

7 Randomized Algorithms

All the algorithms described so far in this article are deterministic: if they are run repeatedly on the same data, they produce the same result every time. Some algorithms make random choices and so generally produce a different result every time they are run. For example, we might approximate

$$\int_0^1 f(x) dx \approx \frac{1}{n} \sum_{i=1}^n f(x_i),$$

where the x_i are independent random numbers uniformly distributed in $[0, 1]$. The standard deviation of the error in this approximation is of order $n^{-1/2}$. This is an example of a *Monte Carlo algorithm*; such algorithms have a deterministic run time and produce an output that is correct (or has a given accuracy) with a certain probability. Of course, there are much more efficient ways to estimate a one-dimensional integral, but Monte Carlo algorithms come into their own for multidimensional integrals over complicated domains.

A *Las Vegas algorithm* always produces a correct result, but its run time is nondeterministic. A classic example is the quicksort algorithm for sorting a list of numbers, for which a randomized choice of the partition element makes the algorithm much faster on average than in the worst case ($O(n \log n)$ running time versus $O(n^2)$, for n numbers).

Randomized algorithms can be much simpler than deterministic alternatives, they may be more able to exploit modern computing architectures, and they may be better suited to large data sets. There is a wide variety of randomized algorithms, and they are studied in mathematics, computer science, statistics, and other areas.

One active area of research is randomized algorithms for numerical linear algebra problems, based on random sampling and random projections. For example, fast algorithms exist for computing low-rank approximations to a given matrix. The general framework is that random sampling is used to identify a subspace that captures most of the action of the matrix, the matrix is then compressed to this subspace, and a low-rank factorization is computed from the reduced matrix.

Examples of randomized algorithms mentioned in this book are the Google PAGERANK ALGORITHM [VI.9], with its use of a random surfer, the k -MEANS ALGORITHM [IV.17 §5.3] for clustering, and MARKOV CHAIN MONTE CARLO ALGORITHMS [V.11 §3].

8 Some Key Algorithms in Applied Mathematics

Table 2 lists a selection of algorithms mentioned in this book. Very general methods such as preconditioning and the finite-element method, which require much more information to produce a particular algorithm, are omitted. The table illustrates the wide variety of important algorithms in applied mathematics, ranging from the old to the relatively new.

A notable feature of some of the algorithms is that they are iterative algorithms, which in principle take an infinite number of steps, for solving problems that can be solved directly, that is, in a finite number of operations. The conjugate gradient and multigrid methods are iterative methods for solving a linear system of equations, and for suitably structured systems they can provide a given level of accuracy much faster than Gaussian elimination, which is a direct method. Similarly, interior point methods are iterative methods for linear programming, competing with the simplex method, which is a direct method.

Further Reading

A classic reference for algorithms and their analysis is Donald Knuth's *The Art of Computer Programming*. The first volume appeared in 1968 and the development is ongoing. Current volumes are *Fundamental Algorithms* (volume 1), *Seminumerical Algorithms* (volume 2), *Sorting and Searching* (volume 3), and *Combinatorial Algorithms* (volume 4), all published by Addison-Wesley (Reading, MA).

- Bentley, J. L. 1986. *Programming Pearls*. Reading, MA: Addison-Wesley.
- Brassard, G., and P. Bratley. 1996. *Fundamentals of Algorithms*. Englewood Cliffs, NJ: Prentice-Hall.
- Cormen, T. H., C. E. Leiserson, R. L. Rivest, and C. Stein. 2009. *Introduction to Algorithms*, 3rd edn. Cambridge, MA: MIT Press.
- Higham, N. J. 2002. *Accuracy and Stability of Numerical Algorithms*, 2nd edn. Philadelphia, PA: SIAM.

1.5 Goals of Applied Mathematical Research

Nicholas J. Higham

A large body of existing mathematical knowledge is encapsulated in theorems, methods, and algorithms, some of which have been known for centuries. But applied mathematics is not simply the application of

Table 2 Some algorithms mentioned in this book.

Algorithm	Reference	Key early figures
Gaussian elimination	IV.10 §2	Ancient Chinese (ca. 1 C.E.), Gauss (1809); formulated as LU factorization by various authors from 1940s
Newton's method	II.28	Newton (1669), Raphson (1690)
Fast Fourier transform	II.10	Gauss (1805), Cooley and Tukey (1965)
Cholesky factorization	IV.10 §2	Cholesky (1910)
Remez algorithm	IV.9 §3.5, VI.11 §2	Remez (1934)
Simplex method (linear programming)	IV.11 §3.1	Dantzig (1947)
Conjugate gradient and Lanczos methods	IV.10 §9	Hestenes and Stiefel (1952), Lanczos (1952)
Ford-Fulkerson algorithm	IV.37 §7	Ford and Fulkerson (1956)
k -means algorithm	IV.17 §5.3	Lloyd (1957), Steinhaus (1957)
QR factorization	IV.10 §2	Givens (1958), Householder (1958)
Dijkstra's algorithm	VI.10	Dijkstra (1959)
Quasi-Newton methods	IV.11 §4.2	Davidon (1959), Broyden, Fletcher, Goldfarb, Powell, Shanno (early 1960s)
QR algorithm	IV.10 §5.5	Francis (1961), Kublanovskaya (1962)
QZ algorithm	IV.10 §5.8	Moler and Stewart (1973)
Singular value decomposition	II.32	Golub and Kahan (1965), Golub and Reinsch (1970)
Strassen's method	I.4 §4	Strassen (1968)
Multigrid	IV.10 §9, IV.13 §3, IV.16	Fedorenko (1964), Brandt (1973), Hackbusch (1977)
Interior point methods	IV.11 §3.2	Karmarkar (1984)
Generalized minimal residual method	IV.10 §9	Saad and Schulz (1986)
Fast multipole method	VI.17	Greengard and Rokhlin (1987)
JPEG	VII.7 §5, VII.8	Members of the Joint Photographic Experts Group (1992)
PageRank	VI.9	Brin and Page (1998)
HITS	I.1	Kleinberg (1999)

existing mathematical ideas to practical problems: new results are continually being developed, usually building on old ones. Applied mathematicians are always innovating, and the constant arrival of new or modified problems provides direction and motivation for their research.

In this article we describe some goals of research in applied mathematics from the perspectives of the ancient problem of solving equations, the more contemporary theme of exploiting structure, and the practically important tasks of modeling and prediction. We also discuss the strategy behind research.

1 Solving Equations

A large proportion of applied mathematics research papers are about analyzing or solving equations. The

equations may be algebraic, such as linear or nonlinear equations in one or more variables. They may be ordinary differential equations (ODEs), partial differential equations (PDEs), integral equations, or differential-algebraic equations.

The wide variety of equations reflects the many different ways in which one can attempt to capture the behavior of the system being modeled. Whatever the equation, an applied mathematician is interested in answering a number of questions.

1.1 Does the Equation Have a Solution?

We are interested in whether there is a unique solution and, if there is more than one solution, how many there are and how they are characterized. Existence of solutions may not be obvious, and one occasionally

hears tales of mathematicians who have solved equations for which a proof is later given that no solution exists. Such a circumstance may sound puzzling: is it not easy to check that a putative solution actually is a solution? Unfortunately, checking satisfaction of the equation may not be easy, especially if one is working in a function space. Moreover, the problem specification may require the solution to have certain properties, such as existence of a certain number of derivatives, and the claimed solution might satisfy the equation but fail to have some of the required properties. Instead of analyzing the problem in the precise form in which it is given, it may be better to investigate what additional properties must be imposed for an equation to have a unique solution.

1.2 Is the Equation Well-Posed?

A problem is *well-posed* if it has a unique solution and the solution changes continuously with the data that define the problem. A problem that is not well-posed is *ill-posed*. For an ill-posed problem an arbitrarily small perturbation of the data can produce an arbitrarily large change in the solution, which is clearly an unsatisfactory situation.

An example of a well-posed problem is to determine the weight supported by each leg of a three-legged table. Assuming that the table and its legs are perfectly symmetric and the ground is flat, the answer is that each leg carries one-third of the total weight. For a table with four legs each leg supports one-quarter of the total weight, but if one leg is shortened by a tiny amount then it leaves the ground and the other three legs support the weight of the table (a phenomenon many of us have experienced in restaurants). For four-legged tables the problem is therefore ill-posed.

For finite-dimensional problems, uniqueness of the solution implies well-posedness. For example, a linear system $Ax = b$ of n equations in n unknowns with a nonsingular coefficient matrix A is well-posed. Even so, if A is nearly singular then a small perturbation of A can produce a large change in the solution, albeit not arbitrarily large: the CONDITION NUMBER [I.2 §22] $\kappa(A) = \|A\| \|A^{-1}\|$ bounds the relative change. But for infinite-dimensional problems the existence of a unique solution does not imply that the problem is well-posed; examples are given in the article on INTEGRAL EQUATIONS [IV.4 §6].

The notion of well-posedness was introduced by Jacques Hadamard at the beginning of the twentieth

century. He believed that physically meaningful problems should be well-posed. Today it is recognized that many problems are ill-posed, and they are routinely solved by reformulating them so that they are well-posed, typically by a process called REGULARIZATION [IV.15 §2.6] (see also INTEGRAL EQUATIONS [IV.4 §7]).

An important source of ill-posed problems is INVERSE PROBLEMS [IV.15]. Consider a mathematical model in which the inputs are physical variables that can be adjusted and the output variables are the result of an experiment. The *forward problem* is to predict the outputs from a given set of inputs. The *inverse problem* is to make deductions about the inputs that could have produced a given set of outputs. In practice, the measurements of the outputs may be subject to noise and the model may be imperfect, so UNCERTAINTY QUANTIFICATION [II.34] needs to be carried out in order to estimate the uncertainty in the predictions and deductions.

1.3 What Qualitative Properties Does a Solution Have?

It may be of more interest to know the behavior of a solution than to know the solution itself. One may be interested in whether the solution, $f(t)$ say, decays as $t \rightarrow \infty$, whether it is monotonic in t , or whether it oscillates and, if so, with what fixed or time-varying frequency. If the problem depends on parameters, it may be possible to answer these questions for a range of values of the parameters.

1.4 Does an Iteration Converge?

As we saw in METHODS OF SOLUTION [I.3], solutions are often computed from iterative processes, and we therefore need to understand these processes. Various facets of convergence may be of interest.

- Is the iteration always defined, or can it break down (e.g., because of division by zero)?
- For what starting values, and for what class of problems, does the iteration converge?
- To what does the iteration converge, and how does this depend on the starting value (if it does at all)?
- How fast does the iteration converge?
- How are errors (in the initial data, or rounding errors introduced during the iteration) propagated? In particular, are they bounded?

To illustrate some of these points we consider the iteration

$$x_{k+1} = \frac{1}{p}[(p-1)x_k + x_k^{1-p}a], \quad (1)$$

with p a positive integer and $a \in \mathbb{C}$, which is Newton's method for computing a p th root of a . We ask for which a and which starting values x_0 the iteration converges and to what root it converges. The analysis is simplified by defining $y_k = \theta^{-1}x_k$, where θ is a p th root of a , as the iteration can then be rewritten

$$y_{k+1} = \frac{1}{p}[(p-1)y_k + y_k^{1-p}], \quad y_0 = \theta^{-1}x_0, \quad (2)$$

which is Newton's method for computing a p th root of unity. The original parameters a and x_0 have been combined into the starting value y_0 .

Figure 1 illustrates the convergence of the iteration for $p = 2, 3, 5$. For y_0 ranging over a 400×400 grid with $\text{Re } y_0, \text{Im } y_0 \in [-2.5, 2.5]$, it plots the root to which y_k from (2) converges, with each root denoted by a different grayscale from white (the principal root, 1) to black. Convergence is declared if after fifty iterations the iterate is within relative distance 10^{-13} of a root; the relatively small number of points for which convergence was not observed are plotted white. For $p = 2$ the figure suggests that the iteration converges to 1 if started in the open right half-plane and -1 if started in the open left half-plane, and this can be proved to be true. But for $p = 3, 5$ the regions of convergence have a much more complicated structure, involving sectors with petal-like boundaries.

The complexity of the convergence for $p \geq 3$ was first noticed by Arthur Cayley in 1879, and an analysis of convergence for all starting values requires the theory of Julia sets of rational maps. However, for practical purposes it is usually principal roots that need to be computed, so from a practical viewpoint the main implication to be drawn from the figure is that for $p = 3, 5$ Newton's method converges to 1 for y_0 sufficiently close to the positive real axis—and it can be proved that this is true.

We see from this example that the convergence analysis depends very much on the precise question that is being asked. The iteration (1) generalizes in a natural way to matrices and operators, for which the convergence results for the scalar case can be exploited.

2 Preserving Structure

Many mathematical problems have some kind of structure. An example with explicit structure is a linear system $Ax = b$ in which the $n \times n$ matrix A is a TOEPLITZ

MATRIX [I.2 §18]. This system has $n^2 + n$ numbers in A and b but only $3n - 1$ independent parameters. On the other hand, if for the vector ODE $y' = f(t, y)$ there is a vector v such that $v^T f(t, y) = 0$ for all t and y , then $(d/dt)v^T y(t) = v^T f(t, y) = 0$, so $v^T y(t)$ is constant for all t . This conservation or invariance property is a form of structure, though one more implicit than for the Toeplitz system.

An example of a nonlinear conservation property is provided by the system of ODEs

$$\begin{aligned} u'(t) &= v(t), \\ v'(t) &= -u(t). \end{aligned}$$

For this system,

$$\frac{d}{dt}(u^2 + v^2) = 2(u'u + v'v) = 2(vu - uv) = 0,$$

so there is a quadratic invariant. In particular, for the initial values $u(0) = 1$ and $v(0) = 0$ the solution is $u(t) = \cos t$ and $v(t) = -\sin t$, which lies on the unit circle centered at the origin in the uv -plane. If we solve the system using a numerical method, we would like the numerical solution also to lie on the circle. In fact, one potential use of this differential equation is to provide a method for plotting circles that avoids the relatively expensive evaluation of sines and cosines. Consider the following four standard numerical methods applied to our ODE system. Here, $u_k \approx u(kh)$ and $v_k \approx v(kh)$, where h is a given step size, and $u_0 = 1$ and $v_0 = 0$:

$$\begin{aligned} \text{Forward Euler} & \quad \begin{cases} u_{k+1} = u_k + hv_k, \\ v_{k+1} = v_k - hu_k, \end{cases} \\ \text{Backward Euler} & \quad \begin{cases} u_{k+1} = u_k + hv_{k+1}, \\ v_{k+1} = v_k - hu_{k+1}, \end{cases} \\ \text{Trapezium method} & \quad \begin{cases} u_{k+1} = u_k + h(v_k + v_{k+1})/2, \\ v_{k+1} = v_k - h(u_k + u_{k+1})/2, \end{cases} \\ \text{Leapfrog method} & \quad \begin{cases} u_{k+1} = u_k + hv_k, \\ v_{k+1} = v_k - hu_{k+1}. \end{cases} \end{aligned}$$

Figure 2 plots the numerical solutions computed with 32 steps of length $h = 2\pi/32$. We see that the forward Euler solution spirals outward while the backward Euler solution spirals inward. The trapezium method solution stays nicely on the unit circle. The leapfrog method solution traces an ellipse. This behavior is easy to explain if we write each method in the form

$$z_{k+1} = Gz_k, \quad z_k = \begin{bmatrix} u_k \\ v_k \end{bmatrix},$$

where $G = \begin{bmatrix} 1 & h \\ -h & 1 \end{bmatrix}$ for the Euler method, for example. Then the behavior of the sequence z_k depends on the

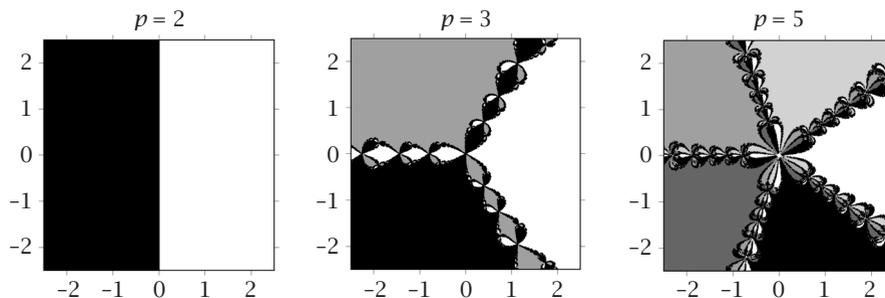


Figure 1 Newton iteration for a p th root of unity. Each point y_0 in the region is shaded according to the root to which the iteration converges, with white denoting the principal root, 1.

eigenvalues of the matrix G . It turns out that the spectral radius of G is greater than 1 for forward Euler and less than 1 for backward Euler, which explains the spiraling. For the trapezium rule G is orthogonal, so $\|z_{k+1}\|_2 = \|z_k\|_2$ and the trapezium solutions stay exactly on the unit circle. For the leapfrog method the determinant of G is 1, which means that areas are preserved, but G is not orthogonal so the leapfrog solution drifts slightly off the circle.

The subject of GEOMETRIC INTEGRATION [IV.12 §5] is concerned more generally with methods for integrating nonlinear initial-value ODEs and PDEs in a way that preserves the invariants of the system, while also providing good accuracy in the usual sense. This includes, in particular, SYMPLECTIC INTEGRATORS [IV.12 §1.3] for Hamiltonian systems.

3 Modeling and Prediction

AS WHAT IS APPLIED MATHEMATICS? [I.1 §1] explains, modeling is the first step in solving a physical problem. Models are necessarily simplifications because it is impractical to incorporate every detail. But simple models can still be useful as tools to explore the broad consequences of physical laws. Moreover, the more complex a model is the more parameters it has (all of which need estimating) and the harder it is to analyze.

In their 1987 book *Empirical Model-Building and Response Surfaces*, Box and Draper ask us to

Remember that all models are wrong; the practical question is how wrong do they have to be to not be useful.

Road maps illustrate this statement. They are always a simplified representation of reality due to representing a three-dimensional world in two dimensions and displaying wiggly roads as straight lines. But road maps

are very useful. Moreover, there is no single “correct” map but rather many possibilities depending on resolution and purpose. Another example is the approximation of π . The approximation $\pi \approx 3.14$ is a model for π that is wrong in that it is not exact, but it is good enough for many purposes.

It is difficult to give examples of the modeling process because knowledge of the problem domain is usually required and derivations can be lengthy. We will use for illustration a very simple model of population growth, based on the *logistic equation*

$$\frac{dN}{dt} = rN \left(1 - \frac{N}{K} \right).$$

Here, $N(t)$ is a representation in a continuous variable of the number of individuals in a population at time t , $r > 0$ is the growth rate of the population, and $K > 0$ is the carrying capacity. For $K = \infty$, the model says that the rate of change of the population, dN/dt , is rN ; that is, it is proportional to the size of the population through the constant r , so the population grows exponentially. For finite K , the model attenuates this rate of growth by a subtractive term rN^2/K , which can be interpreted as representing the increasing effects of competition for food as the population grows. The logistic equation can be solved exactly for $N(t)$ (see ORDINARY DIFFERENTIAL EQUATIONS [IV.2 §2]). Laboratory experiments have shown that the model can predict reasonably well the growth of protozoa feeding on bacteria. However, for some organisms the basic logistic equation is not a good model because it assumes instant responses to changes in population size and so does not account for gestation periods, the time taken for young to reach maturity, and other delays. A more realistic model may therefore be

$$\frac{dN(t)}{dt} = rN(t) \left(1 - \frac{N(t - \tau)}{K} \right),$$

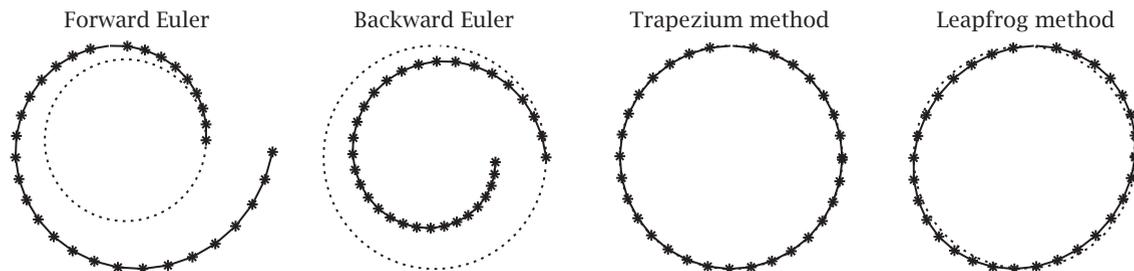


Figure 2 Approximations to the unit circle computed by four different numerical integrators with step size $h = 2\pi/32$. The dotted line is the unit circle; asterisks denote numerical approximations.

where $\tau > 0$ is a delay parameter. At time t , part of the quadratic term is now evaluated at an earlier time, $t - \tau$. This delay differential equation has oscillatory solutions and has been found to model well the population of lemmings in the Arctic. Note that in contrast to the PREDATOR-PREY MODEL [I.2 §10], the delayed logistic model can produce oscillations in a population without the need for a second species acting as predator. There is no suggestion that either of these logistic models is perfect, but with appropriate fitting of parameters they can provide useful approximations to actual populations and can be used to predict future behavior.

3.1 Errors

A lot of research is devoted to understanding the errors that arise at the different stages of the modeling process. These can broadly be categorized as follows.

Errors in the mathematical model. Setting up the model introduces errors, since the model is never exact. These are the hardest errors to estimate.

Approximation errors. These are the errors incurred when infinite-dimensional equations are replaced by a finite-dimensional system (that is, a continuous problem is replaced by a discrete one: the process of discretization), or when simpler approximations to the equations are developed (e.g., by MODEL REDUCTION [II.26]). These errors include errors in replacing one approximating space by another (e.g., replacing continuous functions by polynomials), errors in FINITE-DIFFERENCE [II.11] approximations, and errors in truncating power series and other expansions.

Rounding errors. Once the problem has been put in a form that can be solved by an algorithm implemented in a computer program, the effects of the rounding errors introduced by working in finite-precision arithmetic need to be determined.

Analysis of errors may include looking at the effects of uncertainties in the model data, including in any parameters in the model that must be estimated. This might be tackled in a statistical sense using techniques from UNCERTAINTY QUANTIFICATION [II.34]—indeed, if the model has incompletely known data then probabilistic techniques may already be in use to estimate the missing data. Sensitivity of the solution of the model may also be analyzed by obtaining worst-case error bounds with the aid of CONDITION NUMBERS [I.2 §22].

3.2 Multiphysics and Multiscale Modeling

Scientists are increasingly tackling problems with one or both of the following characteristics: (a) the system has multiple components, each governed by its own physical principles; and (b) the relevant processes develop over widely different time and space scales. These are called *multiphysics* and *multiscale* problems, respectively. An example of both is the problem of modeling how *space weather* affects the Earth, and in particular modeling the interaction of the solar wind (the flow of charged particles emitted by the sun) with the Earth's magnetic field. Different physical models describe the statistical distribution of the plasma, which consists of charged particles, and the evolution of the electric and magnetic fields, and these form a coupled nonlinear system of PDEs. The length scales range from millions of kilometers (the Earth-sun distance) to hundreds of meters, and the timescales range from hours down to 10^{-5} seconds. Problems such as this pose challenges both for modeling and for computational solution of the models. The computations require HIGH-PERFORMANCE COMPUTERS [VII.12], and a particular task is to present the vast quantities of data generated in such a way that users, such as forecasters of space

weather, can explore and interpret them. More on the issues of this section can be found in the articles on COMPUTATIONAL SCIENCE [IV.16] and VISUALIZATION [VII.13].

3.3 Computational Experiments

The step in the problem solution when computational experiments are carried out might seem to be the easiest, but it can in fact be one of the hardest and most time-consuming, for several reasons. It can be hard to decide what experiments to carry out, and it may be necessary to refine the experiments many times until useful or satisfactory results are obtained. The computations may have a long run time, even if executed on a high-performance computer.

Many pitfalls can be avoided by working to modern standards of REPRODUCIBLE RESEARCH [VIII.5], which require that programs, data, and results be recorded, documented, and made available in such a way that the results can be reproduced by an independent researcher and, just as importantly, by the original author.

3.4 Validation

The process of *validation* involves asking the question, “Have we solved the right equations?” This is not to be confused with *verification*, which asks whether the equations in the model have been solved correctly. Whereas verification is a purely mathematical question, validation intimately involves the underlying physical problem. A classic way to validate results from a model is to compare them with experimental results. However, this is not always feasible, as we may be modeling a device or structure that is still in the design phase or on which experiments are not possible (e.g., the Earth’s climate).

Validation may not produce a yes or no answer but may instead indicate a range of parameters for which the model is a good predictor of actual behavior.

Validation may be the first step of an iterative refinement procedure in which the steps in figure 1 are repeated, with the second and subsequent invocations of the first step now comprising adjustments to the current model. Assuming it is feasible to carry out refinement, there is much to be said for starting with the simplest possible model and building gradually toward an effective model of minimal complexity.

4 Strategies for Research and Publishing

Analysis of the research literature in applied mathematics reveals some common features that can be built into a list of strategies for doing research.

- (i) Solve an open problem or prove the truth or falsity of a conjecture that has previously been stated in the literature.
- (ii) Derive a method for solving a problem that occurs in practice and has not been effectively solved previously. Problems of very large dimension, for which existing techniques might be impractical, are good hunting grounds.
- (iii) Prove convergence of a method for which the existing convergence theory is incomplete.
- (iv) Spot some previously unnoticed phenomenon and explain it.
- (v) Generalize a result or algorithm to a wider class of problems, obtaining new insight in doing so.
- (vi) Provide a new derivation of an existing result or algorithm that yields new insight.
- (vii) Develop a new measure of cost or error for a problem and then derive a new algorithm that is better than existing algorithms with respect to that metric. For example, instead of measuring computational cost in the traditional way by the number of elementary arithmetic operations, also include the cost of data movement when the algorithm is implemented on a parallel computer.
- (viii) Find hidden assumptions in an existing method and remove them. For example, it may seem obvious that multiplying two 2×2 matrices requires eight multiplications, but Strassen showed that only seven multiplications are needed (see [I.4 §4] for the relevant formulas), thereby deriving an asymptotically fast method for matrix multiplication.
- (ix) Rehabilitate an out-of-favor method by showing that it can be made competitive again by exploiting new research results, problem requirements, or hardware developments.
- (x) Use mathematical models to gain new insight into complex physical processes.
- (xi) Use mathematical models to make quantitative predictions about physical phenomena that can lead to new procedures, standards, etc., in the target field. Here it will probably be necessary to work with researchers from other disciplines.

Nowadays, publishing the results of one's research is more important than ever. Funding bodies expect to see publications, as they provide evidence that the research has been successful and they help to disseminate the work. Assessment of researchers and their institutions increasingly makes use of metrics, some of which relate to publications, such as the number of citations a paper receives and the ranking of the journal in which it is published. There is therefore a tension between publishing prolifically (which, taken to the extreme, leads to breaking research up into "least publishable units") and publishing fewer, longer, more-considered papers.

In addition to the traditional journals and conference proceedings that publish (usually) refereed articles, nowadays there are many outlets for unrefereed manuscripts, including institutional eprint servers, the global arXiv eprint service, and personal Web pages. Blogs provide yet another venue for publishing research, usually in the form of shorter articles presented in a more accessible form than in a regular paper. The nonjournal outlets provide for instant publication but have varying degrees of visibility and permanence.

The balance between publication of journals and books in print only, in both print and electronic form, and purely in electronic form has been changing for the past decade or more, and the advent of handheld devices such as smartphones and tablets has accelerated developments. Equally disruptive is the movement toward open-access publishing. Traditionally, the publishers of mathematics journals did not charge an author to publish an article but did charge institutions to subscribe to the journal. In recent years a new model has been introduced in which the author pays to publish an article in a journal and the article is freely available to all.

While we can be sure that there will always be outlets for publishing research, it is difficult to predict how the forms that these outlets take will evolve in the future.

I.6 The History of Applied Mathematics

June Barrow-Green and Reinhard Siegmund-Schultze

And as for the Mixt Mathematickes I may only make this prediction, that there cannot faile to bee more kinds of them, as Nature growes furdre disclosed.

Francis Bacon: *Of the Proficiency and Advancement of Learning* (1605)

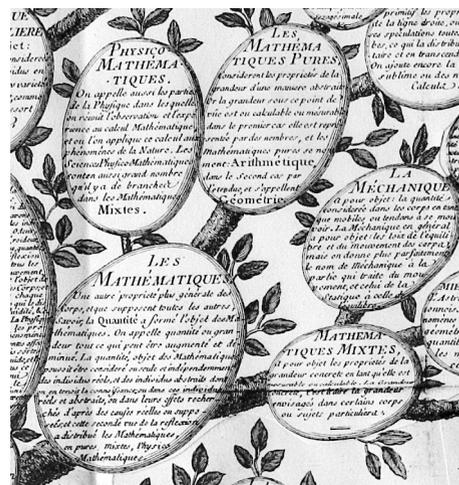


Figure 1 Detail from a pictorial representation of J. le R. d'Alembert's "Système des Connoissances Humaines" in the supplement to volume 2 of *l'Encyclopédie* (1769), mentioning "mixed mathematics," for which Francis Bacon had predicted a great future.

1 Introduction

What is applied mathematics? This is a difficult question—one to which there is no simple answer. The massive growth in applications of mathematics within and outside the sciences, especially since World War II, has made this question even more problematic, the increasing overlap with other disciplines and their methods adding further to the difficulties, creating problems that border on the philosophical.

Given the fact that almost every part of mathematics is potentially applicable, there are mathematicians and historians who consider the term "applied mathematics" primarily as a term of social distinction or a matter of attitude. One such was William Bonnor, the mathematician and gravitational physicist, who in 1962 in a lecture on "The future of applied mathematics" said the following:

Applied mathematics, as I should like it to be understood, means the application of mathematics to *any* subject, physical or otherwise; with the proviso that the mathematics shall be interesting and the results nontrivial. An applied mathematician, on this view, is somebody who has been trained to make such applications, and who is always prepared to look for situations where fruitful application is possible. As such, he is not a physicist *manqué*. I therefore see applied mathematics as an activity, or attitude of mind, rather than as a body of knowledge.