Technologies for College-level Quantitative Education

Daniel T. Kaplan

2026-04-09

Abstract: It is a truism that today's students will go out into a highly quantitative and data-oriented world. The sequence leading to and through calculus provides hardly any preparation for this; it is designed around the technology of algebraic manipulation that predates the modern world or modern capabilities for expressing mathematical ideas. Similarly, the standard elementary statistics course does not cover quantitatively issues of contemporary importantance such as causality, confounding, and covariates. It does not strongly engage the important technologies of the modeling process, of the analysis of the interconnected components of systems, or the quantitative representation of hypotheses and beliefs.

This paper proposes goals for quantitative education that better align with modern needs. I mark those goals as college-level, defined in a way that avoids the traditional pattern in which each course's goal is to prepare for the next in the long-established sequence. The paper then identifies several mathematical technologies relevant to serving the modern goals and discusses some often-made objections to innovation.

Introduction

Mainstream secondary and college-level quantitative education is divided into two separate themes. The older of the two is the GATC sequence that starts in high school: geometry, algebra, trigonometry, and calculus. The sequence continues within calculus: Calc I is mainly about differential calculus, Calc II features integrals, and Calc III deals with multi-variable calculus. Then, changing naming conventions, differential equations, and linear algebra. (It is not uncommon to intermix topics from these five canonical calculus courses.)

The other theme is statistics. The dominant organization is an elementary statistics course that draws only on the formula reading and on squares and square roots introduced in early algebra. There is usually no standard follow-up course. Sometimes the algebra needed is minimized by the use of technology; for instance, software for t-tests or "simulation" techniques such as resampling and permutation testing.

There is no clear dividing line between high school and college-level quantitative education. Large numbers of students take calculus or statistics in high school; larger numbers take much the same statistics as well as algebra and pre-calculus in college.

The framework of the GATC mathematics curriculum has been established for decades, even centuries. Technical innovations are made on the margins, for example, using spreadsheets to calculate Riemann sums or computer algebra systems to calculate integrals and derivatives and draw graphs. But students are still factoring second-order polynomials, symbolically manipulating logarithms and exponentials, calculating limits, applying the chain rule, product rule, and quotient rule for derivatives, and using u-substitution and trigonometric substitution for integration. Functions of one variable remain the norm even through the early college level.

The material in mathematics textbooks from 1900 is easily recognizable to today's instructors. This contrasts with almost every other field, where textbooks must eventually keep up with developments in emphasis, techniques, and areas of application. To do so in statistics involved side-stepping the mathematics curriculum. But, even in statistics, where there has been steady innovation and introduction of new concepts and techniques, the content of elementary courses dates mainly to 1940 and before.

In this paper, I look at the prospects for innovation in early college-level quantitative education. I use the term "quantitative education" rather than "mathematics" to escape from the common notion that the topics of mathematics education are and ought to be fixed and unchanging. I regard the traditional, long-standing math topics as technologies for achieving educational goals rather than the goals themselves. Technology changes. Old technologies can be compared to new ones and, if appropriate, as it so often is outside of mathematics, the old can be replaced with the new.

The focus on "early college level" education is not meant to discount the potential need for education at the primary and secondary levels. However, I lack expertise in developing number, geometric, and other mathematical senses in young students. I also see very strong, institutional, conservative forces that can easily stymie proposals for innovation in the primary and secondary domains. At the college level, however, decentralization leads to much greater scope for innovation. Similarly, with few exceptions, no expectations are imposed on most college graduates to fulfill the mathematical prerequisites for the next stage of education (if any).

Defining college-level

In institutional practice, the shorthand for "college-level" amounts to "calculus." Statistics is often seen as an alternative for students who are not prepared for calculus. But defining "college-level" in these ways makes it difficult to think comprehensively about quantitative education.

I think "college-level" would be better defined in operational terms: the kind of quantitative thinking students should be exposed to, rather than traditional topic categories. I offer three interrelated perspectives on this: closed-ended versus open-ended investigations; deductive versus inductive reasoning; "model" as an object versus "modeling" as a process of investigation.

Open-ended investigations engage judgment, domain expertise, appraisal, testing, and revision in response to identified flaws: a process. Conducting, interpreting, and evaluating the quality of open-ended investigations is key to addressing many important social, medical, scientific/technical, and design questions. This implies that at some point in quantitative education, educators should focus on the techniques, questions, practices, and attitudes that serve open-ended investigation.

In contrast, "closed-ended" problems are those for which there is a correct answer, all other answers being incorrect in some more or less grievous way. In mathematics education, the hallmark of closed-ended problems is an emphasis on deduction, constructing a set of irrefutably logically connected statements.

Open- and closed-ended do not stand in absolute opposition to one another. Among other things, the technologies that address open-ended problems, such as those listed in **?@sec-technology**, themselves need to be learned. Posing closed-ended problems in the use of those technologies can be an effective pedagogy.

Goals of quantitative education

Ideally, educators would come together to form a consensus statement of goals that embraces important ideas and needs from many disciplines. Unfortunately, consensus often comes from using a high level of abstraction in expressing goals. This allows different disciplines to avoid conflict and the need to set priorities and resolve trade-offs. To illustrate the situation, consider these two (of ten) bullet points produced by the vague AI prompt, "Goals for studying calculus."

- "Problem-Solving & Modeling: Develop the ability to translate real-world scenarios into mathematical models and solve them, often using technology to gain insights."
- "Critical Thinking: Rigorously analyze complex problems and evaluate the reasonableness of solutions (checking signs, units, and accuracy)."

Perhaps AI identified these since the phrases "problem solving" and "critical thinking" appear so often in discussions of educational priorities and goals. Without other justification, I'll treat them as components of a pan-disciplinary, consensus view. The problem arises when trying to assign a specific meaning to such highlighted phrases. In my experience, educators differ markedly in interpreting "real world," "solve," "rigorously," "complex,"often," "reasonableness," and "solutions. In most calculus texts, for example,"real-world scenarios" include finding the optimal, metal-saving shape of a cylindrical can or calculating the volume of a "lake" whose shape profile is a solid of revolution of a parabola. The phrase “using technology” is particularly vexing, as it is used to avoid distinguishing between mouse-driven graphical apps, computing syntax, and even the graphing calculators that persisted for too long in mathematics classrooms.

Institutional segregation of disciplines hides conflicts of meaning. A statistical educator might point to “accuracy” as meaningful mainly in counterpoint to statistical “precision,” whereas mathematics teachers might focus on the correctness of an answer, the number of digits reported, or the virtues of using a limiting process rather than finite elements. Similarly, “modeling” is a popular word across disciplines for framing quantitative education goals. Properly, modeling should be an iterative, comparative process that considers several aspects of the domain and the intended uses of the analysis. For instance, tin-can modeling should incorporate the factors that lead to cans of tuna being short, cans of soup being taller, and cans of sardines being utterly different, as well as the costs of shaping and crimping or seaming the metal. In contrast, the common educational expression of quantitative “modeling” is writing a precise analytic expression that will be said to serve as an approximation to whatever the unimportant realities might be.

For the purpose of evaluating what technologies might be most helpful in learning quantitative reasoning, I will work with a set of educational goals that have emerged from my career experiences of teaching and textbook writing in several disciplines: engineering, statistics, computing, epidemiology and medicine, applied mathematics, economics, and quantitative literacy. [GIVE CITATIONS TO TEXTBOOKS] These are not meant to be comprehensive and exclusive, but they have shaped my thinking about evaluating quantitative technologies.

Most readers of this journal will be sympathetic to the idea that reasoning about information presented statistically is essential to effective participation in modern society, whether that be interpreting medical screening test results or evaluating the quality of claims about causal connections. Possibly fewer readers would see calculus as similarly essential. Perhaps this is because calculus, as offered in college, is mainly about algebraic manipulations. Even calculus students rarely understand the actual uses of calculus in the non-mathematical world. As a brief argument, consider this excerpt from C.P. Snow’s famous 1956 lecture, The Two Cultures:

“[In response to] an even simpler question – such as, What do you mean by mass, or acceleration, which is the scientific equivalent of saying, Can you read? – not more than one in ten of the highly educated would have felt that I was speaking the same language. So the great edifice of modern physics goes up, and the majority of the cleverest people in the western world have about as much insight into it as their neolithic ancestors would have had.”

In proposing the following list of goals for college-level quantitative education, I have kept in mind the importance of both calculus ideas and statistical reasoning. Mathematicians might argue that my non-inclusion of the “infinitesimal” means that I am not addressing calculus.

1. Learn elements of a genuine process of modeling appropriate to the scientific method.
2. Learn to quantify uncertainty and risk and to distinguish between accuracy and precision.

3. Learn strategies for working with systems. By "system," I mean a collection of components connected by relationships.
4. Become familiar with several important ways to represent and quantify relationships. Among these are ratios, functions, units and dimension, alignment, and rates of change as in the differentiation and accumulation operations of calculus. The "training" paradigm is also highly relevant.
5. Learn how to state hypotheses and judge to what extent data can serve to identify likely hypotheses. "Likely" here is a code word, intended to generalize the loaded phrase "statistical inference." For instance, one expression of "likely" involves identifying and interpreting Bayesian reasoning. The Frequentist framework also uses likelihood.
6. Learn how to make and evaluate responsible statements about causal relationships and avoid irresponsible claims. This has come to attention in the last 30 years owing to the advent of formal techniques for "causal inference." [CITE Book of Why] But it has for much longer been an important issue for people who need to make decisions and take action. The viewpoints expressed by Austin Bradford Hill in the 1960s, then president of the Royal Statistical Society, remain highly germane in epidemiology and medicine, even if they are not widely known among statistics educators. [Hill, 1965] This goal is closely related to goals (3) & (4).

"Modeling process" versus "modeled by"

The first goal in the previous chapter relates to the modeling process. It is important to emphasize the suffix "…ing" and the word "process" mark an important distinction. On the one hand is the sort of activity that supports open-ended investigations, a process of iteratively constructing and revising the mathematical representation of the situation at hand. On the other hand is a closed-ended statement of a mathematical problem that is intended to evoke a sentiment that the problem is related in a useful way to a practical application.

Consider this evocation that uses the phrase "modeled by":

> [A] pond is nearly circular, is 20 feet deep at its center, and has a radius of 200 feet. The bottom of the pond can be modeled by
>
> $$y = 20\left((0.005x)^2 - 1\right).$$
>
> How much water is in the pond? — Source: (R Larson and BH Edwards (2006) *Calculus: An Applied Approach* (7e)]

The point is to exercise the disk method of calculating volumes, as evident by the title of the chapter where it is found. This closed-ended problem with an "exact" answer. The pond setting is merely decorative. The phrase "can be modeled by" carries the weight of the move from the world of things to the deductive world of mathematics. Grammatically, it is in the passive rather than the active voice.

Another widely-seen, calculus-level example of the passive use of "model" involves the shape of a tin can. Here, the formula given as the model is more plausible; the tin can is treated as the surface of a unit-volume cylinder. The task is to find the "optimal" ratio of height h to radius r that minimizes the surface area. This involves writing equations for the surface area and volume of the cylinder, solving the volume equation to express h as a function of r for unit volume, substituting that solution into the area equation, and differentiating to find the argmax of the surface area. The textbook answer is that the optimal can has shape $h = 2r$; a lovely, simple result.

"Modeling," in the active voice, is different. A proper modeling process should start with a notion of the goals (e.g., understanding the general shape of cans or deciding what shape is best for a particular product), consider a range of relevant factors (e.g., shipping efficiencies, ability to be grasped by the human hand, whether the surface area of the can can effectively stand for the cost of material that goes into the can, whether the material is the dominating factor in the economics of producing a can, and so on. (See Ngo Van Giang 2025 for a discussion). There's also the need to evaluate the model's suitability for addressing the questions at hand, compare its results with observations to judge how well it has captured the realities of food storage, potentially revise it to better suit the purpose, and consider the required precision. Ironically, even accepting the cylinder model and the context of optimization, the textbook answer is misleading. As all calculus books point out, the function is flat at the argmin. Consequently, the surface area will be within 10% of the optimum for a wide range of $0.93r < h < 5r$, that is, from a stubby tuna can to a thin energy drink container.

Returning to the pond volume context of the previous example, consider the problem posed to me by an aquatic ecologist colleague[^Dan Hornbach, personal communication]: Here is a map of Pokegama Lake (in Minnesota, USA), published by the state's Department of Natural Resources. What is the volume of the Lake?

Even with the contour map in hand, this is an open-ended problem. An important issue in practice is the level of precision required. This depends on the motivation behind the work, which raised the question of whether volume was the appropriate quantity to estimate. And, as a matter of technique, it's important to figure out how to assess whether this precision has been achieved or even can be.

I have posed the Pokegama Lake problem at conferences involving hundreds of mathematics educators. The approach proposed by the large majority of participants is to overlay the map with a square grid, measure the depth of each square, multiply its area by its depth, and sum the results to find the total volume. I inquire, what will be the accuracy and precision of a result found in this way? "In the limit as the square size goes to zero," the participants say, "the result is exact." In actuality, the ecologist thought that a result within about 30% would suffice. This can be accomplished by throwing a few dozen darts at the map, a method that could include the marshy areas around the lake for which no contours are provided.

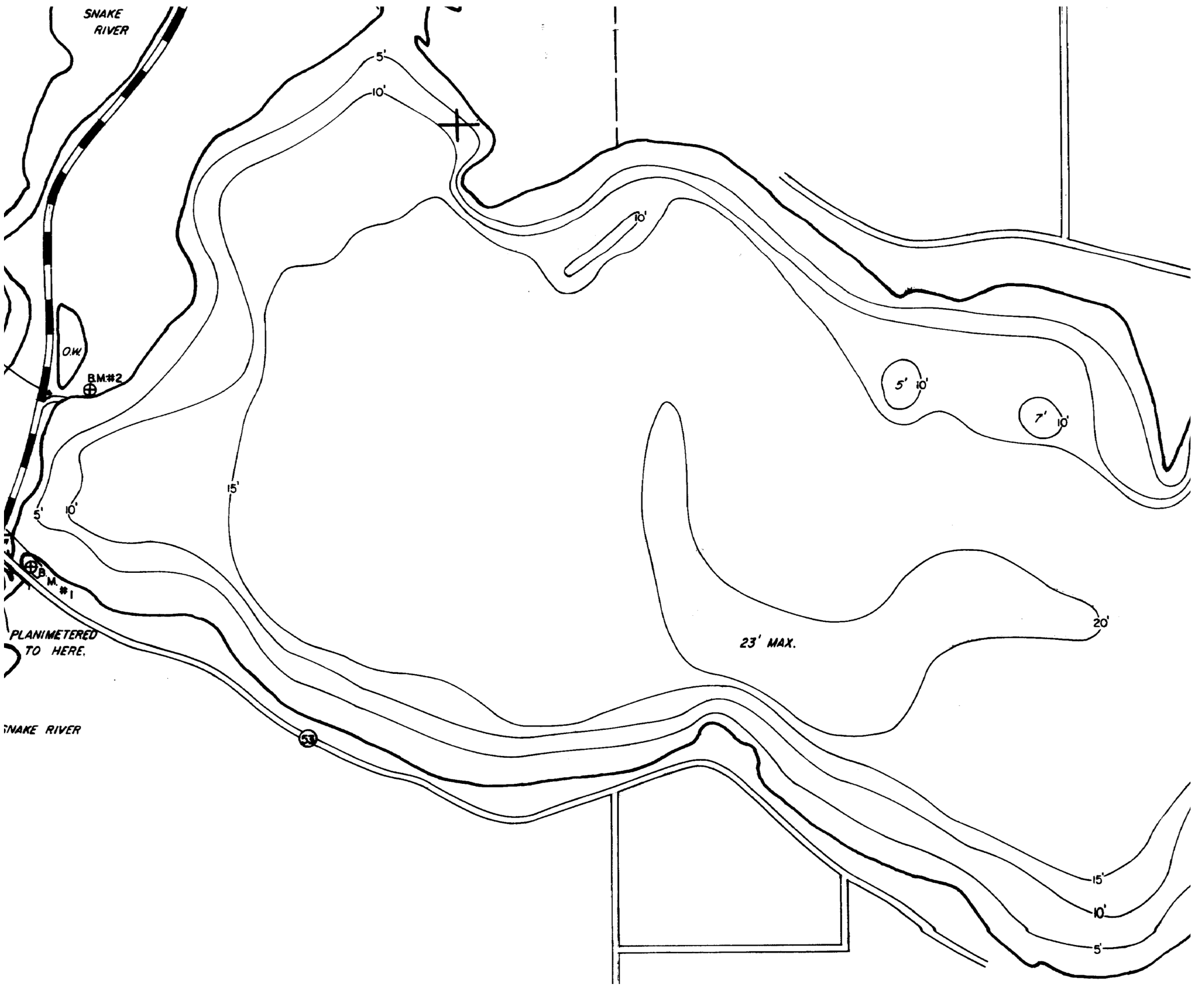

Figure 1: A detail from the Pokegama Lake map.

Quantitative technologies

The first issue of TISE appeared 20 years ago. Some of the technologies highlighted in volume 1 issue 1 were drag-and-drop URLs, Wikis, videotape, and a computer notation that avoids "code." These are in keeping with the conventional view of education technologies as centered on display, communication, and calculation. The very first article in TISE 1(1), George Cobb's "The Introductory Statistics Course: A Ptolemaic Curriculum?", is famous for inspiring now-widely adopted changes in the logical foundation of teaching statistical inference. The invocation of Ptolemy referred to a technology for calculating the orbits of planets based on Earth-centered circles and epicycles. The mechanisms also served as the basis for conceiving the motion of the planets, in vogue from about 100 to 1600 AD, until they were replaced by Kepler's alternative mechanism—sun-centered elliptical orbits—further simplified by Newton's Laws and calculus to the consequences of the inverse-square law of gravitation.

For Cobb, the innovation in the technology applied to planetary orbits was a metaphor. Cobb encouraged the replacement of the dominant statistical inference technology of tests, formulas, and tables with a more general and accessible set of technologies. He called these the "three R's of statistical inference": 1) randomization, 2) repetition (or, perhaps better, "iteration of trials"), and 3) rejection (by which he meant counting the proportion of successful trials).

Doing statistical inference with the three-Rs technology is now common in statistics education. It is far from universal, however. One reason is the reluctance of many instructors to embrace another technology: computing. As stated in the introduction, the interfaces between technologies are important.

The following set of quantitative technologies involves many interconnections and interfaces.

Quantities and rates

Of course, the use of numbers in mathematics education is hardly an innovation. However, an important aspect of "quantity", dimensionality, has not been integrated into the traditional curriculum. Quantities such as position, velocity, and acceleration are not just numbers; they refer to different physical entities. The same is true for a wide variety of entities. Examples include the relationship between deficit and debt, between birth rate and population, and the many different ways of accounting for fertility, mortality, prevalence, and incidence of disease.

In science and engineering, there are standards such as the *Systèm Internationale* that cover both unit and dimension. Different types of quantity, for instance, velocity versus position, can be closely related. New quantities are formed by the multiplication or division of existing quantities. Division of two quantities creates a *rate*, while subtraction is the fundamental way to calculate the *change* in the amount of a quantity, rather than in its dimension. *Rates* of *change* are fundamental to calculus. Many mistakes of comparison or interpretation stem from careless or, more commonly, missing consideration of dimension.

Functions

The traditional curriculum makes extensive use of *equations* and has for more than a millenium. Al-Khwarizmi (c. 780 – c. 850) is considered the father of algebra due to his treatise, *The Compendious Book on Calculation by Completion and Balancing.* (The shorthand Arabic title of the treatise, *Al-Jabr*, is the source of the word "algebra" whereas "algorithm" comes from the name of the treatise's author.) In conventional notation, the = symbol—notation invented by Robert Recorde (1510 – 1558)—is used for three distinct purposes: a. assignment of a value to a name, for instance $x = 4$; b. as the point of balance in an expression, supporting algebra techniques that involve moving quantities from one side to the next or performing the identical operation on each side as in $3x + 2 = 5$ implying $3x = 3$ and $x = 1$; and c. as a typographic separator between equivalent formulations of the same expression (e.g. $(x-3)(x+2) = x^2-x-6$). The equation formalism is so basic to mathematics education that it seems indispensable. But if learning algebra *per se* is not the ultimate goal, there are alternative technologies that can be advantageous. For example, mainstream programming languages do not have built-in support for equations; students learning to program have to be weaned from equations. (Computer expressions like `x == 3` are not equations, they are questions with a true or false answer. Statements like `3*x = 9` are not allowed in most computer languages; the `=` means assignment although other symbols are also used, e.g. `:=` or `<-`.)

Much later, Euler (1707-1783) invented the modern idea of *functions* and the notation $f(x)$. A function translates one or more inputs into an output. This directionality, inputs $\longrightarrow$ output, is not a feature of equations. Functions are a basic structure in computer languages.

Much of what is done with equations in the secondary curriculum through calculus can be done naturally with functions. The tasks performed on functions are various and have names: to find an input that produces a given output is to "solve"; to calculate the output given the input is to "evaluate"; the input that produces the largest output is "argmax", whereas finding that largest output is to "maximize". Other tasks include differentiation, integration, and iteration. One builds new functions from existing ones mainly using a few strategies: multiplication, division, linear combination, and composition. Much of machine learning involves constructing functions by training on data.

Functions play a key role in constructing models. Operations on functions are used to extract latent information from models.

Graphical presentation

Graphics provides a valuable technology both for communicating results and developing intuition. There are many types of graphs: function outputs versus inputs, paths and trajectories, vectors and vector fields, in addition to the many kinds encountered when working with data. A contour plot provides a simple way to visualize functions of two variables. (In my view, this should be placed near the beginning of any curriculum addressing the goals in **?@sec-high-level-goals**.)

Beyond the graphic modes already mentioned, there are other types of graphs relating to the goals. Directed graphs are helpful for tracking the components of a system and their mutual relationships; directed acyclic graphs are useful for stating causal hypotheses.

Seeing graphics as a technology reminds us that graphic modes can become obsolete. For instance, the dot plots and stem-and-leaf plots often found in statistics texts stem from an era before computing and the practicability that computer graphics systems provide for transparency, color, and movement/interactivity. For student success, I have found it useful to embrace these capabilities and adopt the simple principle that statistical commentary should be built on data. (Interpreting a regression line, even one with a confidence band, is prone to misinterpretation unless the underlying data are shown; transparency makes this easy.)

Magnitude

Mathematicians use the word "magnitude" in two different senses: to refer to a quantity as if measured by a ruler and in the phrase "order of magnitude." This dual use can lead to confusion and uncertainty. I prefer to use "magnitude" only in the second sense, and "size" for the first. The reader may think this is silly or contrarian, but there is a good reason to unambiguously separate the two senses.

The mathematics curriculum has included logarithms for hundreds of years. When invented around 1600, the logarithm was an aid to calculation, making multiplication and division easier and enabling the raising of numbers to non-integer powers. For this purpose, human calculators were equipped with telephone-book-sized tables of "common," that is, base-10 logarithms; the first such table was published in 1617.

With the replacement of logs in everyday calculations by digital computation, students encounter logarithms in their algebra course, which emphasize the algebraic properties of logarithms: different bases, conversion between bases, and the prestige assigned to the base of the "natural" logarithm. Students often find logarithms difficult. The closeness of the words "logarithm" and "algorithm" (whose etymologies are unrelated) adds to the confusion. Even among well-educated professionals in fields that are not primarily quantitative, there is anxiety around the word and concept. In contrast, "magnitude" is not so anxiety-provoking. ("Order of magnitude" of a number is the base-10 log of the number's size, rounded downward to an integer. The implementation, counting digits around the decimal point, is not mysterious, unlike ln().)

Consider two important contexts in which the idea of magnitude is important that relate to the goals in **?@sec-high-level-goals** reasoning. One is when displaying data, like the population or GDP of countries, which spans more than one order of magnitude. The other is when thinking about growth, the two main models of growth—"arithmetic" or "geometric" in traditional terms—being distinguished by whether the the *quantity* is changing steadily or the *magnitude* is changing steadily.

Randomization, simulation, and trials

In the 1950s through 1970s, statistics was present in the mathematics curriculum mainly as a third-year class called variously "mathematical statistics" or "probability and statistics." It was well placed in the third year since prerequisites often included multivariable calculus and combinatorics. Also beginning in the 1950s, a method developed for the design of nuclear bombs—the Monte Carlo method—started to find other applications. Today, such methods are known as simulation coupled to random-number generation. Calculus techniques such as multiple integration are not central. Indeed, it's common for theoreticians using calculus techniques to confirm their results using simulation. The same goes for graduate-level techniques from stochastic calculus.

Simulations can be built by connecting random number generators to deterministic calculations such as arithmetic. Complete mastery requires experience and insight, but the starting cost is low enough for undergraduate students. Students can start with pre-packaged simulations and then move on to modifying simple simulations to their own settings.

It's common to run many trials of a simulation. Doing so does *not* require a student to learn enough software to write a loop and implement an accumulator to collect the results. Instead, simple software interfaces such as`do(100) *` in the `{mosaic}` R package or the `|> Trials(100)` in the more recent `{LSTbook}` package.

Simulation technology can contribute to several of the goals in **?@sec-high-level-goals**. They provide implementations of models that students can manipulate; they handle and illuminate many important aspects of uncertainty, risk, and statistical reasoning; and they provide a way to instantiate systems with multiple components.

Hypotheses, likelihood, and belief

The scientific method is a great ally of quantitative thinking. A simple definition will do for our purposes here.

> *Principles and procedures for the systematic pursuit of knowledge involving the recognition and formulation of a problem, the collection of data through observation and experiment, and the formulation and testing of hypotheses.* — Merriam Webster Dictionary

Using the scientific method in quantitative reasoning requires the quantitative representation of *hypotheses.* (Note the plural form: hypotheses.) The constellation of modeling, systems, and simulation provides a good basis for this. In addition, it provides a concrete realization of the concept of mathematical likelihood, central to both Frequentist and Bayesian perspectives. With simulation technology available, calculating a likelihood or a likelihood function does not require calculus.

Teaching and using likelihood does not require the mastery of conditional probability notation. Indeed, it is arguably better in an introduction to avoid the confusion of notation such as $p(A|B) = p(B|A)\ p(B)/p(A)$. A and B are too abstract; the symmetry makes it hard to keep track of what's going on. Instead, emphasize that an observation is the measurement of a quantity. A likelihood is the probability of that measured quantity taking a hypothesis as a given. When a hypothesis is implemented as a simulation, the likelihood is derived from the distribution obtained from many simulation trials.

As for quantifying belief, this is simply a non-negative number assigned to a hypothesis. When, as is usually the case, there are multiple or many hypotheses, the level of belief in each is indicated by the size of the corresponding non-negative number. This does not need to be stated as a probability.

The Bayesian update of belief can occur whenever a new observation becomes available. The updated level of belief in a hypothesis is a matter of multiplying the likelihood of the observation (under that hypothesis) by the original level of belief. The non-negative numbers that result when multiple hypotheses are under consideration become the updated levels of belief.

Spaces, vectors, basic vector operations

The Cartesian plane is a standard part of the high-school curriculum. Vectors, however, are not. For instance, the Common Core educational standards place vectors in the same terms as precalculus and calculus, advanced math not required for all students. Possibly the reason for this placement is the idea that vectors are mainly needed in physics. This attitude neglects the strong connections between vectors, statistics, and machine learning/AI.

When working with functions with two inputs, students can intuit the idea of a gradient, even without calculus. (Everybody knows what "uphill" means!) The gradient is, of course, a vector. Knowing about gradients provides a way to engage the paradigm of *training*, which is so important to the uses of data. (We'll return to the concept of training with data in Section .)

Even the most basic modern introduction to data science should include the concept of a data frame, a rectangular array of rows and columns, as in spreadsheet software. Each *variable* is a column. For quantitative variables, the column of numbers amounts to a vector.

The vector viewpoint is helpful when working with data. At the elementary level, the correlation coefficient is the (cosine of) the angle between the two variables/vectors. Linear regression is the *projection* of a vector onto sets of vectors. What's important here is not the possibility of teaching elementary statistics with vectors. Moving beyond elementary statistics, learning to think about *systems* and the connections between multiple variables, is facilitated by having a means to visualize multiple variables and the mechanism of model fitting.

Looking at vector alignment is a basic technique not only in statistics but also in machine learning and even in the AI of large language models. In LLMs, for instance, text is tokenized: each token is a number. Chunks of text (the "context") are transformed into vectors in high-dimensional spaces. Similarities in meaning are assessed by the alignments among the high-dimensional embeddings of multiple text chunks. This is not to say we should be teaching first-year college students how to train LLMs, but it is beneficial to reveal the aspects that can be made accessible.

Vectors also provide an accessible introduction to dynamics (Section) in addition to their obvious relevance in STEM fields. Taken along with the connections of vectors to statistics, understanding systems, and machine learning, this is a substantial benefit that comes at a *very low cost*: protractors, rulers, the dot product, scalar multiplication, and vector addition/subtraction are the essential techniques. It would be a very ambitious instructor who would add QR decomposition to this list—leave it for the students taking calculus—but using QR software to find the projection of a vector onto a subspace is both accessible to students and highlights ideas like linear combination, minimization of residuals, and statistics such as R^2.

Eulerian framework for dynamics

Many readers with extensive mathematics backgrounds will have encountered the fourth or fifth course in a university-level calculus sequence: differential equations. This course, typically a buffet of solution techniques for equations involving derivatives, is well placed at the end of a calculus sequence since it involves not only differentiation and integration, but also linear algebra.

The differential equations course is, naturally, about solving equations. The equations involved, however, come from a much more widely applicable problem: how to think quantitatively about systems whose state changes over time. An epidemic provides an example, with the state including both the number of susceptible people in the population and the number of infectives. Newton's development of calculus was, in large part, motivated by his need to provide mathematical principles for working with his laws of motion, for example, force equals mass times acceleration, where acceleration is the rate of change of velocity and velocity is the rate of change of position.

However, solving differential equations is hardly the only approach to working with systems whose state changes in time, that is, *dynamical systems*.

I use the name "Eulerian framework" to refer to a straightforward technique for working with dynamical systems. First, the system has a state that can change from instant to instant. The rate of change of the state with respect to time is a function—often a simple function—of the state itself. The framework provides a step-by-step process—the Euler method—for tracing the trajectory of the state. Start at the initial state, calculate the rates of change at that state point (a vector), then multiply by a small interval of time. This yields a change vector, which is added to the state vector to trace the trajectory over a small time interval. Continue with this process, using the location from the previous step at each step.

The Euler method is an accessible way to enable introductory students to work with dynamical systems. As the name "dynamical systems" suggests, students are learning ways to represent interactions among multiple quantities, serving Goal 3 in Section . (There are classic dynamical system models of contexts that students find compelling: e.g., epidemics, sleep and wakefulness,and ecosystems.)

As well, dynamical systems provide a valuable perspective on a major way of extracting information from models: optimization. The differentiate-then-solve method that is standard in textbooks is undoubtedly appropriate when only the technology of algebraic expressions is used to represent the objective function. But unlike in Newton's time, we now have other technologies to represent functions, such as graphics and computer expressions. Once a student learns to read a graph, finding maxima and minima can be done at a glance. Moreover, students can easily judge from the graph the range of function inputs that gives effectively the same output as the argmin/max.

The methods of dynamical systems provide an equally intuitive approach to optimization: Walk uphill! This is everyday-language shorthand for "follow the gradient," with the Euler framework providing the arithmetic structure for this gradient ascent.

Other technologies from Section are also relevant to realistic optimization questions. We often have multiple objectives in mind when working with a real-world situation, for instance, save lives while keeping expenses tolerably small. A standard quantitative approach is to pick one objective to serve as the function to be optimized, treating the other objectives as "constraints." In such problems, an algorithm can be easily formulated as a dynamical system: "Walk in the steepest uphill direction consistent with the constraints."

Discussion

I have presented these ideas in presentations and workshops over many years, to both sympathetic and unsympathetic listeners. These are some of the most frequent reactions to the ideas in this paper, along with my responses.

- Many of these topics are too advanced for students who might otherwise be taking pre-calculus, or Calc I, or elementary statistics. For example, gradients are covered in Calc III, for which such students are not prepared.

RESPONSE: The conventional sequence is symbolic derivatives in Calc I, partial derivatives in Calc III, construction of the vector field of partial derivatives, and algebraic demonstration that each vector in this field points in the most uphill direction. There are other ways. A sequence supported by the technologies in Section goes like this. Introduce functions of two variables using concrete examples of familiar quantities, such as the meteorological "heat index," use contour plots to visualize such functions, and topographical maps as an example of a contour plot. Take a field trip to a place with a hill (bring a compass) and have students point due east and north, then point uphill in the steepest direction. (Also instructive: close your eyes and figure out from short steps which direction is uphill.) Show the (vector) representation of gradients as an annotation on a contour plot, and have students figure out the relationship between the contour direction and the gradient. Many concept reinforcement activities using contour diagrams are possible. One I particularly like is to hand out a landscape-like contour map of a hill, basin, or amphitheater. (Keep it simple!) Mark two points in the latitude/longitude space. Ask students to find the shortest path between the two points that doesn't exceed a specified steepness. Later, when teaching vector decomposition, plot out the steepness in the due east and north direction, representing these (magnitude and direction) as vectors. Show that they are a decomposition of the gradient vector.

- Students have a hard enough time getting through college algebra or Calc I. We have to stabilize their understanding of these topics before they can move on.

RESPONSE: How well has that strategy been working over the last 50 years? The usual word for "stabilize" is "remediate." Remedial courses have a minuscule success rate.

- Students learned about measurement units in middle school. Why bring this up again at a college level?

RESPONSE: I doubt we need to teach more about cups, pints, and gallons. But we do need to explain that there are different types of quantities and that they are often related by multiplication and division. Units can provide an entryway to the idea of dimension, but units are not a substitute for dimension. And dimension can provide access to important ideas from science and calculus. For instance, as many people know, distance has dimension L, velocity is change in distance divided by change in time, a rate of change with dimension L / T. If Snow's observations about the general ignorance of acceleration is true, perhaps it will help to mention that acceleration is also a rate of change: change in velocity divided by change in time, having dimension L / T^2.

- How do you know that the approach will work, either in general or at my school?

RESPONSE: "Will work" might mean either "students will succeed" or "students will achieve the goals of quantitative education." I've introduced curricula incorporating the technologies at three institutions. Not every idea worked. The technologies presented here are the ones that students found successful. As for "achieve the goals," the goals in Section emerged from an iterative process of testing and refining technologies and seeing what students could achieve with them.

We can ask the same question about the conventional curriculum. I'll suggest two metrics for establishing success. 1) What fraction of students succeed at the level of the conventional marker for college-level mathematics, that is, Calc I. Nationally, that is a very small fraction, at least from the perspective of a student trying to progress through the sequence. For many instructors who work with the survival-biased sample of students who enroll in Calc I, the success rate is short of 50%. 2) What fraction of students continue to use the methods of Calc I in their further studies?

- How would this approach serve students, such as engineers, who need to satisfy calculus prerequisites for their engineering (or physics or chemistry) classes?

RESPONSE: Naturally, I accept the right of engineering and science disciplines to set the content and timing of prerequisites. In general, deferring symbolic calculus to the second year of college would not work with many programs. So, for students who face genuine Calc I/II prerequisites, careful thought and design would be needed before adopting the program I advocate here for early college-level quantitative education. Some things to keep in mind. Many engineering programs have statistics requirements that might be better satisfied by quantitative reasoning courses along the lines described here. Additionally, many traditional sequence courses are not designed with engineers in mind and devote considerable time to methods and concepts that are not needed in downstream engineering courses. In fact, many engineering programs offer an alternative pathway that sidesteps math department courses. This is done both for efficiency and to gain the benefits of teaching calculus within the context of engineering.

- Requiring students to learn computer programming will not be practicable due to time constraints, limited instructional resources, or student interest.

RESPONSE: I did not list "computer programming" among the technologies to build on. While some instructors might prefer to go that route, and while the use of computers is a powerful amplifier for gaining understanding in quantitative education, computer services can be supplied in many non-programming ways. Mouse- and menu-driven apps can serve as the foundation and are already available. My preferred style is to integrate computational capabilities into the textbook documents and exercises, creating an interactive notebook that uses a highly focused syntax. Commands can scaffold students so they can do something meaningful merely by changing the names of functions or parameter values. Programming constructs such as explicit loops, array indexing, and accumulators are never needed. See, for example, these packages for R: `{mosaic}`, `{LSTbook}`, and `{mosaicCalc}` .

- The content of our statistics and mathematics curricula is subject to state-level mandates and to two-year college transfer and credit agreements with four-year universities.

RESPONSE: Change is always hard. Some people are better situated to accomplish it than others. Often, such mandates are written so abstractly that compliance is possible even when making substantial changes. For example, "statistical inference" is much broader than t-tests and p-values, and even if these are explicitly dictated, you can cover them in a much more efficient way in a modeling-based curriculum. If mandates interfere with your doing what's best for your students, we as a community need to work to change them. Part of doing this is establishing a body of instructors, not subject to the mandates, who can innovate and report to legislators the reasoning behind and the success of the innovations.

A prevalent obstacle to success is the practice of selecting individual words from the goals or technologies and translating them directly into traditional categories. In other words, paying lip service to the goals while changing little. For example, comparing proportions in two populations is a reasonable topic to include when teaching about risk. Conventional elementary courses regularly do this. Calculating expectation values in the context of gambling is also related. But a more pressing context for our students and their families is medical risk. For example, "risk factors" are a concept that everybody encounters. What do they mean, and how are they quantified and combined? Risk often involves multiple factors, so understanding "risk as a function of" needs to be in the curriculum. Another example: is it sufficient to warn about confounding and lurking variables when addressing causality, or do we need to engage in adjustment for covariates and such? My view is that students need to be able to engage with the uses of statistics as they appear in professional research. Covariates, adjustment, risk, the presentation of data for survival analysis, and causal networks are all important.

The verb "grok," meaning "to understand intuitively or by empathy, to establish rapport with", and "to empathize or communicate sympathetically (with); also, to experience enjoyment" [Source], is helpful to describe attitudes and practices that support meaningful, effective innovation. To engage with the innovations outlined in this article, the instructor has to grok the reasons for change. Our training, specialization, and teaching experience do not always make it easy to do so.

- Are there textbooks and other materials that can give a concrete example of the uses of the technologies?

RESPONSE: The University of Austin has a core "Intellectual Foundations" curriculum that includes two courses that are required for all first-year students: QR-1 and QR-2 offered in 2024-2025 and 2025-2026. Each course is three credit hours. Goals 1-5 in Section are treated in both QR-1 and QR-2 but from different perspectives: mathematical modeling in the first, statistical modeling in the second. Goal 6 is covered in QR-2. The textbook for QR-1, *Quantitative Reasoning: An Analysis Approach* (Kaplan, 2026) is available free online. Owing to the greater level of innovation in statistics curricula, QR-2 could draw on many existing materials. The UATX course uses David Spiegelhalter's commercially published introductory overview *The Art of Statistics: How to Learn from Data* as well as Hans Rosling's remarkable *Factfulness* which provides ten short case studies in the context of international development and health. Computing support for projects and exercises is provided by a set of free online tutorials, collected as the *Data to Information Workbook* (Kaplan, 2025).