

The Making of An Academic Programme in Modeling and Simulation

Mihir Arjunwadkar^{1,2} and Abhijat M. Vichare³

¹ Centre for Modeling and Simulation, University of Pune, Pune 411007, India

² National Centre for Radio Astrophysics, Pune 411007, India

³ Computational Research Laboratories Ltd., Pune 411004, India

Abstract

While mathematical modeling is not a new enterprise, the ubiquity of computing technologies has made it all-pervasive. One may argue that this is precisely the reason why we need a multi-disciplinary, broad-based academic programme in mathematical modeling and computational simulation (M&S) methods, with a problem-centric approach at its heart.

Designing and implementing such a multi-disciplinary academic programme without precedent in the arena of conventionally stratified academic disciplines can be a challenging task. It requires a strong sense of vision coupled with commitment, endurance coupled with flexibility, and team work coupled with wise leadership and organizational skills.

In this paper, we describe the design and development of a Masters-level academic programme in Modeling and Simulation at the Centre for Modeling and Simulation, University of Pune, and first implemented in the academic year 2008-09. The long and arduous process that culminated into this programme began with brain-storming within a core group, a rigorous design exercise, followed by informal expert reviews that led to an initial one-year diploma implementation. Gradual and incremental experience accumulated from running this diploma programme, together with design-redesign cycles, readily available academic and organizational help – and some luck – eventually brought the programme to its current form, i.e., a full-length two-year M.Tech. degree programme in modeling and simulation. We also share our outlook and some of the wisdom gained in the process.

Perhaps the single most important take-home message of our exercise is that curriculum design and implementation in any discipline is a serious endeavour that requires commitment on part of the faculty, long-term organizational commitment, will, and support, and a focus on the overall objective by all.

Keywords: Modeling and Simulation – Computational Science – Problem-Centric Approach – Applied Mathematics, Applied Statistics, and Computing Education – Multidisciplinary Academics – Curriculum Design – Higher Education

This article was contributed to the [NIME National Conference on Mathematics Education, Homi Bhabha Centre for Science Education \(TIFR\), Mumbai, India \(2012\)](#), and formally published as part of its [proceedings](#).

1 Mathematical Models, Computing, Simulation

To set stage for what we are going to discuss here, let us first consider a few “real-life” problems:

1. At the current epoch, India is getting urbanized at an ever-growing pace. Given the fast growth in privately-owned automobiles, this has invariably led to dense, chaotic, and often dangerous traffic, a prime example of which is the city of Pune, India. The design of efficient urban transport systems that

maximize human throughput while minimizing costs (including technological and environmental) is clearly a need of the time.

2. Since the “liberalization” of the Indian economy, an Indian chocolate manufacturer has the choice, at least in principle, of buying sugar, amongst other raw materials, from around the world. Depending on the state of the world economy and the supply-demand forces, prices of commodities fluctuate over time. Based on past trends in the pricing of the raw materials and the expected demand of chocolates in the near future, the manufacturer needs to make a decision now as to when and where to buy sugar, and how much. Clearly, the manufacturer’s end goal is to maximize profit.
3. The spatio-temporal dynamics of infectious diseases such as bird or swine flu, malaria, rabies, tuberculosis, etc., has profound implications to human societies across the world. Policy frameworks for dealing with disease outbreaks on a mass scale necessarily require a detailed understanding of this dynamics.

These illustrative examples above (and, say, a range of similarly-spirited examples in [Steen, 2000](#)) come from domains that apparently have nothing to do with one another. However, many such real-life problems and their reasonable solutions share many common features. First and foremost, there is a need to understand the system, be able to predict its behaviour, and be able to control it, each at an appropriate level of precision. The first two goals are primarily associated with the “science” part of the process, while the third is predominantly in the technological premises. “Understanding” usually implies identifying prominent patterns of behaviour of the system; in other words, discovering or inferring the “laws” that govern the system. Second, in each of these examples, there is a need to go beyond the qualitative, and understand the problem in a quantitative manner. We therefore need to be able to describe patterns in a flexible, economical, and quantitative manner. We may also need the ability to formulate descriptions the same system at varying levels of detail and precision. Perhaps the only “language” that allows describing patterns in this fashion and making inferences about them without actually having to observe the system endlessly is what we call *Mathematics*.

The fact that mathematics is being used for solving problems in almost every domain of human activity is not a new revelation; in fact, [Bickley \(1964\)](#), e.g., pointed out some five decades ago that “Mathematics is relentlessly seeping into the very foundations of the civilisation in which we live and the community of which we are part” Moreover, this game of using mathematics for describing and “understanding” patterns is even more ancient, as evidenced by its use in the natural sciences; in fact, we could go as far as saying that mathematics was perhaps *born* this way in the depths of time.

A mathematical description of a system under scrutiny, referred to as the *mathematical model* of the system, is built by extracting features of the system that are most relevant to the problem at hand. This is the process of abstraction: the actual system is now replaced with its model which tries to capture the essential features of the system that are most relevant for the problem at hand. Understanding the system with reference to the problem at hand is now equivalent to understanding the mathematical model. Mathematical models therefore need to mimic the system sufficiently accurately, which implies that models need to be validated against reality. Better understanding of the system, arrived at through observation, invariably leads to changes in its mathematical model, some times through generalization of the model to incorporate finer nuances of the system’s behaviour, and some times through outright rejection of the current model. This process is perhaps best exemplified by the how mathematical models of the physical world (popularly referred to as theories of physics) changed in the light of new observational evidence. The fact that mathematical models can at best be accurate, but never exact, representations of the reality, is most perceptively expressed in the oft-quoted maxim “Essentially, all models are wrong, but some are useful” and its variant, namely, “Remember that all models are wrong; the practical question is how wrong do they have to be to not be useful.” ([Box and Draper, 1987](#); see also [Box, 1976](#)).

The mathematical structure of a model gets invariably more and more involved as finer and finer nuances of the system’s behaviour get incorporated into the model. This is where pure analytical reasoning become

more and more difficult, if not impossible, and greater mathematical sophistication is required to deal with the greater complexity of the model. This is where computation and simulation make themselves indispensable. From a certain perspective, theory of computation itself could be considered a branch of mathematics.

A simulation attempts to represent a model of the real-life system under study using some other well-understood system, the *simulation system*. A natural and versatile choice for a simulation system is the contraption called a computer that performs computation. To quote [Bickley \(1964\)](#) again, “Similar things were happening in the pre-computer age. ... But the computer *has* increased our powers and widened our scope. In particular, it enables us to conduct mathematical experiments and to construct mathematical models on a vast scale ...” The purpose of a simulation system is to somehow mimic the behaviour of the real-life system under study, and more often than not, this involves the use of a mathematical model of the system. A simulation is thus built on three components, namely, know-how from the problem domain, a (mathematical) modeling formalism, and methodologies and technologies specific to the simulation system.

Another important feature of this problem-centric approach is that we take the view that a problem is a problem in need of a reasonable solution; it doesn’t care how we choose to classify it. While traditional knowledge domains have evolved for a good reason into the current organization as sciences, humanities, engineering and technology, etc., the fact that problems arising in disparate knowledge domains could have common mathematical structure necessarily advocates cutting across traditional knowledge domains. For example, what is common to the spread of a disease, the spread of a rebellion, and chemical reaction kinematics? It turns out that commonly-used mathematical models of these systems all turn out to be systems of coupled ordinary or partial differential equations. A breakthrough in one field can thus benefit a completely disparate field because of this underlying similarity of description at the mathematical level.

2 Why Create An Academic Programme in M&S

While mathematical modeling is conceivably as ancient as mathematics itself, what makes it all-pervasive in modern times – especially the last few decades – is the ubiquity of computing technologies and the ever-increasing availability of comparatively inexpensive computing power. Computation and simulation, in particular, is often seen as a partial substitute for expensive experimentation with the actual system. This has created somewhat specialized niches for people with adequate mathematical background, together with reasonable analytical and computational skills, in almost every area of human activity that benefits from mathematical modeling. Classic examples of this kind include drug discovery, design of automobiles and aircraft, mechanics of engineering structures such as a bridge, pattern discovery in biological sequences, and so on and so forth.

One may argue that this is precisely the reason why we need a multi-disciplinary, broad-based academic programme in mathematical modeling and computational simulation (M&S) methods, with a problem-centric approach at its heart.

3 The Centre’s M.Tech. Programme in M&S

The Centre for Modeling and Simulation, University of Pune, was formally established in mid-2003 with a vision to “promote, support, and facilitate academic and research activities related to mathematical modeling and computational simulation and, in particular, the use of computation as the third scientific methodology (besides theory and experiment)”; “to aggressively promote a problem-centric outlook to real-life problems, and highly multidisciplinary approaches that transcend traditional boundaries separating individual scientific disciplines”; “to keep up with the state-of-the-art in computing and, in particular, develop strong expertise computing technologies such as high-performance computing, grid computing, etc.”; and “to create excellent, versatile minds that are capable of learning by themselves, of thinking deeply, of questioning dogma and authority, and of seeing beyond the immediate.” ([Arjunwadkar et al., 2008](#)).

The Centre became functional with the appointment of its first two faculty members in December 2003. An informal core group consisting of local computational scientists (Arjunwadkar et al., 2005, see History and Credits) in addition to the Centre’s faculty became operational almost immediately, and engaged itself in brainstorming about trends in scientific research, engineering, technology, and beyond, with focus on computation, modeling and simulation. Many spirited discussions convinced us that designing broad-based curricula in M&S was not an unreasonable thing to do, even if no such programme existed at that point in time.

The Centre’s curriculum design exercise was inspired by the methodical and comprehensive outlook reflected in *Computing Curricula 2001* (The Joint Task Force on Computing Curricula, 2001). Several months of brainstorming and design-redesign cycles eventually led to the creation of the *Advanced Diploma Programme in Modeling and Simulation* (Arjunwadkar et al., 2005). It was a conscious and practical decision to first create a one-year programme so as to get better experience in an unprecedented academic territory, implementation logistics, etc. Experience gathered from running this one-year diploma programme for three consecutive batches (2005-08) made us confident enough to dive into the greater complexities of a full-length (i.e., two-year) academic programme.

The full-length programme in M&S eventually came to be called the *M.Tech. Programme in Modeling and Simulation* (Arjunwadkar et al., 2007). We requested many practising experts to review the programme simply to see where it stands in the eyes of both the academic/research community as well as the industry/corporate sector, and found that it was positively received by all with genuine interest. In fact, we received more positive suggestions than what is logistically possible (e.g., suggestions for courses on topics relevant to specific domains). Most of the feedback thus obtained was incorporated in the structure and curriculum prior to its deployment in the academic year 2008-09. This voluntary review initiative on part of the Centre was over and above the University’s formal requirements for approval of an academic programme.

The documents cited above that describe both these programmes are quite detailed and publically available. In the rest of this paper, apart from providing a crisp overview of the M.Tech. programme in Sec. 4, we will therefore try to avoid duplicating the extensive material contained in these documents, and instead focus on salient features of the structure of the full-length programme, design considerations that went into its making, and the lessons learnt in the process. We recommend treating the present article as a companion to these extensive documents.

4 Structure and Curriculum: An Overview

The M.Tech. Programme in M&S designed and implemented by the Centre for Modeling and Simulation, University of Pune, is full-time 2-year post-Masters (or post-Bachelors in Engineering) programme. The only other pre-requisite for admission in the programme is adequate mathematics background approximately at the first-year level of a typical Indian science Bachelors programme. Core curriculum of the programme is by-and-large focused on methodologies, and is founded on the four principle pillars of applied mathematics, applied statistics, computing, and M&S in practice.

The first three semesters of the programme are devoted to coursework, the fourth to a rigorous full-time project involving hands-on M&S work in any knowledge domain. The training and outlook of the incoming students can vary substantially because of their diverse backgrounds and maturity levels. As such, the first semester of the programme is primarily geared to building a uniform background in essential mathematics and computing. Second and third semesters, in addition to the core courses, also include a selection of elective courses on somewhat specialized topics such as computational fluid dynamics, machine learning, etc.

5 Rationale for the Structure and Curriculum

The rationale for the structure and curriculum could be viewed from three perspectives; namely, (a) a student's perspective, (b) the academic view, and (c) the overall education system. The considerations presented below, which were realized through trial-and-error and continuous introspection, offer some rationale for how our M&S programme came to be in its present form. We present a loose collection of specific insights in the next section.

From a Student's Perspective. At the post-Masters level, most prospective students tend to evaluate degree programmes from the point of view of job prospects and value addition. Employment opportunities for a graduate of a M&S programme are likely to come from:

- R&D centres and analysis organizations (in the industrial, defense, academic, government, and other sectors), that use modern computational and simulation methodologies in their design, development, and research/analysis initiatives. Such initiatives include all areas of engineering, science and technology including (but not limited to) materials science and engineering, nanotechnology, bioinformatics and biotechnology, computational fluid dynamics, molecular modeling and drug design, process engineering, finance, etc.
- Research- and computing-oriented *support* positions in R&D or research-and-analysis organizations.
- Research programmes leading to an advanced degree such as Ph.D.

Although some the students may find careers in the IT-related and conventional software industry after undergoing M&S training, that should not be the focus of the programme. However, we do envisage that the skill set and attitude developed through the programme – specifically, problem-solving and programming skills, ability to learn on the fly, and versatility – will enable a student to migrate easily to such careers.

How does this consideration dictate the design and the content of the curriculum? The diversity of problem domains where M&S methodologies can be employed meaningfully argues in favour of a broad-based programme that is primarily focused on methodologies, together with moderate specialization, and flexibility to respond to technological advances as well as “market” needs. This is what led to the inclusion of elective courses in our M&S curriculum.

From an Academic Perspective. Quite independent of M&S considerations, the idea that the academic goals of an academic programme may be characterised in terms of *competence* and *awareness* is useful for curriculum design and content. The curriculum content for an M&S programme is therefore dictated by a set of concepts and skills that are central and necessary for *competence* in M&S, along with another set of concepts and skills that a student is expected to be *aware* of. As mentioned before, a simulation is built on domain expertise, (mathematical) modeling formalisms, and methodologies/technologies specific to the simulation system. Of all the possible combinations of these three components and the two goal levels (competence, awareness), our curriculum is primarily geared towards training students to be skilled (as opposed to conversant/aware) in the last two components, but be aware of/conversant with a few techniques from the first. This choice was made primarily because all possibilities that require building additional domain expertise into the curriculum quickly become impractical for such a programme.

From a Systems Perspective. At the risk of over-generalization, we may say that our education system is organised in the following hierarchy of levels: school, college/university, and the research establishment. Our M&S programme is a part of this whole, and therefore must take into account the inter-dependencies with other levels, as well as with the greater socio-economic reality. This is important because, in our system, failures at one level propagate unhindered to the next. In the context of any programme in higher education, this needs

to be accounted for as much as purely academic considerations about the required preparation at the time of entry to a programme such as ours.

For example, we tend to expect that a student be mature and independent at the post-Masters level. This expectation fails often. This suggests that there should be reasonable corrective mechanisms built into curricula as well as in the organizational ethos and environment. Our M&S curriculum therefore includes a small fraction of modules on written and oral communication/presentation skills. (Incidentally, we have also experimented – sporadically and not methodically – with arranging soft-skills workshops, including unconventional themes such as time management and creativity training, with some positive effect on the students.)

A subtler issue is, however, also a better-known one: “learning to learn” versus “learning by rote”. It is well-known that most seriously minded places of higher learning need to spend significant effort to help students make this transition to “learning to learn”. In our context, for example, the technological component of the programme is the one that is changing at the fastest rate. This implies, among other things, that having programming skills in a given language is not enough; what matters is the ability to learn new computing paradigms and languages as and when required.

The best solution for this is to build an organizational ethos that encourages this in a genial but serious manner, and design in-class and out-of-class activities to this end. Excellent faculty and staff, together with a research-oriented environment, is one clear route to this end.

6 Assorted Notes

Why a Post-Masters Programme? Post-Masters (or post-Bachelors in Engineering) is perhaps the right time for specialized training in M&S for the following reason: by then, a student has hopefully acquired sufficient knowledge from a specific knowledge domain, together with some level of mathematical, problem-solving, critical thinking, and algorithmic thinking skills, and has matured enough to understand the utility of M&S in his or her knowledge domain.

Hierarchical Content Organization. The entire content of the programme is organized on three levels of hierarchy: a *programme* consists of *courses* which in turn comprise of *modules*. A *module* is defined here as an indivisible/logical unit of content/instruction that can be meaningfully handled by a single instructor. This hierarchical organization has many advantages; specifically, it makes it possible to share modules across different programmes, allowing for sharing of instructor resources. Also, for an organization that is forced to rely heavily on visiting teachers, finding a teacher for a single module is much easier compared to finding a teacher for an entire semester-long course.

Syllabi and Student Assessment. In the original conception, the syllabi included in the programme document (Arjunwadkar et al., 2007) are considered indicative of the overall scope and focus of a module, and not as rigid, sacrosanct entities that cannot be touched or altered. The actual module content can be decided by the instructor, explicitly trusting his or her expertise and judgement. In the best-case scenario, this helps alleviate the problem of dead, outdated syllabi. Similarly, evaluation and assessment have been left to the discretion of a competent teacher. With this well-intended freedom for the instructor, the entire onus then rests on finding an able person with appropriate expertise and knowledge, pedagogic skills, maturity, and academic instincts to teach a module.

The Core Courses. Developing survival skills in applied mathematics and statistics is quite rigorously included into the curriculum. Commonly-required areas of analysis (calculus, complex analysis, linear algebra, vector analysis, etc.), and probabilistic reasoning (probability theory, statistical inference, stochastic “simulation” methods) are the focal points of the first year of the curriculum (apart from formalisms expressly useful

for modeling, such as differential equations, numerical analysis, etc.). These are handled in the first year with the dual goal of developing a strong base in modeling as well as to iron out the differences in the preparation levels of students at entry. Computing skills are predominantly geared toward developing algorithmic thinking, and large-scale code reading and writing capabilities. A mathematical formalism is expected to be presented in a three-fold fashion: domain contexts and application in which it is used, mathematical results related to the formalism with focus on concept and visualization rather than mathematical rigour, and related computational methods.

Teaching the “Art” of Mathematical Modeling. By “art” aspects, we mean, e.g., the difficult-to-teach process of arriving at mathematical models in the context of a specific problem. In our experience, the “art” aspects of mathematical modeling and the problem-centric approach are best conveyed and emphasized by exposing the students to a diverse range of problems from many domains. M&S practice ultimately rests on domain knowledge. Given the wide variety of domains in which M&S can be employed, the challenge is primarily in breadth-versus-depth trade-off optimized for the finite time available for instruction. This is best done by organizing colloquia, informal talks, and interactive case-study sessions conducted by practising experts from academics and industry alike, where the expert attempts to illustrate the advantages, techniques, and limitations of M&S in the his/her specific context.

Concept Versus Mathematical Rigour. Perhaps too much emphasis on mathematical rigour in conventional mathematics programmes, together with distancing of mathematics in typical curricula from the “reality” that mathematics can often so well describe, is believed to have harmed mathematics more than anything else: A scathing expression of this extreme view asserts that “mathematics is far too important to be left to mathematicians” (Bickley, 1964). We do not subscribe to this view; in fact, all those involved in the development of this programme have always had a deep respect for mathematics and mathematicians. However, keeping in mind that this programme is not a programme in mathematics, we find it useful to take (and propagate) the view that we are users of mathematics at varying levels of mathematical sophistication. What matters from a practitioner’s point of view is the ability to grasp a mathematical result or concept in its essence, the ability to visualize mathematical constructs, and a reasonable judgement on whether one should trust one’s own mathematical instincts or get help from a real expert. This implies, among other things, putting emphasis on concept instead of proof, and diving into the nitty-gritties of a proof only if it helps understand concepts better. This also puts a lot of burden on an instructor to find innovative ways of illustrating and conveying a concept. (In our limited view of literature, this outlook has been emphatically stated in Wasserman, 2004.)

Learning At a Comfortable Pace. One serious confound in traditional degree programmes, from a student’s perspective, is the expectation that it is “not good” to take longer than the stipulated time period for the programme. In the rapidly changing socio-economic circumstances of this country, we are seeing students who need to support themselves financially in ever-increasing numbers. Although it is in the interest of a student to complete a degree programme in the least possible time, traditional expectations and value judgements often take a toll on a student who is serious about what (s)he wishes to learn, but cannot devote full time to education. As a standard practice at the Centre, we recommend (and help) such students to be realistic about the time they can put in, and take a judgement on how much course work they can possibly take in any semester, without putting any psychological stigma on completing the programme in more than two years. The expectation, in return, is that they excel in whatever they choose to learn.

A Wish List for the Future. First, we had realized very early on that a large compendium of M&S case studies would be immensely useful for this programme. In particular, each case study should begin in the problem domain, explore possible modeling alternatives and their relative merit with respect to the questions

that need to be answered, whether analytical treatment of the chosen model is feasible, and if not, how does it boil down to computation/simulation. Such a “Handbook of M&S” would be a useful resource for the programme. Second, possible additions to the curriculum could usefully include a module on Fermi-like approximate and opportunistic reasoning (see, e.g., [Mahajan, 2010](#)). Third, the sector that has consistently shown the greatest amount of interest in the programme consists of working professionals who wish to further their knowledge in M&S out of necessity or interest. The logistics of making the programme available in a part-time or a web-based virtual mode is clearly enormous. This virtual mode has been on the Centre’s wish list ever since the regular programme was deployed in 2008-09.

M&S in the Greater Context of Mathematics Education One of the challenges in mathematics education has been to debunk the common myth that mathematics and formal sciences are “abstract” and “difficult”, and worse still, not “applicable”. We feel that this is one problem that must be addressed right from the school level. The modeling contexts associated with mathematical constructs and formalisms, and the M&S enterprise in general, is a potent candidate to dispel these myths. It might also be useful, e.g., to consider including elementary Logic at school and graduate levels. The “PQ–” logic ([Hofstadter, 1979](#)) could be a good vehicle around which concept levels can be designed for instruction at school and graduation levels. A significant barrier to mathematics could then possibly be overcome.

Acknowledgments

The M&S programme described in this paper could not have materialized if we did not have Prof. D.G. Kanhere as the founder director of the Centre for Modeling and Simulation, University of Pune. In particular, his uncanny knack to offer the right help and practical advice at the right time, his trust for a (junior) colleague’s academic instincts, his wisdom in dealing with a variety of University situations, and his overall outlook on education have been an integral element in the making of this programme. We would like to thank the organizers of the 2012 National Conference on Mathematics Education (R. Ramanujam and Aaloka Kanhere in particular) for their patience and encouragement. We also take this opportunity to thank our current and former colleagues at the Centre, namely, Sukratu Barve, Prashant Gade, and Kavita Joshi, for intense discussions. Abhay Parvate, another former colleague who was deeply involved in this endeavour and contributed to the M&S programme and the Centre in innumerable ways, should have been a co-author on this paper. He, however, felt that he has not contributed to this paper enough to warrant authorship and, consequently, declined to be a co-author out of modesty and a strong sense of professional ethics.

AMV is required to include the following statement on behalf of his employer: “The opinions expressed in this article are the authors’ personal opinions and views, and are in no way endorsed by Computational Research Laboratories Ltd.”

About the Authors

Mihir Arjunwadkar, a computational physicist by training, is currently on leave from the Centre for Modeling and Simulation, University of Pune, and is a visiting professor at the National Centre for Radio Astrophysics, Pune.

Abhijat Vichare, a computational physicist with deep interest in computation and computer science and a member of the core team that designed the Centre’s academic programmes, is now a scientist at the Computational Research Laboratories Ltd., Pune.

References

- Mihir Arjunwadkar, Dilip G. Kanhere, and P.M. Gade. Advanced Diploma Programme in Modeling and Simulation. Public Document CMS-PD-20050502, Centre for Modeling and Simulation, University of Pune, Pune 411007, India, 2005. Available at <http://cms.unipune.ac.in/reports/pd-20050502/>.
- Mihir Arjunwadkar, Abhay Parvate, Sukratu Barve, P.M. Gade, and Dilip G. Kanhere. Master of Technology (M.Tech.) Programme in Modeling and Simulation. Public Document CMS-PD-20070223, Centre for Modeling and Simulation, University of Pune, Pune 411007, India, 2007. Available at <http://cms.unipune.ac.in/reports/pd-20070223/>.
- Mihir Arjunwadkar, Abhay Parvate, and Dilip G. Kanhere. Centre for Modeling and Simulation: Excellence in Academics, Research, and Outreach (2003-08 Final Report Submitted to the UGC). Public Document CMS-PD-20080218, Centre for Modeling and Simulation, University of Pune, Pune 411007, India, 2008. Available at <http://cms.unipune.ac.in/reports/pd-20080218/>.
- W. G. Bickley. Mathematics for Engineering Students. *The Mathematical Gazette*, 48(366):379–383, 1964. Available at <http://www.jstor.org/stable/3611695>.
- George E. P. Box. Science and Statistics. *Journal of the American Statistical Association*, 71(356):791–799, 1976. Available at <http://www.jstor.org/stable/2286841>.
- George E. P. Box and Norman R. Draper. *Empirical Model-Building and Response Surfaces*. Wiley, New York, 1987.
- Douglas Hofstadter. *Gödel, Escher, Bach: an Eternal Golden Braid*. Basic Books, 1979.
- Sanjoy Mahajan. *Street-Fighting Mathematics: The Art of Educated Guessing and Opportunistic Problem Solving*. MIT Press, 2010.
- Lynn Arthur Steen. Revolution by Stealth: Redefining University Mathematics. In Derek Holton, editor, *Teaching and Learning of Mathematics at the University Level (Proceedings of the ICMI Study Conference, Singapore, 1998)*, pages 303–312. Kluwer Academic Publishers, 2000. Available as <http://www.stolaf.edu/people/steen/Papers/stealth.html>.
- The Joint Task Force on Computing Curricula. Computing Curricula 2001. *Journal on Educational Resources in Computing (JERIC)*, 1(3):1–240, 2001. doi: 10.1145/384274.384275.
- Larry Wasserman. *All of Statistics: A Concise Course in Statistical Inference*. Springer, 2004.