Continuity as Organized Discrete Data

The finitary work behind infinity

Henry Shin · First published 15 November 2026 · Revised 18 November 2026

Refactor · web edition

Every row of this infinite table is eventually zero. Its main diagonal is one all the way down.

$$x(m,n) = \begin{cases} 1, & n \le m, \\ 0, & n > m. \end{cases}$$

The main diagonal picks the last one in every row, so $x(m,m) = 1$ for every m. Move one column farther right in each row and the entries $x(m,m+1)$ are all zero.

All these rows represent the same number. Once we pass to their limit, we can forget how long each row's initial run of ones lasted. Choosing an entry from each row brings that detail back into play. How far along the row must we go before the entry is accurate enough for the next calculation?

In the Cauchy construction of the real line, each approximation is a fraction we can write with two integers. What makes a point is agreement across arbitrarily fine scales. One condition governs how the late terms of a sequence must agree with one another. An equivalence rule says when two such sequences count as the same point. That settles what the point is. Computing with a particular sequence raises a further question. How far must we go for the accuracy we need? The point can be well defined even when the sequence has no computable rate of convergence.

Choosing how far to look

A rational Cauchy sequence is a sequence whose terms eventually become arbitrarily close to one another. Two such sequences represent the same real number when their difference tends to zero.

Where does $\sqrt{2}$ first enter the mathematics? Start with rational numbers only. Bisect $[1, 2]$, square the midpoint, and keep the half whose endpoint squares bracket 2. The first cut keeps $[1, 3/2]$. The next keeps $[5/4, 3/2]$. After n cuts the endpoints differ by 2^{-n}. Their squares still bracket 2 and differ by at most $4 \cdot 2^{-n}$. Every decision so far could be made by someone who had never heard of $\sqrt{2}$.

Now comes the declaration. We will count a whole class of equivalent rational Cauchy sequences as one number. The shrinking widths prove that the two endpoint sequences belong to the same class. The bound on their squares proves that squaring this class gives the class of the constant sequence $2, 2, 2, \ldots$. The declaration tells us what counts as a number. The estimate tells us that this one has square two.

I take the finite rule, with its guarantees and up to that equivalence, to constitute $\sqrt{2}$. We can now look back and say that $\sqrt{2}$ lay inside every bracket all along. **We made the point from the brackets before we could say the brackets contained the point.**

Take the second bracket, from $5/4$ to $3/2$. Its endpoints are a quarter of a unit apart. Keep those very same fractions and try a distance rule that records only equality. Two fractions with the same value have distance zero. Any two with different values have distance one. Now the second bracket

is one unit wide. So is the hundredth. Under this rule, for all the late terms of a Cauchy sequence to lie within half a unit of one another, they must all have the same value. Completing the rationals adds no new points.

Return to ordinary distance. The gaps between the endpoints are a half, a quarter, an eighth, and keep halving. Completing the rationals with this distance gives the real line. Look at the step that gives us $\sqrt{2}$ in particular. The two endpoints differ at every cut. Their two histories name one point because, for every tolerance however small, their gap eventually falls below it and stays there.

Now suppose (X_m) is itself a Cauchy sequence of real numbers. For each m, present X_m by a rational Cauchy row

$$x(m,1),\ x(m,2),\ x(m,3),\ldots.$$

There are two independent demands. Terms within a row must settle toward the real represented by that row. The reals represented by later rows must settle toward one another. The opening table shows why a common index supplies insufficient control over those demands.

Choose k_m far enough along row m that

$$|x(m,n) - x(m,k_m)| < 2^{-m}$$
$$\text{for every } n \geq k_m,$$

and set $q_m = x(m,k_m)$. Regarding a rational as the real represented by its constant sequence, we obtain

$$|q_m - X_m| \leq 2^{-m}.$$

For two selected entries, the triangle inequality gives

$$|q_p - q_m| \leq 2^{-p} + |X_p - X_m| + 2^{-m}.$$

Given a positive tolerance, take both indices large enough that each of the three terms is less than one third of it. The outer Cauchy hypothesis supplies control of the middle term. Our choices within the rows supply control of the other two. Thus (q_m) is a rational Cauchy sequence.

Let Y be the real represented by that sequence. The embedded rationals q_m converge to their own class Y, a direct consequence of the definition of the Cauchy construction. Therefore

$$|X_m - Y| \leq |X_m - q_m| + |q_m - Y|,$$
$$|X_m - q_m| + |q_m - Y| \longrightarrow 0.$$

The sequence of reals has a limit among the reals. We have established completeness by returning a two-level description to a single rational sequence.

The schedule $m \mapsto k_m$ is the essential connective information. It makes the accuracy obtained inside a row answer to the accuracy required when that row joins the outer calculation. Different

schedules can work, and their outputs represent the same limiting real. The point remains independent of the particular presentation used to reach it.

This is also why completing a metric space a second time adds no further points once its first completion is complete, in the classical setting and up to canonical isometry. A Cauchy sequence in an already complete space has a unique limit there. Sending its equivalence class to that limit preserves distances and inverts the constant-sequence embedding. In the rational case, the estimate above displays the absorption explicitly.

Calling the result a point lets later arguments use this coordinated construction as one object. We can add it to another real, compare distances, or apply a continuous function because those operations respect the identification of equivalent approximations. The simple notation rests on relations among the finite stages. A chosen representative exposes the rational data again, and quantitative estimates determine how accurately we must obtain them. The completed point summarizes this organization. Establishing and using the underlying relations is the substantive mathematical work here.

A useful way to expose the role of quantitative control is to compare consecutive approximations. Suppose $|q_{n+1} - q_n| \leq e_n$, where $e_n \geq 0$ is a proved error allowance for that step. Repeated application of the triangle inequality gives, for $m > n$,

$$|q_m - q_n| \leq \sum_{j=n}^{m-1} e_j.$$

If the remaining allowances have a summable tail

$$E_n = \sum_{j=n}^{\infty} e_j \longrightarrow 0,$$

then every later approximation lies within E_n of q_n. Completeness supplies a limit x, with $|x - q_n| \leq E_n$. For $e_n = 2^{-n-1}$, the geometric-series identity gives $E_n = 2^{-n}$. A rule governing individual steps and an estimate of their remaining total control an arbitrarily long continuation of the calculation.

The distinction between local information and global control is explicit here. A bound on each step becomes useful for convergence through a bound on how the steps accumulate. Merely knowing that the steps tend to zero allows the harmonic partial sums, whose steps are $1/(n+1)$ and whose values diverge. Even summability can leave the effective tail bound inaccessible, as the computable example below will show.

Summable step bounds give one sufficient presentation. A Cauchy sequence can also approach its limit with infinite total variation, as $q_n = (-1)^n/n$ does for $n \geq 1$. Its convergence is established by a different estimate. The useful organization depends on the sequence. What matters is finding a relation that makes the passage from a finite observation to an unbounded family of later observations something we can justify.

What a finite description promises

The proof chose sufficiently late entries. An executable version must obtain them when asked.

Suppose a single procedure $Q(m, p)$ returns a rational satisfying

$$|Q(m, p) - X_m| \le 2^{-p}.$$

The word *single* matters. Given the row number and precision, the procedure must reach the corresponding approximation. Separate assertions that every row has some algorithm leave this uniform access to be established.

Suppose we also have an outer Cauchy modulus N, meaning

$$i, j \ge N(p) \quad \Longrightarrow \quad |X_i - X_j| \le 2^{-p}.$$

Passing to the limit in the second index shows that $|X_{N(p)} - Y| \le 2^{-p}$. Hence the explicit request

$$q_p = Q\big(N(p+1), p+1\big)$$

returns a rational with

$$\begin{aligned} |q_p - Y| &\le |q_p - X_{N(p+1)}| \\ &\quad + |X_{N(p+1)} - Y| \\ &\le 2^{-(p+1)} + 2^{-(p+1)} = 2^{-p}. \end{aligned}$$

Half the error allowance goes to the outer approximation and half to evaluating it. If Q and N are computable, so is this procedure.

In computable analysis, a rapidly converging rational sequence of this kind is a standard way to give access to a real. The request is a precision, and the response is an approximation with a specified error bound. Brattka's tutorial develops the representation and the importance of uniform access when passing from individual numbers to sequences and functions. [1]

A finite program can answer indefinitely many such requests, completing each in finite time. A finite proof can establish its error bound for every requested precision. Those two finite objects have different jobs. The program produces values. The proof justifies the relationship between every produced value and the real it represents.

Zeilberger makes this computational demand explicit in Opinion 146. His account of the decimal truncations $0.9, 0.99, 0.999, \ldots$ asks for an algorithm that takes a symbolic positive rational tolerance ε and returns a threshold $N(\varepsilon)$, beyond which the truncations stay within ε of one. He interprets the assertion entirely through finite symbols. In the classical account used here, that procedure also supplies an effective rate of convergence. The computational content gives the two interpretations substantial common ground. [16]

The equivalence relation defining the real deliberately ignores how fast a representing sequence converges. Two sequences with very different rates can represent the same point. This makes

the point independent of how we first approached it. Calculating with that point requires access through some particular presentation, whose speed and available error bounds affect what we can afford to do. Mathematical equality leaves these operational differences intact.

Classical real numbers have a wider scope. There are only countably many finite programs and uncountably many reals. A general Cauchy name can therefore contain information that no finite program generates. A procedure may also receive a name as input, querying its approximations without possessing a program that created it. Representing an arbitrary real and computing a particular real are separate mathematical questions.

Foundational assumptions also affect the available data. Selecting representatives for all rows can invoke countable choice. Classically, canonical rational approximations give a completeness proof without it. Constructive settings make the availability of representatives and moduli consequential, as Lubarsky's countermodels show. The explicit procedure above states its required names and modulus at the input. [2]

A computable sequence with an inaccessible rate

The missing modulus can carry decisive information.

Number computer programs by the nonnegative integers. At stage s, run each of the first $s+1$ programs for s steps, with its own index as input. Let H_s be the indices observed to halt, and form

$$a_s = \sum_{e \in H_s} 4^{-(e+1)}.$$

Each stage involves a finite computation and returns an exact rational. The sequence is nondecreasing and bounded by $1/3$, so it converges to a real a.

An effective procedure for approximating a to every requested accuracy would decide which programs eventually halt. To see why, suppose the halting status of indices below e has already been recovered and subtract their contributions from a. Put $w = 4^{-(e+1)}$. If program e halts, the remainder is at least w. If it never halts, the remainder is at most

$$\sum_{j>e} 4^{-(j+1)} = \frac{w}{3}.$$

Approximating that remainder to error less than $w/6$ distinguishes the two cases. Beginning at $e = 0$ and proceeding inductively would decide the halting problem. Thus a is noncomputable. A computable Cauchy modulus for (a_s) would compute a, so such a modulus is unavailable.

Its computable increments $e_s = a_{s+1} - a_s$ are nonnegative and summable. An effective bound tending to zero for their tails would compute a, so such a bound is unavailable.

The increments themselves converge to zero. Even for this familiar limit, the sequence has no computable rate. Whenever a program j first appears in H_{s+1} after being absent from H_s, that one increment is at least $4^{-(j+1)}$. Suppose we could compute a stage M such that every increment e_s with $s \geq M$ is smaller than this weight. We could then decide whether program j ever halts by checking whether it belongs to H_M. The constant zero sequence and these increments represent

the same real number. Only the first gives a computable deadline for its terms to stay within a requested distance of zero.

The cumulative sequence (a_s) is a Specker sequence, a computable nondecreasing bounded rational sequence whose limit is noncomputable. Powell discusses what such sequences imply for extracting quantitative information from convergence theorems. [3]

The obstruction resides in the future contributions. A long quiet stretch cannot certify that a particular program will never halt later. Every term is accessible, and the sequence has a perfectly definite classical limit. The demand for a justified stopping point introduces information absent from the rule producing its terms.

There is, however, a revealing finite question that this sequence can answer. Supply a positive rational accuracy ε and a computable rule $g : \mathbb{N} \to \mathbb{N}$ prescribing how far ahead to inspect from any starting point N. Ask for a window

$$[N, N + g(N)]$$

throughout which all the values stay within ε of one another.

For any computable nondecreasing rational sequence in $[0, 1]$, there is a finite search. Start at $N_0 = 0$, and set $N_{k+1} = N_k + g(N_k)$. A failed window has

$$a_{N_{k+1}} - a_{N_k} > \varepsilon.$$

More than $\lfloor 1/\varepsilon \rfloor$ successive failures would consume more than the total available increase of one. Testing endpoints therefore finds a successful window among the first $\lfloor 1/\varepsilon \rfloor + 1$ candidates. Monotonicity controls every intermediate value. The indices requested by g may be enormous, so the number of tests alone leaves their execution cost open.

The guarantee covers the requested window. A later change can occur beyond it. The corresponding classical formulation quantifies over every horizon function, including noncomputable ones,

$$\forall \varepsilon > 0 \, \forall g : \mathbb{N} \to \mathbb{N} \, \exists N$$
$$\forall i, j \in [N, N + g(N)], \quad |a_i - a_j| \leq \varepsilon.$$

This property is called *metastability* and is classically equivalent to Cauchy convergence. Its computational requests differ from obtaining one starting point that controls the entire infinite tail. For computable g and rational ε, the finite search just given supplies an executable response. This change in the form of the question can expose finite information that a global convergence rate conceals. [3,4]

The classical equivalence makes the role of the horizon precise. A Cauchy sequence satisfies every requested window once its whole tail is sufficiently close. If the sequence fails to be Cauchy, some fixed positive tolerance has a violating pair beyond every starting point. Define $g(N)$ to extend to one such pair, selecting the first in a fixed enumeration. Every requested window then fails. The quantification over all g detects exactly this possibility, even when the detecting horizon cannot be computed.

Infinity enters here through the range and order of the requests we must answer. The same sequence supports a finite search for every supplied observation window while withholding an effective guarantee about its whole future. The logical form of a theorem affects what a calculation can obtain from it.

Extracting such information from proofs is an established research programme. Proof mining investigates the effective bounds, algorithms, and dependence on parameters that can be recovered from qualitative arguments. Kohlenbach's survey develops this approach in nonlinear analysis. The form of the conclusion and the available quantitative hypotheses govern what can be extracted. [17]

How accuracy passes through a function

Completion gives a point to a Cauchy description. Continuity governs what happens when another operation acts on that description.

This can even define an operation on the completed space. Let D be dense in a metric space X, and let $f : D \to Y$ be uniformly continuous, with Y complete. Given $q_n \in D$ converging to x, uniform continuity makes $(f(q_n))$ Cauchy, so it has a limit in Y. Interleaving two sequences approaching the same x shows that both yield the same limit. Defining $F(x) = \lim_n f(q_n)$ therefore gives a well-defined extension, and the same uniform control proves its continuity. The sequence disappears from the notation $F(x)$ because the function values converge to the same answer along every sequence approaching x. [15]

The hypothesis matters. The function $q \mapsto 1/q$ is continuous on the positive rationals, but its values diverge along $q_n = 1/n$, approaching the missing endpoint zero. Uniform continuity is a sufficient condition that prevents this failure when extending to the closure.

For a function already defined at x, continuity there gives $q_n \to x \Rightarrow f(q_n) \to f(x)$. A quantitative version asks how close the input must be to make the output sufficiently close. Suppose a modulus μ satisfies, on the relevant domain,

$$\begin{aligned} |x - y| \le 2^{-\mu(p)} \quad &\Longrightarrow \\ |f(x) - f(y)| &\le 2^{-p}. \end{aligned}$$

Assume that, from the supplied name of x, we can obtain rational approximations in the domain with any requested error. To approximate $f(x)$ within 2^{-p}, obtain a rational q in the domain within $2^{-\mu(p+1)}$ of x, then evaluate $f(q)$ within $2^{-(p+1)}$. Again, two controlled errors add to the requested tolerance.

The modulus translates output accuracy into input accuracy. An evaluation procedure must then deliver the second bound. Ordinary continuity asserts the existence of suitable local control. A computable implementation requires effective access to enough control and evaluation information. These requirements belong to the representation of the function as well as the representation of its argument. [1]

A Lipschitz bound gives a particularly simple translation. If $|f(x) - f(y)| \le L|x - y|$, an input error of at most $\varepsilon/(2L)$, for $L > 0$, and an evaluation error of at most $\varepsilon/2$ suffice. The arithmetic cost of obtaining either accuracy remains a further question.

Even an elementary continuous function can demand much more accurate input than its output. For nonnegative x, y,

$$|\sqrt{x} - \sqrt{y}| \le \sqrt{|x - y|}.$$

At zero this estimate is attained. If the available information is only $x \in [0, \delta]$, then $\sqrt{x}$ can range over $[0, \sqrt{\delta}]$. Making that range at most ε wide requires $\delta \le \varepsilon^2$.

After d successive square roots, the function is

$$f_d(x) = x^{1/2^d}.$$

The same input enclosure allows an output range of width $\delta^{1/2^d}$. To limit that range to 2^{-p}, require

$$\delta \le 2^{-p2^d}.$$

In this enclosure model near zero, the required number of input bits doubles with every added square root. A compact formula for the composition preserves that dependence. If the input is known symbolically to equal zero, the additional information settles the evaluation immediately. The cost belongs to the function together with the way its input is available.

Thus a hierarchy can collapse to one expression while retaining an expensive demand for precision. Evaluation is already sensitive to how information arrives. Asking a function to reveal one of its roots introduces a further difficulty.

What it takes to find a zero

The intermediate value theorem brings these ideas together. Let f be continuous on $[a, b]$, with $f(a) < 0 < f(b)$. Classical bisection either encounters a zero or produces nested brackets $[a_n, b_n]$ of length

$$b_n - a_n = 2^{-n}(b - a)$$

whose endpoint values retain opposite signs. Nestedness makes both endpoint sequences Cauchy. Completeness supplies their limits, and the vanishing lengths make those limits equal to a point c. Continuity then gives

$$f(c) \le 0 \le f(c),$$

so $f(c) = 0$.

Each ingredient has done a particular job. The bracket rule preserves opposite signs at the endpoints. The length estimate controls location. Completeness produces the point, and continuity connects the endpoint signs to its value. If the endpoints themselves arrive through rational names, their evaluation can be scheduled by the earlier diagonal construction.

For an algorithm receiving a general computable continuous function, exact sign decisions introduce a difficulty. Increasing the accuracy of an evaluation can eventually certify a strictly positive or strictly negative value. An exact zero may keep every enclosing interval touching both signs. The classical branch that asks whether the midpoint value equals zero has no general terminating test.

The difficulty also has a geometric form. Define on $[0, 1]$

$$g(x) = \begin{cases} x - \frac{1}{3}, & x < \frac{1}{3}, \\ 0, & \frac{1}{3} \le x \le \frac{2}{3}, \\ x - \frac{2}{3}, & x > \frac{2}{3}, \end{cases}$$

Set $f_t(x) = g(x) + t$.

For $|t| < 1/6$, every f_t has opposite signs at its endpoints and is 1-Lipschitz. For $t > 0$, its unique root is $1/3 - t$. For $t < 0$, its unique root is $2/3 - t$. At $t = 0$, every point of $[1/3, 2/3]$ is a root.

As t approaches zero from its two sides, the forced choices of root approach different points. Arbitrarily close input functions therefore demand separated answers from any rule selecting one root. The continuity of evaluation and the continuity of solution selection are distinct properties.

Computable analysis states the consequence precisely. In the standard representation of continuous functions on $[0, 1]$, there is no uniform algorithm selecting an exact root from every computable continuous function with a strict endpoint sign change. Each individual such function nevertheless has some computable root. Restricting the input to functions with a unique zero restores a uniform computable operation, with computational cost still depending on the input. [5]

For the unique-zero case on $[0, 1]$, divide the current rational bracket into thirds. Evaluate the function at both interior division points in parallel, refining until one value has a certified nonzero sign. Uniqueness guarantees that both points cannot be zeros, so this search terminates. Keep the interval joining that point to the endpoint of opposite sign. Its length is at most two thirds of the previous length. Repetition yields a computable name for the root while allowing each successful sign test to take as much precision as it needs.

Individual computable roots exist in the general case for another reason. If the zero set contains a nondegenerate interval, it contains a rational root. Otherwise every open subinterval contains a rational point with nonzero value, by continuity. A parallel search through rational points in the middle third then supplies a nonzero sign and the same shrinking-bracket construction. These two cases establish existence, while choosing the appropriate case uniformly remains part of the obstruction.

An approximate answer must specify what is small. Consider two requests, each for a single rational point. For any positive rational ε, we can find q with $|f(q)| < \varepsilon$. Search through rational points, refining their evaluations in parallel until one certifies the inequality. Continuity gives a neighborhood of a zero throughout which the test succeeds.

Now ask for a rational point within $1/8$ of some zero. Even this fixed location tolerance has no uniform procedure for the plateau family f_t. Supply t by rational approximations with requested errors, so that the known function g and this input give access to $f_t = g + t$. At $t = 0$ we can answer every precision request with zero. By the time the procedure returns its point, it has made only

finitely many requests. The same replies are valid for sufficiently small shifts of either sign, whose unique roots are more than $1/3$ apart. No one point can meet the location tolerance for both.

A small value of $|f(x)|$ therefore needs an additional estimate to certify location. For $f_\delta(x) = \delta(x - 1/2)$ with $\delta > 0$, the residual at zero is $\delta/2$, while the distance to its root is always $1/2$. A known inverse estimate

$$|f(x) - f(c)| \geq m|x - c|, \qquad m > 0,$$

converts residual into location error through $|x - c| \leq |f(x)|/m$.

A short bracket known to contain a zero gives a location guarantee directly. A uniformly computable nested sequence of such rational brackets, with widths tending to zero, already supplies a name for an exact root. To answer a precision request, search until the rational width is sufficiently small. Consequently, producing compatible brackets at all accuracies carries the solution-selection obligation with it.

Zeilberger's Opinion 160 asks for an atomic sign-changing interval at a fixed symbolic mesh h. His earlier account of discrete analysis treats $h = 0$ as formal degeneration. On the mesh, either a sampled value is zero or two adjacent values have strictly opposite signs. That answers the fixed-mesh question. [6] [7]

If the samples come from a continuous function, a sign-changing interval contains at least one zero. Suppose a later calculation needs a root to greater precision. What ensures that successive answers are closing in on the same zero of the original function? In the nested-bracket argument, each refinement stays inside the previous bracket, and the shrinking widths tell us how accurately we have located their common point.

These distinctions let mathematical purposes remain explicit. A finite location tolerance may completely meet one application. A symbolic certificate may establish a property for every precision at once. The question being asked determines which relationship has to be proved.

A continuum problem inside finite systems

Numerical analysis encounters the same structure when it replaces a differential equation by a family of finite problems. Mesh refinement supplies one index. Solving the resulting algebraic system supplies another. Their interaction can be calculated exactly in a simple example.

Consider

$$\begin{aligned} -u''(x) &= 1 \quad (0 < x < 1), \\ u(0) &= u(1) = 0. \end{aligned}$$

The exact solution is $u(x) = x(1 - x)/2$. Measure an error $w \in H_0^1(0, 1)$, with square-integrable weak derivative and zero endpoint values, by its energy norm

$$\|w\|_E^2 = \int_0^1 |w'(x)|^2 \, dx.$$

Divide the interval into $n \geq 2$ equal cells of length $h = 1/n$. Let V_h consist of continuous functions that are linear on each cell and vanish at the endpoints. The Galerkin solution $u_h \in V_h$ satisfies

$$\int_0^1 u_h' v_h' \, dx = \int_0^1 v_h \, dx \qquad (v_h \in V_h).$$

In this example, u_h is the nodal interpolant of u. Indeed, the interpolation error vanishes at the ends of every cell, while v_h' is constant there. Its product with the derivative of that error therefore integrates to zero. The interpolant satisfies the displayed Galerkin equations.

On the cell $[jh, (j+1)h]$,

$$u'(x) - u_h'(x) = (j + \tfrac{1}{2})h - x.$$

Its squared integral is $h^3/12$. Summing over the n cells gives the exact discretization error

$$\|u - u_h\|_E^2 = \frac{h^2}{12}.$$

Now let $v_{h,k} \in V_h$ be the result after k iterations of an algebraic solver. Galerkin orthogonality makes the discretization error orthogonal to every function in V_h, including $u_h - v_{h,k}$. Expanding the squared norm gives

$$\|u - v_{h,k}\|_E^2 = \frac{h^2}{12} + \|u_h - v_{h,k}\|_E^2.$$

The two indices from the rational table have returned as two sources of error. Refinement controls the first term. Solver progress controls the second. An exact solve on a fixed mesh leaves the error $h/\sqrt{12}$. Refining the mesh while maintaining insufficient algebraic accuracy can also leave a persistent error. Choose $k = k(h)$ so that

$$\|u_h - v_{h,k(h)}\|_E \leq \frac{h}{\sqrt{12}}.$$

Then

$$\|u - v_{h,k(h)}\|_E \leq \frac{h}{\sqrt{6}}.$$

This is a controlled diagonal through the family of numerical calculations. It includes a rate, and it states which norm measures success. The exact equalities are specific to this example. The underlying projection argument is standard Galerkin theory, developed in Schöberl's account of coercive variational problems. [8]

For a general bounded coercive bilinear form a on a real Hilbert space V, assume

$$|a(v, w)| \leq C_a \|v\|_V \|w\|_V,$$
$$a(v, v) \geq \alpha_a \|v\|_V^2, \qquad \alpha_a > 0.$$

With a bounded load and a closed conforming Galerkin subspace $V_h \subseteq V$, Céa's estimate gives

$$\|u - u_h\|_V \le \frac{C_a}{\alpha_a} \inf_{w_h \in V_h} \|u - w_h\|_V.$$

Convergence requires the best-approximation error on the right to vanish. A rate needs corresponding approximation and regularity estimates. Parameter-dependent problems also require attention to the constants. Each computation can use a finite space, while the theorem connects its answer to the infinite-dimensional target and controls what happens as the spaces change. [8]

Making the error affordable to measure

The stopping rule above uses an error relative to u_h, which is precisely the solution still being computed. We must obtain a usable bound without making its evaluation as expensive as finishing the original problem.

In the interior nodal basis of our example, the exact system is

$$A_h U = b_h, \qquad b_h = h\mathbf{1},$$
$$A_h = \frac{1}{h}\,\mathrm{tridiag}(-1, 2, -1),$$

This matrix is the Laplacian of a weighted path with its endpoint values fixed. Extend an interior vector z by $z_0 = z_n = 0$, and let D record neighboring differences, $(Dz)_i = z_{i+1} - z_i$. Direct multiplication gives

$$A_h = D^T (h^{-1} I) D,$$
$$z^T A_h z = \frac{1}{h} \sum_{i=0}^{n-1} (z_{i+1} - z_i)^2.$$

The incidence matrix records which nodal values to subtract. The weight on each edge comes from the length of the cell. For a linear field its slope is the change in value divided by that length. Square the slope and integrate over the cell, and one factor of the length cancels. The displayed sum is therefore exactly the continuum energy of the piecewise-linear field. A graph solve can control error within that finite space. The mesh estimate from the preceding section bounds the part of the continuum solution that the space cannot represent. Together they tell us what an answer on the path means on the interval. [8] [9]

For an approximate coefficient vector V, put $r = b_h - A_h V$. Take $V = 0$ on every mesh. Then $v_h = 0$ and $r = b_h$, giving

$$\|r\|_2^2 = (n-1)h^2 = h - h^2 \longrightarrow 0.$$

Yet $\|u - v_h\|_E^2 = \int_0^1 (1/2 - x)^2 \, dx = 1/12$ on every mesh. A fixed positive absolute tolerance on this residual eventually accepts an approximation that has made no progress. The Euclidean norm here measures the vector of residuals against the nodal test functions. Both those functions and their

number change with the mesh. The relative residual $\|r\|_2/\|b_h\|_2$ stays equal to one. Interpreting the absolute residual requires a relation to the mesh and the intended error norm.

For a general approximation V, the coefficient error is $A_h^{-1}r$, so

$$\|u_h - v_h\|_E^2 = r^T A_h^{-1} r.$$

As with the root-location estimate, the inverse response connects the residual to the error in the chosen norm. This relation also survives a change of basis. If old coefficients equal T times new coefficients, with T invertible, the same approximation has transformed matrix $A' = T^T A_h T$ and residual $r' = T^T r$. Direct substitution gives

$$(r')^T (A')^{-1} r' = r^T A_h^{-1} r.$$

Rescaling the basis can make the raw residual norm arbitrarily small while leaving the represented function unchanged. The inverse operator compensates for that coordinate change, preserving the algebraic energy error inside V_h. A zero matrix residual means that the equations hold for every test function in that space. The mesh estimate accounts for the continuum error that remains. Computing the inverse action may itself require a substantial solve.

Suppose instead that a symmetric positive-definite matrix M_h admits a cheaper inverse action, with proved bounds

$$\alpha M_h \preceq A_h \preceq \beta M_h, \qquad 0 < \alpha \le \beta.$$

Here $B \preceq C$ means $z^T B z \le z^T C z$ for every vector z. Reversing the order under inversion gives

$$\frac{1}{\beta} r^T M_h^{-1} r \le r^T A_h^{-1} r \le \frac{1}{\alpha} r^T M_h^{-1} r.$$

Thus

$$\|u - v_h\|_E^2 \le \frac{h^2}{12} + \frac{1}{\alpha} r^T M_h^{-1} r.$$

For a target ε, one sufficient allocation is

$$h \le \sqrt{6}\,\varepsilon, \qquad r^T M_h^{-1} r \le \frac{\alpha \varepsilon^2}{2}.$$

Half the squared-error allowance goes to discretization and half to the algebraic solve. Other allocations can be more economical. Their validity and cost can be investigated because the error identity makes their consequences explicit.

Uniform control adds another requirement. If α and β remain fixed throughout a mesh family, they bound the condition number of the symmetric transformed operator $M_h^{-1/2} A_h M_h^{-1/2}$ by β/α. This is the spectral comparison used in positive-definite preconditioning. Its computational value also depends on constructing M_h and applying its inverse economically. [9]

Finite solvability alone leaves substantial room for deterioration across a family. Apply the tridiagonal matrix to the vectors with entries $z_i = \sin(ij\pi/n)$, for $1 \le i \le n-1$. The sine addition formula gives

$$\lambda_j(A_h) = \frac{4}{h}\sin^2\left(\frac{j\pi}{2n}\right),$$
$$1 \le j \le n-1,$$

so

$$\kappa_2(A_h) = \cot^2\left(\frac{\pi}{2n}\right) \sim \frac{4}{\pi^2 h^2}.$$

Every matrix is finite and positive definite. Its condition number nevertheless grows without bound as the mesh is refined. The family presents a mathematical problem that no inventory of the properties of one fixed matrix resolves. The present tridiagonal system still admits a direct solve in $O(n)$ arithmetic operations, using bounded work per row in elimination and back substitution. Conditioning and the work required by a particular solution method are separate questions. [9]

This example also admits a certificate using only the residual norm. Concavity places sine above its endpoint chord on $[0, \pi/2]$, giving $\sin(\pi h/2) \ge h$ for $0 < h \le 1$. The smallest eigenvalue therefore satisfies $\lambda_{\min}(A_h) \ge 4h$, and

$$\|u - v_h\|_E^2 \le \frac{h^2}{12} + \frac{\|r\|_2^2}{4h}.$$

The stopping rule $\|r\|_2 \le h^{3/2}/\sqrt{3}$ makes the algebraic error at most $h/\sqrt{12}$, as requested earlier. A finite residual calculation now certifies the chosen balance. Its threshold changes with the mesh, making the relationship between the two levels explicit.

The displayed error identities use exact matrices and residuals. They apply to any represented coefficient vector, including one produced by floating-point arithmetic. Certifying their numerical evaluation requires bounds on rounding errors. More general discretizations can also introduce quadrature, geometric, and coefficient errors that must enter the final estimate. Ball arithmetic, as developed in Arb, supplies enclosures through finite computations, while truncation and approximation bounds still require the mathematics appropriate to the problem. [10]

The one-dimensional example makes all these relations visible because its solution is explicit. In a difficult problem, obtaining the analogue of its error identity, a sharp bound, or a cheap inverse action can be the research contribution.

Finite proofs of infinite obligations

Exact cancellation supplies another way for a finite relation to control an unbounded family. Wilf and Zeilberger's theory of summation certificates begins with functions satisfying

$$F(n+1,k) - F(n,k)$$
$$= G(n,k+1) - G(n,k).$$

Here n is a nonnegative integer and k is an integer. Summing over a finite interval leaves only its boundary terms,

$$\sum_{k=L}^{K} \big(F(n+1,k) - F(n,k)\big) = G(n,K+1) - G(n,L).$$

Suppose each fixed-n row of F and G has finite support. Choose the endpoints beyond the relevant supports. The boundary terms vanish, so $S(n) = \sum_k F(n,k)$ satisfies $S(n+1) = S(n)$. Its initial value determines the sum for every n. For infinite support, absolute summability of the F-rows and $G(n,k) \to 0$ as $k \to \pm\infty$ suffice for this same conclusion. The hypotheses give the local identity its global force. [18]

In the hypergeometric setting, a rational function can encode G relative to F, making the displayed relation reducible to a polynomial identity after denominators are cleared, wherever the ratios are defined. Exceptional indices and boundary behaviour still require justification. The certificate controls a whole family because its algebra and those conditions work together. The earlier summable-step estimate controlled accumulation by bounding it. Here the relation arranges exact cancellation.

A finite certificate can also control a tail whose terms remain present. Analytic estimates make this possible even when a computation never stores those terms.

In one computation for the three-dimensional Swift–Hohenberg equation, Gameiro and Lessard obtained a candidate equilibrium from a finite Fourier expansion in 5.85 seconds. To verify it, they checked inequalities for a finite box of Fourier modes and used analytic estimates to bound every mode outside it. Together these bounds showed that a map sent a ball around the candidate into itself and contracted there. The contraction theorem gave a unique equilibrium of the full equation within that ball in the even-periodic class. [11]

Running their verification algorithm took another 15.73 seconds. Recomputing the candidate at twice the resolution in each direction took 3,084.07 seconds, just over fifty-one minutes. Agreement between the two finite computations would still leave open whether an equilibrium lay nearby. In this run it cost less to account for every omitted mode by a proved estimate than to compute the next, finer approximation. [11]

A specific relation supplies the control in each example. The Cauchy estimate combines accuracy within a row with convergence across rows. The WZ identity converts a change in the sum into boundary terms whose vanishing must be justified. In the Poisson problem, Galerkin orthogonality separates mesh error from solver error, and a spectral bound converts a residual into a solution-error estimate. The Fourier tail bound covers all omitted coefficients. Each relation specifies what the finite work establishes and which hypotheses support its reach beyond the terms explicitly examined.

Finding the relation that supports such a justification can be the difficult part of the mathematics. A sharper tail estimate can make a previously inconclusive computation decisive. A less expensive certificate can make verification affordable. Once established, such a relation becomes equipment for later calculations, carrying its hypotheses and error bounds with it.

The numerical research bet

The sixteen-second Fourier check began after the researchers had found a way to bound every omitted mode. An agent could generate a hundred more candidate equilibria and still owe an argument that an exact equilibrium lies near any one of them. Once the estimate is known, anyone can run the check wherever its hypotheses hold. Suppose the estimate is too weak for a new equation. Could a different choice of coordinates give control of the tail at a cutoff we can afford?

The same demand arises in adaptive sparse computation. Even with every matrix entry in hand, one still has to prove which errors reach the physical quantity of interest and how accurately to solve for it. A bound sharp enough to stop the calculation early can repay the work spent finding it on every later instance where its hypotheses hold. [11] [12]

Two fields can give the same new cell masses after a mesh transfer and still disagree on a later regional mass. If the requested regions share a geometric relation, one extra measurement saved from the old field can answer a whole family of questions.

Zeilberger's computational philosophy already joins representation to these questions of cost. His *Enquiry* urges mathematicians to make computational frameworks explicit and teach them to computers, while observing that a short program can take an impractically long time. *Theorems for a Price* examines the expense of proof certificates and proposes accepting weaker assurance when full verification costs too much. The numerical programme here specifies the required accuracy and assurance through the declared task, then includes the cost of meeting them in the comparison. [19] [20]

Established numerical methods already automate substantial parts of this process. Inexact Newton methods tie the accuracy of an inner linear solve to progress in the outer nonlinear problem. Eisenstat and Walker study forcing terms η_k satisfying

$$\|F(x_k) + F'(x_k)s_k\| \le \eta_k \|F(x_k)\|,$$
$$0 \le \eta_k < 1.$$

Their choices aim to avoid oversolving a local model whose prediction may be poor far from the solution. Their convergence statements have local regularity and initial-point hypotheses, while their experiments address the practical expense of the choices. [13]

Automatic construction also has a substantial history. Brandt's algebraic coarsening work uses compatible relaxation to assess proposed coarse variables and guide improvements when slowly reduced errors reveal inadequate representation. The decision about what to retain can itself become an algorithm. [14]

Halve the cell width in the Poisson calculation above. The mesh error in the energy norm falls by half. To certify the same balance between mesh error and solver error, our stopping rule lowers the absolute residual tolerance to about 35 percent of its former value. Keep any fixed positive tolerance and refinement alone will eventually let the zero answer pass, with its full energy error intact. I would look for rules of this kind on families whose exact solutions are out of reach. Count what it costs to find the bound, including the attempts that fail, and what it costs to check its hypotheses on a new mesh. Where those hypotheses hold, the next investigator can decide when to stop on a mesh the inventor never tried.

Returning from infinity

The opening table contains nothing difficult at any individual entry. Its failure arises from how entries are selected across the table. The repaired diagonal succeeds because an accuracy requirement governs that selection.

The question we intend to answer changes which information matters. The Specker sequence admits computable finite windows of stability while withholding a computable rate for its whole tail. A continuous function can be easy to evaluate while selecting one of its roots is unstable. The inverse-energy relation preserves the meaning of algebraic error across coordinate changes, while the mesh estimate controls the continuum error left by an exact finite solve.

Even the distance between two points on the completed line comes from the fractions. If (q_n) and (r_n) represent the points, the rational distances $|q_n - r_n|$ themselves form a Cauchy sequence. Its class is the distance between the points. Replace either representative by an equivalent one and the triangle inequality shows that the change in those distances tends to zero. The organization that made the points also determines which points are close.

The relations do different mathematical work in these examples. The shrinking brackets constitute a point. A uniformly continuous rule on a dense set extends to the missing points when its target is complete. The Fourier bounds certify an equilibrium of an equation already specified. In each case, ask what the finite data establish, which relation carries that conclusion across all later stages, and what further information a calculation needs. A finite proof can cover infinitely many precision demands by establishing the relationship they share.

To make a construction elementary is to make that relationship available for further thought. The point, the estimate, and the procedure then become equipment for the next problem. Infinity enters the work as a structure of obligations that mathematics can sometimes organize, transfer, and discharge with remarkable economy.

Notes and sources

1. Vasco Brattka, *Computability and Complexity in Analysis*, tutorial, 2005–2006, especially the sections on rapidly converging rational names, uniformity, and computable real functions. The explicit two-level approximation formula in this essay is derived on the page. Tutorial ↩ ↩

2. Robert S. Lubarsky, *On the Cauchy Completeness of the Constructive Cauchy Reals*, Mathematical Logic Quarterly 53, 396–414, 2007. The introduction distinguishes representatives, moduli, and equivalence classes. For the classical proof without countable choice, enumerate the rationals and select the first within 2^{-m} of each X_m. Density guarantees a candidate, and classical least-number selection determines one uniquely. The three-term estimate in the text then gives the limiting class. This is a set-theoretic definition whose computability requires further information. Section 7 explicitly observes that classical logic suffices for the positive results. The arXiv deposit dates from 2015. Full text ↩

3. Thomas Powell, *A note on the finitization of Abelian and Tauberian theorems*, Mathematical Logic Quarterly 66, 300–310, 2020, introduction and §3. Discusses Specker sequences, absent computable convergence rates, and metastability. The halting-weight construction and finite-window calculation above provide explicit versions of the mechanisms used here. The linked preprint dates from 2019. Full text ↩ ↩

4. Terence Tao, *Soft analysis, hard analysis, and the finite convergence principle*, 2007. Develops finite convergence principles and the role of a prescribed observation horizon. Author's essay ↩

5. Vasco Brattka and Guido Gherardi, *Effective Choice and Boundedness Principles in Computable Analysis*, Bulletin of Symbolic Logic 17, 73–117, 2011, §6. Distinguishes uniform IVT root selection, computable solutions to individual computable instances, and computability under the unique-zero restriction. Full text ↩

6. Doron Zeilberger, *Opinion 160*, July 2, 2017. States the fixed symbolic mesh and atomic sign-changing interval interpretation of the intermediate value question. Opinion ↩

7. Doron Zeilberger, *"Real" Analysis Is a Degenerate Case of Discrete Analysis*, 2001. Develops symbolic discretization and the proposed formal degeneration at zero mesh. Paper ↩

8. Joachim Schöberl, *Interactive Finite Elements*, chapter 10, *Coercive variational problems and their approximation*. Gives the variational assumptions, Galerkin orthogonality, and Céa estimate. The uniform one-dimensional Poisson error formulas are derived explicitly in this essay. Chapter ↩ ↩ ↩

9. Joachim Schöberl, *Interactive Finite Elements*, chapter 76, *Preconditioning*. Develops spectral comparison, the symmetric transformed operator, residuals, and the cost of inverse application. The zero-vector residual and change-of-basis identities are derived here from the displayed system. Tobin A. Driscoll and Richard J. Braun, *Fundamentals of Numerical Computation*, §2.9, explains the arithmetic savings from banded factorization, including the tridiagonal case. Preconditioning · Banded factorization ↩ ↩ ↩

10. Fredrik Johansson, *Arb — Efficient Arbitrary-Precision Midpoint-Radius Interval Arithmetic*, IEEE Transactions on Computers 66, 1281–1292, 2017, especially §§2 and 5. Describes enclosure arithmetic and the management of working precision and propagated error. Author's paper ↩

11. Marcio Gameiro and Jean-Philippe Lessard, *Efficient Rigorous Numerics for Higher-Dimensional PDEs via One-Dimensional Estimates*, SIAM Journal on Numerical Analysis 51, 2063–2087, 2013. Lemma 1.1 in the author manuscript gives the contraction criterion. Section 3.1, Table 1, reports the times for computing and verifying the approximation at cutoff 9 in each direction and for recomputing it at cutoff 18. Published paper · Author manuscript · Timing comparison ↩ ↩ ↩

12. Henry Shin, *Abstraction is no refuge from AI*, Refactor, revised November 18, 2026, especially §§1, 4, 6, 8, and 9. Compares what representations preserve and proposes a test of what specialist participation adds to their discovery. Essay ↩

13. Stanley C. Eisenstat and Homer F. Walker, *Choosing the Forcing Terms in an Inexact Newton Method*, SIAM Journal on Scientific Computing 17, 16–32, 1996. Relates forcing terms to local convergence and examines their role in avoiding oversolving. Author's paper ↩

14. Achi Brandt, *General Highly Accurate Algebraic Coarsening*, Electronic Transactions on Numerical Analysis 10, 1–20, 2000, especially §3. Develops compatible relaxation as a way to assess and improve coarse-variable selection. Paper ↩

15. Jean-François Babadjian, Didier Smets, and Franck Sueur, *Basic Functional Analysis*, course notes, October 18, 2011, §1.2.4, Theorem 1.2.3. Establishes extension of a uniformly continuous function from a dense subset to a complete target. Notes ↩

16. Doron Zeilberger, *Opinion 146*, September 12, 2015. States the decimal-convergence claim through an algorithm from symbolic rational tolerance to a symbolic threshold, while rejecting quantification over infinite sets. Opinion ↩

17. Ulrich Kohlenbach, *Proof-theoretic Methods in Nonlinear Analysis*, 2018 ICM survey, especially §1. Discusses extracting effective bounds, algorithms, parameter dependence, and weaker hypotheses from proofs. The present essay brings established methods into a connected account of approximation, certification, and computational cost. Author manuscript ↩

18. Herbert S. Wilf and Doron Zeilberger, *Rational Functions Certify Combinatorial Identities*, Journal of the American Mathematical Society 3, 147–158, 1990, especially §2. The WZ relation, telescoping proof, and boundary hypotheses support the summation argument above. The linked manuscript is dated January 25, 1989. Author manuscript ↩

19. Doron Zeilberger, *An Enquiry Concerning Human (and Computer!) [Mathematical] Understanding*, in Cristian S. Calude, ed., *Randomness & Complexity, from Leibniz to Chaitin*, World Scientific, 383–410, 2007, especially §§1.14–1.18. Develops the relation among computational frameworks, representations, and execution cost. Author manuscript ↩

20. Doron Zeilberger, *Theorems for a Price: Tomorrow's Semi-Rigorous Mathematical Culture*, Notices of the American Mathematical Society 40, 978–981, 1993. Discusses the cost of obtaining proof certificates and proposes semi-rigorous alternatives. Author manuscript ↩

Related essays

Abstraction is no refuge from AI

How accuracy, representation, and total cost enter the numerical research programme and its test of specialist contribution.

What a New Proof Makes Possible

How another proof can expose a procedure, an error estimate, or a useful change of representation.

When a Description Becomes a Verdict

How the question determines which information a representation must retain.