

Incorporating Computational Challenges into a Multidisciplinary Course on Stochastic Processes*

Mark Jayson Cortez[†]
Alan Eric Akil[‡]
Krešimir Josić[§]
Alexander J. Stewart[¶]

Abstract. Quantitative methods and mathematical modeling are playing an increasingly important role across disciplines. As a result, interdisciplinary mathematics courses are increasing in popularity. However, teaching such courses at an advanced level can be challenging. Students often arrive with different mathematical backgrounds, different interests, and divergent reasons for wanting to learn the material. Here we describe a course on stochastic processes in biology delivered between September and December 2020 to a mixed audience of mathematicians and biologists. In addition to traditional lectures and homework, we incorporated a series of weekly computational challenges into the course. These challenges served to familiarize students with the main modeling concepts and provide them with an introduction on how to implement the concepts in a research-like setting. In order to account for the different academic backgrounds of the students, they worked on the challenges in small groups and presented their results and code in a dedicated discussion class each week. We discuss our experience designing and implementing an element of problem-based learning in an applied mathematics course through computational challenges. We also discuss feedback from students and describe the content of the challenges presented in the course. We provide all materials, along with example code for a number of challenges.

Key words. stochastic processes, problem-based learning, small-group learning, computational challenges

MSC codes. 97M10, 92-08

DOI. 10.1137/21M1445545

1. Introduction. Quantitative methods are of growing importance across many fields, as access to big data and powerful computational tools become the norm, and

*Received by the editors September 10, 2021; accepted for publication (in revised form) January 23, 2023; published electronically November 7, 2023.

<https://doi.org/10.1137/21M1445545>

Funding: This work was funded by the National Science Foundation through grants DMS-1662305 (first and third authors), MCB-1936770 (third author), and DBI-1707400 (second and third authors) and by the National Institutes of Health grant RO1 GM117138 (first and third authors).

[†]Department of Mathematics, University of Houston, Houston, TX 77204 USA. Current address: Institute of Mathematical Sciences and Physics, University of the Philippines Los Baños, Laguna, 4031, Philippines (mvcortez1@up.edu.ph).

[‡]Department of Mathematics, University of Houston, Houston, TX 77204 USA (alan.akil@yahoo.com).

[§]Department of Mathematics, Department of Biology and Biochemistry, University of Houston, Houston, TX 77204 USA (kresimir.josic@gmail.com).

[¶]School of Mathematics and Statistics, University of St Andrews, Fife, KY16 9SS, UK (ajs50@st-andrews.ac.uk).

so familiarity with mathematical modeling is increasingly valuable for researchers from all backgrounds. Mathematical modeling requires a combination of skills, in order to successfully bring together discipline-specific knowledge of the processes being described and the mathematical knowledge required to properly formulate a model. Researchers need access to mathematical tools to determine whether a model is well-posed, to analyze its behavior, and to determine solutions when possible. Just as important, interpretation and analysis of a model often require familiarity with computational methods to implement it and statistical methods to understand and analyze its output.

A course on mathematical modeling thus requires introducing students to a variety of mathematical ideas and domain-specific concepts. This becomes particularly challenging when teaching an advanced undergraduate or graduate modeling course aimed at an interdisciplinary audience. Given the range of tools and ideas required to formulate, implement, and analyze all but the simplest of toy models, such courses need to be organized differently from more classical applied mathematics courses.

Here we describe an approach to delivering such a course, which was implemented between September and December 2020 at the University of Houston and titled “Stochastic Processes in Biology.” We developed an introductory graduate course on stochastic processes in biology for an interdisciplinary audience of biologists and applied mathematicians, incorporating an element of problem-based learning [25, 26, 44, 51] in the form of computational challenges. We designed a series of nine sets of such challenges, which were assigned to groups of students to tackle in Python. These challenges were integrated into the course alongside more traditional lectures and homework. The challenges in each set were thematically related to each other and to the lectures that immediately preceded them. Each challenge required the implementation of a mathematical model making use of techniques taught in class, but which required students to explore the models in a more open-ended way than in homework. Students tackled the challenges in groups and presented their code and results in a dedicated class following submission. The ensuing discussions allowed us to revisit aspects of the material covered in lectures, talk about concrete implementation questions and bug fixing, and connect the material to current questions and methods in biology and to the practical challenges of implementing mathematical models in a research setting. Translating algorithms discussed in class into models of concrete biological processes also offered students a view of the role of stochastic processes in biology, allowing them to develop intuitions into the biological relevance of concepts such as a stationary distribution or a first passage time. We focused on practical aspects of model development and implementation, but also introduced the distinction between extrinsic and intrinsic noise. Fortunately, stochastic models are often useful even when the origin of noise is not known or is not completely specified.

The more open-ended nature of these programming challenges allowed us to go beyond what students typically learn in an introductory course. Indeed, as the complexity of the challenges increased, it was impossible to fully specify the modeling approach. This allowed us to discuss how some results depend on the modeling choices made by the student or researcher. An important consequence of such open-endedness is that the computational challenges did not come with a simple set of right answers. While this was uncomfortable for some students, it also reflects a difficulty faced by many students as they transition from undergraduate courses to independent research and problem solving in graduate school or outside of academia. In contrast to a typical applied mathematics course, the computational challenges helped introduce

students to simple heuristics for model validation when no analytical solution is available, and it encouraged them to think more carefully about whether computational results made sense given the modeling assumptions, and whether the assumptions made sense given the observed results.

Students tackled the computational challenges in small groups, which were assigned at the start of each challenge. These groups necessarily contained individuals with different levels of programming skill and mathematical and biological knowledge. We discuss the efficacy of using this kind of small-group learning [32, 47] as part of an applied mathematics course, in particular, the challenges of maintaining equitable division of labor and maintaining successful group interactions under the conditions of remote learning imposed by the COVID-19 pandemic.

In a repository¹ accompanying this article we provide a total of nine challenges as supplementary material, which can be used in other courses as written, or used as a basis for the development of similar challenges in other disciplines. Each challenge is broken into four interrelated group assignments, and we provide complete sample solutions to three challenges, including Python code. We also provide a summary of feedback from the students and discuss planned points of improvement when the course is given in the future.

Our approach of using small-group learning and computational challenges to bridge the gaps among an interdisciplinary audience, and to introduce students to some of the practical realities of applied mathematics research, could be replicated by others and can be easily adapted to fit alternative sets of research interests.

2. Course Overview. Biological processes are inherently stochastic [16, 50]. While deterministic models can offer valuable insights into the function and behavior of living systems, they do not capture the effects of randomness and variability that characterize and often drive many biological processes. Hence, a familiarity with stochastic processes is vital for anyone seeking to build models of biological systems and simulate their behavior numerically. The goal of our course is to teach students how to use the tools of the mathematical theory of probability and stochastic processes to develop, analyze, and implement models of living systems.

Our course met for 80 minutes twice a week for 14 weeks, but our approach could be easily adjusted to fit different schedules. Due to the COVID-19 pandemic, our class took place remotely via Zoom, and so the need to maintain student engagement was an important issue. To address this, we broke the class into two 25 minute periods of instruction, separated by 30 minutes devoted to student presentations and discussions of computational challenges.

We chose the mathematical topics for the course to be of use to a wide audience of graduate students in applied mathematics, biology, biomedical engineering, and related disciplines. Our audience in Fall 2020 consisted of 16 Ph.D. students (11 in applied mathematics, 4 in biology, and 1 in biomedical engineering). The prerequisites included differential equations and linear algebra, as well as undergraduate level probability. We also assumed that the students had some experience programming in MATLAB or Python.

We asked students to use a Git repository to facilitate sharing of code related to computational challenges. Two months before the course began we checked with students individually to see if they had the requisite background. We pointed students with a weaker background in programming to introductory online Python courses [18]

¹https://github.com/josic/stochastic_process_bio

and asked them to complete the courses before the start of our course. In particular, we recommended a tutorial on agent-based modeling, which introduced a number of ideas used in the course [41]. We set up a Slack team to communicate with students before and during the course, to facilitate group discussion and keep a clear record of problems that arose and their solutions.




Fig. 1 Summary of the course structure. Lectures (blue) were divided into five sections plus an introduction. In each section, key models and concepts related to the topic were introduced alongside the tools required to analyze and simulate them. Each homework (orange) covered six lectures' worth of material. Group computational challenges (green) occurred either every week (challenges 0–5) or every two weeks (challenges 7–9) as the challenges became more advanced, and thus covered either two or four lectures worth of material.

An overview of the course is provided in Figure 1, and a detailed syllabus with a complete list of topics covered in the lectures and homework can be found in the GitHub repository accompanying this article.

The course contained two types of material (Figure 1), with core material mostly delivered first, covering a review of probability, followed by a discussion of Markov processes with discrete and continuous space variables, diffusion processes, stochastic differential equations, Wiener and Ornstein–Uhlenbeck processes, and point processes. The discussion of these mathematical concepts was kept at a practical level, driven by

examples and the ultimate goal of simulating stochastic models of living systems. We were also inspired by D. Gillespie's dictum that "one's knowledge of any dynamical system is deficient unless one knows a valid way to numerically simulate that system" [21]. Thus, as far as possible, we used an algorithmic approach in introducing stochastic processes motivated by D. Higham's articles in this journal [23, 24] and D. Gillespie's classic book [21]. While this approach is, of course, no substitute for a rigorous course on the theory of stochastic processes, it allowed us to present the main ideas of the subject succinctly and in a way that was understandable to our interdisciplinary audience.

Following on from this, we focused on applying the core material to research-based topics from across biology. These topics were chosen to reflect the interests and expertise of the two instructors and covered models from systems biology, evolutionary game theory, and neuroscience. In particular we introduced examples based on biochemical reaction networks, gene regulatory systems, and neuronal networks, as well as models of epidemics and evolutionary processes and finally stochastic games and agent-based models.

This approach allowed us, for example, to motivate continuous Markov processes using the birth-death process, the Gillespie algorithm using biochemical reaction networks, and point processes using neuronal spike trains.

Because of the wide range of topics covered in the course, we did not follow a single textbook. However, we offered suggested reading from a number of texts covering different topics addressed in the lectures [1, 2, 5, 6, 15, 20, 36, 49]. In addition, the computational challenges drew on models published in research papers related to the topic at hand. As such, we introduced stochastic modeling using a variety of examples from across biology and provided interested students with groundwork that they could build on in subsequent courses or their own research.

3. Computational Challenges. We implemented small-group learning in the course by assigning a sequence of computational challenges of increasing difficulty to groups composed of four students: Tackling each challenge required the students to work together to understand the background material describing biological and mathematical concepts associated with the problem, combining both to formulate a model. This model then needed to be translated into code contained in a single Jupyter notebook (submitted to a Git repository), with results that could be presented and interpreted in a class discussion. As such, the challenges required the students in a group to share their knowledge and ideas and learn from each other [47]. Once an assignment was completed, each group's code was made available to the whole class.

Group presentations took place during class, with two groups presenting for 15 minutes each. Each group presented their solution, giving an overview of the challenge and an explanation of how they approached the problem, implemented the solution, and interpreted the results of their simulations. Guidelines for presentations were as follows:

- a brief (1–2 minute) description of the problem;
- 5 minutes explaining how they approached and implemented the solution;
- 5 minutes on results and interpretation; and
- a few minutes for questions.

Each group received a grade based on the quality of their code, their solution and analysis of the results, and the presentation. We provided feedback on the presenta-

tions, the results, and the code. However, the time reserved for in-class presentations was too short to give substantive feedback. We therefore held optional, weekly 60 minute Q&A sessions to discuss the week's material, including the challenges. These sessions were well attended. We recommend that students be given written feedback and be incentivized to attend the discussion sections (see section 6 for more details).

At the end of the course, students gave a short (15 minute) presentation either on a paper related to the subjects discussed in the course or about their own research, in the style of a short conference talk. If choosing to present a paper, students were asked to verify the results. These final presentations were intended to provide the opportunity for students to practice delivering professional presentations.

3.1. Choice of Programming Language. We required students to implement their models in Python and use Jupyter notebooks for their presentations. Python is widely used by many researchers as well as in industry, it is flexible and has appropriate libraries for building stochastic models, and it is also free to download and use. In addition, example code for standard procedures is widely available, and we encouraged students to make use of such resources as this is an important skill when writing code for research and other real-world settings. Similarly, by asking students to use a Git repository for their code, we were able to introduce them to a widely used tool for version control and code sharing.

We provided the following instructions to students about coding:

- You do not have to develop all the code yourself, but you need to understand it. If we don't say explicitly to implement an algorithm, you can use a package or code that others have written. For example, there are plenty of implementations of the Gillespie algorithm out there. However, make sure you understand what they do. Sometimes the safest thing to do is implement it yourself.
- Find a way to communicate and share code among yourselves. Git is the best for version control, but you may prefer Dropbox. We leave this up to you. You will need to use Git to share the code with us.
- Use Slack. We are here to help you. If you need help understanding something or need a pointer, send us a message. You will likely get a reply on the same day, often within an hour, if not from us, then from other students in the course.

Our choice of programming language did create some problems as some students were not familiar with Python. We tried to mitigate this problem by communicating our expectations to students and providing resources to help them learn Python well before the start of the course, as explained above. We believe that Python is a good choice for most modeling courses in applied mathematics.

4. Example Computational Challenges. Here we provide a sample of the computational challenges we assigned, along with solutions. Further details about these assignments, model details, and parameters used in the simulation, as well as commented code showing the solutions, can be found in the accompanying repository.

We note that as the complexity of the mathematical ideas covered in the lectures increased, the associated challenges became more open-ended. In particular this meant that there were multiple valid modeling approaches to the same challenge, and that there was no single "correct" solution that could be provided in an answer sheet. As such, the class presentations and discussions were the primary form of assessment.

This gradual transition to open-ended problems with more than one acceptable solution was an important part of the course. Beginning graduate students in applied mathematics often find it challenging to accept that in many research problems there may be no single right answer. This problem is compounded by the fact that in most applied mathematics classes, the problems presented do have a single right answer. The computational challenges allowed us to emphasize the importance of making good modeling choices, looking at limiting cases, and asking whether the results of a model make biological sense.

A good challenge problem should

- clearly illustrate a mathematical concept discussed in class, in the context of a concrete biological problem;
- have a solution that offers insight into the behavior of a real biological system;
- have a solution that can be arrived at and implemented within 10 hours by an average programmer; and
- have identifiable limiting and test cases where the system's behavior can be predicted either intuitively or using some straightforward analysis.

The three challenges presented below reflect these criteria and illustrate the progression from simple questions with a single concrete solution toward more open-ended complex problems.

4.1. Rejection Sampling (Challenge I). Generating random numbers and samples from an arbitrary distribution is fundamental to stochastic modeling. We began the course with a discussion of the numerical techniques for quasi-random number generation and their limitations. We used linear congruential generators (LCGs) to illustrate the ways in which random number generators can fail [39] and introduced the Mersenne Twister as a tool appropriate for the class.

Our first computational challenge was to use such random number generators to implement rejection sampling and illustrate how random samples from a distribution $f(x)$ can be used to generate samples from a different distribution, $g(x)$ [38] (see Figure 2). We chose four different combinations of univariate distributions for $f(x)$ and $g(x)$, to produce four different group challenges. We asked the students to visualize the results and make a convincing case that the samples were generated from the correct distribution. We posed the following questions for discussion: When does the rejection sampling method become inefficient? Do you think that rejection sampling can be extended to higher dimensions? How?

Students solved this challenge easily, and dug deeper: For instance, some looked at how the probability of acceptance of a proposed sample depends on the scaling parameter in the algorithm (see Figure 2(c)). This scaling parameter serves to ensure that the candidate density has heavier tails than the target, and it is optimal at the supremum of the densities ratio. Others implemented the algorithm in higher dimensions (see Figure 2(d)), allowing us to discuss in more detail how the efficiency of the algorithm depends on the shape of the proposed and target distributions.

While this challenge is simple, it introduced the students to the small-group learning aspect of the class—to work together to solve the problem and plan a presentation of their results. The simplicity of the problem and its solution allowed us to concentrate on presentation structure and to communicate expectations. As the format of the course is different from most others students have attended, the first discussion concentrated on clearly defining what constitutes an acceptable solution to the challenges and how the solutions should be communicated to the class.




Fig. 2 In this example of rejection sampling, the normal distribution was used to generate samples from the t -distribution. (a) A comparison of the probability densities of the proposal and target distributions. (b) Kernel density estimate (blue) obtained using 10,000 samples closely matched the target density (orange). (c) Acceptance rates are inversely proportional to the scaling constant, c , in the algorithm. (d) In an extension to two dimensions, students obtained samples from the bivariate t -distribution with parameters $(\mu = [0, 0], \Sigma = I_2, df = 2)$, using samples from a normal distribution with parameters $(\mu = [0, 0], \Sigma = 2I_2)$. The contour map obtained using the generated samples (right) shows that the target distribution is well approximated (left).

4.2. The Genetic Switch and the Repressilator (Challenge 4). The construction of the first synthetic genetic toggle switch [19] and genetic oscillator [13] were landmark events in modern biology. The two original papers are lucidly written and explain how the design of both circuits was inspired by mathematical models. We therefore assigned these two papers as background for this challenge, with optional references for those who wanted to learn more [9].

Although the models presented in these papers are deterministic, genetic circuits often operate at small molecular numbers and are thus inherently noisy [12]. The aim of this challenge was to illustrate how to model noisy genetic circuits and show how molecular noise can shape their dynamics by inducing state transitions and oscillations. Stochastic versions of the genetic toggle switch and oscillator are described in a pair of papers by A. Loinger and coauthors [30, 31].

Two groups were assigned the genetic switch model, while the other two worked on the repressilator. To illustrate the relationship between the deterministic and stochastic descriptions of the system, one group in each pair implemented the deterministic model of the system, while the second group implemented its stochastic counterpart. To simulate the model numerically, the second group used the Gillespie algorithm [22], which was the subject of a previous challenge. We therefore also asked the pair of groups assigned to each project to share and discuss their results with one another.

The first group analyzed the deterministic version of the genetic switch system described in section III in [31]. They first constructed a Petri net representation of the system as an exclusive switch and were asked to explain the reason for the term “exclusive.” They were then asked to solve the ODE numerically and show that for a set




Fig. 3 The stochastic exclusive genetic switch of Loinger and Biham [30] exhibited bistability, while its deterministic analogue was monostable. (a) A phase plane analysis of the continuous exclusive switch shows a single equilibrium state (black star marker) with a large basin of attraction (circles represent initial conditions). (b) The stochastic model displayed stochastic switching between two complementary states in which one, but not the other protein was highly expressed. Whether the promoter is bound by either A or B, $r_A = 1$ or $r_B = 1$, respectively, is also indicative of the dominant species as both proteins negatively regulate each other's synthesis. (c) The average residence time in the two states increases with repression strength $k = a_0/a_1$.

of initial conditions the system approached a single equilibrium (see Figure 3(a)) using phase plane analysis and computing the equilibrium and determining its stability.

The second group implemented a stochastic model of the same exclusive switch, starting with the master equation for the system. They first showed that for the same parameters used in the deterministic model, the stochastic model is bistable and exhibits noise driven transitions between the two states (see Figure 3(b)). They next changed a parameter that governs the strength of repression between the two genes in the switch. As repression increased, so did the average time in each state (see Figure 3(c)). Both groups were asked to refer to the model equations and the Petri net to explain these observations.

The assignment handed to group three paralleled that of the first group: They implemented and analyzed two different versions of a deterministic model of the repressilator [13, 30]. The first version of the model included mRNA dynamics (corresponding to equation 1 in [30]), while the second did not (corresponding to equation 3 in [30]). The inclusion of mRNA level in the model introduced an effective delay in the production of mature proteins, which can be important for generating oscillations. The third group was then asked to show that oscillations in the deterministic system only appear for a high enough value of the Hill coefficient, n , which measures the sensitivity of expression (or cooperativity) of each gene to changes in concentration of the regulatory protein (see Figure 4(a),(b)).




Fig. 4 The Michaelis-Menten rate equation models of the repressilator circuit. The circuit consists of three genes that negatively regulate each other and can exhibit oscillatory behavior with the three genes expressed in alternation. (a) Simulations show that oscillations appear for the Hill coefficient, $n \geq 2$, in the deterministic model that includes mRNA concentrations (equation 1 in [30]). (b) However, when mRNA concentration is not modeled explicitly (equation 3 in [30]), oscillations occur only when $n = 3$. (c) Simulations of the stochastic model including mRNA and using $n = 3$ result in oscillations. (d) FFT-based denoising of one of the trajectories showed that a filtered signal allows for an estimate of the oscillation period.

The fourth group worked in parallel with group three and developed a stochastic model of the repressilator based on equation 1 in [30], showing that oscillations appear (see Figure 4(c)) consistently for high Hill coefficients, while smaller amplitude, random fluctuations are typical at low Hill coefficient values. We also asked the fourth group to develop a method to detect oscillations and their period by using the FFT of one of the protein concentrations (see Figure 4(d)).

At this point in the course, students were familiar with the Gillespie algorithm. The goal of this assignment was to apply this knowledge to model biological systems that stimulated the development of the field of synthetic biology, as well as to show the sometimes nontrivial relationship between stochastic models and their deterministic counterpart. This challenge was somewhat different from the others we assigned, since two pairs of groups worked on different models of the same systems. It is necessary to clearly communicate that the two groups need to meet before the assignment is due to compare results and coordinate their presentations. This could be done by requiring that the two deliver a joint conclusion to their projects that compares and contrasts the results of the two modeling approaches.

4.3. Agent-Based Models (Challenge 9). The last challenge of the semester consisted of four group assignments aimed at demonstrating the diverse applications of agent-based simulations in biology. These assignments were the most open-ended of the semester and required students to apply much of what they had learned throughout the course. We asked that they animate their simulations using, for example, `FuncAnimation` from `matplotlib` in Python. For this challenge, each group was given a separate topic based on a classic paper or concept in which space plays a central role in determining the behavior of the system.

Group 1: Schelling Segregation. This example of segregation is a classic example of the insight that can be gained by implementing an agent-based model. We provided students with both the original paper introducing the model [45] and other more recent discussions of the topic (e.g., see page 108 in [11]). In the model, two types of agents reside on a lattice and decide whether to move or stay put according to the number of neighbors of the opposite type in the adjoining spaces. There are several choices that can be made in implementing these rules, but the results are largely independent of those choices. Students were asked to examine the dependence of the final state of the system on the initial “empty ratio,” the fraction of the domain that is initially unoccupied, and the similarity threshold that determines the fraction of neighbors of opposite type that cause an agent to move.

Although the Schelling model illustrates a sociological (rather than a biological) phenomenon—segregation in housing—similar lattice agent-based models are now used widely in biology [28, 48].

Group 2: A Spatial Moran Model of Cancer. This is an abstract lattice model of carcinogenesis [29]. Cells are arranged in a regular grid at locations $i = 0, 1, 2, \dots, N$. The total number of cells, N , is fixed, as each cell that dies is replaced by a new cell, with probability determined by the fitness of these cells. Here is a simplified version of the algorithm:

1. A cell is chosen for death and is removed from the population. All cells are equally likely to die.
2. One of the two neighboring cells is chosen for reproduction. If the fitness values of the two neighboring cells are r_{left} and r_{right} , the probability that the left will reproduce is $r_{\text{left}}/(r_{\text{left}} + r_{\text{right}})$, and the probability that the right will reproduce is $r_{\text{right}}/(r_{\text{left}} + r_{\text{right}})$.
3. The descendant of the dividing cell fills the empty spot created by the removal of the cell in step 1.

Students were asked to start with different initial positions for a single mutant cell and compute the fixation probabilities (the probability that the mutant cell will take over the entire population) for different fitness and initial position values. They showed the fixation probability against initial position to check for edge effects and determined the effects of domain size on fixation probability. A master equation for the first version of the model can be found in [29].

Group 3: A Model of Leadership and Decision-Making in Animal Groups. The goal of this challenge was to examine how the decision of informed individuals can impact the behavior of animals in a collective using the agent-based model described in [10]. In contrast to the previous two models, the domain is continuous and agents can occupy any position in space. The movement of each individual is determined by the motion of its close neighbors. Some of the agents have information about a

source of food or a predator and adjust their movement by balancing their direction preference and the influence of social interaction. Others are naive, do not have a preferred direction, and only follow their neighbors.




Fig. 5 A model of information transfer in moving groups introduced by Couzin et al. [10] shows that the strength of social interactions, ω , and the fraction of informed individuals, p , are central in a group's success, be it moving to a desired fixed location (a)–(b), like food and other resources, or escaping from a predator (c)–(d). Students simulated $N = 50$ agents in a square domain showing the following: (a) When agents (gray) are informed and social interaction weight is sufficiently high, the group is able to reach a stationary target (green). (b) Group accuracy, quantified by the probability to move to within 10 units of the target destination within a given time, increases with the fraction of informed individuals p and social interaction strength ω , but decreases with the increase in population size N . (c) Students also simulated the escape of the group from a single, fast-moving predator (red). Each agent tries to escape the predator by coordinating with its neighbors, but only some agents have information about the predator's location. Neighboring agents coordinate motion as indicated by their direction (orange arrows). (d) Survival time increases with both interaction strength and the number of informed individuals.

In the first part of the challenge, students examined the ability of a collective to find a source of food. They started their simulations with food located at the origin, and placed the group in a small circle far from the food (see Figure 5(a)). They then showed that the collective reaches the origin in a time that depends on the number of informed individuals. In the second part of the challenge, students started with populations of different sizes and different numbers of informed individuals, as well as various social interaction strengths. These simulations verified that group accuracy increased with the number of informed agents and the ability of these agents to

influence their neighbors, but decreased as population was increased (see Figure 5(b)).

The second part of the assignment departed from the original model, as the individuals in the population were attempting to escape a predator rather than find food (see Figure 5(c)). The predator started at some distance from the population and always moved to the closest agent, regardless of distance. If the predator came sufficiently close to an agent, that agent was consumed and removed from the population. The predator's speed was slightly faster than the speed of each agent. Figure 5(d) shows the survival time of the population as a function of the fraction of informed individuals and the strength of interactions between the agents in the population.

Group 4: A Spatial Rock-Paper-Scissors Game. In this simulation, agents were placed on a lattice with periodic boundary conditions, and each agent used one of the three eponymous strategies [42]. Each agent thus effectively belonged to one of three species. Each cell (location) in the lattice was initially either empty or occupied by an agent.

Students were asked to start by populating the lattice with agents of each type and leave a fraction of cells empty. Simulations proceeded by picking a cell uniformly at random in the lattice and then picking a random neighboring cell. If both cells are empty or occupied by agents of the same species, nothing happens. Otherwise, three things can happen:

1. If only one of the cells is occupied, the agent in the occupied cell reproduces. The descendant belongs to the same species, i.e., uses the same strategy.
2. If the two cells are occupied by different species, then with probability p they fight. The agent with the losing strategy dies and their cell is vacated.
3. With probability $1 - p$, the agents in the two cells swap places.

The three strategies can coexist in a sufficiently large domain (see Figure 6(a)) and the dynamics of the system changes significantly with the parameter p (see Figure 6(b)). Mobility, modeled through agents swapping positions with probability $1 - p$, is crucial to population diversity, as species form spiraling spatial patterns that grow in size with mobility (see Figure 6(a),(b)). Over a wide range of parameters, the fraction of each species oscillates through time at a period that depends on p (see Figure 6(c)). A possible interesting extension of the model, which we left as a bonus challenge, was to simulate five species by using the rules of the rock-paper-scissors-lizard-Spock game.

These agent-based model challenges were the most complex we assigned. Their implementation required that students use much of what they had learned up to this point in the course. We noticed that, in some groups, this resulted in students with weaker programming skills not contributing much. This could be addressed by asking the students to subdivide the programming and data analysis tasks.

5. Group Interactions. The benefits of collaborative learning are well documented [27, 40]. Problem-based learning is widely used in some sectors such as medical education [3, 33, 35, 46], but it is less common in mathematics and STEM education more generally. Solving problems in groups allows students to learn inductively, as well as develop communications skills. The ability to work as a member of a team is useful, regardless of their career path. A succinct and extremely valuable introduction to this topic is provided by Richard M. Felder and collaborators [17, 37], and there are many other excellent books on the subject [34]. Here we focus on the issues we encountered implementing small-group learning in a graduate course on mathematical modeling aimed at an interdisciplinary audience.




Fig. 6 Lattice size and mobility strongly affect population dynamics in a spatial rock-paper-scissors game. In all simulations, a cell in the lattice was initialized to be empty (black) with probability 0.1 and belonging to one of the species (red, blue, yellow) with probability 0.3, respectively. (a) Games on square lattices of different sizes, 30×30 , 75×75 , and 150×150 , with mobility probability $1 - p$ for $p = 0.2$, showed that a single species is likely to dominate in a small domain. Coexistence is more likely in larger domains where spiral patterns developed, consistent with the observation of Reichenbach, Mobilia, and Frey [42]. (b) In a lattice of size 75×75 , students observed oscillation in population fractions with frequency decreasing and cluster size increasing with p .

For each challenge we randomly assigned students to teams of four and asked them to choose a person who would present their work. However, to ensure all students participated in presenting, only students with the least number of presentations to date were able to present in any given week. The strategy for choosing a presenter worked well, allowing all students to present about the same number of times during the semester. However, we found that reshuffling groups for each assignment was disruptive and made it harder to assess contributions and ensure equitable division of labor in the long run. In future, we plan to keep group composition constant throughout the semester with initial group assignments informed by responses to a short questionnaire about skills and background. The questionnaire will be given to individual students and include questions about their familiarity with different mathematical ideas, concepts in biology, and programming skills. The questionnaire can be accompanied by an elementary computational challenge that students need to tackle

individually. Responses will allow us assign teams of students with heterogeneous skill sets. Assigning groups prevents homogeneity in which math students seek out other math students, etc. Such homogeneity is not desirable in an interdisciplinary group since it puts some groups at a disadvantage and prevents knowledge exchange, in addition to the loss of the other benefits of interdisciplinary groups [37].

The biggest challenge we faced with small-group learning was maintaining a fair division of labor in the computational challenges. Students were asked to fill out feedback forms after each challenge, with many noting this problem repeatedly. However, students also found providing feedback weekly to be onerous. In future we will collect more detailed feedback after every three challenges.

There are several collective learning strategies that can be implemented to improve the success of small-group learning in our course: Providing a *Team Policy Statement* and *Team Expectation Agreement* at the outset of the course would provide a better basis for collaboration, regardless of the course topic [37]. In particular, rather than just the presenter, assigning roles to all members of the group would make sure that everyone was accountable and less able to free ride: e.g., assigning a *coordinator* to organize meetings and keep the group on task, a *recorder* and *presenter* to prepare the final presentation that will be delivered to the class, a *monitor* who checks that everyone agrees with the final results, and a *checker* who checks the final results. Finally, if meeting in person, we would institute a group office hour *before* the assignment is due to go over the solution and the presentation with the students.

6. Assessment. Assessment was based on the homework (20%), the final presentation (20%), and the group challenges (60%). As a result, the overall grade was based 40% on individual performance and 60% on group work. The grade for the individual challenges was further broken down so that Group Challenges 1–6 were worth 5% each, while Group Challenges 7–9 were worth 10% each, reflecting the greater difficulty of the final three challenges (Figure 1). Because grades were awarded at the level of the group for each challenge, with a single individual presenting for the group in a given week, there was a risk of free riding in which some members of a group contribute little, but still receive a high grade. In order to mitigate this, we required that the role of presenter was rotated between group members. We also introduced feedback forms in which group members could comment on one another’s contribution to each challenge [27, 34]. However, we did not attempt to adjust individual grades based on this feedback, and some students continued to indicate concern about free riding at the end of the course (see below). Free riding could be further reduced in future by asking that presentations be given in “tag-team” form, or by choosing a presenter or presenters at random from the group. We often asked a number of questions during the presentation. We suggest directing these not only to the person presenting, but to other team members.

Since different groups tackled different problems, grades had to take into account problem difficulty. Verbal feedback was provided after each presentation based on both the results presented and a review of the submitted code. Sometimes there were clear problems with the solutions. During the presentations we tried to get students to recognize these issues, and then explained how we recognized them. In these situations, we recommend allowing the students to correct and resubmit their code, providing an improved grade as the motivation. We reserved part of our online office hour for discussion of the revised code. In future we also suggest providing feedback via written comments with marks based on a rubric provided to the students at the start of the course [7, 43].

7. Student Feedback. At the end of the course we administered an exit survey, the results of which we summarize here. All questions and full results are available in the repository accompanying this article.

The 14 students who responded overwhelmingly agreed that the computational challenges helped them understand the course material. They were somewhat more divided on whether what they learned was directly applicable to their research, although the majority agreed. Most students who answered (8/14) indicated that they agreed or strongly agreed with the statement that they would take a course with computational challenges again. Most students strongly agreed with the statement that the course helped improve their coding skills.

Students were more critical about the division of labor within the groups. There was a weak majority who thought that there were too many challenges. This is really something that depends on the composition of the class. We asked them to self-report the time they spent on the challenge (three reported 2–4 hours, six reported 4–6 hours, and seven reported more than 6 hours per week). We note that all groups completed all challenges. Although the quality of the results varied, they generally exceeded our expectations. A better guide to division of labor, as discussed above, would likely improve this. Students were largely positive about small-group learning, suggesting that problem-based learning can be successful in the context of an applied mathematics course.

8. Conclusions. Designing and integrating a computational component into a course on applied mathematics can be time consuming; however, we observed clear benefits to both the students and the instructor. Small-group learning, particularly with an interdisciplinary student cohort, also provided clear benefits despite the challenges with ensuring equitable division of labor, and we plan to continue using this approach in future. This type of course is especially relevant for graduate students. Science is becoming increasingly collaborative. Whether our students continue in academia or industry, they will likely have to work in teams, and these teams are likely to include far greater diversity of skill and knowledge than the groups in this course. For many, their future success will thus depend on how well they can work as part of such groups.

While our course covered many topics, several were only touched upon. While we asked students to provide a well-motivated analysis of the results, and we provided feedback, we only spent the last week of the course on fitting models to data [49]. A companion course, including an element of problem-based learning and focused on model fitting, validation, and data driven model discovery and assimilation, would be an excellent complement to the one described here. Such a course could be taught using publicly available, curated datasets. As these are areas of high interest and intense research activity, there are now excellent resources that could be used to develop such a course [8].

Throughout the semester, we effectively provided four 25 minute lectures per week that we could easily have prerecorded. Alternatively, we could have used existing resources, such as the videos [14] accompanying the introductory book to stochastic processes in biology [15]. This would have allowed us to “flip” the classroom [4] and spend less time in class on instruction. Our optional discussions/office hours were well attended, and students were most engaged when we discussed their modeling and implementation choices and helped them interpret their results. It was difficult to provide extensive feedback in class after the presentations, and we therefore suggest providing written comments and using rubrics. Students also benefit from the oppor-

tunity to correct their mistakes and to reexamine their results. We therefore believe that an additional hour devoted to further discussions of the challenges is essential for the course, and attendance should be encouraged if not required. We expect that an online discussion session would be better attended, even if the lectures are delivered face-to-face. In a four hour course, attendance could be required.

A course making use of examples from current research can make effective use of crowdsourcing. We welcome help and suggestions to expand on the set of challenges we have made available. Developing the challenges, seeing the different, often innovative ways in which students solved them, and discussing the solutions was very rewarding for the instructors. We therefore hope to spur the development of similar courses across different branches of applied mathematics.

Appendix A. Repository of Resources. We provide a detailed syllabus, homework, and the text of all computational challenges, along with code for the examples we discussed here, at https://github.com/josic/stochastic_process_bio. This repository will be maintained.

Acknowledgment. We would like to thank the students in the course for their enthusiasm and willingness to participate in an experimental course, all during a pandemic.

REFERENCES

- [1] L. J. ALLEN, *An Introduction to Stochastic Processes with Applications to Biology*, CRC Press, Boca Raton, FL, 2010. (Cited on p. 1156)
- [2] U. ALON, *An Introduction to Systems Biology: Design Principles of Biological Circuits*, CRC Press, Boca Raton, FL, 2019. (Cited on p. 1156)
- [3] H. S. BARROWS AND R. M. TAMBYN, *Problem-Based Learning: An Approach to Medical Education*, Springer Ser. Med. Educ. 1, Springer, New York, 1980. (Cited on p. 1164)
- [4] D. BERRETT, *How ‘flipping’ the classroom can improve the traditional lecture*, *Chronicle Higher Educ.*, 12 (2012), pp. 1–3. (Cited on p. 1167)
- [5] J. M. BOWER AND H. BOLOURI, *Computational Modeling of Genetic and Biochemical Networks*, MIT Press, Cambridge, MA, 2001. (Cited on p. 1156)
- [6] P. C. BRESSLOFF, *Stochastic Processes in Cell Biology*, *Interdiscip. Appl. Math.* 41, Springer, New York, 2014. (Cited on p. 1156)
- [7] S. M. BROOKHART AND F. CHEN, *The quality and effectiveness of descriptive rubrics*, *Educ. Rev.*, 67 (2015), pp. 343–368. (Cited on p. 1166)
- [8] S. L. BRUNTON AND J. N. KUTZ, *Data-Driven Science and Engineering: Machine Learning, Dynamical Systems, and Control*, Cambridge University Press, New York, 2019. (Cited on p. 1167)
- [9] S. B. CAHN AND S. G. J. MOCHRIE, *Biologic: Gene circuits and feedback in an introductory physics sequence for biology and premedical students*, *Amer. J. Phys.*, 82 (2014), pp. 412–421. (Cited on p. 1159)
- [10] I. D. COUZIN, J. KRAUSE, N. R. FRANKS, AND S. A. LEVIN, *Effective leadership and decision-making in animal groups on the move*, *Nature*, 433 (2005), pp. 513–516. (Cited on pp. 1162, 1163)
- [11] D. EASLEY AND J. KLEINBERG, *Networks, Crowds, and Markets: Reasoning about a Highly Connected World*, Cambridge University Press, New York, 2010. (Cited on p. 1162)
- [12] A. ELДАР AND M. B. ELOWITZ, *Functional roles for noise in genetic circuits*, *Nature*, 467 (2010), pp. 167–173. (Cited on p. 1159)
- [13] M. B. ELOWITZ AND S. LEIBLER, *A synthetic oscillatory network of transcriptional regulators*, *Nature*, 403 (2000), pp. 335–338. (Cited on pp. 1159, 1160)
- [14] R. ERBAN AND S. J. CHAPMAN, *Stochastic Modelling of Biological Processes*, Online Course, <http://people.maths.ox.ac.uk/erban/cupbook/>. (Cited on p. 1167)
- [15] R. ERBAN AND S. J. CHAPMAN, *Stochastic Modelling of Reaction–Diffusion Processes*, Cambridge University Press, New York, 2019. (Cited on pp. 1156, 1167)
- [16] A. A. FAISAL, L. P. SELEN, AND D. M. WOLPERT, *Noise in the nervous system*, *Nat. Rev. Neurosci.*, 9 (2008), pp. 292–303. (Cited on p. 1154)

- [17] R. M. FELDER AND R. BRENT, *Effective strategies for cooperative learning*, J. Cooperation Collaboration College Teaching, 10 (2001), pp. 69–75. (Cited on p. 1164)
- [18] M. GALARNYK, *Python Tutorials*, <https://github.com/mGalarnyk/Python.Tutorials> (accessed 2021-08-29). (Cited on p. 1154)
- [19] T. S. GARDNER, C. R. CANTOR, AND J. J. COLLINS, *Construction of a genetic toggle switch in escherichia coli*, Nature, 403 (2000), pp. 339–342. (Cited on p. 1159)
- [20] W. GERSTNER, W. M. KISTLER, R. NAUD, AND L. PANINSKI, *Neuronal Dynamics: From Single Neurons to Networks and Models of Cognition*, Cambridge University Press, New York, 2014. (Cited on p. 1156)
- [21] D. GILLESPIE, *Markov Processes: An Introduction for Physical Scientists*, Academic Press, Cambridge, MA, 1991. (Cited on p. 1156)
- [22] D. T. GILLESPIE, *Exact stochastic simulation of coupled chemical reactions*, J. Phys. Chem., 81 (1977), pp. 2340–2361. (Cited on p. 1159)
- [23] D. J. HIGHAM, *An algorithmic introduction to numerical simulation of stochastic differential equations*, SIAM Rev., 43 (2001), pp. 525–546, <https://doi.org/10.1137/S0036144500378302>. (Cited on p. 1156)
- [24] D. J. HIGHAM, *Modeling and simulating chemical reactions*, SIAM Rev., 50 (2008), pp. 347–368, <https://doi.org/10.1137/060666457>. (Cited on p. 1156)
- [25] C. E. HMELO-SILVER, *Problem-based learning: What and how do students learn?*, Educ. Psychol. Rev., 16 (2004), pp. 235–266, <https://doi.org/10.1023/B:EDPR.0000034022.16470.f3>. (Cited on p. 1153)
- [26] W. HUNG, D. H. JONASSEN, AND R. LIU, *Problem based learning*, Handb. Res. Educ. Commun. Technol., 3 (2008), pp. 485–506. (Cited on p. 1153)
- [27] D. W. JOHNSON, R. T. JOHNSON, AND M. B. STANNE, *Cooperative Learning Methods: A Meta-Analysis*, University of Minnesota Press, Minneapolis, 2000. (Cited on pp. 1164, 1166)
- [28] B. R. KARAMCHED, W. OTT, I. TIMOFEYEV, R. N. ALNAHHAS, M. R. BENNETT, AND K. JOSIĆ, *Moran model of spatial alignment in microbial colonies*, Phys. D, 395 (2019), pp. 1–6. (Cited on p. 1162)
- [29] N. L. KOMAROVA, *Spatial stochastic models for cancer initiation and progression*, Bull. Math. Biol., 68 (2006), pp. 1573–1599. (Cited on p. 1162)
- [30] A. LOINGER AND O. BIHAM, *Stochastic simulations of the repressilator circuit*, Phys. Rev. E Stat. Nonlin. Soft Matter Phys., 76 (2007), art. 051917. (Cited on pp. 1159, 1160, 1161)
- [31] A. LOINGER, A. LIPSHAT, N. Q. BALABAN, AND O. BIHAM, *Stochastic simulations of genetic switch systems*, Phys. Rev. E Stat. Nonlin. Soft Matter Phys., 75 (2007), art. 021904. (Cited on p. 1159)
- [32] Y. LOU, P. C. ABRAMI, AND S. D'APOLLONIA, *Small group and individual learning with technology: A meta-analysis*, Rev. Educ. Res., 71 (2001), pp. 449–521. (Cited on p. 1154)
- [33] D. MANSUR, S. KAYASTHA, R. MAKAJU, AND M. DONGOL, *Problem based learning in medical education*, Kathmandu Univ. Med. J., 10 (2014), pp. 78–82, <https://doi.org/10.3126/kumj.v10i4.11002>. (Cited on p. 1164)
- [34] B. J. MILLIS AND P. G. COTTELL, JR., *Cooperative Learning for Higher Education Faculty. Series on Higher Education*, Oryx Press, Phoenix, AZ, 1997. (Cited on pp. 1164, 1166)
- [35] A. J. NEVILLE, *Problem-based learning and medical education forty years on*, Med. Princ. Pract., 18 (2009), pp. 1–9, <https://doi.org/10.1159/000163038>. (Cited on p. 1164)
- [36] M. A. NOWAK, *Evolutionary Dynamics: Exploring the Equations of Life*, Harvard University Press, Cambridge, MA, 2006. (Cited on p. 1156)
- [37] B. OAKLEY, R. M. FELDER, R. BRENT, AND I. ELHAJJ, *Turning student groups into effective teams*, J. Student Centered Learning, 2 (2004), pp. 9–34. (Cited on pp. 1164, 1166)
- [38] R. D. PENG, *Advanced Statistical Computing*, <https://bookdown.org/rdpeng/advstatcomp/rejection-sampling.html> (accessed 2021-07-26). (Cited on p. 1158)
- [39] W. H. PRESS, S. A. TEUKOLSKY, W. T. VETTERLING, AND B. P. FLANNERY, *Numerical Recipes: The Art of Scientific Computing*, 3rd ed., Cambridge University Press, New York, 2007. (Cited on p. 1158)
- [40] M. J. PRINCE AND R. M. FELDER, *Inductive teaching and learning methods: Definitions, comparisons, and research bases*, J. Eng. Educ., 95 (2006), pp. 123–138. (Cited on p. 1164)
- [41] L. RAVIV AND B. THOMPSON, *A Short Tutorial on Agent Based Modeling in Python*, <https://github.com/Limor-Raviv/Tutorial.Agent.Based.Models> (accessed 2021-07-26). (Cited on p. 1155)
- [42] T. REICHENBACH, M. MOBILIA, AND E. FREY, *Mobility promotes and jeopardizes biodiversity in rock-paper-scissors games*, Nature, 448 (2007), pp. 1046–1049, <https://doi.org/10.1038/nature06095>. (Cited on pp. 1164, 1165)
- [43] G. REYNDERS, J. LANTZ, S. M. RUDER, C. L. STANFORD, AND R. S. COLE, *Rubrics to assess*

- critical thinking and information processing in undergraduate stem courses*, Internat. J. STEM Educ., 7 (2020), pp. 1–15. (Cited on p. 1166)
- [44] J. R. SAVERY, *Overview of problem-based learning: Definitions and distinctions*, in Essential Readings in Problem-Based Learning: Exploring and Extending the Legacy of Howard S. Barrows, A. Walker, H. Leary, C. Hmelo-Silver, and P. Ertmer, eds., Purdue University Press, West Lafayette, IN, 2015, pp. 5–15. (Cited on p. 1153)
- [45] T. C. SCHELLING, *Dynamic models of segregation*, J. Math. Sociol., 1 (1971), pp. 143–186. (Cited on p. 1162)
- [46] P. B. A. SMITS, J. H. A. M. VERBEEK, AND C. D. DE BUISONJÉ, *Problem based learning in continuing medical education: A review of controlled evaluation studies*, Brit. Med. J., 324 (2002), pp. 153–156, <https://doi.org/10.1136/bmj.324.7330.153>. (Cited on p. 1164)
- [47] L. SPRINGER, M. E. STANNE, AND S. S. DONOVAN, *Effects of small-group learning on undergraduates in science, mathematics, engineering, and technology: A meta-analysis*, Rev. Educ. Res., 69 (1999), pp. 21–51. (Cited on pp. 1154, 1156)
- [48] P. VAN LIEDEKERKE, M. PALM, N. JAGIELLA, AND D. DRASDO, *Simulating tissue mechanics with agent-based models: Concepts, perspectives and some novel results*, Comput. Part. Mech., 2 (2015), pp. 401–444. (Cited on p. 1162)
- [49] D. J. WILKINSON, *Stochastic Modelling for Systems Biology*, Chapman and Hall/CRC, Boca Raton, FL, 2006. (Cited on pp. 1156, 1167)
- [50] D. J. WILKINSON, *Stochastic modelling for quantitative description of heterogeneous biological systems*, Nat. Rev. Genet., 10 (2009), pp. 122–133. (Cited on p. 1154)
- [51] D. F. WOOD, *Problem based learning*, Brit. Med. J., 326 (2003), pp. 328–330, <https://doi.org/10.1136/bmj.326.7384.328>. (Cited on p. 1153)