Professional statement — Alex H. Barnett, October 2016

Numerical simulation and data analysis have always been key to scientific and technological progress. Their role is growing, but so is the scale of complexity of the problems. This motivates my work as an applied mathematician to develop the computational algorithms that will enable the challenging problems of today and tomorrow to be solved efficiently. My main area is solving the partial differential equations of wave propagation and viscous fluid flow, which have applications to acoustics, electromagnetics (communications, internet switches, radar, imaging), quantum physics, and the design of devices such as solar cells and microfluidic channels that sort cancerous from healthy cells. I also create algorithms for reconstructing underlying signals from noisy measurements, in bio-medical arenas such as electrical recording inside the brain, and protein imaging. My work, and that of my group, spans analysis (proving theorems about numerical methods), to the invention and testing of new numerical algorithms, to interdisciplinary collaborations (e.g. with neuroscientists, engineers, and pure mathematicians). I release documented software implementations used by others—I believe that this is crucial for dissemination of ideas, and for scientific progress and reproducibility. In the period 2011–2016, I will have given 36 invited research talks, received external research funding from the NSF for three years, and published or submitted at least 19 new papers, almost all in top-level international journals.

As an educator I strive to ignite a passion for "doing mathematics:" this means not only understanding, but also contributing to, possibly the most beautiful edifice ever created by humans. It also means getting hooked on problem solving, especially solving numerical problems via the smart use of computer code. Students also leave my classes appreciating the huge impact of mathematics in today's world. As a mentor, I have guided several young researchers into research careers, and work to build a thriving applied mathematics community on campus and internationally. In this statement I overview my academic contributions over the last six years.

1 Scholarship

Background for the non-specialist. A partial differential equation (PDE) relates the *rates of change* of fields in space, such as their curvature, to their values. Each PDE has its own character: the Laplace and Stokes equations have smooth solutions away from obstacles, whereas the Helmholtz equation, while also smooth on the smallest scale, can oscillate rapidly like a wave, as in Fig. 1a. If there are many oscillations, this is called *high frequency*.

My approach to computational PDEs often involves *integral equations*, since they enable the fastest and most accurate solutions in a variety of applications, in particular when the medium is *uniform* in each of its geometric regions. Examples include radio or light waves traveling in air then bouncing off a metal obstacle, or

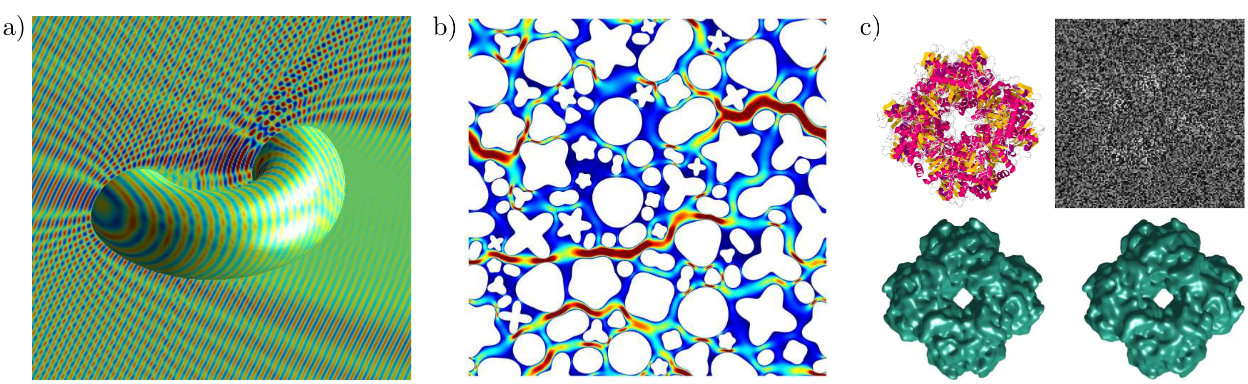

Figure 1: a) 3D acoustic scattering from a torus 50 wavelengths across, computed via QBX ($N = 145000$ unknowns, 5-digit accuracy, 2 hours). b) Pressure-driven viscous fluid flow through 2D doubly-periodic array of no-slip particles—one unit cell is shown, containing 100 particles ($N = 80000$, 12-digit accuracy) [38]. c) Cryo-electron microscopy: a known protein (top left), simulated density isosurface (bottom left), one of 50000 simulated noisy projections at SNR 0.1 (top right), reconstruction from noisy images (bottom right, $N = 67000$ unknowns, 2 hours) [36].

the flow of blood (a viscous fluid containing vesicles) through a capillary. An everyday example of an integral equation is as follows. Consider when you look at the outside world through a frosted bathroom window: the blurred scene you see is the true scene acted on by an *integral operator* (here the optics of the glass). Solving the integral equation means, given only this blurred scene, finding the (unknown) original scene in all its clarity. Obviously, in this case, such an equation is very hard to solve! (this is an example of a "first kind" equation). Now, when applied to solving PDEs, what we called "scenes" are, mathematically speaking, functions that instead live on the geometric boundaries (surfaces). A good physical analogy becomes: the unknown function is an electric charge "sprayed" onto the surface (as when you rub a balloon against your sweater in dry weather), and the integral operator now generates the surrounding electrostatic field. (Such a field is felt easily, for instance, by the hairs on back of your hand.) With a careful set-up, one can create "second kind" equations that are much easier to solve than first kind (in essence they are like bathroom glass that blurs, but also overlays a crystal clear copy of the original scene).

Because all of the unknown quantities live on the surfaces, integral equation methods are computationally more efficient than methods that treat the electrostatic field throughout the volume as unknown. In fact, they enable the solution of scientific problems that essentially cannot be solved by any competing methods. Integral equations naturally involve integrals (loosely, the area under a curve). Computers can rarely evaluate such integrals exactly, since that would involve knowing the function at an infinite number of points; instead one seeks good *quadrature* schemes which approximate the integral to *many digits* using a small number N of discrete samples. The captions of the figures show that for challenging problems, N, which is also the number of unknowns, varies from thousands to millions. Careful design of such schemes, as well as cunning algorithms to solve the resulting huge sets of equations, enables reliable, rapid solution, often to one-part-in-a-billion accuracy.

1.1 Frequency-domain wave scattering

A major contribution of the work of myself and my group has been developing better tools for integral equation methods to solve the boundary value problem (BVP) for scattering for the Helmholtz equation $(\Delta + k^2)u = 0$. Integral equations have well known advantages over conventional volumetric discretization, especially when the frequency k is high, but in this regime, or when boundaries are not smooth, the large N brings special challenges (see Figure 1(a)). We often exploit "fast" solutions for the linear system, whose cost grows much more slowly than the prohibitive $\mathcal{O}(N^3)$ of dense linear algebra. In this we build upon the work of Rokhlin, Greengard, Martinsson, Tygert, and others. Creating schemes compatible with fast algorithms—such as the fast multipole method (FMM) or fast direct solvers—requires in depth understanding of technical details such as quadrature. The ultimate goal is to develop adaptive solvers, i.e. new algorithms and accompanying software implementations, that, with minimal user adjustment, are accurate and efficient for most geometries relevant to engineering and science problems.

Periodic problems. Diffraction gratings and other periodic geometries remain essential for the control of electromagnetic waves. Building on an idea in [16],[1] we bypass the traditional use of the periodic Green's function (which does not always even exist, and can require complicated lattice sums to compute), and instead combine the *free space* Green's function for nearby lattice neighbors with an auxiliary set of free space solutions (such as proxy sources) that accurately capture the rest of the lattice. Periodicity and radiation conditions are applied directly to form an expanded linear system. With JWY instructor Adrianna Gillman [21], we showed how a fast direct solver makes this extremely efficient (600× faster than a "fast" iterative FMM scheme) when multiple incident angles are needed, as in many applications. Then, with IACM instructor Min Hyung Cho [26], we used auxiliary solutions in each layer of a dielectric grating, and by exploiting block tridiagonal structure to handle huge numbers of layers, as occur in microchips (Fig. 2a). A similar scheme created with collaborators at NYU [30] allows thousands of inclusions to inhabit one layer, as occurs in composite optical media, such as those designed to enhance absorption in thin-film solar cells. In each of these three 2D works the cost was $\mathcal{O}(N)$. Moving to 3D, with student Yuxiang Liu we developed efficient solvers for doubly-periodic arrays of bodies of revolution for acoustics [32], and even full Maxwell [39] (Fig. 2b). In this axially symmetric setting, we prefer the method of fundamental solutions to second-kind integral equations. We also exploit skeletonization via the interpolative decomposition to accelerate the application of obstacle interactions by an

[1]Numbered citations are to the works enumerated in the publications section of my CV.

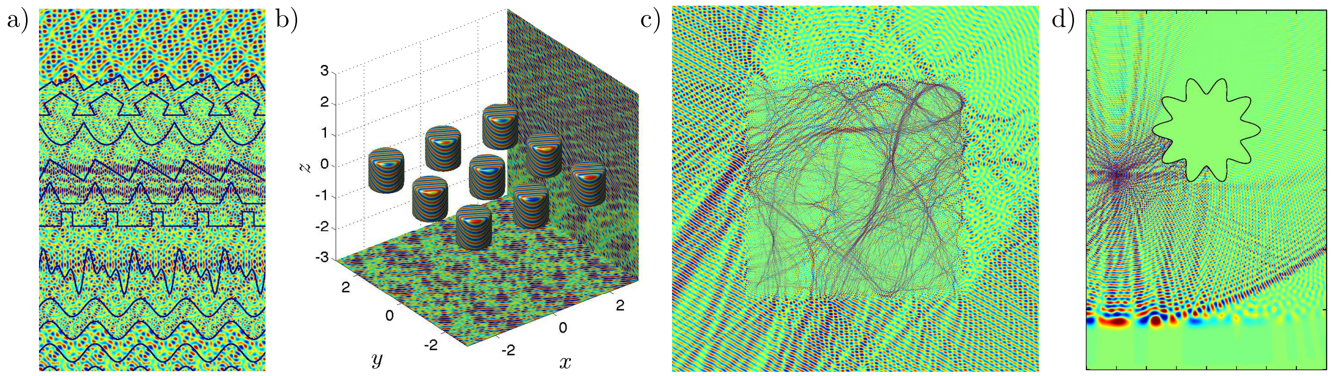

Figure 2: a) Scattering from 1000-layer 2D dielectric grating ($N = 6 \times 10^5$ unknowns, 8-digit accuracy, 6 minutes) [26]. b) 3D electromagnetic scattering from a 13-wavelength doubly-periodic grating of dielectric rounded cylinder shapes ($N = 83000$, 8-digit accuracy, 2 hours) [39]. c) 2D acoustic scattering from variable medium 200 wavelengths across, using hierarchical Poincaré–Steklov direct solver ($N = 1.4 \times 10^7$, 9-digit accuracy, 20 minutes) [25]. d) 2D scattering from Dirichlet obstacle in vertically-varying medium ($N = 2000$, 12-digit accuracy, 1 hour) [28].

order of magnitude.

Variable-coefficient PDEs. In applications such as tissue imaging with ultrasound or microwaves, a piecewise-constant model is inappropriate. With Gillman and Gunnar Martinsson, we created a smoothly-varying-coefficient scattering solver based on the hierarchical merging of provably robust impedance-to-impedance maps [25] (Fig. 2c). This solver has already had impact, enabling 10^6 PDEs to be solved efficiently in the recent inverse scattering scheme of Borges–Gillman–Greengard. In a special case of sound speed variation in only the vertical coordinate, as might occur in underwater density gradients, we show in [28] how a novel steepest-descent quadrature enables Green's function evaluation with cost independent of k, hence high-order scattering from embedded inclusions (Fig. 2d).

Quadrature. High-order quadrature schemes that can handle singular kernels are crucial for accurate elliptic PDE solutions, hence I have devoted much effort to this, and plan to continue. We review and provide a first systematic comparison of log-singular Nyström schemes on curves in [23]. However, since the solution is represented as a convolution of the Green's function with a density on a curve or surface, when the *evaluation* point approaches the surface, conventional smooth quadratures lose all accuracy—this is the "close evaluation" problem. In [24] I provide the first rigorous analysis of this, for the commonly-used periodic trapezoid rule, then present a new scheme to cure it, based on *local expansions* with centers near the boundary. This became the key idea in *quadrature by expansion* (QBX), which included on-surface evalution, in joint work with Leslie Greengard's group at NYU [20]. QBX has led to works by several other researchers including Epstein and Racch, and has enabled simulation of densely-packed ellipsoids in a Stokes fluid by Klinteburg–Tornberg. Because it needs only smooth surface quadratures, QBX has the potential to become a standard tool, especially in 3D (see Fig. 1a, unpublished). Finally, we have generalized QBX to arbitrary constant-coefficient elliptic PDEs in 2D, exploiting kernel-independent proxy sources [37].

Achieving many accurate digits for integral equations on surfaces with edges and corners poses a challenge. Automated medium-to-high order quadrature of such surfaces is one of my future goals. I am also in the process of building documented software libraries for 2D and 3D quadrature schemes for Laplace, Helmholtz and Stokes.

1.2 Stokes flow, eigenvalue problems, and computational biology

Stokes flow in complex geometries. In 2014 I started a collaboration with Shravan Veerapaneni and his group at Michigan, applying some of the advances described above to Stokes BVPs. These model viscous flows in a range of important applications, including porous media, sedimentation, microfluidic devices, and flows of massive numbers of deformable vesicles such as red blood cells. In the case of vesicles, large numbers of *timesteps* are needed, which motivates our work to develop more efficient solvers for each timestep. I started out by generalizing an elegant 2D Laplace double-layer spectral close-evaluation quadrature scheme of Helsing to the single-layer, opening up its use in Stokes applications [29]. This required complexifying a log-singular spectral scheme of Kress. Periodic boundary conditions are common in applications, so, with Gillman we combined a Stokes variant of the periodizing idea to solve flows of thousands of vesicles through a complicated channel, with a fast direct solver that solves the channel problem "for free", accelerating the timestepping [33]. For porous media and microfluidic separators, the effective permeability of a complicated periodic geometry is desired. However, applying a pressure drop with arbitrary unit cells is complicated in lattice-sum based Stokes periodic solvers. We recently invented and analyzed a simple scheme for 2D doubly-periodic flows [38], exploiting rank-1 or rank-3 perturbations to make well-conditioned systems compatible with the adaptive FMM, allowing very complicated geometries to be handled with linear complexity (see Fig. 1b). This motivates future plans to release a set of periodized FMMs, with all types of periodicity and unit cells in 2D and 3D, for use by the Stokes fluid community. This should aid applications in bio-fluids, active swimmers, sedimentation, and shearing flows.

The high frequency spectral problem for the Laplacian. The "drum problem" (eigenvalues of the Laplacian in a domain with Dirichlet boundary conditions) is a paradigm for many resonance problems in acoustics, optics and electromagnetics. The behavior of eigenfunctions at high eigenvalue is also of ongoing interest in pure mathematics (quantum chaos). Building on work of physicist Vergini, I pioneered pre-tenure the "scaling method" which, for high eigenvalues of star-shaped domains, is hundreds of times more efficient than any competing scheme. Yet I had been unable rigorously to analyze the method. Recently, in a 57-page work with Andrew Hassell at ANU, we re-expressed it in terms of the flow (with frequency k) of the spectrum of the domain's Helmholtz Neumann-to-Dirichlet map, enabling the first error analysis a potential theoretic formulation, and high-order accurate variants [22]. To find lower eigenvalues (e.g. the first 10^3), graduate student Lin Zhao and I combined the Fredholm determinant—analytic with respect to k—with Boyd's efficient rootfinder, and showed that, due to the possibility of exterior resonances, a combined-field formulation is key for robustness [27]. Finally, Hassell, Melissa Tacy and I have tightened the best inclusion bounds on Neumann eigenvalues by a factor of $\lambda^{5/6}$, where $\lambda = k^2$ is the eigenvalue, and present a numerical implementation that shows that accuracy can be improved from 9 to 14 digits without extra effort. I built the numerical methods presented in the above three works into a software package that I created, the `MPSpack` toolbox for MATLAB, enabling their dissemination and testing. My future plans include scaling-type interpolation-based rootfinding methods for the doubly-layer operator that have no restriction on domain shape, and user-friendly software releases of 2D and 3D eigenvalue solvers that supercede `MPSpack`.

Computational biology and data analysis. Since Fall 2014, have also been working at the new Simons Center for Data Analysis. This has led to two major new applied threads in my research: spike sorting, and cryo-electron microscopy (cryo-EM).

The goal in spike sorting is to process terabytes of electrical recordings from electrode arrays embedded in the brain (usually rats or mice), extracting the identities and firing times of as many distinct neurons as possible, as reliably as possible. Such electrode technology is enabling an explosion in understanding the brains of behaving animals. Challenges include "noise" (largely due to distant firings), the scale of the data, modeling the judgments made by expert human neuroscientists, and the lack of any *ground truth* in almost all experiments. With collaborators Jeremy Magland and Leslie Greengard, we first proposed new *metrics* based on cross-validation [31]. Magland and I also invented a clustering algorithm that handles skew and non-Gaussian clusters better than competing methods, without user-adjustable parameters [34]. This is built into our new C++ based software platform for spike sorting and visualization, which is in beta, and is being tested by collaborating neuroscience labs at UCSF, Stanford, Columbia, and NYU. My future plans include developing statistical inference based solutions to the challenge of overlapping spikes, and comparing ours and other software packages with standardized metrics.

The goal in cryo-EM is to reconstruct the 3D electron density function of a protein from of order 10^5 noisy projection images taken at unknown orientations. The images come from transmission electron microscopy of protein molecules embedded in a layer of ice; the fact that no crystallization is needed for cryo-EM is part of why it is revolutionizing protein structure determination. The signal-to-noise ratio is *very* poor (typically 0.1), and yet, due to the large number of images, structures accurate down to a few angstrom resolution can be found. However, existing reconstruction algorithms can take weeks of CPU time, which is a bottleneck to progress. With Spivak, Greengard and Pataki, we showed that combining accelerations enabled by spherical harmonic transforms with *marching in spatial frequency* (an idea of Yu Chen in inverse imaging), leads to reliable reconstructions in a couple of CPU hours (see Fig. 1c) [36], a huge improvement in solution time. I am now working on creating a turn-key software solution for labs to use, which requires more work on low-resolution initialization, in-plane translations, and better noise models.

Other future projects. In each topic above I have sketched directions in which I intend to continue. I also have ongoing projects in time-domain integral equations for wave scattering applications, in real-world acoustic phenomena such as racquetball court echoes, and in the topology of nodal domains of random waves (a pure mathematical question). In the long term I also would love to write a hands-on book (with codes) on integral equation methods for PDEs, as well as one for non-science majors on the mathematics of music and sound.

2 Teaching

My teaching practice has continued to evolve after tenure. I have co-taught our graduate student teaching seminar three times, and created two new courses, bringing my total of courses created at Dartmouth to six.

Inspirational teaching is, of course, more than merely explaining the material—although joyful presentation of well-edited material is an essential part. For each student, the act of learning is to build, in their own brain, a secure framework of new ideas and techniques (in learning theory this is called *constructivism*). Thus an effective professor must diagnose and model what is inside each individual's mind, be a motivational coach and guide, and be attuned to emotions that affect learning. With this in mind, I also bring mathematics alive through memorable/humorous physical demonstrations (often with unlikely household objects: e.g. I twist sushi mats to illustrate mixed partials in calculus, twang rulers to illustrate oscillatory ODEs. . .), through questions that hook one's interest (e.g. having students vote on what a code will do *before* running it), and through in-class worksheet activities that I design to enable the students (often in groups, via *peer instruction*) to make the right mental leap. I give frequent pointers to real-world applications (proven in studies to increase minority engagement), and motivate why the course can help their future careers.

My most significant new venture has been teaching, and helping shape, the **graduate student teaching seminar** (Math 147). This 2-course-equivalent program is required of all our graduate students, and is a key reason that they are sought out for high-quality teaching positions. (In our department External Review it was recognized as an asset "unique in mathematics programs across the country".) It includes practical daily exercises, cognitive psychology, affective aspects, diversity training, and the graduate students creating two fun week-long "math camps" for middle- and high-school students that include a deep pratice of evaluation and observation. This requires intense involvement for 6 weeks of the summer, but the rewards in instructor competence are huge. The effect on my own teaching has also been profound. For instance, a conventional work flow is "choose material → decide how to teach it → decide how to evaluate (homework, finally exams)." However, now I find *goal-oriented design* more effective: "choose tasks you want students to be able to do → design evaluations for those tasks → finally decide how to teach it." This "flipped design" is central to Math 147. For this, and many other insights, I am indebted to co-instructors Marcia Groszek and Rosa Orellana.

Dartmouth provides a special opportunity to develop **new courses**, and post-tenure I have developed two:

- Math 56: *Computational and Experimental Mathematics*, surveying a wide range of efficient methods for high-accuracy numerical computing in both applied and pure math contexts. I introduce key topics in numerical analysis and applied math (rounding, stability, numerical linear algebra, fast Fourier methods, deconvolution) but also algorithms of interest in pure math (factoring huge integers, cryptography, computation of π to millions of digits). We program in *two* languages with different strengths (MATLAB and python), since such flexibility is essential in quantitative careers. Student project work, culminating in often-entertaining research presentations, replaces a final exam. I have taught this twice.

- Math 126: *Great Papers in Numerical Computation*, a graduate-level survey of key papers and algorithms over the decades that underlie the tremendous power of modern numerical computing. Presentation skills (clarity, organization, and numerical demonstrations) were taught, and evaluated: students were required to give two lectures on the paper of their choosing, and were given anonymous feedback by their peers, and by myself. I have taught this once.

I have also been lucky to be able to continue to teach courses that I previously created at Dartmouth, namely Math 126: *Numerical Analysis for PDEs and Wave Scattering* (graduate level), Math 5: *The Mathematics of Music and Sound*, Math 46: *Introduction to Applied Mathematics*, and Math 53: *Chaos!* The last two have also been successfully taught by colleagues, and form part of the standard curriculum for applied majors.

3 Service to Department, College, and Wider Community

Mentoring. Since being granted tenure, I have mentored eleven undergraduate research projects, including three senior theses, two summer research projects funded from my NSF grant, one WISP student, and several Presidential Scholars. Two of the three senior thesis students went on to prestigious graduate programs in applied mathematics, inspired and enabled, I believe, by their undergraduate research experience. Kathleen Champion '11 is now at U. Washington, and Brad Nelson '13 at Stanford.

Over this same period I trained and supervised two graduate students to completion of their Ph.D.s. Lin Zhao (Ph.D '15, mathematics) solved problems integral equations and eigenvalues, and landed a job in a mathematically-oriented investment company. Yuxiang Liu (Ph.D '16, physics) solved periodic problems in acoustic and electromagnetic wave scattering, and just started postdoctoral research in applied physics at Yale. I have just started mentoring another graduate student.

I have also mentored two postdoctoral instructors, thanks to the generosity of the JWY and IACM programs in our department. Both instructors pushed their research to impressive levels, collaborated intensely with myself, brought strength to my research group, and landed excellent jobs in an academic market that has never been tougher. Adrianna Gillman, whose field is fast direct solvers, is now tenure-track at Rice, while Min Hyung Cho, whose field is electromagnetics and layered media, is tenure-track at UMass Lowell. I have four joint publications with them, and more on the way.

Organizing and outreach. Building community is essential for an academic field to flourish. Post-tenure I have devoted more time to this. The largest endeavor was organizing a NSF-CBMS conference, hosted by our department, on the exciting new field of fast direct solvers. With 57 participants, 5 days of videoed tutorial lectures by Gunnar Martinsson, tutorials by myself and other experts, hands-on computer sessions, and 2 days of talks by young researchers, this conference placed Dartmouth on the map in computational PDEs. I am currently one of three organizers for a 100-person workshop at Yale in 2017, have co-organized four research "minisymposia" at the top conferences in our field (eg, ICIAM, SIAM), co-organized a conference at the Banff International Research Station, and I continue to run the Applied and Computational Math Seminar (which has hosted around 90 talks since I started it in 2006). Since 2012 I have been an associate editor for Advances in Computational Mathematics. I was heavily involved in our department's submissions for NSF RTG grants to enhance graduate training in our department.

Music and its connection to mathematics continues to be a love of mine, and a compelling way to engage a wider audience on the power and ubiquity of applied mathematics. In collaboration with graduate student Megan Martinez (a GK12 participant), we developed an 8-hour module *Periods, Pitches and Pipes* connecting music and mathematics at the middle-school level, and field-tested them at the Newton school in Strafford, VT. Recently I gave two 2-hour long sessions to high-school teachers through the *Math for America* program, and a lecture at Hanover high school.

Service. My most significant service to the College has been my elected position on the Committee on Organization and Policy, including as Chair in Fall 2015. Prior to this I served for three years on the Committee on Student Life, and continue to take an active role on campus issues such as student social life and the problem of sexual assault. On a departmental level, I have served as advisor to graduate students, have served on the graduate program committee for several years, have been involved in applied mathematics faculty recruiting, and am the faculty advisor to our student chapter of SIAM (the chief professional society in applied mathematics).