

On the numerical solution of a semilinear elliptic eigenproblem of Lane–Emden type, II: numerical experiments

F. J. Foss, Roland Glowinski, Ronald H. W. Hoppe

Angaben zur Veröffentlichung / Publication details:

Foss, F. J., Roland Glowinski, and Ronald H. W. Hoppe. 2007. "On the numerical solution of a semilinear elliptic eigenproblem of Lane–Emden type, II: numerical experiments." *Journal of Numerical Mathematics* 15 (4). <https://doi.org/10.1515/jnum.2007.013>.

Nutzungsbedingungen / Terms of use:

licgercopyright

Dieses Dokument wird unter folgenden Bedingungen zur Verfügung gestellt: / This document is made available under these conditions:

Deutsches Urheberrecht

Weitere Informationen finden Sie unter: / For more information see:

<https://www.uni-augsburg.de/de/organisation/bibliothek/publizieren-zitieren-archivieren/publiz/>



On the numerical solution of a semilinear elliptic eigenproblem of Lane–Emden type, II: Numerical experiments

F. J. FOSS, II*, R. GLOWINSKI* and R. H. W. HOPPE*

Received February 22, 2007

Received in revised form May 5, 2007

Abstract — In this second part of our two-part article, we present and discuss the corresponding numerical results from implementations of the numerical algorithms described in the first part. With these results, we observed that

- operator splitting applied to the associated time-dependent problem is suitable for solving only the first eigenproblem,
- among those tried, the perturbation and arclength continuation approach was the sole effective and robust approach for solving higher eigenproblems,
- on the eigenproblems for which (undamped or damped) Newton’s method converged, it was without question the most efficient.

Keywords: numerical method, Lane, Emden, semilinear, elliptic, eigenproblem, operator splitting, finite element, arclength continuation, least-squares, control, Newton’s method

1. Introduction

In this second part of our two-part article, we present the results of implementing the numerical algorithms detailed in the first part. For simplicity, we take the computational domain Ω to be a right triangle whose sides have the ratios 3:4:5 (a ‘Fermat triangle’), anticipating that with such a symmetry-breaking choice, all computed eigenpairs will be simple (although we do not attempt to prove this). Figure 1 shows this domain along with the finest (uniform) triangulation used in computations.

Starting with the principal eigenproblem, in Section 2 we present numerical results from the operator splitting method applied to the time-dependent problem discussed in Part I, Section 2. Next, in Section 3, we present the numerical results

*Department of Mathematics, University of Houston, 4800 Calhoun Rd, Houston, TX 77204-3008
This work was supported by NSF grant DMS 0412267.

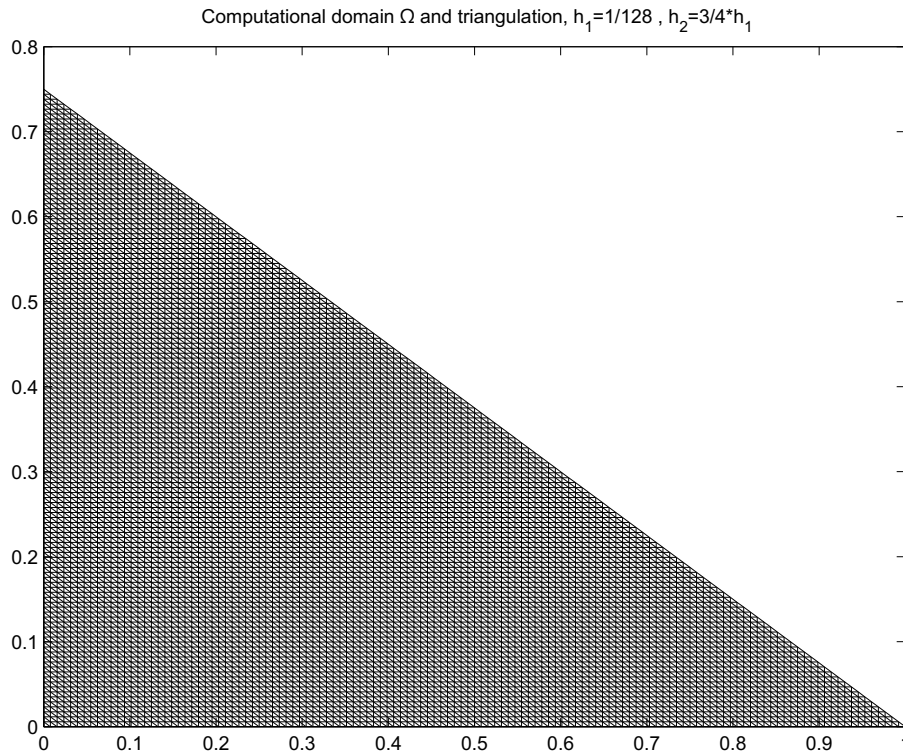


Figure 1. Domain Ω with finest triangulation.

from the perturbation and arclength continuation methods discussed in Part I, Subsection 3.1. Finally, in Section 4, we present some numerical results from the application of Newton's method discussed in Part I, Subsection 3.2.

2. The principal eigenproblem

Figure 2 shows the solutions of the initializing linear principal eigenproblem (2.21)–(2.23) (top) and the nonlinear principal eigenproblem (1.1)–(1.3) (bottom). Figure 3 shows the convergence results with an initial guess of $\{u^0, \lambda_0^1\} = \{w_1, 0\}$ (top), the evolution of the $H_0^1(\Omega)$ norms of the minimizing sequence elements generated by the computational scheme (middle), and graphs of the *a posteriori* necessary upper bounds on τ computed using conditions (2.30) (solid) and (2.32) (dashed), plotted as functions of the time step n (bottom). From these figures, we can see that the reduction in norm occurs mostly before about the 70th iteration. Notice also that the former condition on τ is always more restrictive. Attempts to use even a slightly larger τ than the smallest necessary condition value resulted in failure.




Figure 2. Principal linear eigenpair $\{w_1, \mu_1\}$ solving problem (2.21)–(2.23) (top) and nonlinear eigenpair $\{u_1, \lambda_1\}$ solving problem (1.1)–(1.3) (bottom).




Figure 3. Convergence results with an initial guess of $\{u^0, \lambda_0^1\} = \{w_1, 0\}$ (top), evolution of the $H_0^1(\Omega)$ norm of the minimizing sequence elements (middle), and *a posteriori* necessary upper bounds on operator splitting time step τ (bottom).

3. Higher eigenproblems

Despite the relative ease with which we were able to reproduce results for some validation problems, such as the classical *Bratu problem* (cf. [4,5]) and a model problem with a $\sin u$ nonlinearity from the NETLIB software package PLTMG (cf. [2]), our initial experiences with the pseudo-arclength continuation computational framework applied to the perturbed formulation (3.18) of our main problem were rather puzzling. Since we already found a solution of the 2nd (unperturbed) nonlinear eigenproblem using other methods, it was natural and desirable (as an independent verification of these earlier results) to first attempt to find the nontrivial solution branch for the 2nd perturbed nonlinear eigenproblem with the ultimate goal (after sufficient continuation) of ‘jumping’ to the corresponding unperturbed branch. Initial efforts to continue this nontrivial solution branch met with failure.

To explain this failure, we refer to Fig. 4. Each beginning branch segment shown in this figure was initialized with the trivial branch point $\{u_0, \lambda_0\} = \{0, \mu_2/\delta\}$, where μ_2 is the 2nd eigenvalue solving the linear eigenproblem (2.21)–(2.22) (for simplicity we fixed the perturbation parameter $\delta \equiv 1$). Three choices for the initial tangent were used in turn. Because the nontrivial nonlinear solution branches for all eigenmodes are symmetric about the trivial branch, our first and theoretically-consistent choice for the initial tangent was $\{\dot{u}_0, \dot{\lambda}_0\} = \{w_2/\sqrt{\mu_2}, 0\}$. This choice produced the dashed curve shown in Fig. 4. The computed solution branch is returning to the trivial solution branch! Initially thinking that this unfortunate development had something to do with the fact that we were initializing with $\dot{\lambda}_0 = 0$, we next derived a first-order approximation (in Δs) of the solution along the nontrivial branch, which turned out to be given by $\{u_1^0, \lambda_1^0\} \approx \{w_2/|w_2|_{1,2}, \mu_2/\delta - \Delta s/(\delta^2|\Omega|)\} \Delta s$. We used this approximation in the correction step and for the first-order tangent approximation (3.11) at the corrected solution branch point. The resulting computed solution branch essentially retraced the one resulting from our first tangent choice (dotted curve overlaying the dashed curve in Fig. 4). The third initial tangent was chosen using (3.24) with $n = 2$, taking $\dot{\lambda}_0 = -1/\sqrt{1 + \mu_2}$ which in turn gives $c_2 = \pm \dot{\lambda}_0$. This choice produced the dashed-dotted curve shown in Fig. 4. Again, the computed solution branch eventually starts returning to the trivial solution branch.

The remedy for this undesirable behavior turned out to be the renormalization of the tangent pair $\{\dot{u}_1, \dot{\lambda}_1\}$ (via solution of the Davidenko equations (3.28)) when, after enough updates using the tangent approximation (3.11), a simple monitor, testing whether or not the computed arclength fell within a small neighborhood of unity, was violated. The manifestation of this continuation with the tangent approximation followed by tangent renormalization is shown in the solid curve of Fig. 4. This computed solution branch was initialized as with the dashed curve (so they initially overlay), but at some point the tangent renormalization criteria was satisfied and the renormalization performed, giving the subsequent abrupt change in the direction of the solution branch seen in the figure. Subsequent renormalizations were required as well, although the directional corrections were substantially less dramatic than the first. This renormalization was enough to admit the following of the nontrivial

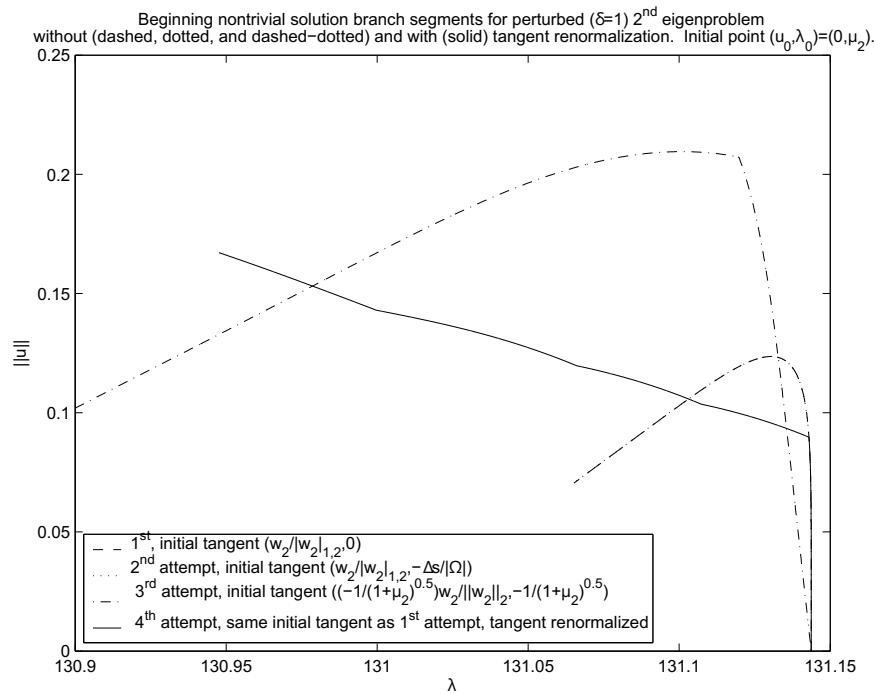


Figure 4. Beginning solution branch segments for 2nd perturbed nonlinear eigenproblem resulting from three different initializations without (—, ···, and — · —) and one with (—) tangent renormalization.

solution branch instead of returning to the trivial branch.

After resolving the mystery of the computed nontrivial solution branch returning to the trivial branch, we successfully continued nontrivial solution branches for the perturbed ($\delta = 1$) eigenproblem (3.18) as well as the ‘jump’ to the corresponding unperturbed ($\delta = 0$) solution branches for seven eigenproblems corresponding to the eigenpairs $\{u_n, \lambda_n\}$, $n \in \{1, 2, 3, 4, 5, 10, 20\}$. For brevity, we only include the results for the 3rd and 20th eigenproblems. Since none of the other methods attempted were successful at finding the 3rd eigenpair, we focus our discussion on these results shown in Figs. 5–7. The top graphic of Fig. 5 shows the continued branch $\{(\lambda_1, |u_1|)_n\}$, $n = 0, 1, \dots$, of perturbed solutions as well as the ‘jump’ to the corresponding unperturbed branch for the 3rd semilinear eigenproblem. The bottom graphic shows the corresponding conjugate gradient convergence history, i.e., the number of CG iterations required for convergence to a solution on the branch at each step of the continuation.

Since the initializing trivial solution along the solution branch of the 3rd perturbed eigenproblem is rather uninteresting, we show for comparison purposes the linear eigenpair $\{w_3, \mu_3\}$ solving (2.21)–(2.23) in the top graphic of Fig. 6. We then

show, in the bottom graphic of this same figure, the last solution obtained via continuation along the perturbed ($\delta = 1$) solution branch, starting from the bifurcation point $\{0, \mu_3\}$, where μ_3 is the 3rd linear eigenvalue solving the linear eigenproblem (2.21)–(2.22). The top graphic of Fig. 7 shows the unconstrained eigenpair $\{u_3, \lambda_3\}$ obtained by ‘jumping’ from this last perturbed solution to the unperturbed ($\delta = 0$) solution branch. The bottom graphic of this figure shows the unit $L^4(\Omega)$ norm constrained eigenpair $\{u_3, \lambda_3\}$ solving the original constrained semilinear eigenproblem (1.1)–(1.3).

Finally, Figure 11 superimposes the continued perturbed, ‘jump’, and continued unperturbed branches for all eigenmodes investigated with arclength continuation. This shows the relative orientation of the solution branches for all eigenmodes revealed in this study. Additionally, the ‘ \times ’s mark the unit $L^4(\Omega)$ norm constrained solution points along the unperturbed branches.

It is important to note that, compared with other solution methods, the perturbation and arclength continuation approach proved to be the most computationally reliable, predictable, and effective in capturing eigenmodes of the model problem, particularly the higher ones. In fact, at the expense of some global efficiency (locally, at each continuation step, it is extremely efficient), it was the only method we tried that was able to capture eigenmodes beyond the 2nd one.

To complete the analysis and discussion of the results obtained thusfar, we include Fig. 12, each graphic of which highlights, with contrasting symbols, the regions in the domain Ω within which the respective eigenfunctions have one sign. The boundaries between these regions constitute the so-called *nodal sets*, defined as the closure of the sets of interior zero crossings of the eigenfunctions. The geometry of these nodal sets within Ω as well as their proximity to $\partial\Omega$ for various domain geometries continue to attract a fair amount of attention (cf. [1]). It is evident from the figures that, for the eigenfunctions computed herein, the nodal sets that intersect $\partial\Omega$ do so orthogonally (within computational resolution).

4. Newton’s method results

Our experience with the implementation of the damped Newton algorithm discussed in Part I, Subsection 3.2, as applied to problem (1.1)–(1.3), proved rather fascinating despite being only partially successful. Specifically, we found that although we were able to successfully solve for the first two nonlinear eigenpairs with breathtaking efficiency, we failed in our many attempts to solve for the 3rd nonlinear eigenpair. Even the successful outcomes in the cases of the first two eigenpairs were obtained through nothing short of pure luck in ‘stumbling on’ good initial guesses $\{u^0, \lambda^0\}$. We were not so lucky in the case of the 3rd eigenpair. Indeed, it was our experience that the performance of the method was extremely sensitive to the choice of initial guess, even among the ‘natural’ although relatively small set of choices provided by the one-dimensional linear eigenmanifolds. Other than the trivial solution manifold, which is useless as a source of initial guesses, the linear eigenmanifolds are really the only ‘natural’ sources of initial guesses that readily present themselves




Figure 5. Continued perturbed solution branch with jump to unperturbed branch (top) and conjugate gradient convergence behavior (bottom) for 3rd semilinear eigenproblem.

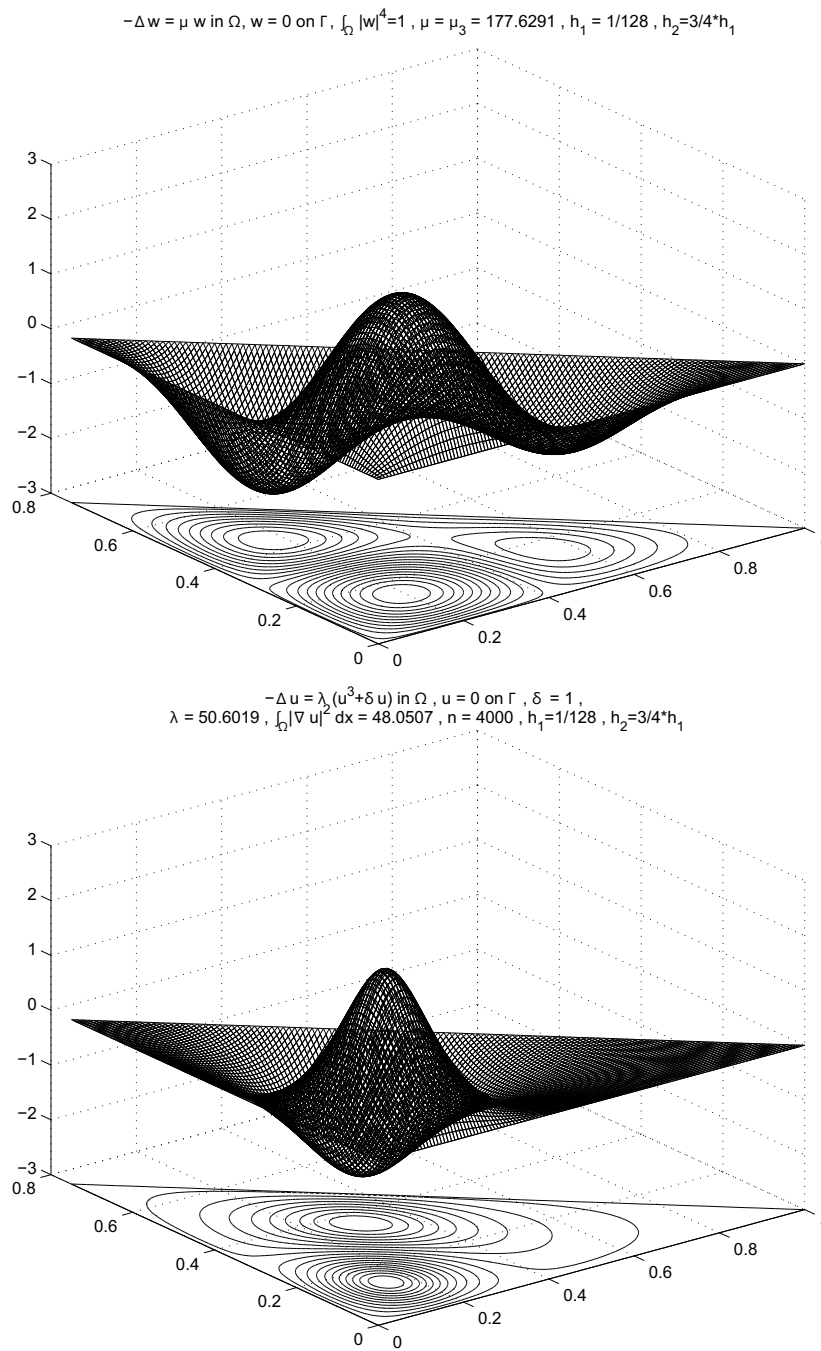


Figure 6. 3rd eigenpair $\{w_3, \mu_3\}$ solving the linear eigenproblem (2.21)–(2.23) (top) and continued 3rd eigenpair $\{u_3, \lambda_3\}$ solving the perturbed ($\delta = 1$) semilinear eigenproblem (3.18)–(3.19) (bottom).

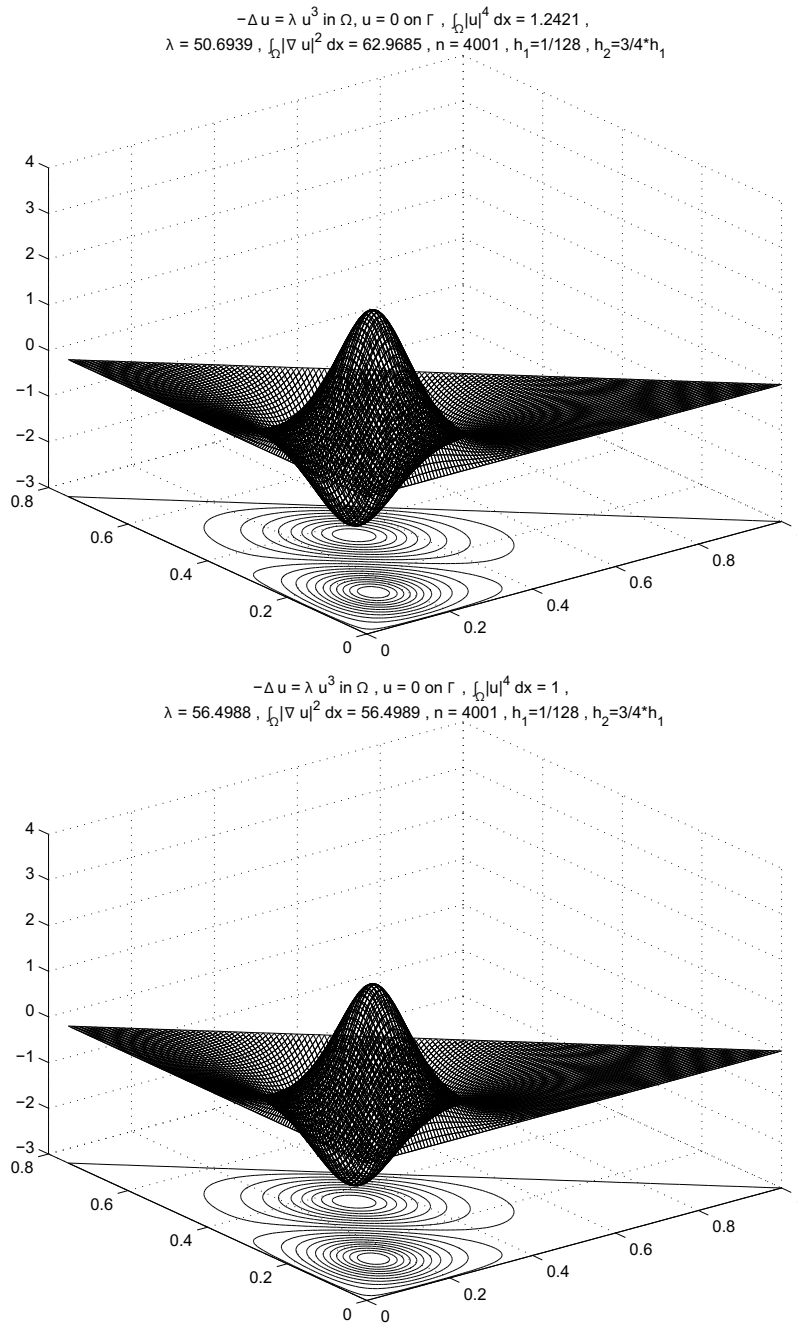


Figure 7. Continued 3rd eigenpair $\{u_3, \lambda_3\}$ solving the unperturbed ($\delta = 0$) semilinear eigenproblem (3.18)–(3.19) (top) and normalized 3rd eigenpair $\{u_3, \lambda_3\}$ solving the original constrained semilinear eigenproblem (1.1)–(1.3) (bottom).

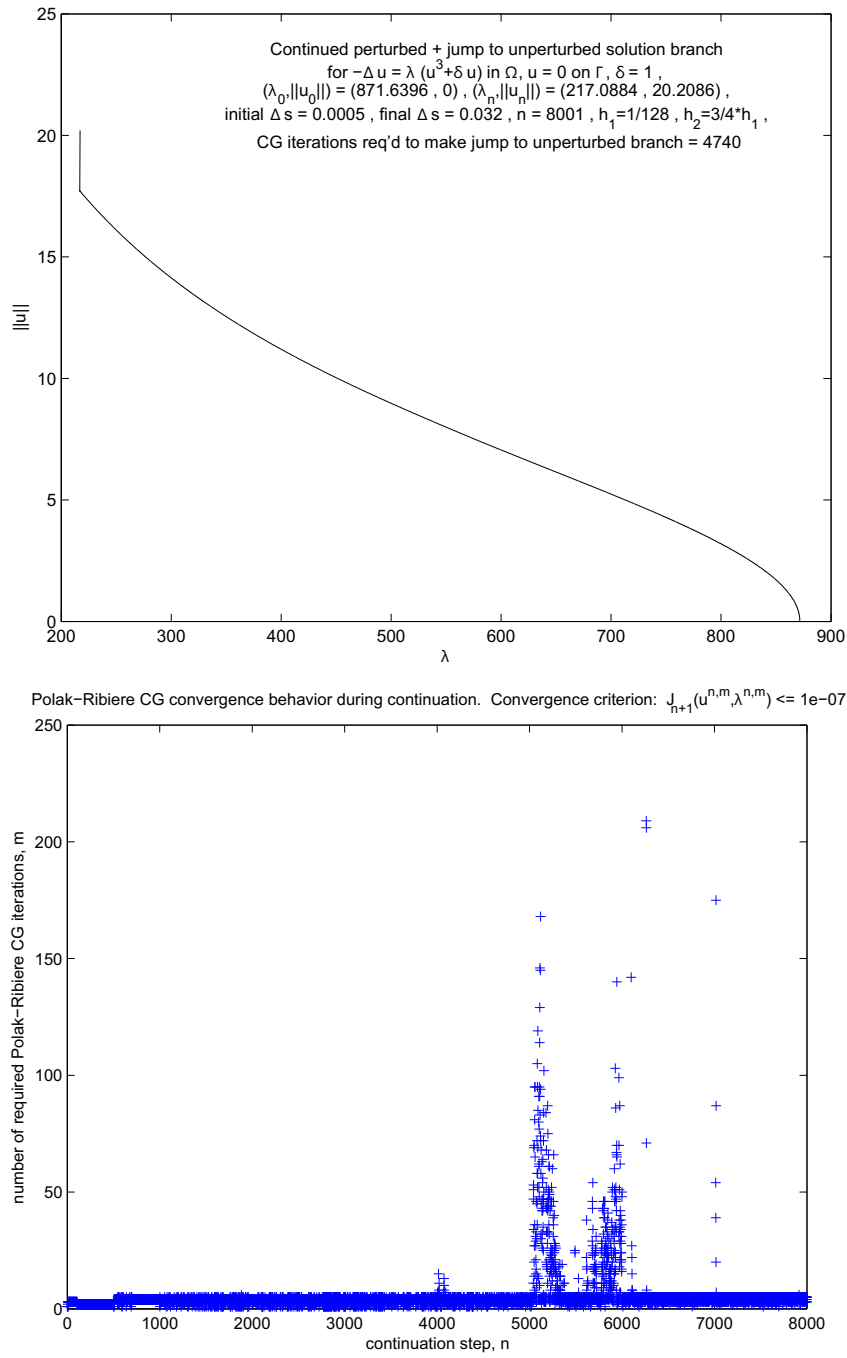


Figure 8. Continued perturbed solution branch with jump to unperturbed branch (top) and conjugate gradient convergence behavior (bottom) for 20th semilinear eigenproblem.

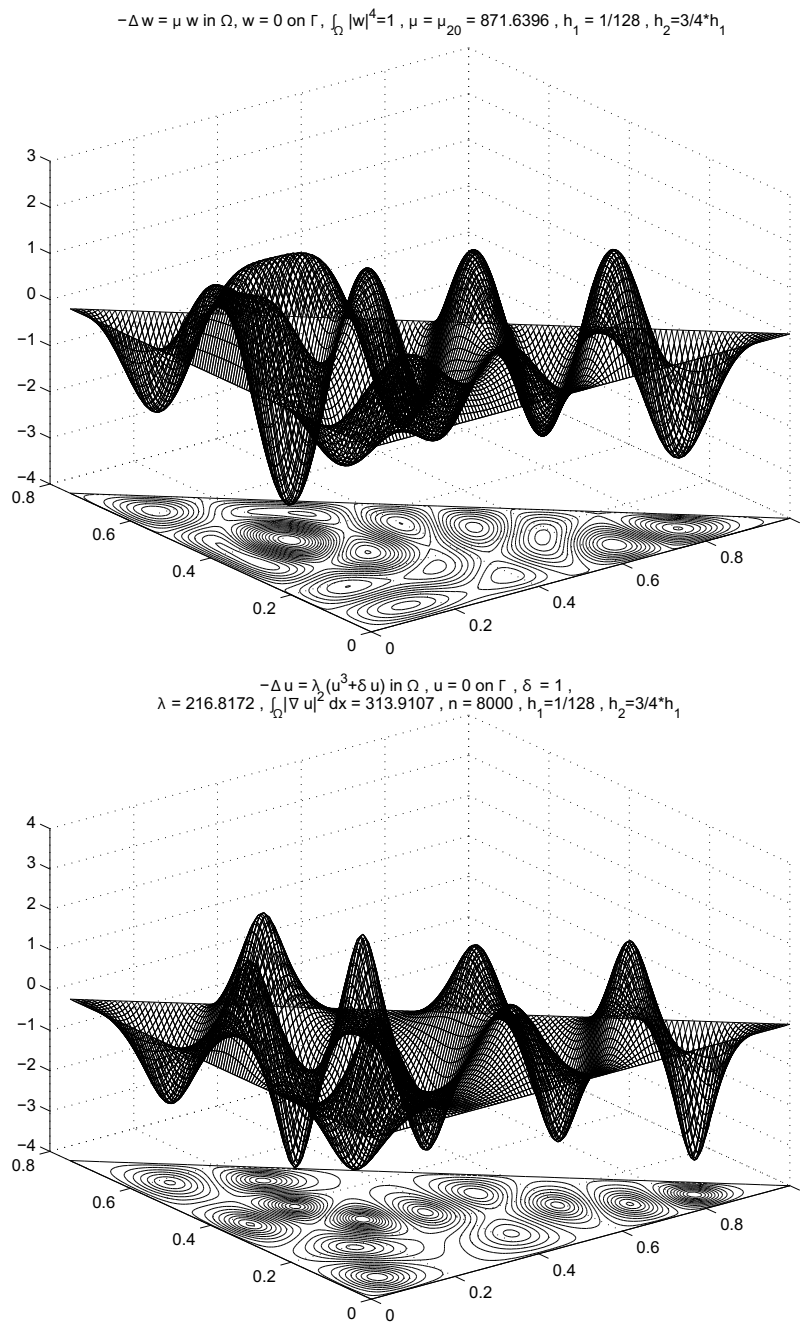


Figure 9. 20th eigenpair $\{w_{20}, \mu_{20}\}$ solving the linear eigenproblem (2.21)–(2.23) (top) and continued 20th eigenpair $\{u_{20}, \lambda_{20}\}$ solving the perturbed ($\delta = 1$) semilinear eigenproblem (3.18)–(3.19) (bottom).

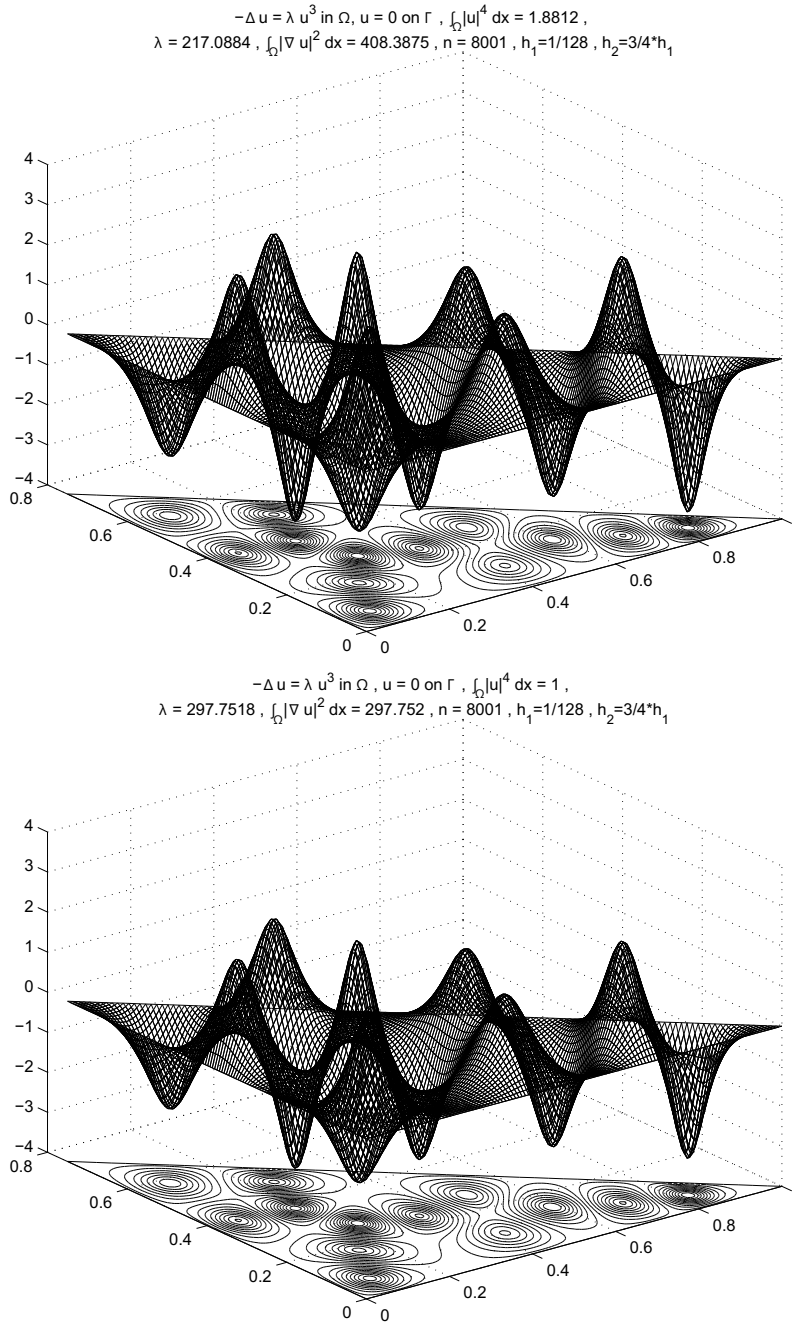


Figure 10. Continued 20th eigenpair $\{u_{20}, \lambda_{20}\}$ solving the unperturbed ($\delta = 0$) semilinear eigenproblem (3.18)–(3.19) (top) and normalized 20th eigenpair $\{u_{20}, \lambda_{20}\}$ solving the original constrained semilinear eigenproblem (1.1)–(1.3) (bottom).




Figure 11. Continued perturbed solution branches with jump to, and additional continuation along (excluding 1st and 20th), unperturbed branches for 1st–5th, 10th, and 20th semilinear eigenproblems. ‘x’ marks the solution points of the respective unit $L^4(\Omega)$ norm constrained semilinear eigenproblems.

a priori in this case. If ‘good’ initial guesses cannot be found easily within these sets, then without further modification (some suggestions for which can be found in [3], Chap. 8), the method is of limited practical use for this particular semilinear problem.

To support this empirical assessment, we provide computational results for the 1st–3rd eigenproblems. We first applied the method to, and thus focus our discussion on, the 2nd eigenproblem. After many unsuccessful tests of damped (using various damping factors) and undamped Newton’s methods, taking the eigenpair $\{w_2, \mu_2\}$ solving the unit $L^4(\Omega)$ norm constrained linear eigenproblem (2.21)–(2.23) as initial guess (the most ‘natural’ choice), we literally stumbled upon a ‘good’ initial guess by scaling the linear eigenfunction by a factor of 0.5. With this lucky scaling, the method converged in 12 iterations without damping and 9–10 iterations with damping (depending on the damping factor used)! Compared with the overall convergence rates of any of the other previously used methods, this is a truly remarkable benefit of either method.




Figure 12. Signed regions of 3rd (top) and 20th (bottom) eigenfunctions solving the original constrained semilinear eigenproblem (1.1)–(1.3). '+' marks regions where $u > 0$ and '.' marks regions where $u < 0$.

The costs of this remarkable performance come in the forms of (1) a fairly narrow window of choices for ‘good’ initial guesses as well as (2) extreme sensitivity of the behavior of the Newton sequence to the choice of initial guess. As we mentioned, we got lucky with the choice of $\{0.5w_2, \mu_2\}$ as our second initial guess, which resulted in the spectacular convergence behavior indicated above for the corresponding damped and undamped Newton sequences. However, the sensitivity of the Newton sequence behavior to a simple scaling of the linear eigenfunction initial guess is extreme. Without damping, convergence occurs when the scale factor falls within a rather small neighborhood of 0.5, e.g., one just containing the interval $[0.4875, 0.50625]$. With damping, this neighborhood is broadened to one just containing $[0.39375, 0.525]$. Furthermore, the delineation of these intervals of ‘good’ scale factors was only possible because of our good fortune in guessing the value 0.5 within the intervals to begin with as opposed to a systematic or algorithmic search. Without a ‘good’ initial guess, the damped Newton method stagnates without converging (due to the ultimate failure of the monotonicity test even with vanishingly small damping factors) and the undamped Newton method simply diverges. Unfortunately, ‘good’ initial guesses for the 3rd nonlinear eigenproblem were very elusive indeed (we did not find any).

Figures 13–15 show the results for the 2nd eigenproblem. Figure 13 shows several converging Newton sequences $(|u^k|, \lambda^k)$ (note that λ^k is plotted on the horizontal axis and $|u^k|$ on the vertical axis, marked by ‘×’s), along with the corresponding undamped (dotted line) and several damped (solid line) approximate Newton path segments connecting the iterates. A circle marks the solution point $(|u^*|, \lambda^*)$. The top graphic shows the entire sequences while the bottom graphic shows a zoomed subset in a neighborhood of the solution point. The intersection of the dotted and solid segments marks the initial guess $(|u^0|, \lambda^0)$. Note the dramatic difference between the undamped and damped approximate Newton paths. The damping has the effect of changing both the magnitude and direction of the first and subsequent Newton increments to improve convergence. In the undamped case, it is interesting that the first Newton step, taken from the initial guess within the linear eigenmanifold, is consistently to a point further away from the solution point, from which subsequent iterates then approach the solution point. This already suggests that the linear eigenmanifold may be a problematic source of initial guesses, but what other ‘natural’ choices are there?

Figure 14 shows the damped Newton solution $\{u^*, \lambda^*\}$ of the 2nd semilinear eigenproblem (1.1)–(1.3) (bottom) compared with the initial guess $\{u^0, \lambda^0\} = \{0.5w_2, \mu_2\}$ (top), where $\{w_2, \mu_2\}$ is the 2nd eigenpair solving the unit $L^4(\Omega)$ norm constrained linear eigenproblem (2.21)–(2.23). The solution agrees with that obtained using the arclength continuation method.

Figure 15 shows converging undamped (top) and various damped (bottom) Newton sequences resulting from various initial guesses generated by different scalings αw_2 of the 2nd linear eigenfunction w_2 together with the linear eigenvalue μ_2 , i.e., various points in the 2nd linear eigenmanifold $\{\{\alpha w_2, \mu_2\} | \alpha \in \mathbb{R}\}$. With such scalings, the initial guesses have $H_0^1(\Omega)$ norms scatter along the vertical line $\lambda = \mu_2$,

as can be clearly seen in the damped case and barely seen in the undamped case. In the damped case, solid Newton sequences converge while dotted ones do not. In the undamped case, only solid converging Newton sequences are plotted (plotting diverging ones spoil the axes scaling for the converging ones). Note that we attempted to find nearly extremal scalings admitting convergence so as to isolate windows of ‘good’ initial guesses for both the damped and undamped methods. As can be seen in the graphics, and as discussed previously, the ‘good’ initial guess window in the undamped case is extremely narrow indeed, with some significant widening in the damped case.

Although not shown here, we repeated our experiments using both the damped and undamped Newton methods on the 1st nonlinear eigenproblem, anticipating and confirming that it performs even better in this case. Not only was the solution of the 1st nonlinear eigenproblem found by both damped and undamped Newton methods in fewer iterations than that of the 2nd, the windows of ‘good’ initial guesses within the 1st linear eigenmanifold are significantly wider in both cases.

Unfortunately, despite a fair amount of effort (even with the benefit of knowing the desired outcome in advance thanks to arclength continuation), we failed to produce a damped or undamped Newton sequence converging to the solution of the 3rd semilinear eigenproblem, as can be seen by examining Fig. 16 showing the stagnating damped Newton sequences resulting from a large number of initial guesses in the 3rd linear eigenmanifold. These represent only some of the effort made to find a damped Newton sequence converging to the solution point (marked with a circle). Even the stagnation point for one of the best results obtained (i.e. nearest to the solution point) is still a poor approximation of the actual solution produced via arclength continuation.

In conclusion, based on the experience with the damped and undamped Newton’s methods in this case, it would appear that despite the outstanding performance of these methods when ‘good’ initial guesses are used, the methods are not practical if such initial guesses are difficult to find. This is somewhat disappointing, since the purpose of employing a damped Newton method is to globalize the ordinary Newton method in the sense of removing any restrictions on the initial guess. Further investigation as to why the globalization fails in this case could provide valuable insight into both the methods themselves as well as the current problem to which they are applied.

5. Conclusions

We have presented some numerical techniques for solving a particular semilinear elliptic eigenproblem, each with its own strengths and weaknesses with respect to algorithmic complexity, robustness for solving higher eigenproblems, and computational efficiency.

Operator splitting applied to the associated time-dependent problem proved effective and fairly efficient for solving the principal (1st) eigenproblem.

For higher eigenproblems, we have shown that, among the methods used herein,

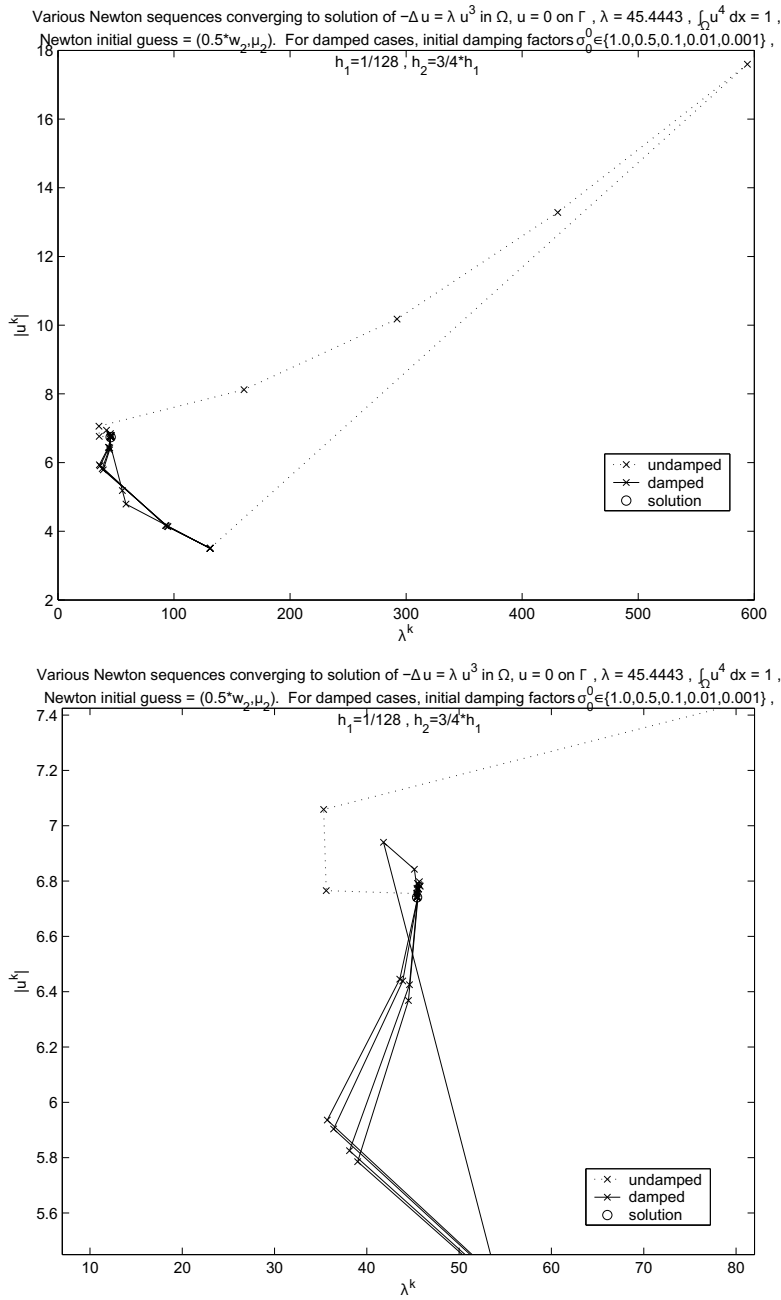


Figure 13. Various undamped and damped Newton sequences converging to the 2nd eigenpair $\{u_2, \lambda_2\}$ solving the original semilinear eigenproblem (1.1)–(1.3) resulting from the initial guess $\{0.5w_2, \mu_2\}$ (top), where $\{w_2, \mu_2\}$ solves the linear eigenproblem (2.21)–(2.23), along with a zoom of same in a neighborhood of the solution point (bottom).

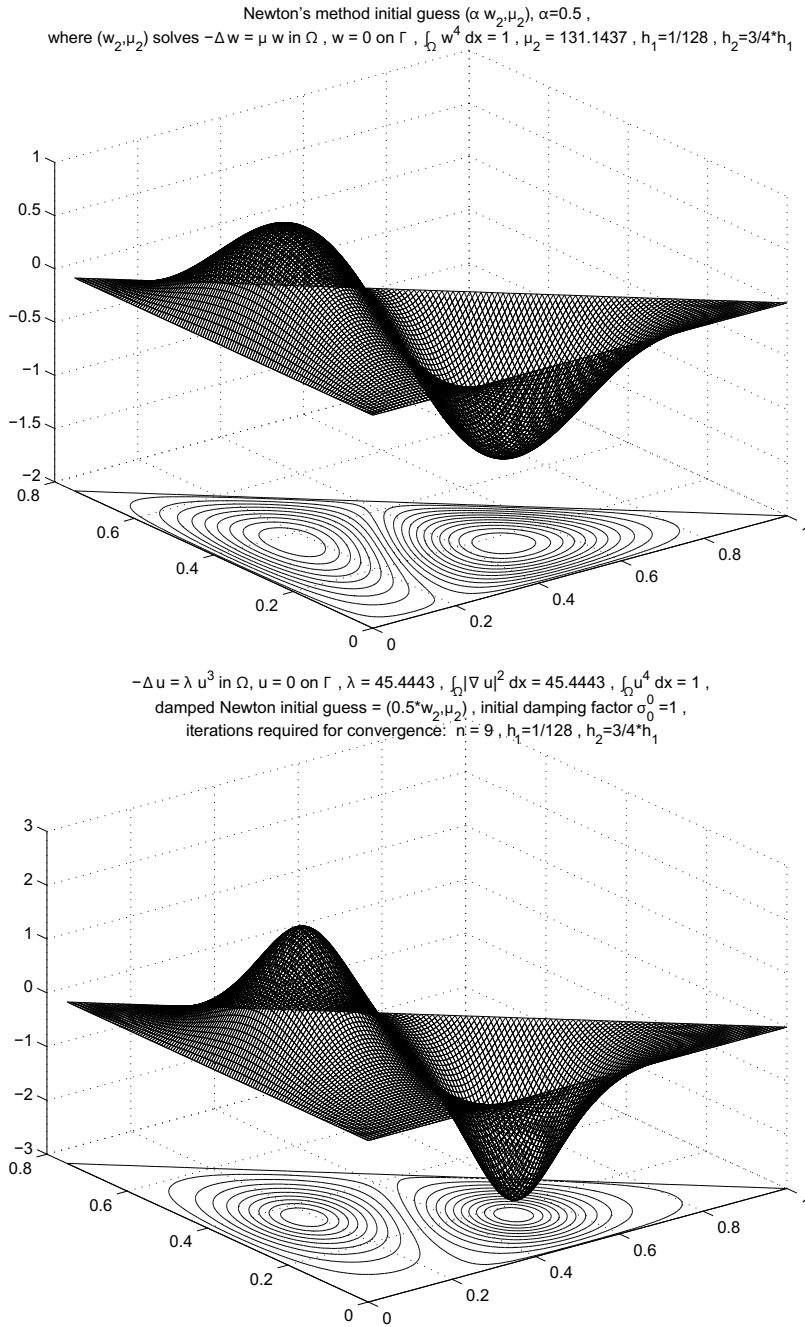


Figure 14. Comparison of 2nd eigenpair $\{u_2, \lambda_2\}$ solving the original semilinear eigenproblem (1.1)–(1.3) (bottom) resulting from a Newton's method initial guess $\{0.5w_2, \mu_2\}$ (top), where $\{w_2, \mu_2\}$ solves the linear eigenproblem (2.21)–(2.23).

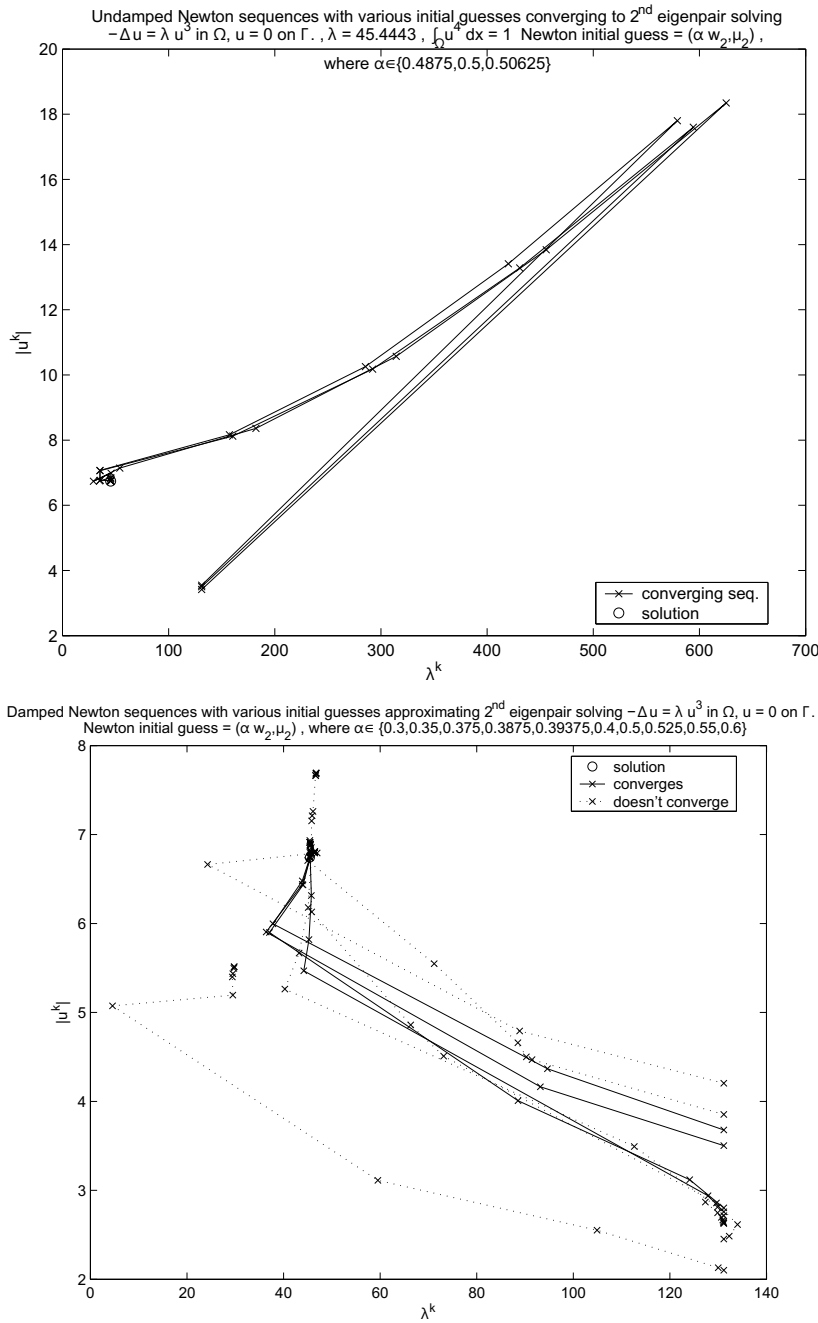


Figure 15. Converging undamped (top) and various damped (bottom) Newton sequences resulting from initial guesses $\{\alpha w_2, \mu_2\}$, where $\{w_2, \mu_2\}$ solves the linear eigenproblem (2.21)–(2.23), for various $\alpha \in [0.3, 0.6]$.

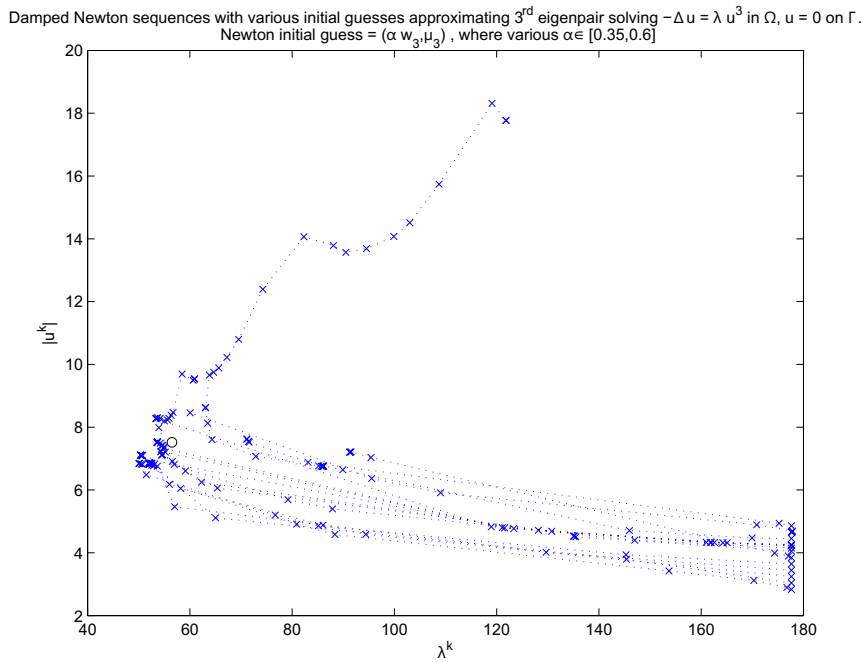


Figure 16. Various stagnating damped Newton sequences resulting from initial guesses $\{\alpha w_3, \mu_3\}$, where $\{w_3, \mu_3\}$ solves the linear eigenproblem (2.21)–(2.23), for various $\alpha \in [0.35, 0.6]$.

the most robust in terms of resolving the most eigenmodes is the so-called arclength continuation method. We successfully applied this method to a perturbation of the original unconstrained nonlinear eigenproblem for the purpose of generating segments of nontrivial solution branches bifurcating from a trivial solution branch whose endpoints serve to initialize ‘jumps’ to corresponding solution branches of the original eigenproblem. To increase the overall efficiency of the method, the arclength continuation process can be optimized via automatic updating of the arclength stepsize Δs and tangent renormalization, based on second-order information in $S(\cdot, \cdot)$ and the conditioning of the Davidenko system (cf. [6, pp. 88–89]), and/or adjusting/continuing in the perturbation parameter δ so that the jump to the unperturbed branch can be done after a lesser number of continuation steps.

Finally, we have demonstrated that an error-oriented, or affine covariant, implementation of Newton’s method, undamped or damped, applied directly to the original constrained nonlinear eigenproblem is extremely effective and efficient for finding the first two eigenmodes, but without further modification fails, for example, to find the 3rd eigenmode. We showed that the key ingredient in the successful cases is proper initialization.

Acknowledgments. For inspiring this work, we would like to thank Mónica Clapp.

For their interest in, and comments on, various aspects of this work, we would like to thank Herb Keller and Peter Deuffhard. Finally, for funding this work, we would like to thank NSF.

References

1. A. Aftalion and F. Pacella, Qualitative properties of nodal solutions of semilinear elliptic equations in radially symmetric domains. *Tech. Report No. 1*, C. R. Acad. Sci. Paris, 2004.
2. R. E. Bank, *PLTMG: A Software Package for Solving Elliptic Partial Differential Equations*. Department of Math., University of California at San Diego, La Jolla, CA, 2004.
3. P. Deuffhard, *Newton Methods for Nonlinear Problems: Affine Invariance and Adaptive Algorithms*. Number 35 in Springer Series in Comp. Mathematics, Springer-Verlag, Berlin-Heidelberg, 2004.
4. R. Glowinski, H. B. Keller, and L. Reinhart, Continuation-conjugate gradient methods for the least squares solution of nonlinear boundary value problems. *SIAM J. Sci. Stat. Comp.* (1985) **6**, No. 4, 793–832.
5. J. W. He and R. Glowinski, Neumann control of unstable parabolic systems: Numerical approach. *J. Optim. Theory Appl.* (1998) **96**, No. 1, 1–55.
6. H. B. Keller, Global homotopies and newton methods. In: *Recent Advances in Numerical Analysis* (Eds. C. de Boor and G. H. Golub), Math. Research Center, University of Wisconsin at Madison, Academic Press, Inc., 1978, No. 41, pp. 73–94.