh in 2 corners is clearly too poor, and before continuing for smaller λ we should refine the mesh.

Thus, some (automatic) mesh adaptation may be vital for reliable continuation. The `pdetoolbox` comes with mesh refinement based on an a posteriori error estimator as follows. For the scalar Poisson problem $-\Delta u = f$, $u|_{\partial\Omega} = 0$, let u_h be the FEM solution and u the PDE solution. Then, with $\alpha, \beta > 0$ some constants independent of the mesh,

$$||\nabla(u - u_h)||_{L^2} \leq \alpha||hf|| + \beta D_h(u_h), \tag{3.5}$$

where $h = h(x)$ is the local mesh size, and $D_h(v) = \left(\sum_{\tau \in E_i} h_\tau^2 (\partial_{n_\tau} v)^2\right)^{1/2}$. Here $\partial_{n_\tau} v$ is the jump in normal derivative of v over the edge τ, h_τ the length of the edge, and E_i the set of all interior edges. For equations $-\nabla(c \otimes \nabla u) + au = f$ this suggests the error indicator function

$$E(K) = \alpha||h(f - au)||_K + \beta\left(\frac{1}{2}\sum_{\tau \in \partial K} h_\tau^2 (n_\tau \cdot c\nabla u_h)^2\right)^{1/2} \tag{3.6}$$

for each triangle, which is calculated by the `Matlab` routine `pdejmps`. For convenience we provide the interface routine `err=errcheck(p)`.* Calling, e.g., `err=errcheck(p)` yields `err=0.273` which somewhat overestimates the error in Fig. 3(a). In `cont`, for `errchecksw>0` we call `errcheck` after each successful step and store the result in `p.err`.

The (basic) mesh refinement strategy is to introduce new triangles where $E(K)$ is large.† This is done by the `pdetoolbox` routine `refinemesh`, but we provide the interface

* With `a,f,b` returned from `p.f`, this also re-calculates f (via `fb=bgradu2f(p,f,b,u)`) to include $b \otimes \nabla u$ into `fb`, as `pdejmps` does not take $b \otimes \nabla u$ into account directly. Similarly, the mesh refinement below always recalculates f to include $b \otimes \nabla u$. Moreover `errcheck` also contains our settings for some tunable parameters of `pdejmps`. Finally, at the end of most `*cmds.m` files there are templates how to (re)evaluate $E(K)$ at any given solution.

†It is not a priori clear if this is also suitable for systems, but it works well for all our examples.

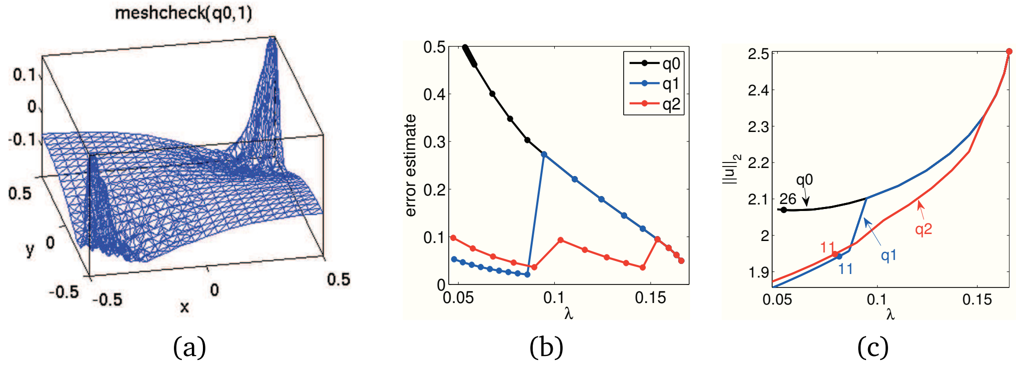

(a) (b) (c)

Figure 3: (a) Error in u at q0=point 20 on the q branch from Fig. 2 obtained from calling meshcheck(q0,1). (b) p.err over λ for continuing from point 10 on the q branch, without mesh-adaptation (labeled q0), with strategy (i) (amod=10, labeled q1) and with strategy (ii) (p.errbound=0.1, labeled q2). (c) The bifurcation diagrams belonging to (b). In q1 a rather large jump appears from refinement after the 10^{th} step.

routine p=meshref(p, varargin).* Since meshref also interpolates the tangent τ to the new mesh, we can continue with cont immediately after mesh refinement. However, instead of mesh refinement, which means introduction of new points into the mesh, we rather need mesh adaptation, which means refinement where necessary, but coarsening where possible, to limit the problem size. In pde2path, mesh-adaptation is implemented in an ad hoc way in the function p=meshadac(p,varargin) by first interpolating a given solution to a (typically somewhat coarse) "base-mesh" or "background-mesh" and then refining. (Base-mesh given by p.bpoints, p.bedges, p.btria, varargin as above).

During continuation runs there are basically two strategies for mesh-adaptation, which can also be mixed:

(i) call meshadac every p.amod$^{\text{th}}$ step, for p.amod>0, or

(ii) call meshadac whenever p.err>p.errbound (choose p.errchecksw>1 for this), where we refine until err<p.errbound/2 in order to allow some margin for the next steps.

See Figs. 3(b), (c) for a comparison of the different approaches.

In summary, mesh adaptation strategies and error bounds are highly problem dependent, and moreover, may not be rigorously justified for the system case or general BC. Thus, for a given problem we recommend to first experiment with mesh adaptation to, e.g., achieve a given error bound ε, i.e., call p_refined=meshref(p,'eb',ε). Mesh adaptation works well for all examples we considered, and is essential in some, in particular in Section 4.2 and Section 5.1. See also Remark 3.5 for another simple scalar example with details on mesh sizes and error estimates.

*where varargin takes pairs 'maxt',maxt=number of triangles aimed at, or 'ngen', ngen=number of refinement steps, or 'eb',eb=error bound. Calling for instance q0r=meshref(q0, 'eb', 0.0025, 'maxt', 50000, 'ngen',20) shows that to achieve an estimated error$\leq$ 0.0025 we need about 30.000 triangles.

Remark 3.4. For bifurcation from trivial branches, another good strategy is to prepare a *finer* base mesh than the starting mesh, for instance if the trivial branch consists of spatially homogeneous solutions, but the bifurcating solutions develop sharp gradients. For convenience we also provide the functions `p=newmesh(p)`, which interpolates the current (u,τ) to a new mesh generated after user input in the form `hmax` or `nx,ny`, and `p=setbmesh(p)` which sets the base mesh to the current mesh. This should be called if it is expected that the current mesh is a good base for adaptation in the next steps.

3.1.6. Eigenvalues and `findbif`

To obtain stability information of u as a stationary solution of $\partial_t u = -G(u,\lambda)$ we provide the function `[ineg,muv]=spcalc(Gu,p)` which by default uses `eigs.m` to calculate the set Σ_0 of `p.neig` smallest eigenvalues of G_u, and returns the number `ineg` of eigenvalues with negative real-part and the (vector of) eigenvalues `muv`. The implicit assumption, based on ellipticity of G_u, is that `p.neig` is sufficiently large such that all eigenvalues with negative real part are always contained in Σ_0. *

A similar method can also be used to calculate sign(det$A$) by calculating `p.neigdet` eigenvalues η_i of A closest to 0, `p.deig` > 0, and then use

$$\text{sign}(\det A)=\text{sign}\left(\sum_{i=1}^{d_{\text{eig}}}\text{Re}\eta_i\right). \tag{3.7}$$

This, however is usually much slower (though it also gives more information) than using the LU decomposition of A, which we use as standard by setting `p.neigdet=0`.

For problems with symmetries, often multiple eigenvalues of A and G_u cross the imaginary axis simultaneously. Even if such a symmetry is artificially broken, for instance by a small "detuning" of the domain, often still a rather large number of eigenvalues of A and G_u crosses the imaginary axis on rather short parts of branches. To find the associated bifurcation points would then require unpractically small `p.ds`. Therefore we also provide a routine `findbif` which scans a branch for changes of the signature of G_u, i.e., the number `ineg` of unstable eigenvalues of G_u, to detect bifurcations, where again `ineg` is based on calculating `p.neig` eigenvalues as above. This in particular can be used to detect if an even number of eigenvalues crosses the imaginary axis in a given step, and then start a bisection. In detail, `p=findbif(p)` continues a branch until the signature changes and then tries to detect the successive points between which it changes by ± 1. `p=findbif(p,m)` with $m \geq 1$ tries to detect m such index changes. Besides the bifurcation points, the first point, each point just ahead of and behind a bifurcation, and the last point on the considered branch are saved. Thus `findbif` can be considered as a version of `cont` with bifurcation detection based on the eigenvalues of G_u. This can also be extended

*Use `p.eigref='sa'` for alternatively calculating eigenvalues with smallest real parts (typically somewhat slower, even if used with small `p.neig`). Our standard setting is $n_{\text{eig}} = 50$, but of course this is highly problem dependent and should be adapted by the user when needed. We give a warning if $|\mu_1| > |\mu_{\text{neig}}|/2$ since then eigenvalues might wander out of Σ_{G_u} to the left in the next steps, which might give false stability information.

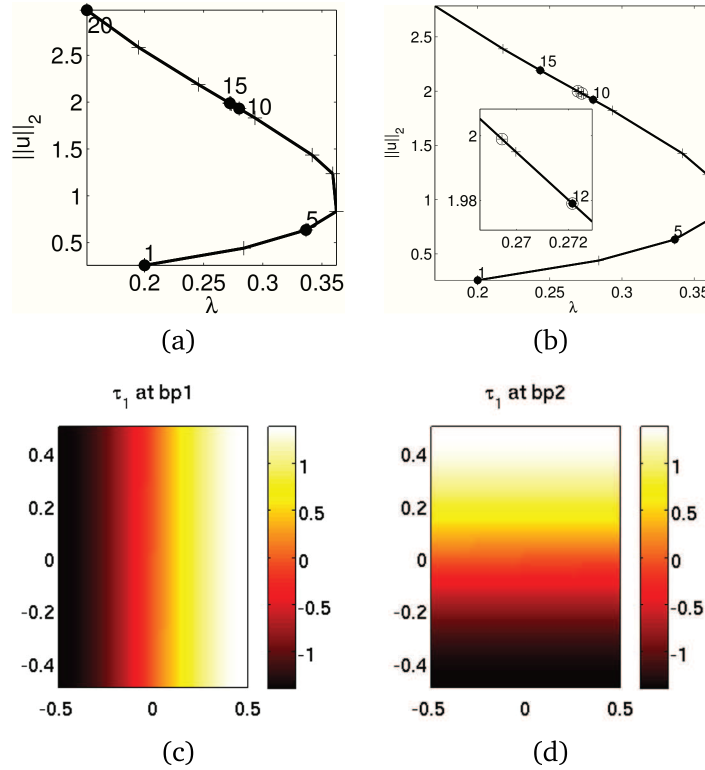

Figure 4: Finding bifurcation points with `findbif`. (a) On a square domain $\Omega = (-0.5, 0.5)^2$, the double bifurcation point at $\lambda \approx 0.2724$ is initially detected by `findbif` but not localized. (b) Breaking the square symmetry by setting $\Omega = (-0.5, 0.5) \times (-0.495, 0.495)$ this bifurcation unfolds and `findbif` conveniently localizes the resulting two simple bifurcation points, although quite close. Associated eigenvectors in (c), (d).

to detect and localize bifurcations with a prescribed signature change $\delta = \pm 2, \pm 3, \cdots$, i.e., double/triple/$\cdots$ bifurcation points. However, as already said, presently we stick to simple bifurcation points.

In principle, a similar algorithm could be set up based on the eigenvalues of A, but as the signature of G_u also immediately gives the stability information, here we opted for G_u. In any case, the final localization of the bifurcation points presently again uses bisection based on sign changes of det$A$, without calculation of eigenvalues. In future versions we plan to improve this to some regula falsi or Anderson-Björk-King method based on the actual eigenvalue crossing, and also to try some more sophisticated methods such as continuation of invariant subspaces (CIS) [3, 4].

See Fig.4 for examples of using `findbif` from the end of `bratucmds.m`. In some of the examples below we use `findbif` to locate bifurcations from "the main trivial" branches, on which the solutions do not change (or at least not significantly); since the stepsize adaptation in `findbif` is based on the signature of G_u, in this case `findbif` can be run with rather large `p.ds` and thus the bifurcations can be found quickly.

3.1.7. The linear system solvers

Recall that after discretization with n_p points we have nodal values $u \in \mathbb{R}^p$ with $p = Nn_p$ large, and

$$G_u \in \mathbb{R}^{p\times p} \quad \text{and} \quad A = \begin{pmatrix} G_u & G_\lambda \\ \xi\dot{u} & (1-\xi)\dot{\lambda} \end{pmatrix} \in \mathbb{R}^{(p+1)\times(p+1)} \tag{3.8}$$

are large, but sparse (block) matrices. The question is how to best solve $G_u v = r$ and the bordered systems such as $A\tau = z$, respectively. In all the examples that we considered, our experience is that the highly optimized matlab solver `z=A\b` of $Az = b$ works remarkably well, but for easy customization of the code we never call \ directly but use two interface routines:

1. `v=p.lss(M,r,p,lam)` to solve $Mv = r$ with $M = G_u \in \mathbb{R}^{p\times p}$;
2. `z=p.blss(A,b,p,lam)` to solve $Az = b$ with $A \in \mathbb{R}^{p+1\times p+1}$.

Here `blss` and `lss` stand for `(bordered)linear system solver`.

The default solvers `lss` and `blss` just contain one command, namely `v=M\r` resp. `z=A\b`. Nevertheless, for large systems or for some special classes of problems iterative solvers might work better, and as templates we provide the two routines `ilss` and `iblss`, using `gmres` with (incomplete `ilu`) LU factorization as preconditioners. These should, of course, be reused as long as gmres converges quickly, and here (and in `resinj.m`) we thus introduce some *global* variables, namely `global L U;` resp. `global bL bU`. Thus, when using, e.g., `ilss` the user must also issue `global L U; L=[]; U=[];` from the command line. The reason for this construction is that we do not want to make L, U a part of `p` since this needs a lot of disk space when saving `p`: typically, we get fill-in factors for L, U of 10 and larger.

It turns out that for scalar problems these outperform the direct solvers `lss` and `blss` for large n_p, $n_p > 10^5$, say. On the other hand, for systems, \ beats `gmres` with LU preconditioning even for very large n_p. In summary, the iterative solvers `ilss` and `iblss` should only be regarded as template files to create problem specific iterative solvers when needed. See also Section 3.5 for adaptations of `lss` and `blss` to some special situation.

Finally, various approaches have been proposed for the solutions of the bordered systems $Az = b$, see, e.g., [15, 19]. As alternatives to `blss` we provide `bellss` ("bordered elimination") and `belpolss` ("bordered elimination plus one"). To use these, simply set, e.g., `p.blss=@belpolss`. In our tests the performances of `bellss` and `belpolss` are roughly the same as `blss`.

3.1.8. Screen output, plotting, convergence failure, auxiliary functions

The screen output during runs is controlled by the two functions `p.headfu` (headline) and the function `p.ufu`. These are preset in `stanparam` as `p.headfu=@stanheadfu, p.ufu=@stanufu` to first print a headline and then, after each step, some useful information. To print some other information the user should adapt `stanheadfu` and `stanufu`

to a local copy, say `myhead.m`, and set `p.headfu=@myhead`, and similar for `stanufu` and `p.ufu`. The bifurcation diagram and solution plots are also generated during continuation runs, but in general it is more convenient to postprocess via `plotbra, plotsolf` etc.

The files `p*.mat` and `bp*.mat` contain the complete data of the respective point on a branch, including the mesh, which is necessary if, e.g., some mesh refinement occurred during continuation. To save disk space, however, we deliberately chose to *not* make G_u a part of p. Thus, a run which is no longer in memory can be simply reloaded by, e.g., `q=loadp(pre,pname,'q')`, where `pre, pname` is the name data of a previously saved point, and the third argument is used to set the directory name for the newly created struct. The loaded point will often be either the last one or the first; in the latter case, to change direction of the branch, use, e.g., `q=loadp('p','p1','q'); q.ds=-q.ds; q=cont(q);`

If the Newton-loop does not converge even after reducing `ds` to `dsmin` then `cfail.m` is called. The standard option is to simply abort `cont`, but we offer a number of alternatives, e.g., to change some parameters like `dsmin` or `imax`, or to try, e.g., some mesh refinement or adaptation. Clearly, the choice here is strongly problem dependent, and thus we recommend to adapt `cfail.m` if needed; see Section 3.6 for remarks on such "customization without function handles".

Besides those already mentioned we provide further auxiliary functions, see [26, m2html] for a complete documented list.

3.2. The Allen-Cahn equation with Dirichlet boundary conditions (ac)

In our second example we use Dirichlet boundary conditions (DBC), and explain some ad hoc parameter switching, and time integration. We consider a cubic-quintic (to have folds) Allen-Cahn equation

$$-\mu\Delta u-\lambda u-u^3+u^5=0 \quad \text{on } \Omega=[-L_x,L_x]\times[-L_y,L_y], \quad u|_{\partial\Omega}=0, \tag{3.9}$$

with two parameters $\mu>0$ and $\lambda\in\mathbb{R}$. We use `[p.geo,bc]=recdbc1(lx,ly,1e3)` to approximate the DBC, and set $L_x=1$ and $L_y=0.9$ to break the square symmetry present in `bratu` in order to have only simple bifurcations, namely at $\lambda_{kl}=\mu\pi^2((k/L_x)^2+(l/L_y)^2)$. First we fix $\mu=0.25$ which yields $\lambda_{11}=1.3784$, $\lambda_{21}=3.2289$, $\lambda_{12}=3.6630,\cdots$, and continue in λ, which yields Figs. 5(a)-(c). After branch switching we turn on mesh-adaptation after each 5 steps, see Remark 3.5, and see `acdemo.m` or `accmds.m` for more details, which also contain an example of perturbing a solution and subsequent time integration.

3.2.1. Parameter switching

Unlike AUTO, `pde2path` (currently) has no switches or presets for multi-parameter continuation. However, switching to a new parameter for continuation can be achieved in a simple and flexible way by modifying the structure p from the command line. As an example we want to continue in μ from point 10 on q to $\mu=0.1$. This is achieved by the commands in Table 6, and yields Fig. 5(d). The basic idea is to copy q to w (this is

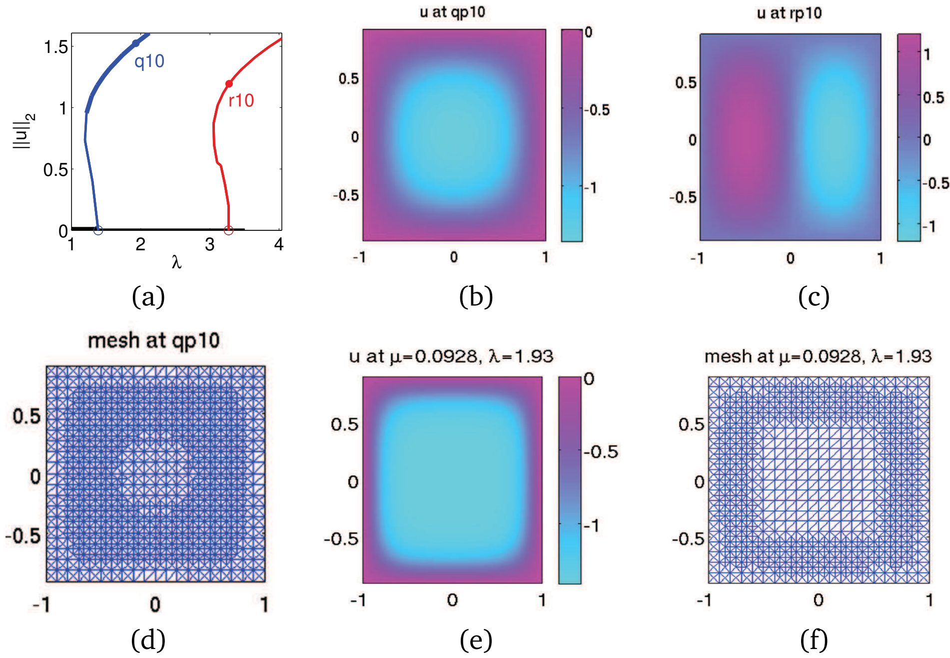

Figure 5: (a) Elementary bifurcation diagram for (3.9) with $\mu = 0.25$. No secondary bifurcations occur, and the mode-structure on each branch is completely determined at bifurcation. (b), (c), (d) points resp. mesh on branches as indicated. (e), (f) Solution and mesh after continuation in μ from (b) to $\mu \approx 0.093$.

not strictly necessary) and then reset `w.f` and `w.jac`. We make our life simple by setting `w.jsw=1` such that we do not need G_λ, and set $\xi = 10^{-6}$ since the dependence on μ is quite sensitive. We introduce a new parameter `w.up1` ('user parameter 1', but any name will be fine) which is used to pass the current λ to `acfmu`, see Table 6.

Table 6: Switching to continuation in μ, commands, and modified coefficient and Jacobian functions.

```
w=q;w.up1=w.lam;w.lam=0.25;w.lammin=0.05;w=setfn(w); w.ds=-0.01;
w.f=@acfmu;w.f=@acjacmu;w.jsw=1;w.parasw=0;w.xi=1e-6;w.restart=1;w=cont(w);
```

```
function [c,a,f,b]=acfmu(p,u,lam)  % AC for cont in mu
u=pdeintrp(p.points,p.tria,u); c=lam; a=0; f=p.up1*u+u.^3-u.^5; b=0;
```

```
function [cj,aj,glam,bj]=acjacmu(p,u,lam) % Jac for cont in mu; use jsw=1
u=pdeintrp(p.points,p.tria,u);cj=lam;glam=0;bj=0;aj=-p.up1-3*u.^3+5*u.^4;
```

Remark 3.5. The base mesh in Fig. 5 consists of 800 triangles, and the mesh adaptation produces the following numbers: if we continue q without adaptation, then from (3.6) we obtain an error estimate $E(K) \approx 0.0047$ at q10, while with adaptation to the mesh in Fig. 5(d) with $\mathtt{nt} = 2484$ triangles we obtain $E(K) \approx 0.0013$. Similarly, if we continue with this mesh and no further adaptation to $\mu = 0.1$ we get $E(K) \approx 0.0016$, while the

error estimate with adaptation to the mesh in Fig. 5(f) with `nt` $= 2112$ is $E(K) \approx 0.0001$. Together with the quite reasonable meshing obtained, we take this as a further illustration of efficiency of the mesh adaptation.

3.2.2. Time integration

For time integration of (1.3) using the struct `p` we provide a simple semi-implicit Euler method. Writing $u^{(n)}$ for $u(t_n, \cdot)$, choosing a time-step h, approximating

$$\partial_t u(t_n) \approx \frac{1}{h}(u^{(n+1)} - u^{(n)}),$$

where $t_n = t_0 + nh$, and evaluating, e.g., $\nabla \cdot (c \otimes \nabla u)$ as $\nabla \cdot (c(u^{(n)}) \otimes \nabla u^{(n+1)})$ we obtain, on the FEM level,

$$\begin{aligned} &\frac{1}{h} M\left(u^{(n+1)} - u^{(n)}\right) = -K(u^{(n)}) u^{(n+1)} + F(u^{(n)}) \\ \Leftrightarrow\ & u^{(n+1)} = \left(M + hK(u^{(n)})\right)^{-1}\left(Mu^{(n)} + hF(u^{(n)})\right). \end{aligned}$$

Here M is the mass matrix and K is the stiffness matrix on time-slice n. This is implemented in `tint(p,h,nstep,pmod)`, where `nstep` is the number of time steps, and a plot (of component `p.pcmp`) is generated each `pmod`'th step. Thus we may call, e.g., `p.u=p.u+0.1*rand(p.neq*p.np,1); p=tint(p,0.1,50,4,10)` to first perturb a given solution and then time-step. See `accmds.m` for an example, where we perturb a solution on the unstable part of the `q` branch into both directions of the unstable manifold; in the subsequent time integration the solution converges to the stable trivial solution or the stable `q` branch, respectively, as expected. However, the main purpose of `tint` is to generate (stable) initial data for continuation, i.e., after `tint` call `cont`. See also Section 5.2 for an example where `tint` is used in this spirit.

Remark 3.6. In each step in `tint` we assemble $K(u^{(n)})$, and call `p.lss` to solve $(M + hK(u^{(n)}))u^{(n+1)} = g^{(n)}$. Clearly, for special cases this can be optimized: for instance, if c, a, b do not depend on u, then the textbook approach would be to assemble K at the start, followed by some incomplete LU-decomposition of $M + hK$ combined with some iterative solver. However, similar remarks as in Section 3.1.7 apply, and thus we use the very elementary form above, but stress again that `tint` in its present form is not intended for heavy time-integration.

3.3. The Allen-Cahn equation with mixed λ-dependent boundary conditions (`achex`)

We illustrate a few more possibilities with `pde2path` by modifying the Allen-Cahn example (3.9) from above. We consider again (3.9), i.e., $-0.25\Delta u - \lambda u - u^3 + u^5 = 0$, but instead of homogeneous Dirichlet BC on a rectangle, we consider hexagonal domains Ω and parameter dependent mixed Dirichlet-Neumann BC. Fig. 6(a) shows an example

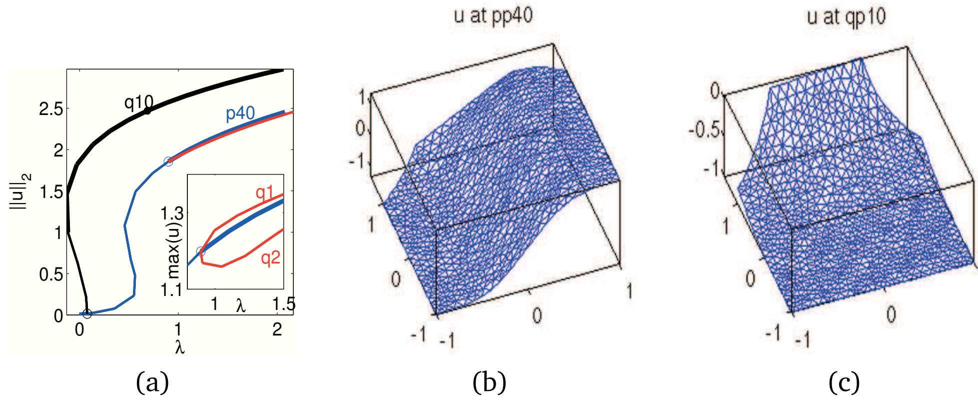

(a) (b) (c)

Figure 6: (a) Bifurcation diagram for (3.9) with BC given by (3.10). (b), (c) selected solution plots. Basic mesh created by `p=stanmesh(p,0.1)` with 1602 triangles and $E(K) \approx 0.007$ at, e.g., q10.

for Ω, which basically consists of a square, with the top boundary shifted by $\delta_y = 0.5$ between $[-l_x, l_x]$, $l_x = 0.5$. Denote this part of $\partial\Omega$ by Γ_D, and set $\Gamma_N = \partial\Omega \setminus \Gamma_D$. To define the domain we could, e.g., use the `pdetool` GUI to draw a polygon composed of six edges, one for each segment and export the geometry. However, usually the function `polygong` [27] is much more convenient. See `geo=hexgeo(lx,dely)`, which also contains a slightly edited output of the GUI for comparison. Second, we want to define the boundary conditions

$$n \cdot \nabla u = 0 \quad \text{on } \Gamma_N, \quad u = \lambda x \quad \text{on } \Gamma_D. \tag{3.10}$$

To implement this we use a stiff spring approximation on Γ_D in via `gnbcs`, i.e.,

```
qd=mat2str(10^4);gd=[mat2str(10^4*lam) '*x']; qn='0'; gn='0';
bc=gnbcs(1,qn,gn,qn,gn,qn,gn,qn,gn,qd,gd,qn,gn).
```
(3.11)

With `pde2path` we perform a continuation starting from the trivial zero solution and obtain the bifurcation diagram plotted in Fig. 6; see `ac6cmds`. Bifurcation detection and branch switching work without problems, and the error estimate $E(K)$ is always well below 0.01. To generate both parts of the `r` branch we first call `r1=swibra('p','bp2','r1',-0.1);r1=cont(r1)` and then `r2=loadp('r1','p1','r2');r2.ds=-r1.ds;r2=cont(r2)` to proceed in the other direction. At the end of `ac6cmds` we also run an example with $u = \lambda$ on Γ_D implemented via `gnbc`.

3.4. A quasilinear Allen-Cahn equation (`acql`)

To give an example of a more complicated Jacobian we modify (3.9) to the quasilinear Allen-Cahn equation

$$\begin{cases} -\nabla \cdot [(0.25 + \delta u + \gamma u^2)\nabla u] - f(u,\lambda) = 0 \quad \text{on } \Omega = [-L_x, L_x] \times [-L_y, L_y], \\ u|_{\partial\Omega} = 0, \end{cases} \tag{3.12}$$

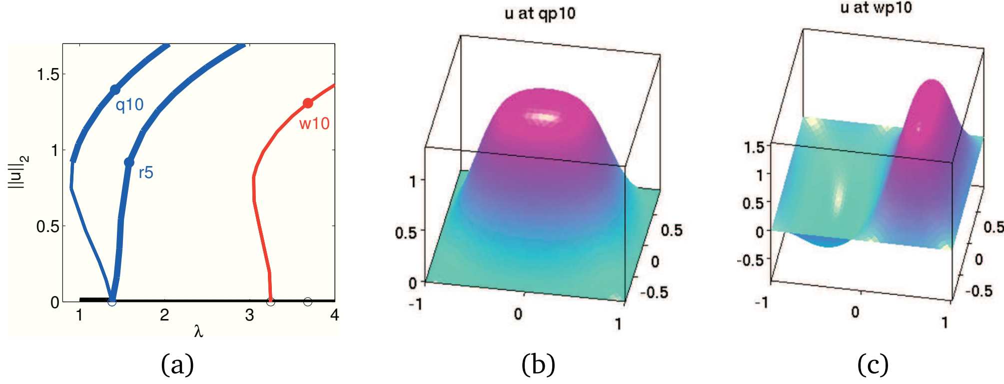

(a) (b) (c)

Figure 7: Elementary bifurcation diagram for (3.12) with $\delta = -0.2$ and $\gamma = 0.05$, and some solution plots. The two blue branches are in fact one branch, and the corner at the transcritical bifurcation from the trivial branch is due to the choice of vertical axis. The symmetry $u \mapsto -u$ no longer holds, and "up" humps are steeper than "down" humps due to $\delta < 0$. The w-branch is still double due to the $x \mapsto -x$ symmetry.

with $f(u,\lambda) = \lambda u + u^3 - u^5$ and $L_x = 1$, $L_y = 0.9$ as before. See `acqlf.m`. The linearization around u gives the linear operator

$$G_u(u,\lambda)v = -\nabla\cdot[(0.25+\delta u+\gamma u^2)\nabla v] + [-f_u(u,\lambda) - \delta\Delta u - 2\gamma(\nabla u\cdot\nabla u + u\Delta u)]v \\ - [(\delta+2\gamma u)\nabla u]\cdot\nabla v.$$

Hence, in `acqljac.m` we now have `fu` $= f_u + \delta\Delta u + 2\gamma(\nabla u\cdot\nabla u + u\Delta u)$, and $\mathtt{b}_{111} = (\delta + 2\gamma u)u_x$ and $\mathtt{b}_{112} = (\delta + 2\gamma u)u_y$, cf. Remark 3.2. To generate (u_x, u_y) and Δu as coefficients in `acqljac.m` we use `pdegrad` resp. `pdegrad, pdeprtni` and `pdegrad` again, see [35].

The term δu in c changes the $u \mapsto -u$ symmetry of the Allen-Cahn equation (3.9). The bifurcation points from the trivial branch $u \equiv 0$ in (3.12) are as in (3.9), but the bifurcations change, see Fig. 7. In particular the first bifurcation changes from pitchfork to transcritical.

3.5. An Allen-Cahn equation with global coupling (`acgc`)

As an example of a "non-standard" elliptic equation we treat an Allen-Cahn equation with a global coupling. We fix $\mu = 0.1$ and $\lambda = 1$ in (3.9), introduce a new parameter (again called λ) and consider

$$G(u,\lambda) := -0.1\Delta u - u - u^3 + u^5 - \lambda\langle u\rangle = 0 \quad \text{on } \Omega = [-\pi/2, \pi/2]^2, \quad u|_{\partial\Omega} = 0, \quad (3.13)$$

where $\langle u\rangle = \int_\Omega u\,dx$. The term $\lambda\langle u\rangle$ is called a global coupling or global feedback, positive for $\lambda > 0$ resp. negative for $\lambda < 0$. Problems with global coupling occur, e.g., in surface catalysis, where global coupling arises through the gas phase [28], in semi-conductors and gas-discharges [34,37], and as "shadow systems" in pattern formation when there is a very fast inhibitor diffusion [18].

Table 7: Sherman-Morrison in customized gclss.m; see also acgcf.m and gcblss.m.

```
function x=gclss(A,b,p,lam) % lss for AC with GC, Sherman-Morrison
global eta nu; y=A\b; z=A\nu; al0=lam*eta*z; al=lam*eta*y/(1-al0); x=y+al*z;
```

The global feedback does not fit into the framework of (1.1) if f is assumed to be local. For the definition of $G(u)$ this is not yet a problem as we may simply define f as, e.g., `f=u+u.^3-u.^5+lam*triint(u,p.points,p.tria)` where `triint(g,points,tria)` is the Riemann sum of $\int g(x)\,\mathrm{d}x$ over the given mesh. However, for continuation we make extensive use of Jacobians, and $G_u(u)$ is now given by

$$[G_u(u)v](x) = -0.1\Delta v(x) - (1+3u(x)^2 - 5u^4(x))v(x) - \lambda \langle v \rangle .$$

As yet we cannot deal with last term, cf. Remark 3.2. The first try would be to simply ignore it in continuation, but this in general only works for small $|\lambda|$ while for larger $|\lambda|$ we loose convergence in the (false) Newton loop. We can express $\langle v \rangle$ on the FEM level via a matrix M such that $G_u(u)v = (K - \lambda M)v$. Essentially, for "natural parametrization" we need to solve

$$G_u(u)v = r, \quad \text{where } \ G_u(u) = (K - \lambda \nu \eta^T) \quad \text{with } \ \nu, \eta \in \mathbb{R}^{n_p}. \tag{3.14}$$

Here $\eta = (aT)^T$, where $(a_1, \cdots, a_{n_t})$ contains the triangle areas, $T \in \mathbb{R}^{n_t \times n_p}$ interpolates $u \in \mathbb{R}^{n_p}$ from nodal values to triangle values (hence $\langle u \rangle$ =`triint(g,points,tria)` =`eta*u`), and $\nu_i = \int_\Omega 1\phi_i\,\mathrm{d}x$ corresponds to adding $\langle v \rangle$ to all nodes with the correct weight. However, $(K - \lambda \nu \eta^T)$ is a full matrix and should never even be formed. Instead we customize `lss` to use a Sherman-Morrison formula which gives (for (3.14))

$$v = K^{-1}r + \alpha(K^{-1}\nu)(\eta^T K^{-1})r, \quad \alpha = \frac{\lambda}{1 - \lambda \eta^T K^{-1} \nu}. \tag{3.15}$$

In `acgcjac.m` we then just ignore the term $\lambda \langle u \rangle$. Similar remarks apply to the bordered systems solved by `blss`.

In the actual implementation we introduce *global variables* `nu,eta`. The idea is that it is sufficient to calculate ν, η once for a given mesh, for instance in `acgcf.m`, as this is always called before the Jacobian `acgcjac` or the linear system solvers `gclss` or `gcblss`. If we set aside mesh-refinement then we could calculate ν, η at startup and store them e.g., as `p.nu,p.eta` but with mesh refinement global variables are more convenient. See Table 7 for the full code of `gclss.m`, and Fig. 8 for the result of the basic continuation runs contained in `acgccmds.m`. We switch off spectral calculations and bifurcation checks by setting `spcalcsw=0; bifchecksw=0;` since out of the box these would be based on the (wrong) local Jacobian, and the two branches were generated by using two different starting points.

3.6. First summary, and some remarks on customization

We end this introductory section based on scalar examples with a first summary and some implementation remarks.

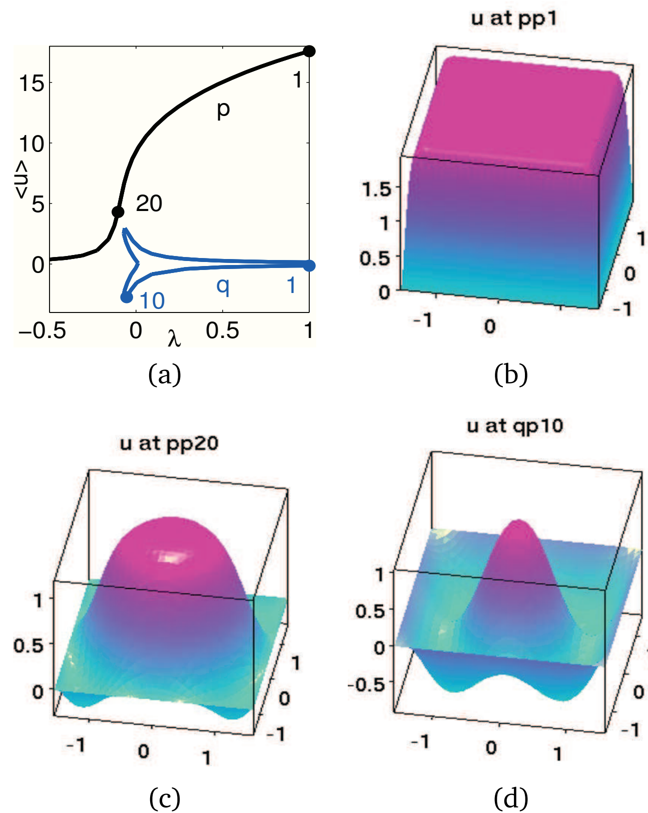

Figure 8: (a) Two solution branches for (3.13), and three selected solutions. By positive global feedback, the plateau in (b) (u around 1.93) is substantially above the zero $(1+\sqrt{5})/2 \approx 1.62$ of $f(u) = u+u^3-u^5$. Here, some mesh refinement near the boundary is also crucial; e.g., at pp1 we have an $E(K) \approx 0.063$ with the base mesh of nt $= 1800$, and $E(K) = 0.009$ after refinement to nt $= 3768$. Decreasing λ to slightly negative values u gets pushed below 0 near the boundary (c). Note that (3.13) is symmetric w.r.t. $(u,\lambda) \mapsto (-u,\lambda)$.

The `p.f=@..` syntax has the advantage that multiple versions of `f` can be maintained and switching can be done by only changing `p.f=@..`. On the other hand, we do not want to overwhelm the user with such options, and thus we restricted the "user-definable" functions to `p.f`,$\cdots$,`p.headfu` from Table 2, where in fact in most cases the user only needs to set up `p.f`, and `p.jac` for `p.jsw`$\leq$ 2. Nevertheless, as outlined above any function of `pde2path` can be customized for a given problem by just copying it from `../p2plib/` to the current directory (where `Matlab` searches first) and then modifying it. Main candidates for customization are, e.g., `plotbra.m`, `plotsol.m` if additional features/options are desired in plotting the bifurcation diagram or/and the solutions. See, e.g., Section 5.1 and Section 5.2.

Most functions of `pde2path` only require a few input/output arguments. An important exception is `plotbra(p,wnr,cmp,varargin)` where `varargin` is a possibly long list of argument/value pairs. See `plotbra.m` for a detailed description, and also the various `plotbra` example calls in the demos.

Table 8: Typical software usage, including pseudo-code of p=cont(p), with main function calls.

Initialization. Declare (user defined) global variables (if any). Initialize structure p, typically by first calling p=stanparam(p), followed by problem dependent calls to define (function handles for the) PDE coefficients, BC and Jacobian, and the geometry, mesh, and starting point.

function p=cont(p)

1. If restart=1 then inistep: generate first two points on branch and (secant) τ_0.
2. Predictor $(u^1,\lambda^1)=(u_0,\lambda_0)+\mathrm{ds}\tau_0$ with stepsize ds.
3. Corrector: depending on parasw and $\dot\lambda_0$ use nlooppde for (2.8) or nloopext for (2.3) or (2.6). This uses getder, getGu resp. getGlam to obtain derivatives, and lss resp. blss as linear systems solver.
4. Call sscontrol to assess convergence (res,iter returned from nlooppde resp. nloopext): If res≤p.tol accept step, i.e., goto 5, (and increase ds if iter<imax/2). If res>p.tol and ds> dsmin then decrease ds and goto 2. If res>p.tol and ds=dsmin then no convergence, hence call cfail.
5. Postprocessing: calculate new tangent τ_1 by (2.5), call spcalc (if spcalcsw=1), bifdetec (if bifchecksw=1). Check for error and mesh adaptation. Update p, i.e., put $u_0=u_1$, $\lambda_0=\lambda_1$, $\tau_0=\tau_1$ into p, call out=outfu(p,u,lam), plot and save to disk. Call p.ufu for printout and further user-defined actions.
6. If stopping criteria met (p.ufu returned 1 or stepcounter>nsteps) then stop, else next step, i.e., goto 2.

Post-processing. Plot bifurcation diagrams via plotbra (plotbraf) and solutions via plotsol (plotsolf). If bifurcations have been found, use swibra (and cont to follow some of these).

As mentioned, by default there are no global variables in pde2path, with the exceptions pj,lamj which are set for numerical differentiation in resinj, and possibly LU preconditioners for iterative linear system solvers, see Section 3.1.7. On the other hand, e.g., p.f, p.jac, p.lss etc. do not return p. This is to have a somewhat clean distinction between functions for specific calculations and function like p=cont(p), p=swibra(···), p=meshref(p) which modify the structure p, including the mesh. As a result, the user might want to introduce some global variables to streamline calculations, see Section 3.5 for an example. These should then be declared before initialization of p.

For convenience, in Table 8 we summarize the typical steps in the usage of the software.

4. Two prototype Reaction-Diffusion Systems

Pattern-formation in Reaction-Diffusion Systems (RDS), in particular from mathematical biology [25], is one of the main applications of path-following and bifurcation software. Here we first consider a quasilinear two-component system with "cross diffusion" from chemotaxis [23] to explain the setup of c in this rather general case, and the setup of general domains in pde2path. We essentially recover the bifurcation diagrams from [23] without special tricks or customization.

Our second example is the Schnakenberg model [32], which is semilinear with a diagonal constant diffusion matrix, and thus in principle simpler than the first example. However, here we are interested in a more complete bifurcation picture, and the Schnakenberg model shows many bifurcations already on small domains. Therefore we need some adaptations of the basic `cont` algorithm to `pmcont` (parallel `m`ulti `cont`inuation), and we introduce `findbif` to locate some first bifurcations from the homogeneous branch.

4.1. Chemotaxis

An interesting system from chemotaxis has been analyzed in [23], including some numerical path-following and bifurcations using ENTWIFE. The (stationary) problem reads

$$0 = G(u,\lambda) := -\begin{pmatrix} D\Delta u_1 - \lambda\nabla\cdot(u_1\nabla u_2) \\ \Delta u_2 \end{pmatrix} - \begin{pmatrix} ru_1(1-u_1) \\ \dfrac{u_1}{1+u_1} - u_2 \end{pmatrix}, \tag{4.1}$$

where $\lambda \in \mathbb{R}$ is called the chemotaxis coefficient and $D > 0$ and $r \in \mathbb{R}$ are additional parameters. In (4.1) we have $b \equiv 0$, may set $a = 0$, and identify the second term with $f(u)$. The linearization reads

$$G_u(u,\lambda)\begin{pmatrix} v_1 \\ v_2 \end{pmatrix} = \left[-\begin{pmatrix} D\Delta & -\lambda\nabla\cdot(u_1\nabla\cdot) \\ 0 & D\Delta \end{pmatrix} + \begin{pmatrix} r(2u_1-1)+\lambda\Delta u_2 & 0 \\ -(1+u_1)^{-2} & 1 \end{pmatrix}\right]\begin{pmatrix} v_1 \\ v_2 \end{pmatrix} + \lambda\begin{pmatrix} \nabla u_2\cdot\nabla v_1 \\ 0 \end{pmatrix}, \tag{4.2}$$

and in the notation from Remark 3.2, the first matrix in (4.2) relates to `cj`, the second to `aj`, and the last term gives $-b\otimes\nabla v$ with $b_{111} = -\lambda\partial_x u_2$, $b_{112} = -\lambda\partial_y u_2$, and $b_{ijk} = 0$ else.

4.1.1. Bifurcation diagram over rectangles (`chemtax`)

Following [23] we first study (4.1) on a rectangular domain $\Omega = [-L_x/2, L_x/2] \times [-L_y/2, L_y/2]$ with homogeneous Neumann BC. Again a number of results can then be obtained analytically. There are two trivial stationary branches, namely $u = (0,0)$ which is always unstable, and $u = u^* = (1,1/2)$. From the BC, the eigenvalue problem $Mv = \mu v$ for the linearization around u^* has solutions of the form $\mu = \mu(m,l,\lambda)$, $v = v(m,l,\lambda;x) = \phi e_{m,l}(x,y)$ with

$$e_{m,l}(x,y) = \cos\left(\frac{m\pi}{L_x}\left(x+\frac{L_x}{2}\right)\right)\cos\left(\frac{m\pi}{L_y}\left(y+\frac{L_y}{2}\right)\right),$$

$(m,l) = (1,0),(0,1),(2,0),(1,1),\cdots$, and $\phi \in \mathbb{R}^2$. To study bifurcations from u^* we solve $\mu(m,l,\lambda) \stackrel{!}{=} 0$ for λ which yields

$$\lambda_{m,l} := 4(Dk^2+r)(k^2+1)/k^2, \quad \text{where } k^2 := \pi^2\left(\frac{m^2}{L_x^2}+\frac{l^2}{L_y^2}\right).$$

Table 9: Bifurcation from $u^* = (1, 1/2)$ in (4.1), $D = 1/4$, $r = 1.52$.

(m,l)	(0,2)	(0,3)	(0,1),	(1,0),(0,4)	(1,1)	(1,2)	⋯
λ_{ml}	12.01	13.73	17.55	17.57	18.15	19.91	⋯

Table 10: `chemf.m` as a prototype for definition of PDE coefficients in case of a (nonsymmetric) c depending on u and λ. See `isoc` and the `assempde` documentation for the order of c_{ijkl} in `c`.

```
function [c,a,f,b]=chemf(p,u,lam) % chemotaxis system with isoc
u=pdeintrp(p.points,p.tria,u);a=0;b=0; v1=ones(1,p.nt);
f1=r*u(1,:).*(1-u(1,:));f2=u(1,:)./(1+u(1,:))-u(2,:); f=[f1;f2];
D=0.25;r=1.52; c=isoc([[D*v1 -lam*u(1,:)];[0*v1 v1]],p.neq,p.nt);
```

As in [23, Fig.3] we choose $D = 1/4$, $r = 1.52$ and the "1×4" domain $L_x = 1$, $L_y = 4$, which yields Table 9.

To encode (4.1) we note that $c_{1111} = c_{1122} = D, c_{1211} = c_{1222} = -\lambda u_1, c_{2211} = c_{2222} = 1$, and all other entries of c are zero. In particular, c is isotropic and thus we may use `isoc.m` to encode it, see Table 10 for `chemf.m`. For convenience and illustration, here by default we first use `p.jsw=3` in `cheminit.m` such that `p.jac` need not be set. With `p=stanmesh(p,0.075)` leading to `p.nt=2376` this still gives quick results, which moreover essentially do not change under mesh refinement. Also, in `cheminit` we introduce `p.vol=`$|\Omega|$; we want to use this quantity in `chembra.m` since the bifurcation diagrams in [23] plot $\|u_1 - 1\|_{L^1}/|\Omega|$ over λ. Again this is a simple example that the user can augment the structure `p` with whatever is useful. The commands in `chemcmds.m` yield the bifurcation diagram in Fig. 9, where the bifurcation values $\lambda_{m,l}$ (except for $\lambda_{10} = \lambda_{04}$) are found with reasonable accuracy, and which agrees well with [23, Fig.3(a)], with one exception: on the $(1,1)$ branch there is a loop near $\lambda = 20.5$ with two bifurcations, during

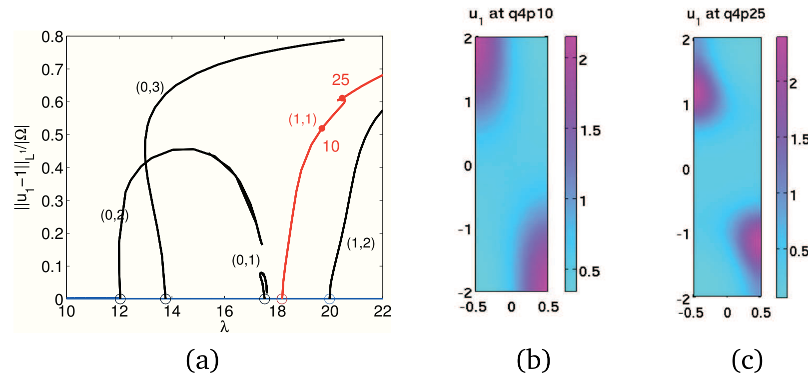

Figure 9: (a) Bifurcation diagram for (4.1) with `jsw=3`, i.e., numerical Jacobians, and `nt=2376`. Of the bifurcating branches only the $(0,2)$-branch is stable in a certain λ range, and a number of secondary bifurcations occur on each branch. (b), (c) The shape of solutions before and after the loop on the $(1,1)$ branch. For `jsw=1` we need finer meshes (`nt` $\approx 10^4$ to $2 \cdot 10^4$ and adaptive refinement), which, while giving smaller error-estimates, also destroy the speed advantage of assembled Jacobians.

which the solution structure changes as detailed in (b), (c). Presumably, this loop was just missed in [23] due to a larger stepsize.

Alternatively, to run (4.1) with `jsw=1` we also provide `chemjac.m` which encodes (4.2). For this, however, we need considerably finer meshes, mainly since the calculation of the coefficient Δu_2 (needed for `jsw<2`) via `pdegrad` and `pdeprtni` does not go together well with Neumann boundary conditions, since the averaging involved in `pdeprtni` produces some error at the boundaries. Therefore, we also replace `p=cheminit(p)` by `p=cheminitj(p)`, which resets a number of switches to (re)run (4.1) with `jsw=1`. See the end of `chemcmds.m` resp. `chemdemo.m`.

4.1.2. Drawing general domains (`animalchem`)

In `animalchem` we consider (4.1) on the animal-shaped domain in Fig. 10, taken from [25], with Neumann BC. To set up Ω we proceed graphically as explained in Section 3.1.4, see `animalgeo.m`, also for the setup of the BC. The plots in Fig. 10 are generated from the commands in `animalcmds.m`. For problems of this type, the bifurcation directions from a trivial branch are often most interesting, and can be generated with `plottauf`.

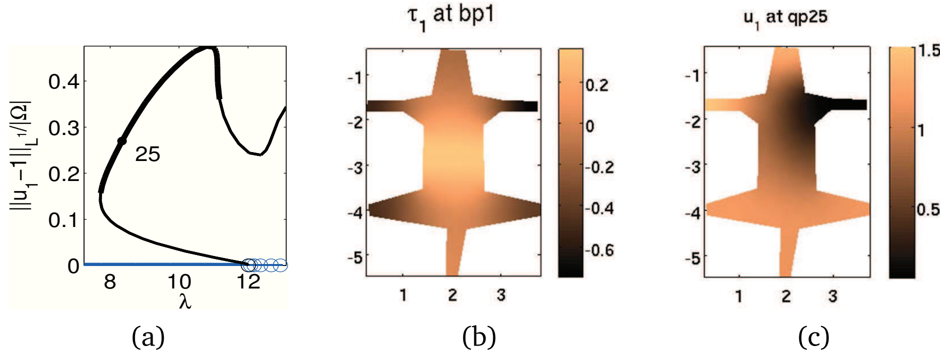

Figure 10: Bifurcation diagram for (4.1) ($\|u_1 - 1\|_{L^1}/|\Omega|$ over λ), first bifurcation direction from the trivial branch, and a stable solution on the bifurcating branch.

4.2. The Schnakenberg model (`schnakenberg`)

We consider the (stationary) Schnakenberg system in the form

$$0 = G(u) = \begin{pmatrix} -\Delta u_1 + u_1 - u_1^2 u_2 \\ -d\Delta u_2 - \lambda + u_1^2 u_2 \end{pmatrix}, \tag{4.3}$$

fix $d = 60$, use λ as bifurcation parameter, and let $(x, y) \in \Omega = [-l_x, l_x] \times [-l_y, l_y]$ with Neumann BC. Over $\Omega = \mathbb{R}^2$ the spatially homogeneous solution $u^*(\lambda) = (\lambda, 1/\lambda)$ becomes Turing unstable [25] when decreasing λ below $\lambda_c = \sqrt{d}\sqrt{3-\sqrt{8}} \approx 3.2085$, with critical

wave number $k_c = \sqrt{\sqrt{2}-1} \approx 0.6436$. In 2D, the most famous Turing patterns are stripes and spots which modulo rotational symmetry and spatial translations can be expanded as

$$u(x,y) = u^* + \left[A\cos(k_c x) + B\cos\left(\frac{k_c}{2}x\right)\cos\left(\frac{\sqrt{3}}{2}k_c y\right)\right]\Phi + \text{h.o.t.},$$

where A, $B \in \mathbb{R}$ are suitable amplitudes, $\Phi \in \mathbb{R}^2$ is the critical eigenvector of the linearization of G around u^* at λ_c, h.o.t. stands for higher order terms (in $A, B, \lambda - \lambda_c$), and we dropped the λ dependence of all terms. For $A \neq 0$ and $B = 0$ we have stripes, for $A = B \neq 0$ hexagonal spots, and there are also so called mixed patterns with $0 \neq A \neq B \neq 0$.

If the domain and BC permit it, both (spots and stripes) bifurcate simultaneously from the trivial branch at $\lambda = \lambda_c$. Here, to make λ_c a simple bifurcation point (for vertical stripes), we choose $l_x = 2m\pi/k_c$ and $l_y = 2n\delta\pi/(\sqrt{3}k_c)$, $m, n \in \mathbb{N}$, where $\delta \approx 1$ is a deformation parameter, such that for $\delta \neq 1$ the multiple bifurcation point splits into simple bifurcation points. The (analytical and numerical) bifurcation diagram for (4.3) with these settings is discussed in detail in [36], including symmetries and a number of new tertiary bifurcations leading so called snaking branches between stripes and spots. However, since the Schnakenberg system is a standard model in pattern formation, and since it motivated us to implement the extension `pmcont` to `cont` described below, here we also include a short discussion.

Fig. 11 shows a bifurcation diagram and some solution plots obtained for $m = n = 2$ and $\delta = 0.99$, which we call a "2 × 2" domain as 2 spots in both directions fit in. To locate the primary bifurcations we use `findbif.m` on the homogeneous branch. Here we use the term "hot" ("cold") spot for spots with a maximum (minimum) of u_1 in the middle of the domain, and in analogy keep these names also for mixed modes. We also plot $|\hat{u}_{kl}|^{1/2}$, where $\hat{u}$ is the discrete Fourier transform of $u_1 - \langle u_1 \rangle$, see `fourierplot.m` in directory `schnakenberg`. These Fourier plots are often interesting for pattern forming systems: for instance $\mathscr{F}$`hhp10` shows that the pattern is essentially still generated from the basic harmonics $\exp(ik_c x)^m$ and $\exp(ik_c(x/2 \pm \sqrt{3}y/2)^n$, $m, n = \pm 1$, while at `hhp30` higher harmonics $\exp(ik_c x)^m \exp(ik_c(x/2 \pm \sqrt{3}y/2))^n$, $m, n \in \mathbb{N}$, contribute more significantly.

The tangents plots for bifurcation in Fig. 12 show that the slightly "detuned" domain yields that the first branch point shows a 3 × 1.5 (rectangular) nodal structure. The stripes are then found as the second bifurcation point, and the third branch point actually corresponds to a "rectangle" branch with a 2 × 2 nodal structure; see also the associated Fourier plot in Fig. 11. However, for this branch the Newton loop immediately takes us on the hexagonal branch, which in this case we happily accept since this is the branch we are interested in. Here, to speed up calculations, after calculating the first bifurcations from `hom` with `findbif`, we mostly switch off bifurcation detection, and instead locate the bifurcations to the bean branches a posteriori via `findbif`. Thus, in the bifurcation diagram we also plot only very few of the branch points that can be found on these branches, namely the branch points to the so called bean branches, and the bifurcation points on the hot beans, discussed further in [36].

Fig. 11 is not generated by `cont` alone, because with `cont` we quickly obtain some

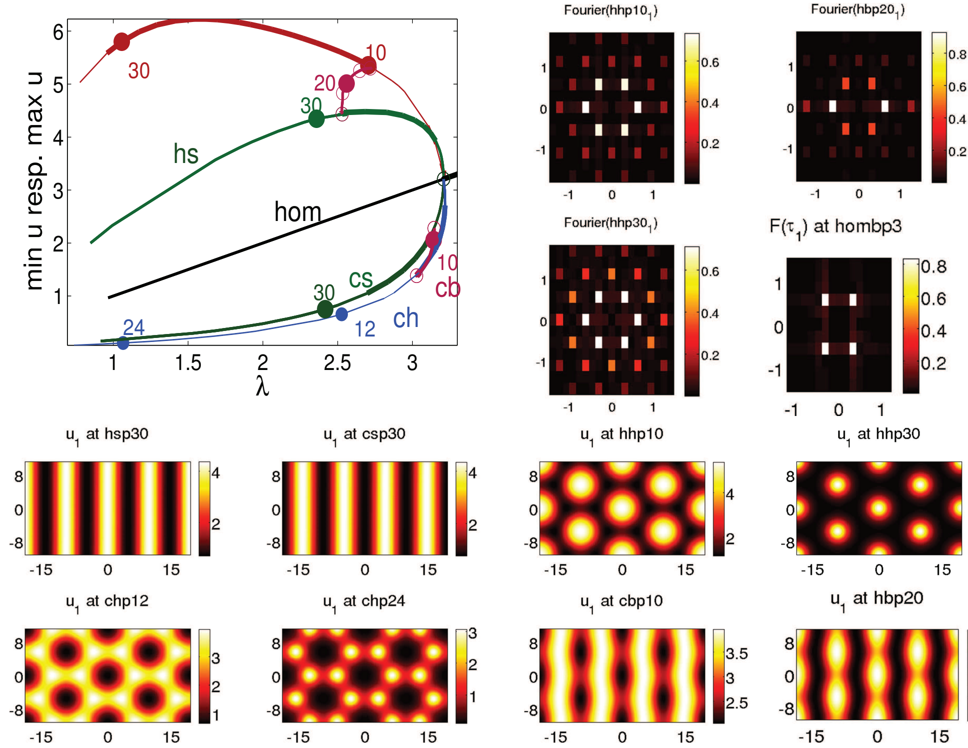

Figure 11: Bifurcation diagram and a selection of patterns for (4.3) over a "2×2"-domain, $\delta = 0.99$, see schnak22cmds.m. The branch hom in the bifurcation diagram are the homogeneous solutions, hs the (hot) stripes, cs the phase shift of hs, hh the hot hexagons, ch the cold spots. hb and cb are mixed modes, also called beans. $\circ$ are bifurcation points. Thick lines mean stable branches and thin lines unstable. For hs, hh and hb we plot the maximum of u_1 and for cs, ch and cb the minimum. As usual, hsp30 stands for the third point of hs.

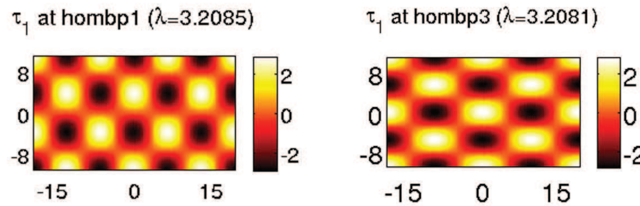

Figure 12: Tangents for bifurcation at the first and the third bifurcation points on hom.

undesired branch switching. For instance when continuing the stripe branch hs with standard settings we switch to the beans branch hb when approaching its bifurcation point. This particular branch switching can be avoided by decreasing ξ to $\xi = 0.1$/p.np, say, but only to the effect that we get branch switching at some later point on the s branch. Such undesired branch switching is a serious problem in all continuation algorithms, see, e.g., [33, section 3]. Here we use a modification pmcont of cont explained in the next section, which also incorporates some parallel computing for speedup. Still, the bifurcation

diagram in Fig. 11 is computationally quite expensive: about 40 minutes on a quad-core desktop PC, with 20.000 triangles in the base mesh and about 60.000 triangles on average during mesh adaptation. This large number of triangles is mainly needed to avoid undesired branch switching. It gives error-estimates $E(K) \le 0.04$ for all solutions calculated. To also provide a cheaper example, the init-function `p=schnakinit(p,m,n,nx,del)` takes the domain sizes m, n, the deformation parameter δ and the startup spatial discretization `nx` as parameters. `schnak11demo.m` (or schnak11cmds.m) then uses $m = n = 1$ and $\delta = 0.97$ and only takes a few minutes for a bifurcation diagram similar to Fig. 11 over the smaller domain.

4.3. `pmcont`

Theorem 4.4 in [19] guarantees that the standard continuation converges to a given branch for "sufficiently small" ds, but near bifurcation points only in cones around the branch. Thus, near a bifurcation point it is often not useful to choose very small ds. To circumvent this and similar problems we provide the function `pmcont`. The basic idea is explained in Fig. 13.

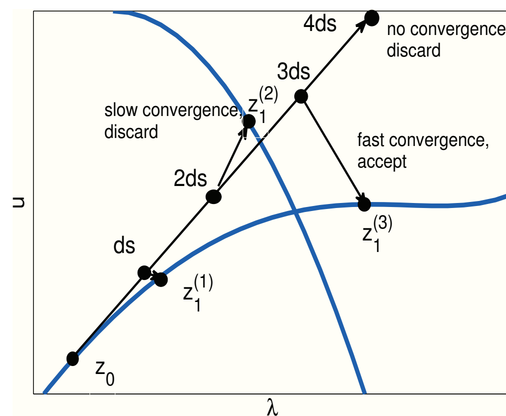

Algorithm `pmcont`

1. Multi-predictors. $(u^i, \lambda^i) = (u_0, \lambda_0) + i$ p.dsτ, $i = 1, \cdots$,p.mst
2. Newton-loops (parallel). Use (4.4) to identify "good points".
3. Tangents (sequentially). Calculate new tangents $\tau_1, \cdots, \tau_m$ at good points, using (2.5).
4. Bifurcation detection and localization (parallel).
5. Postprocess (sequentially). Call `ufu`, save, plot, return to 1.

Figure 13: Sketch of the basic idea of multiple predictors and convergence monitoring, and pseudo code of `pmcont`. The arrows from, e.g. "2ds" to $z_1^{(2)}$ just illustrate the result of the Newton loops, not the hyperplanes $\{(u, \lambda) \in \mathbb{R}^{Nn_p+1} : \langle \tau_0, (u, \lambda) \rangle_\xi = \text{ds}\}$ as in Fig. 1.

Instead of using just one predictor $(u^1,\lambda^1)=(u_0,\lambda_0)+\mathrm{p.ds}\,\tau$, `pmcont` creates in every continuation step the predictors

$$(u^i,\lambda^i)=(u_0,\lambda_0)+i\,\mathrm{p.ds}\tau,\quad i=1,\cdots,\mathrm{p.mst},$$

and starts a Newton loop for each. `pmcont` then monitors the convergence behavior of each loop to decide whether it yields a "good" point, i.e., a point on the present branch. The criterion is that in each Newton step the residual has to decrease by a factor $0<\alpha<1$, i.e.,

$$\|G(u_{n+1},\lambda_{n+1})\|\le\alpha\|G(u_n,\lambda_n)\|, \tag{4.4}$$

otherwise the loop is stopped. The heuristic idea is that if the Newton loop converges slowly, then probably the solution is on a different branch, because the loop has to (slowly) change the solution "shape".

For the crucial parameter α, which describes the desired convergence speed, we recommend trying $\alpha = 0.1$.* Of course, these heuristics in no way guarantee that no branch switching occurs, or anyway that we get convergence for long predictors $(u^i,\lambda^i)=(u_0,\lambda_0)+i\cdot\mathrm{ds}\tau$ with $i>1$, but in practice we find the idea to work remarkably well. See also Section 5.2 for an example of the "unreasonable effectiveness" of `pmcont`, together with an example that long predictors in `pmcont` tend to branch-switching in imperfect bifurcations.

We need three additional parameters: the number of predictors `p.mst`, α = `p.resfac`, and `p.pmimax`. These are used to gain some flexibility for (4.4) via stepsize control. If `p.mst` equals the number of solutions found and `p.ds` is smaller than `p.dsmax/p.dsincfac`, then `p.ds` will be increased by the factor `p.dsincfac`, essentially as in `cont`. If no solution is found and `p.ds` is greater than `(1+p.mst)`·`p.dsmin`, then the step size `p.ds` will be divided by `1+p.mst` in the next continuation step. Finally, if `p.ds` is less than `(1+p.mst)`·`p.dsmin` and `p.pmimax` is less than `p.imax`, then `p.pmimax` will be increased to `p.pmimax+1`, and $\|G(u_{n+\mathrm{p.pmimax}},\lambda_{n+\mathrm{p.pmimax}})\|\le$ `p.resfac`$\|G(u_n,\lambda_n)\|$ is required instead of (4.4). The only new functions used in `pmcont` are `pmnewtonloop.m` and `pmbifdetec.m`.

Another advantage of the `p.mst` predictors is that the Newton loops can be calculated in parallel, which on suitable machines gives substantial speedups (where we use the Matlab Parallel Computing Toolbox in a standard setup). All final Newton iterates with a residual smaller than `p.res` are taken as solutions, and plotted and saved as in `cont`. Next, the tangents $\tau_1,\cdots,\tau_m$ are calculated sequentially, because τ_{l+1} needs τ_l, and afterwards the bifurcation detection and localization is again in parallel. The last solution and its tangent will be used for the next continuation step. Mesh adaptation/refinement is inquired at the start of `pmcont`, i.e., before generating the predictors, but not on the individual correctors.

*However, for instance on the "hot hexagon branch" hh in Fig. 11 we need to use $\alpha=10^{-6}$ to avoid branch-switching, see `schnak22cmds.m`.

In summary, for `p.resfac=1` and `p.mst=1` we have that `pmcont` is roughly equivalent to `cont`, except for slightly less versatile mesh adaptation and error estimates. For `p.mst>1`, `pmcont` takes advantage of parallel computing, and is often useful to avoid convergence problems and undesired branch switching close to bifurcation points. The main reasons why we (currently) keep the two versions and do not combine them into one is that `cont` is simpler to hack and implements Keller's basic, well tested algorithm.

5. Three classical examples from physics

In this section we consider models for Bose-Einstein (vector) solitons, Rayleigh-Bénard convection, and the von Kármán system for buckling of an elastic plate, as examples for systems with more than two components, and with BC implemented via `gnbc` as described in Section 3.1.4. The largest and most complicated system (in the sense of number of components and implementation of BC) here is the von Kármán system. Hence, for this we also explain the coding in `pde2path` in most detail, while for Bose-Einstein solitons and Rayleigh-Bénard convection we mostly refer to the m-files for comments.

5.1. Bose-Einstein (vector) solitons (`gpsol`)

As an example with x, y dependent coefficients, nontrivial advection, and interesting localized solutions we consider (systems of) Gross-Pitaevskii (GP) equations with a parabolic potential that arise for instance as amplitude equations in Bose-Einstein condensates.

5.1.1. The scalar case

First, following [21] we consider the scalar equation

$$\mathrm{i}\partial_t \psi = -\Delta\psi + r^2\psi - \sigma|\psi|^2\psi, \tag{5.1}$$

where $\psi = \psi(x, y, t) \in \mathbb{C}$, $r^2 = x^2 + y^2$, and $\sigma = 1$ (focusing case). This has a huge number of families of localized solutions, aka solitons, which may be time periodic, standing or rotating in space. Going into a frame rotating with speed ω and splitting off harmonic oscillations with frequency μ, i.e.,

$$\psi(x, y, t) = \Phi(r, \phi - \omega t)\mathrm{e}^{-\mathrm{i}\mu t}, \tag{5.2}$$

we obtain

$$\left[\partial_r^2 + \frac{1}{r}\partial_r + \frac{1}{r^2}\partial_\theta^2 - \mathrm{i}\omega\partial_\theta + \mu - r^2\right]\Phi + \sigma|\Phi|^2\Phi = 0. \tag{5.3}$$

A typical ansatz for (approximate) solutions has the form

$$\Phi(r, \theta) = A\phi(r/a)(\cos(m\theta) + \mathrm{i}p\sin(m\theta)), \quad \text{with e.g.,} \quad \phi(\rho) = \rho^m L_n^{(m)}(\rho^2)\mathrm{e}^{-\rho^2/2}, \tag{5.4}$$

where $L_n^{(m)}$ is the n^{th} Laguerre polynomial. Plugging this into (5.3) yields expressions for A, a, p, ω for approximate solutions. The case n=0 and p=0 corresponds to so called (non rotating, since $\omega = 0$) real m-poles, and $|p| = 1$ to a so called radially symmetric vortex of charge m, while the intermediate cases $0 < |p| < 1$ give to so called rotating azimuthons with interesting angular modulations of $|\Phi|$.

Our goal is to calculate these solutions numerically with `pde2path`. Returning to Cartesian coordinates, i.e., setting $\Phi(r,\theta) = u(x,y) + \mathrm{i}v(x,y)$, we obtain the 2-component real elliptic system

$$-\Delta u + (r^2 - \mu)u - |U|^2 u - \omega(x\partial_y v - y\partial_x v) = 0, \tag{5.5a}$$

$$-\Delta v + (r^2 - \mu)v - |U|^2 v - \omega(y\partial_x u - x\partial_y u) = 0, \tag{5.5b}$$

where $|U|^2 = u^2 + v^2$. Our strategy is to use (5.4) for $\omega = 0$ and to continue in $\lambda := \omega$. A measure for the deformation of multi-poles into vortices for the numerical solutions is the "modulation depth" p of the soliton intensity

$$p = \max|\mathrm{Im}\Phi| / \max|\mathrm{Re}\Phi| = \max|v| / \max|u|. \tag{5.6}$$

The result of typical continuation of a quadrupole using a stiff-spring approximation of DBC for u, v on domain $\Omega = [-5,5]^2$ is shown in Figs. 14(a)-(f), see `gpf.m`, `gpjac.m`, `gpcmds.m` and `gpinit.m`, and also `plotsol.m` in directory `gpsol` for the customized of `plotsol`.

The following remarks are in order. First, we switch off stability or bifurcation tracking (`spcalcsw=0`, `bifchecksw=0`), see Remark 5.1 below. Second, the linearization of (5.5) is given by (recall that $\lambda = \omega$)

$$G_u(u,v) = \begin{pmatrix} -\Delta & 0 \\ 0 & -\Delta \end{pmatrix} + \begin{pmatrix} -\mu + r^2 - 3u^2 - v^2 & -2uv \\ 2uv & -\mu + r^2 - 3v^2 - uv \end{pmatrix}$$
$$-\lambda \begin{pmatrix} 0 & x\partial_y - y\partial_x \\ -x\partial_y + y\partial_x & 0 \end{pmatrix}.$$

Thus, the last term is a good example how to use `assemadv` with a relatively complicated b. Third, some problems might be expected from the large number of solutions of (5.3), in particular the phase-invariance: if Φ is a solution, so is $\mathrm{e}^{\mathrm{i}\alpha}\Phi$ for any α, or equivalently, (5.5) is invariant under multiplication with

$$M(\alpha) = \begin{pmatrix} \cos\alpha & -\sin\alpha \\ \sin\alpha & \cos\alpha \end{pmatrix},$$

and indeed we can also phase rotate our numerical solutions without changing the residual, see the end of `gpcmds.m`. Thus, even for all parameters fixed, solutions of (5.5) always come in continuous families, and hence G_u as a linear operator on $[L^2(\mathbb{R}^2)]^2$, say, always has a zero eigenvalue. See also [21] and the references therein for some tricks for

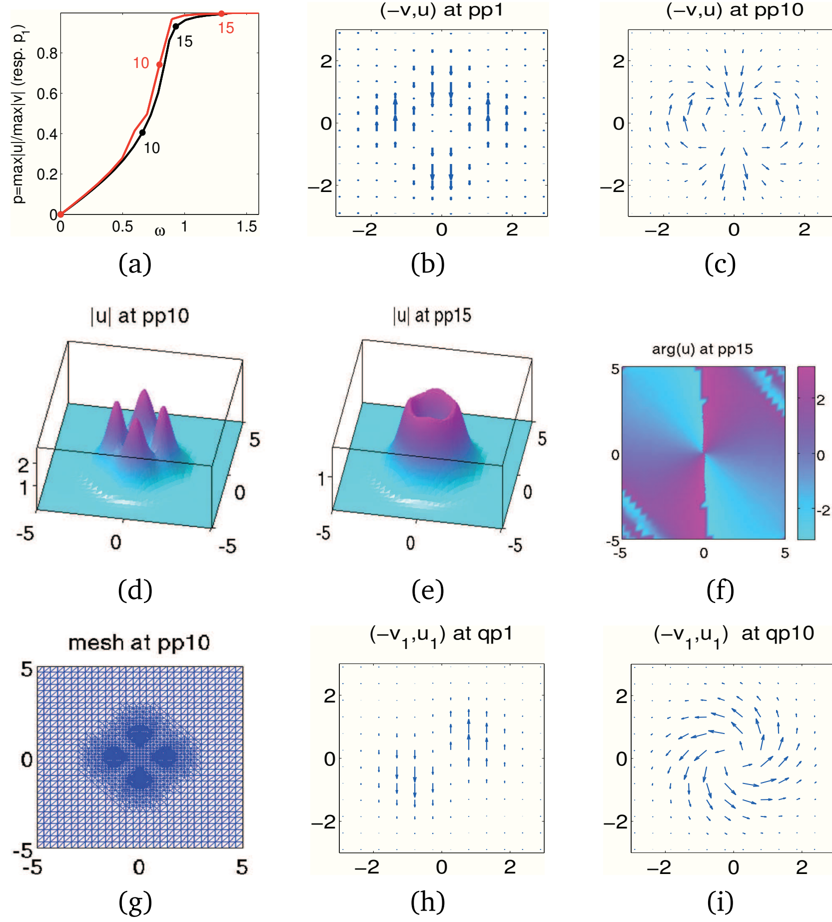

Figure 14: (a) Continuing a quadrupole to an azimuthon to a 2-vortex, $p = \max|u|/\max|v|$ over $\lambda = \omega$, (p branch, black); and continuing a vector dipole to a vector-azimuthon to a vector-vortex, $p_1 = \max|v_1|/\max|u_1|$ over $\lambda = \omega$ (q branch, red), $\mu = 2$ resp. $\mu_1 = 2, \mu_2 = 2.2$. (b), (c) vectorfield plots at $\lambda = 0$ (quadrupole) resp. $\lambda \approx 0.66$ (azimuthon). (d)-(f) $|U|$ and $\arg(u+iv)$ as indicated. (g) mesh for p `p.nt=5208` from mesh-refinement during initialization. (h), (i) vectorfield plots for the first condensate ψ_1 in q as indicated, second condensate similar. Also see the customized `plotsol.m` in gp. The phase plot in (f) is somewhat ragged due to the coarse mesh away from the center. In the q branch we used `q.amod=8` with `q.maxt=9000` which lead to `q.nt=10064` at, e.g., point 10. Error estimates $E(K)$ around 0.025 on the p branch and around 0.007 on the q branch.

the numerical solution of (5.3), which however we do not need, for the following reasons. Analytically, the kernel of $G_u(U)$, $U = (u, v)$, is spanned by

$$\Psi = \left(\frac{d}{d\alpha}M(\alpha)|_{\alpha=0}\right)(u,v) = (-v,u).$$

Then, in the Newton step $U_{n+1} = U_n - G_u(U_n)^{-1}G(U_n)$ we have

$$\langle \Psi, G(U)\rangle = \int \begin{pmatrix} v \\ -u \end{pmatrix} \cdot \begin{pmatrix} \Delta u - (r^2-\mu)u + |U|^2u + \omega(x\partial_y v - y\partial_x v) \\ \Delta v - (r^2-\mu)v + |U|^2v + \omega(y\partial_x u - x\partial_y u) \end{pmatrix} \,\mathrm{d}x$$

$$=\frac{\omega}{2}\int x\partial_y(u^2-v^2)-y\partial_x(u^2-v^2)\,\mathrm{d}x=0, \tag{5.7}$$

due to the Dirichlet boundary conditions. Thus, $G(U_n)$ is orthogonal to the kernel of $G_u(U_n)$. Numerically we find that the eigenvalue of $G_u(U)$ closest to zero typically has modulus 10^{-5} or smaller, and the condition estimate for G_u is on the order of 10^6 to 10^8; however, (5.7) also holds to the order of 10^{-7} or 10^{-8}, and the \ operator accurately solves $G_u(U_n)V_n=G(U_n)$. Again see the end of `gpcmds.m`.

One trick we do use is to start with a coarse mesh of 30×30 points, first take some (rather arbitrary) monopole as dummy-starting guess, use `meshref` to generate a rather fine mesh in the center, define a quadrupole initial guess using (5.4) on that first refined mesh, and then refine again, yielding the (still small) number of 5208 triangles for this continuation, with an error-estimate less than 0.01. On a small laptop computer the whole continuation takes about a minute.

5.1.2. A two-component condensate

The above can be generalized to multi-component condensates [22]. For two components, we then have coupled GP equations of the form, e.g.,

$$\mathrm{i}\partial_t\psi_1=[-\Delta+r^2-\sigma|\psi_1|^2-g_{12}|\psi_2|^2]\psi_1, \tag{5.8a}$$

$$\mathrm{i}\partial_t\psi_2=[-\Delta+r^2-\sigma|\psi_2|^2-g_{21}|\psi_1|^2]\psi_2, \tag{5.8b}$$

where g_{12}, g_{21} are called interspecies interaction coefficients. Physically, it makes sense to use ansätze of the form (5.2) with different μ but equal ω, i.e., $\psi_j(x,y,t)=\Phi_j(r,\phi-\omega t)\mathrm{e}^{-\mathrm{i}\mu_j t}$. Next we can use the form (5.4) for each component Φ_j and classify the thus obtained approximate solutions as soliton-soliton, soliton-vortex, soliton-azimuthon etc pairs. To calculate such solutions numerically we set $\Phi_j(r,\theta)=u_j(x,y)+v_j(x,y)$ and obtain an elliptic system of the form (5.5) but with four real equations. This has been implemented in `vgpf`, with Jacobian `vgpjac`. Figs. 14(a), (g), (h) shows the continuation of a two-dipole obtained from `vgpcmds`. Similar remarks as for the scalar case apply, see the comments in `vgpcmds.m`.

Remark 5.1. clearly was just a very introductory demo of continuation of solutions of (5.1) resp. (5.8); there are *many* more and interesting branches, and further questions, again see, e.g., [21, 22]. Interesting questions concern, e.g., the dependence of vector solitons on $|\mu_1-\mu_2|$ which can for instance be studied by fixing ω and continuing in $\lambda=\mu_2$, or the effect of including a periodic potential, leading to gap solitons [8]. Some of these questions will be considered elsewhere. Also, to study any bifurcations for (5.1) or (5.8) the phase invariance must be broken by adding some constraints.

5.2. Rayleigh-Bénard convection (`rbconv`)

As an example from fluid dynamics we consider two-dimensional Rayleigh-Bénard convection in the Boussinesq approximation in the domain $\Omega=[-2,2]\times[-0.5,0.5]$. In the

streamfunction formulation the stationary system reads

$$-\Delta\psi+\omega=0, \tag{5.9a}$$
$$-\sigma\Delta\omega-\sigma R\partial_x\theta+\partial_x\psi\partial_z\omega-\partial_z\psi\partial_x\omega=0, \tag{5.9b}$$
$$-\Delta\theta-\partial_x\psi+\partial_x\psi\partial_z\theta-\partial_z\psi\partial_x\theta=0, \tag{5.9c}$$

with streamfunction ψ, temperature θ, and the auxiliary $\omega=\Delta\psi$. Moreover, σ is the Prandtl number, set to 1 here, and R the Rayleigh number, which will be the continuation parameter. The implementation of (5.9) in `pde2path` is relatively straightforward, including analytical Jacobians, see `rbconvf.m` and `rbconvjac.m`.

The boundary conditions at the top and bottom plates are taken at constant temperature and with zero tangential stress

$$\psi=\partial_{zz}\psi=\theta=0, \quad \text{at } z=\pm 0.5,$$

which also means $\Delta\psi=0$ at $z=\pm 0.5$. Motivated by the analysis in [17], laterally we consider on the one hand "no-slip" (and perfectly insulating) BC

$$\psi=\partial_x\psi=\partial_x\theta=0, \quad \text{at } x=\pm L, \tag{5.10}$$

and on the other hand "stress free" BC

$$\psi=\partial_{xx}\psi=\partial_x\theta=0, \quad \text{at } x=\pm L. \tag{5.11}$$

See the comments in `rbconvbc_noslip.m` resp. `rbconvbc_stressfree.m` for the implementation (approximation) of these BC based on (1.2).

In both cases it is known that continuation of the trivial zero state for increasing R gives a sequence of bifurcations alternating between even and odd modes. The stability thresholds are plotted in [17], Fig. 1(a) for (5.11) and (b) for (5.10). We use these to choose initial values of $\lambda=R$ for the first two bifurcations, respectively.

For the no-slip case (5.10), the resulting bifurcation diagram is plotted in Fig. 15, which corresponds to the sketch Fig. 2 in [17]. No secondary bifurcations are found up to $R=900$. For stress-free BC we obtain the bifurcation diagram in Fig. 16(a), which corresponds to the case $b'^2>a'^2, b'>0$ in Fig. 3 of [17]. Here the secondary symmetry breaking pitchfork from [17] is turned into an imperfect pitchfork. The $x\to -x$ reflection symmetry is broken by the triangle data of the mesh (here we use `poimesh`) and the stiff-spring approximation of the boundary conditions. We have located the stable branch `s` of the imperfect pitchfork by time-integrating (which is not equivalent to the time integration of the time-dependent Boussinesq equations) with `tint` from the unstable branch in the suitably chosen unstable direction.

The demos run on a rather coarse mesh of 100×25 grid points, because even with assembled Jacobians the calculations are rather slow due to a non-simple structure of the Jacobians. Thus, we use `pmcont` for the bifurcating `q` and `r` branches, which gives a huge speed advantage and works remarkably well even directly after bifurcation from the trivial branch, where long predictors are far off the actual branches, e.g., in Fig. 15 *all* points are

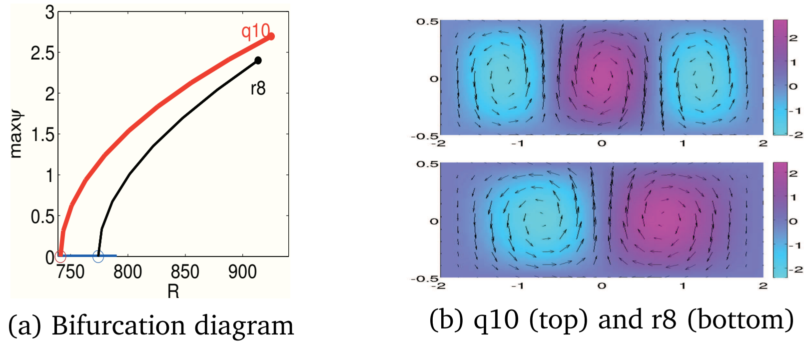

(a) Bifurcation diagram (b) q10 (top) and r8 (bottom)

Figure 15: (a) Bifurcation diagram of (5.9) with (5.10). (b) sample solutions (ψ, and arrows indicating the fluid flow) from (a). See `arrowplot.m`.

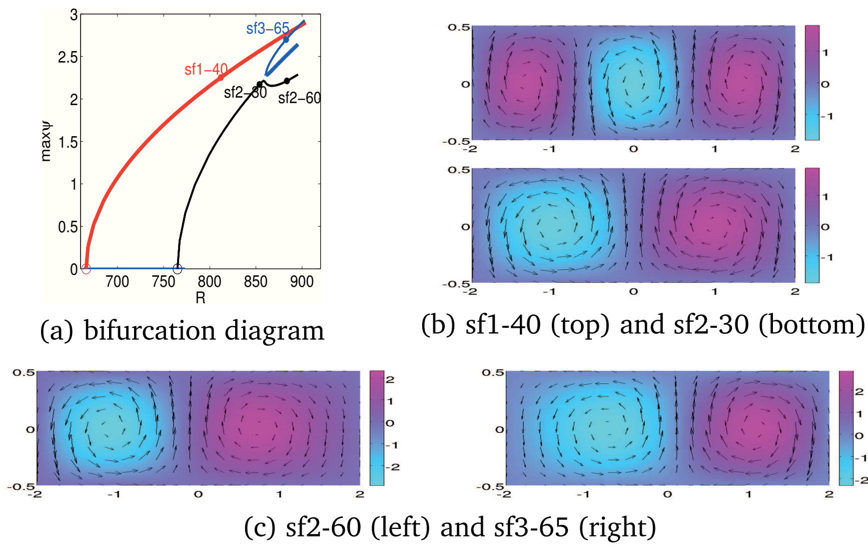

(a) bifurcation diagram (b) sf1-40 (top) and sf2-30 (bottom)

(c) sf2-60 (left) and sf3-65 (right)

Figure 16: Bifurcation diagram of (5.9) with stress-free b.c. and sample solutions (ψ). Here `sf1-40` and `sf2-30` are approximately symmetric, while `sf2-60` and `sf3-65`, generated in a (numerically) imperfect pitchfork around $R = 860$, are not.

calculated from the essentially "vertical" predictors at bifurcation. To check that this does not miss any (possibly imperfect) bifurcation we compared with `cont` with small `ds` and obtained the same branches but much slower. In Fig. 16, however, we switch back to `cont` when approaching the imperfect pitchfork since `pmcont` tends to switch from the `r` to the `s` branch via long predictors.

Error estimates $E(K)$ are around 0.125 for the uniform 100×25 mesh. Here mesh refinement also tends to be quite uniform, and $E(K)$ roughly halves for the double number of triangles. Increasing the number of meshpoints brings the diagram closer to a symmetry breaking pitchfork. In fact, for both boundary conditions, qualitatively the same bifurcation diagram can be found for even coarser meshes, but the location of the branches can be off by order 100 in R.

5.3. Von Kármán description of the buckling of plates (`vkplate`)

The von Kármán equations

$$-\Delta^2 v - \lambda\partial_x^2 v + [v,w] = 0, \quad -\Delta^2 w - \frac{1}{2}[v,v] = 0, \tag{5.12}$$

can be derived to describe the deformation of an elastic (rectangular) plate $\Omega = [-l_x, l_x] \times [-l_y, l_y] \subset \mathbb{R}^2$ under compression. Here $v : \Omega \to \mathbb{R}$ is the out of plane deformation, $w : \Omega \to \mathbb{R}$ is the Airy stress function, $\Delta^2 = (\partial_x^2 + \partial_y^2)^2$ is the squared Laplacian, λ is the compression parameter, and the bilinear form $[\cdot,\cdot]$ is given by

$$[v,w] := v_{xx}w_{yy} - 2v_{xy}w_{xy} + v_{yy}w_{xx}.$$

There are a number of choices for the boundary conditions for (5.12). For v one can choose for instance between (in the notation from [13])

$$\begin{aligned} \text{I(v)}: &\quad v = \Delta v = 0 \text{ on } \partial\Omega \quad \text{(simply supported)},\\ \text{II(v)}: &\quad v = \Delta v = 0 \text{ on } y = \pm l_y, \quad v = \partial_n v = 0 \text{ on } x = \pm l_x,\\ &\quad \text{(simply supported on the sides, clamped at the ends)},\\ \text{III(v)}: &\quad v = \partial_n v = 0 \text{ on } \partial\Omega \quad \text{(clamped on whole boundary)}. \end{aligned}$$

Similarly, for w we may consider, on $\partial\Omega$,

$$\text{I(w)}: \ w = \Delta w = 0, \quad \text{II(w)}: \ \partial_n w = \partial_n(\Delta v) = 0, \quad \text{III(w)}: \ w = \partial_n w = 0.$$

Clearly, for all BC-combinations and all λ the trivial state $v = w = 0$ is a solution. Mathematically, the combination a) (I(v),I(w)) (sometimes as a whole called simply supported) is most simple because it allows an easy explicit calculation of bifurcation points from the trivial branch. However, [30] argues that physically the combinations b) (II(v),I(w)) or c) (II(v),II(w)) are more reasonable, and various combinations and modifications have been studied since, see [6] and the references therein for an overview.

Here we focus on case b) since this yields secondary bifurcations, called "mode jumping" in this field. The other cases can be handled quite similarly and, e.g., a) is in fact slightly simpler. The aim is to show how (5.12) can be put into `pde2path` and thus recover a number of interesting bifurcations.

Clearly, the first idea to set up (5.12) would be to introduce auxiliary variables $\Delta v, \Delta w$ and set

$$u = (u_1, u_2, u_3, u_4) = (v, \Delta v, w, \Delta w)$$

to obtain the (quasilinear elliptic) system

$$\begin{pmatrix} -\Delta & 1 & 0 & 0 \\ -\lambda\partial_x^2 & -\Delta & 0 & 0 \\ 0 & 0 & -\Delta & 1 \\ 0 & 0 & 0 & -\Delta \end{pmatrix} u - \begin{pmatrix} 0 \\ -[u_1,u_3] \\ 0 \\ \frac{1}{2}[u_1,u_1] \end{pmatrix} = 0,$$

for instance in case a) with homogeneous Dirichlet BC $u_1 = u_2 = u_3 = u_4 = 0$. The problem with this formulation in `pde2path` are the derivatives $\partial_x^2 u_1, \cdots, \partial_x \partial_y u_3$ in the nonlinearity. In principle, these can be obtained from calling `pdegrad`, `pdeprtni`, and `pdegrad` again.* However, the first problem is that this introduces some averaging into the second derivatives, in particular at the boundaries. The second problem is that with this approach we have no easy way to generate the Jacobian of G since `pdegrad/pdeprtni` neither fit to matrix assembling nor to numerical differentiation. For the latter the next-next-neighbor effect of `pdegrad/pdeprtni` does not comply with the Jacobian stencil, see Remark 3.2.

Thus, here we choose to introduce additional auxiliary variables, i.e., set

$$u = (v, \Delta v, w, \Delta w, \partial_x^2 v, \partial_y^2 v, \partial_x \partial_y v, \partial_x^2 w, \partial_y^2 w, \partial_x \partial_y w) \in \mathbb{R}^{10}.$$

For instance, $u_5 = \partial_x^2 u_1$ can then be simply added as a linear equation $-\partial_x^2 u_1 + u_5 = 0$ in the `pde2path` formulation (see below for the BC for $u_5, \cdots, u_{10}$). However, since this way we get a number of indefinite equations, in particular the mixed derivatives $-\partial_x \partial_y u_1 + u_7 = 0$ and $-\partial_x \partial_y u_3 + u_{10} = 0$, here we use an ad hoc regularization and set $-\partial_x^2 u_1 + (1 - \delta\Delta) u_5 = 0$ with small $\delta > 0$ (i.e., $\delta = 0.05$ numerically) and similarly for $u_6, \cdots, u_{10}$. Thus, instead of (5.12) we now really treat the problem

$$-\Delta^2 v - \lambda \partial_x^2 v + \left(S v_{xx} S w_{yy} - 2 S v_{xy} S w_{xy} + S v_{yy} S w_{xx} \right) = 0, \tag{5.13a}$$

$$-\Delta^2 w - \left(S v_{xx} S v_{yy} - S v_{xy} S w_{xy} \right) = 0, \tag{5.13b}$$

with the smoothing operator $S = (1 - \delta\Delta)^{-1}$. However, for small δ, comparison of our results with the literature shows that the regularization plays no qualitative or even quantitative role (in the parameter regimes we consider).

Thus, we now have a 10 component system, and to illustrate its implementation in `pde2path` we write it in the form $(-C + A)u - f = 0$ with

$$f = (0, -(u_5 u_9 - 2 u_7 u_{10} + u_6 u_8), 0, u_5 u_6 - u_7^2, 0, 0, 0, 0, 0, 0)^T,$$

and

$$-C+A = \begin{pmatrix}
-\Delta_1 & 1^{11} & & & & & & & & \\
-\lambda \partial_x^2{}_5 & -\Delta_{45} & & & & & & & & \\
 & & -\Delta_{89} & 1^{33} & & & & & & \\
 & & & -\Delta_{133} & & & & & & \\
-\partial_x^2{}_{17} & & & & -\tilde{\mathscr{D}}^{45}_{177} & & & & & \\
-\partial_y^2{}_{21} & & & & & -\tilde{\mathscr{D}}^{56}_{221} & & & & \\
-\partial_x \partial_y{}_{25} & & & & & & -\tilde{\mathscr{D}}^{67}_{265} & & & \\
 & & -\partial_x^2{}_{109} & & & & & -\tilde{\mathscr{D}}^{78}_{309} & & \\
 & & -\partial_y^2{}_{113} & & & & & & -\tilde{\mathscr{D}}^{89}_{353} & \\
 & & -\partial_x \partial_y{}_{117} & & & & & & & -\tilde{\mathscr{D}}^{100}_{397}
\end{pmatrix}.$$

---

* e.g., `[u1xt,u1yt]=pdegrad(p.points,p.tria,u(1:p.np)); u1x=pdeprtni(p.points,p.tria,u1xt); [u1xx,u1xy]= pdegrad(p.points,p.tria,u1x);` could be used to calculate (approximate) $\partial_x^2 u_1$.

Here, (for layout reasons) $\tilde{\mathscr{D}} = \delta\Delta + 1$, and the subscripts $1, 5, 17, \cdots$, denote the starting positions of the respective 2×2 tensor stored in the "400 rows vector" c, i.e., $-\Delta_1$ means $c_1 = [1; 0; 0; 1]$ stored in positions 1 to 4 in c, $\partial^2_{x\,5}$ means $c_2 = [1; 0; 0; 0]$ stored in positions 5 to 8 in c, and so on. The superscripts $11, 33, \cdots$ denote the positions in the "100 rows vector" a, and for $\tilde{\mathscr{D}}$ subscripts refer to $\delta\Delta$ and superscripts to $+1$. See `vkf.m`. Similarly, it is now rather easy to put the linearization f_u into `pde2path`, i.e., the second and fourth row of f_u as a 10×10 matrix read

$$\begin{array}{llllllllllll} f_u, \text{ 2nd row: } & (0 & 0 & 0 & 0 & -u_9{}^{42} & -u_8{}^{52} & 2u_{10}{}^{62} & -u_6{}^{72} & -u_5{}^{82} & 2u_7{}^{92}), \\ f_u, \text{ 4th row: } & (0 & 0 & 0 & 0 & u_6{}^{44} & u_5{}^{54} & -2u_7{}^{64} & 0 & 0 & 0). \end{array}$$

Here again the superscripts give the positions in aj. Of course, the full $\text{aj} = a - f_u$ also contains the constant coefficient terms at positions 11, 33, 45 etc from A; see `vkjac.m`.

It remains to encode the boundary conditions. First, $v = 0$ and $w = \Delta w = 0$ imply

$$\begin{aligned} &u_5 = v_{xx} = 0 \text{ and } u_6 = v_{yy} = 0 \text{ on horizontal edges, and} \\ &u_6 = v_{yy} = 0 \text{ on vertical edges, but no condition for } u_5 = v_{xx}, \text{ and} \\ &u_8 = w_{xx} = 0 \text{ and } u_9 = w_{yy} = 0 \text{ on all edges.} \end{aligned}$$

For u_5 on the vertical edges and u_7, u_{10} on all edges we take homogeneous Neumann boundary conditions. To put this into `pde2path` via (1.2) we thus need two boundary matrices q^h and q^v. For the horizontal boundaries ($y = \pm l_y$), q^h has diagonal $q^h_d = \begin{pmatrix} s & s & s & s & s & s & 0 & s & s & 0 \end{pmatrix}$. For the vertical boundaries ($x = \pm l_x$) q^v has diagonal $q^v_d = \begin{pmatrix} 0 & 0 & s & s & 0 & s & 0 & s & s & 0 \end{pmatrix}$ and additionally $q^v_{2,1} = s$, where $s = 10^3$ stands for the stiff spring constant. Positions 7 and 10 in q^h_d and q^v_d give the Neumann BC for u_7, u_{10}, while the top left 2×2 block $\begin{pmatrix} 0 & 0 \\ s & 0 \end{pmatrix}$ in q_v gives $\partial_n u_1 = 0$ via the first row and $u_2 = 0$ via the second row.

The (analytical) calculation of bifurcation points from $(v, w) = 0$ in case b) is rather tedious, see [30]. There, motivated by mode-jumping, the particular interest is in (the lowest) double bifurcation points, which yields $l = \sqrt{k(k+2)}$ with eigenfunctions

$$w_1(x, y) = \left(\frac{k+2}{k} \sin\left(k\frac{x}{l}\right) - \sin\left((k+2)\frac{x}{l}\right) \right) \sin(y)$$

resp.

$$w_2(x, y) = \left(\cos\left(k\frac{x}{l}\right) - \cos\left((k+2)\frac{x}{l}\right) \right) \sin(y)$$

(over the domain $[0, l\pi] \times [0, \pi]$). The first bifurcation is then obtained for $k = 1$, hence $l = \sqrt{3}$. The idea is to perturb l slightly which may lead to secondary bifurcations between branches coming originally from the same λ.

Putting all these ideas together we indeed get a secondary bifurcation between the first two primary branches, see Fig. 17. A number of further bifurcations from the trivial branch is also detected and can be followed. However, in the tutorial run `vkcmds` we use a rather coarse mesh with 1250 triangles, which should be refined before following higher bifurcations.

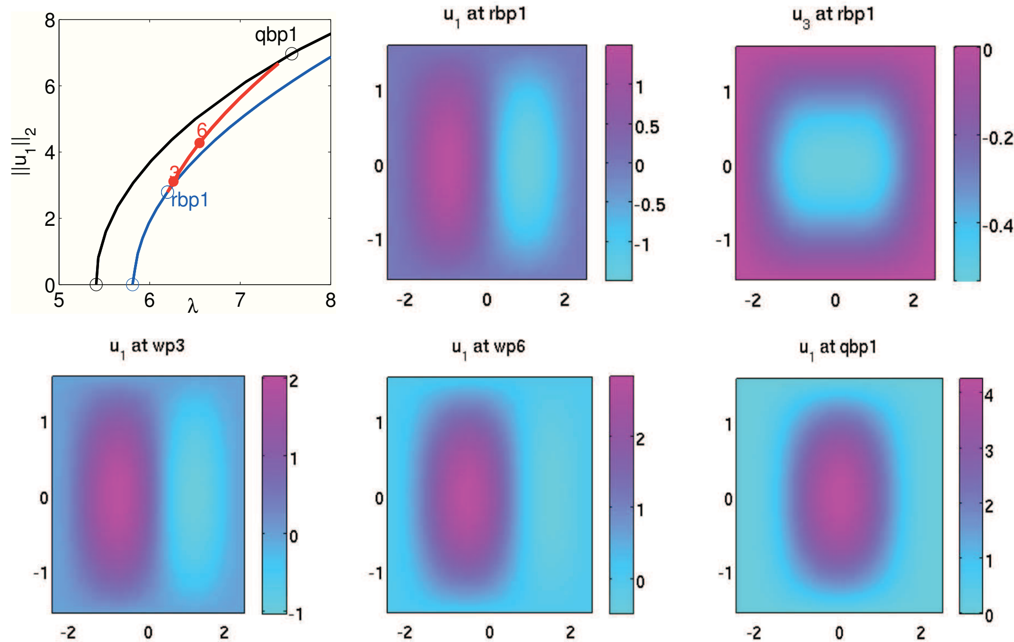

Figure 17: Secondary ("mode jumping") bifurcation (w-branch, red) in the (regularized) partially clamped plate (5.13). Bifurcation diagram and selected plots of u_1 and u_3. $l_y=\pi/2$, $l_x=4\pi/5$, regular mesh with 25×25 points (nt=1250 triangles). Error estimate $E(K)\approx 0.3$ at, e.g., rbp1. By mesh refinement we can obtain, e.g., $E(K)\approx 0.04$ with nt$\approx$ 15000. Then, however, a typical step takes a couple of minutes, where about 80% of the time is spent in blss or lss (standard setting). We expect that this can be optimized considerably, but here we content ourselves with the "proof of principle" setup for the 10 components system for (5.13).

6. Discussion

Clearly, numerical continuation and bifurcation analysis for 2D elliptic systems poses additional challenges compared to algebraic equations or 1D BVP, partly of course due to the more demanding numerics, but in particular also due to the typically very rich solution and bifurcation structure. With pde2path we believe to provide a *general* tool that works essentially out-of-the-box also for non-expert users and allows to start exploring such systems and the rich zoo of their solutions. Of course, in many respects this is just a first step, and probably the main entries on our **to-do-list** are:

1. Implement some more general (stationary) bifurcation handling, in particular bifurcation via non simple eigenvalues as these are quite ubiquitous in 2D systems due to various symmetries. The detection and localization of these is relatively straightforward by a modification of `findbif.m`, see Section 3.1.6, but the implementation of branch-switching will require additional work. Hopf bifurcations and similar phenomena will still be much more demanding.
2. Implement some (genuine) multi-parameter continuation. For instance, the bifurca-

tion to traveling waves generically requires a second parameter γ (the wave speed) to adapt, and consequently we need to further extend the "extended system" (2.1) by one more equation, the "phase condition". More generally, adding some constraints to (2.1) will also be useful to remove degeneracies as, e.g., the phase invariance in Section 5.1.

We believe that our set-up of pde2path is sufficiently modular and transparent such that this and similar adaptations will pose no implementation problems, but for now we confine ourselves to the basic one-parameter continuation and simple bifurcations.

Acknowledgments We thank Uwe Prüfert for providing his extension and documentation [27] of the pdetoolbox. Users familiar with AUTO will recognize that AUTO has been our guide in many respects, in particular concerning the design of the user interface. We owe a lot to that great software. We thank Tomas Dohnal for testing early versions of pde2path and providing valuable hints for making pde2path and this manual more user friendly, and we thank the anonymous referees for their helpful comments on the manuscript. JR acknowledges support by the NDNS+ cluster and the Complexity program of the Dutch Science Fund NWO, as well as his previous employer Centrum Wiskunde & Informatica (CWI), Amsterdam.

References

[1] E. Allgower and K. Georg, Numerical Continuation Methods, Springer, 1990.
[2] R. E. Bank, PLTMG, http://ccom.ucsd.edu/~reb/software.html.
[3] W.-J. Beyn, W. Kless, and V. Thümmler, *Continuation of low-dimensional invariant subspaces in dynamical systems of large dimension,* Fiedler, Bernold (ed.), Ergodic Theory, Analysis, and Efficient Simulation of Dynamical Systems, Springer Berlin, 47–72, 2001.
[4] D. Bindel, J. Demmel, and M. Friedman, *Continuation of invariant subspaces in large bifurcation problems,* SIAM J. Sci. Comput., 30(2) (2008), pp. 637–656.
[5] G. Bratu, *Sur les equations integrales non lineares,* Bull. Soc. Math. France, 42 (1914), pp. 113–142.
[6] C.-S. Chien, S.-Y. Gong, and Z. Mei, *Mode jumping in the von Kármán equations,* SIAM J. Sci. Comput., 22(4) (2000), pp. 1354–1385.
[7] E. J. Doedel, *Lecture notes on numerical analysis of nonlinear equations,* Krauskopf, Bernd (ed.) et al., Numerical Continuation Methods for Dynamical Systems, Path Following and Boundary Value Problems, 1–49, Springer, 2007.
[8] T. Dohnal and H. Uecker, *Coupled mode equations and gap solitons for the 2D Gross–Pitaevsky equation with a non-separable periodic potentia,* Phys. D, 238 (2009), pp. 860–879.
[9] G. Engeln-Müllges and F. Uhlig, Numerical Algorithms with *C*, Springer Berlin, 1996.
[10] B. Ermentrout, *XPP-Aut*, www.math.pitt.edu/~bard/xpp/xpp.htm.
[11] E. Doedel et al., *AUTO: continuation and bifurcation software for ordinary differential equations*, http://cmvl.cs.concordia.ca/auto/.
[12] K. Georg, *Matrix-free numerical continuation and bifurcation*, Numer. Funct. Anal. Optim., 22 (2001), pp. 303–320.
[13] J. Gervais, A. Oukit, and R. Pierre, *Finite element analysis of the buckling and mode jumping of a rectangular plate,* Dyn. Stab. Syst., 12(3) (1997), pp. 161–185.

[14] W. Govaerts, *MatCont*, `http://sourceforge.net/projects/matcont/`.

[15] W. Govaertsm, *Numerical methods for bifurcations of dynamical equilibria,* SIAM, 2000.

[16] M. Heil and A. L. Hazel, *oomph*, `http://oomph-lib.maths.man.ac.uk/doc/html/`.

[17] P. Hirschberg and E. Knobloch, *Mode interactions in large aspect ratio convection*, J. Nonlinear Sci., 7 (1997), pp. 537–556.

[18] D. Iron and M. J. Ward, *A metastable spike solution for a nonlocal reaction–diffusion model*, SIAP, 60(3) (2000), pp. 778–802.

[19] H. B. Keller, *Numerical solution of bifurcation and nonlinear eigenvalue problems,* App. Bifur. Theory Proc. Adv. Semin., Madison/Wis, 1976, 359–384, 1977.

[20] Y. A. Kuznetsov, Elements of Applied Bifurcation Theory, 3rd ed, Springer, 2004.

[21] V. M. Lashkin, *Two–dimensional multisolitons and azimuthons in Bose-Einstein condensates*, Phys. Rev. A, 77 (2008), 025602.

[22] V. M. Lashkin, E. A. Ostrovskaya, A. S. Desyatnikov, and Yu. S. Kivshar, *Vector azimuthons in two-component Bose-Einstein condensates*, Phys. Rev. A, 80 (2009), 013615–6.

[23] P. K. Maini, M. R. Myerscough, J. D. Murray, and K. H. Winters, *Bifurcating spatially heterogeneous solutions in a chemotaxis model for biological pattern formation*, Bull. Math. Biol., 53 (1991), pp. 701–719.

[24] H. D. Mittelmann, *Multilevel continuation techniques for nonlinear boundary value problems with parameter dependence,* Appl. Math. Comput., 19 (1986), pp. 265–282.

[25] J. D. Murray, Mathematical Biology, Springer-Verlag, Berlin, 1989.

[26] pde2path homepage, `www.staff.uni-oldenburg.de/hannes.uecker/pde2path`.

[27] U. Prüfert, *PDE Toolbox,* `tu-freiberg.de/fakult4/iec/mitarbeiter/profil/Uwe_Pruefert_ntfd`.

[28] K. C. Rose, D. Battogtokh, A. Mikhailov, R. Imbihl, W. Engel, and A. M. Bradshaw, *Cellular structures in catalytic reactions with global coupling*, Phys. Rev. Lett., 76 (1996), pp. 3582–3585.

[29] A. Salinger, *LOCA*, `http://www.cs.sandia.gov/LOCA/`.

[30] D. Schaeffer and M. Golubitsky, *Boundary conditions and mode jumping in the buckling of a rectangular plate,* Commun. Math. Phys., 69 (1979), pp. 209–236.

[31] F. Schilder and H. Dankowicz, *coco*, `http://sourceforge.net/projects/cocotools/`.

[32] J. Schnakenberg, *Simple chemical reaction systems with limit cycle behaviour*, J. Theoret. Biol., 81(3) (1979), pp. 389–400.

[33] R. Seydel, Practical Bifurcation and Stability Analysis, 3rd ed., Springer, 2010.

[34] L. Stollenwerk, S. V. Gurevich, J. G. Laven, and H.-G. Purwins, *Transition from bright to dark dissipative solitons in dielectric barrier gas-discharge,* Euro. Phys. J. D, 42 (2007), pp. 273–278.

[35] Matlab PDE Toolbox, online documentation.

[36] H. Uecker and D. Wetzel, *Numerical results for snaking of patterns over patterns in some 2D Selkov-Schnakenberg Reaction-Diffusion systems*, SIAM J. Appl. Dyna. Syst., to appear, 2014.

[37] R. Woesler, P. Schütz, M. Bode, M. Or-Guil, and H.-G. Purwins, *Oscillations of fronts and front pairs in two- and three-component reaction-diffusion systems*, Phys. D, 91 (1996), pp. 376–405.

[38] `www.sercoassurance.com/entwife/introduction.html`, 2001.