

Introducing Constructive Mathematics

Douglas S. Bridges

Department of Mathematics & Statistics,

University of Canterbury,

Christchurch, New Zealand

Constructive vs nonconstructive

Nonconstructive proof: an existence proof-by-contradiction of this schematic form:

Suppose that the desired object x does not exist.

Derive a contradiction.

Claim that x must exist after all.

This proves that it is impossible for x not to exist; but it does not tell us how to find/compute/construct x .

“[Nonconstructive existence proofs] inform the world that a treasure exists without disclosing its location.”
Hermann Weyl

Constructive proof of the existence of an object x : a proof that embodies an algorithm for the construction/computation of the desired object x .

Constructive proof of the existence of an object x : a proof that embodies an algorithm for the construction/computation of the desired object x .

Note: Not all proofs-by-contradiction are nonconstructive. It is perfectly constructive to prove P false by assuming that P is true and deriving a contradiction. This process just captures the constructive meaning of negation.

A nonconstructive proof:

There exists a digit that appears infinitely often in the decimal expansion of the number π .

Note first that the decimal expansion of π is nonterminating and nonrecurring, since π is irrational.

Suppose that each of the digits $0, 1, 2, \dots, 9$ occurs only finitely many times in the decimal expansion of π .

Then there exists a positive integer N such that each of $0, 1, 2, \dots, 9$ appears at most N times in the decimal expansion of π .

So that decimal expansion cannot have more than $10N$ places, which contradicts the “Note first ...” above.

Although the decimal expansion of π has been computed to billions of places, the foregoing proof does not tell us (and nobody knows) which of the digits $0, 1, 2, \dots, 9$ appears infinitely often in the nonterminating, nonrecurring expansion.

All we know is that it is impossible that each of the ten digits appears only a finite number of times.

For another nonconstructive proof, consider the statement:

There exist irrational numbers a, b such that a^b is rational.

Either $\sqrt{2}^{\sqrt{2}}$ is rational or it is irrational.

In the first case, take $a = b = \sqrt{2}$.

In the second case, take $a = \left(\sqrt{2}^{\sqrt{2}}\right)$ and $b = \sqrt{2}$.

Why is this proof nonconstructive?

- 1) It does not tell us which of the two alternatives for $\sqrt{2}^{\sqrt{2}}$ (rational or irrational) actually holds.
- 2) It therefore does not tell us which of the two choices for a and b actually produces irrational numbers with the desired property.

A constructive proof would produce, unambiguously, two irrational numbers a and b and show us that a^b is rational.

Why is this proof nonconstructive?

- 1) It does not tell us which of the two alternatives for $\sqrt{2}^{\sqrt{2}}$ (rational or irrational) actually holds.
- 2) It therefore does not tell us which of the two choices for a and b actually produces irrational numbers with the desired property.

A constructive proof would produce, unambiguously, two irrational numbers a and b and show us that a^b is rational.

Explicit example of irrational numbers a, b such that a^b is rational:

$$a = \sqrt{2}, \quad b = \log_2 9, \quad a^b = 3.$$

In fact, $\sqrt{2}^{\sqrt{2}}$ is transcendental, by the (classical) Gelfand-Schneider theorem:

a^b is transcendental if (i) a is algebraic, (ii) $a \neq 0, 1$ and (iii) b is both algebraic and irrational.

If we want to prove something constructively, then we must **not** use the law of excluded middle,

LEM: For any proposition P , either P is true or else P is false.

Allowing the use of **LEM** is tantamount to allowing nonconstructive existence proofs.

Historical note: Existence proofs-by-contradiction go back at least as far as Gauss (Fundamental Theorem of Algebra, 1799).

They became dominant after Hilbert's proof of his "basis theorem" (1888).

"Das ist nicht Mathematik. Das ist Theologie."

Paul Gordan

The constructive vs nonconstructive controversy goes back at least to Kronecker's attacks on Cantor's set theory (1877–).

It was strengthened by Brouwer's campaign, from 1907 onwards, to convert all mathematicians to the exclusive use of constructive methods, and culminated in the Grundlagenstreit between Brouwer and Hilbert in the 1920s.

“Taking the principle of excluded middle from the mathematician would be the same, say, as proscribing the telescope to the astronomer or to the boxer the use of his fists.”

David Hilbert (1928)

What is constructive mathematics?

Three ways to approach computability in mathematics:

1) Use classical computability theory.

The logic allows “decisions” that cannot be made by any real computer, so we need a clearly specified type of algorithm.

This is the approach of recursive analysis and Weihrauch’s TTE theory.

What is constructive mathematics?

Three ways to approach computability in mathematics:

1) Use classical computability theory.

The logic allows “decisions” that cannot be made by any real computer, so we need a clearly specified type of algorithm.

This is the approach of recursive analysis and Weihrauch’s TTE theory.

2) Use classical *proof mining* (Kohlenbach).

This requires a heavy logical analysis in order to extract (admittedly often good) constructive estimates from classical proofs. Moreover, it is not clear that this technique would work with deep, highly nonconstructive, results of e.g. operator algebra theory.

3) Use intuitionistic logic (Brouwer, Markov, Bishop, Martin-Löf, ...).

This

automatically takes care of the problem of noncomputational “decisions”, and

enables us to work, with any mathematical objects, in the familiar style of the analyst, algebraist, geometer, ...

Bishop-style constructive mathematics (**BISH**) is just

mathematics with intuitionistic logic

and some appropriate foundation such as

- the constructive set theory of Myhill, Aczel, and Rathjen, or
- Martin-Löf type theory.

Using intuitionistic logic, we can

- clarify distinctions of meaning obscured by classical logic, and
- allow results to have a wider range of interpretations (including recursive ones) than their counterparts proved with classical logic.

Using intuitionistic logic, we can

- clarify distinctions of meaning obscured by classical logic, and
- allow results to have a wider range of interpretations (including recursive ones) than their counterparts proved with classical logic.

“Intuitionistic logic is richer than classical logic, since the former makes distinctions that the latter fails to make.”

J.L. Bell & M. Machover

We

- do not restrict to a class of “constructive/computable objects” ;
- use intuitionistic logic to deal with the normal objects of mathematics.

Ishihara's classification:

- ▷ a constructive theory of real numbers: the usual \mathbf{R} studied with intuitionistic logic.

Ishihara's classification:

- ▷ a constructive theory of real numbers: the usual \mathbf{R} studied with intuitionistic logic.
- ▷ a theory of constructive real numbers: the recursive reals studied with classical logic.

Ishihara's classification:

- ▷ a constructive theory of real numbers: the usual \mathbf{R} studied with intuitionistic logic.
- ▷ a theory of constructive real numbers: the recursive reals studied with classical logic.
- ▷ a constructive theory of constructive real numbers: the recursive reals studied with intuitionistic logic.

Ishihara's classification:

- ▷ a constructive theory of real numbers: the usual \mathbf{R} studied with intuitionistic logic.
- ▷ a theory of constructive real numbers: the recursive reals studied with classical logic.
- ▷ a constructive theory of constructive real numbers: the recursive reals studied with intuitionistic logic.

BISH deals with the first of these.

The BHK interpretation

Modern intuitionistic logic is based on the *BHK-interpretation** of the connectives

\vee (or), \wedge (and), \rightarrow (implies), \neg (not)

and quantifiers

\exists (there exists), \forall (for all/each).

*Brouwer-Heyting-Kolmogorov

The BHK interpretation

Modern intuitionistic logic is based on the *BHK-interpretation** of the connectives

\vee (or), \wedge (and), \rightarrow (implies), \neg (not)

and quantifiers

\exists (there exists), \forall (for all/each).

Note that it is **provability**, rather than an a priori notion of truth, that is fundamental to the constructive approach.

*Brouwer-Heyting-Kolmogorov

► $P \vee Q$: either we have a proof of P or else we have a proof of Q .

▶ $P \vee Q$: either we have a proof of P or else we have a proof of Q .

▶ $P \wedge Q$: we have both a proof of P and a proof of Q .

- ▶ $P \vee Q$: either we have a proof of P or else we have a proof of Q .
- ▶ $P \wedge Q$: we have both a proof of P and a proof of Q .
- ▶ $P \rightarrow Q$: by means of an algorithm we can convert any proof of P into a proof of Q .

- ▶ $P \vee Q$: either we have a proof of P or else we have a proof of Q .
- ▶ $P \wedge Q$: we have both a proof of P and a proof of Q .
- ▶ $P \rightarrow Q$: by means of an algorithm we can convert any proof of P into a proof of Q .
- ▶ $\neg P$: assuming P , we can derive a contradiction (such as $0 = 1$); equivalently, we can prove $(P \rightarrow (0 = 1))$.

- ▶ $\exists x P(x)$: we have (i) an algorithm which computes a certain object x , and (ii) an algorithm which, using the information supplied by the application of algorithm (i), demonstrates that $P(x)$ holds.

- ▶ $\exists x P(x)$: we have (i) an algorithm which computes a certain object x , and (ii) an algorithm which, using the information supplied by the application of algorithm (i), demonstrates that $P(x)$ holds.
- ▶ $\forall x \in A P(x)$: we have an algorithm which, applied to an object x and a proof that $x \in A$, demonstrates that $P(x)$ holds.

- ▶ $\exists x P(x)$: we have (i) an algorithm which computes a certain object x , and (ii) an algorithm which, using the information supplied by the application of algorithm (i), demonstrates that $P(x)$ holds.
- ▶ $\forall_{x \in A} P(x)$: we have an algorithm which, applied to an object x and a proof that $x \in A$, demonstrates that $P(x)$ holds.

Note that in the interpretation of the statement $\forall_{x \in A} P(x)$, the proof of $P(x)$ will normally use *both the data describing the object x and the information supplied by a proof that x belongs to the set A .*

Consider the statement:

LPO For each binary sequence $a \equiv (a_n)_{n \geq 1}$ either $a_n = 0$ for all n , or else there exists N such that $a_N = 1$.

This is trivially true under classical logic.

What is its BHK interpretation?

We have an algorithm which, applied to any binary sequence a , either produces a proof that $a_n = 0$ for each n , or else computes N such that $a_N = 1$.

Claim: Such an algorithm is unlikely to be found.

Note: If the conclusion of **LPO** holds for all *increasing* binary sequences, then **LPO** holds.

The *Goldbach conjecture* (GC, 1742):

Every even integer > 2 is a sum of two primes.

Status still unknown.

The *Goldbach conjecture* (GC, 1742):

Every even integer > 2 is a sum of two primes.

Status still unknown.

Define a binary sequence a as follows.

If $2n + 2$ is a sum of two primes, set $a_n = 0$.

If there exists $k \leq n$ such that $2k + 2$ is not a sum of two primes, set $a_n = 1$.

Suppose we have an algorithm as in the BHK interpretation of **LPO**. Applied to this binary sequence, this algorithm

either proves that $a_n = 0$ for all n (i.e., proves GC)

or else computes N such that $a_N = 1$ (i.e., gives a counterexample to GC).

The use of GC here is purely illustrative: we could have used any of a multitude of unsolved problems of a certain type (Riemann hypothesis, ...).

Conclusion: the existence of an algorithm as in the BHK interpretation of **LPO** is highly doubtful.

Moreover, **LPO** is provably false in certain models of constructive mathematics (but it is not provably false in Bishop-style constructive mathematics).

We therefore stay clear of **LPO** as a working constructive principle.

Consequence: we also must avoid using any proposition that constructively implies **LPO**.

In particular, we must avoid using the full law of excluded middle.

This has a serious impact on even elementary analysis.

Consider the classically trivial proposition:

$$\forall_{x \in \mathbf{R}} (x = 0 \vee x \neq 0),$$

where

$$x \neq 0 \Leftrightarrow \exists_{n \in \mathbf{N}} (|x| > 2^{-n}).$$

Suppose we have a constructive proof—that is, an algorithm which, applied to any real number x either proves that $x = 0$ or else computes a positive integer N such that $|x| > 2^{-N}$.

Given an *increasing* binary sequence a , apply this algorithm to the real number

$$x = \sum_{n=1}^{\infty} 2^{-n} a_n = 0 \cdot a_1 a_2 a_3 \dots \text{ (infinite binary expansion).}$$

If $x = 0$, then $a_n = 0$ for all n .

Suppose we have a constructive proof—that is, an algorithm which, applied to any real number x either proves that $x = 0$ or else computes a positive integer N such that $|x| > 2^{-N}$.

Given an *increasing* binary sequence a , apply this algorithm to the real number

$$x = \sum_{n=1}^{\infty} 2^{-n} a_n = 0 \cdot a_1 a_2 a_3 \dots \text{ (infinite binary expansion).}$$

If $x = 0$, then $a_n = 0$ for all n .

If there exists N such that $|x| > 2^{-N}$, then $a_N = 1$: for if $a_N = 0$, then $x = \sum_{n=N+1}^{\infty} 2^{-n} a_n \leq 2^{-N}$, a contradiction.

Thus the proposition

$$\forall x \in \mathbf{R} (x = 0 \vee x \neq 0)$$

implies **LPO** and is therefore essentially nonconstructive!

Here is another essentially nonconstructive principle that is trivially true under classical logic.

LLPO For each binary sequence a with at most one term equal to 1, either $a_n = 0$ for all even n , or else $a_n = 0$ for all odd n .

BHK-interpretation:

We have an algorithm which, applied to any binary sequence a and the data that $a_n = 1$ for at most one n , either proves that all even-indexed terms of the sequence are 0, or else proves that all odd-indexed terms are 0.

Again, it is extremely unlikely that such an algorithm could be produced.

Moreover, **LLPO**, like **LPO**, is provably false in certain models of constructive mathematics (but it is not provably false in **BISH**).

We therefore avoid using **LLPO** as a working constructive principle.

Note that **LLPO** is a consequence of **LPO**; but **LPO** cannot be derived from **LLPO**.

Consider the classically trivial proposition: For each real number x , either $x \geq 0$ or $x \leq 0$.

Suppose we have a constructive proof: that is, an algorithm which, applied to any given real number x , either decides that $x \geq 0$ or else decides that $x \leq 0$.

Given a binary sequence a with at most one term equal to 1, apply this algorithm

$$\begin{aligned} x &= \sum_{n=1}^{\infty} (-1)^{n+1} 2^{-n} a_n \\ &= \frac{a_1}{2} - \frac{a_2}{4} + \frac{a_3}{8} - \frac{a_4}{16} + \dots \end{aligned}$$

If $x \geq 0$, then $a_n = 0$ for all even n ; if $x \leq 0$, then $a_n = 0$ for all odd n .

Conclusion: The statement

$$\forall x \in \mathbf{R} (x \geq 0 \vee x \leq 0)$$

implies **LLPO** and is therefore essentially nonconstructive.

The following elementary classical statements also turn out to be nonconstructive.

- ▷ Each real number x is either rational or irrational (that is, $x \neq r$ for each rational number r). To see this, consider

$$x = \sum_{n=1}^{\infty} \frac{1 - a_n}{n!},$$

where a is any increasing binary sequence. This is equivalent to **LPO**.

- ▷ Each real number x has a binary expansion. Note that the standard interval-halving argument for “constructing” binary expansions does not work, since we cannot necessarily decide, for a given number x between 0 and 1, whether $x \geq 1/2$ or $x \leq 1/2$. In fact, the existence of binary expansions is equivalent to **LLPO**.

- ▷ The intermediate value theorem, which is equivalent to **LLPO**.
- ▷ For all $x, y \in \mathbf{R}$, if $xy = 0$, then either $x = 0$ or $y = 0$. This is equivalent to **LLPO**. The constructive failure of this proposition clearly has implications for the theory of integral domains.

Note: classically valid statements like “each real number is either rational or irrational” that imply omniscience principles are *not false* in constructive mathematics. They cannot be, since **BISH** is consistent with classical mathematics (**CLASS**):

*Every theorem in **BISH** is also a theorem of **CLASS**.*

In fact, we can regard **CLASS** as **BISH** + **LEM**.

Another way of looking at **CLASS**: it is a model/extension of **BISH**.

Brouwer's intuitionistic mathematics (**INT**) and the recursive constructive mathematics (**RUSS**) of the Markov School both use intuitionistic logic, and both are models/extensions of **BISH**:

*Every theorem in **BISH** is also a theorem of **INT** and of **RUSS**.*

Another way of looking at **CLASS**: it is a model/extension of **BISH**.

Brouwer's intuitionistic mathematics (**INT**) and the recursive constructive mathematics (**RUSS**) of the Markov School both use intuitionistic logic, and both are models/extensions of **BISH**:

*Every theorem in **BISH** is also a theorem of **INT** and of **RUSS**.*

Andrej Bauer (Ljubljana) has shown that **BISH** can be interpreted within Weihrauch's Type 2 Effectivity framework for computable analysis.

We use these models of **BISH** to obtain independence results.

Since

INT/CLASS \vdash *Every continuous function $f : [0, 1] \rightarrow \mathbb{R}$ is uniformly continuous*

and

RUSS \vdash *There exists a continuous, real-valued function on $[0, 1]$ that is not uniformly continuous,*

we see that each of the propositions following “ \vdash ” is neither provable nor disprovable in **BISH**. In other words, each of them is independent of **BISH**.

In place of the essentially nonconstructive propositions

$$\forall x \in \mathbf{R} (x = 0 \vee x \neq 0),$$

$$\forall x \in \mathbf{R} (x \geq 0 \vee x \leq 0),$$

we have these constructively valid propositions:

1) If $a < b$, then for each real number x , either $a < x$ or $x < b$.

In place of the essentially nonconstructive propositions

$$\forall x \in \mathbf{R} (x = 0 \vee x \neq 0),$$

$$\forall x \in \mathbf{R} (x \geq 0 \vee x \leq 0),$$

we have these constructively valid propositions:

1) If $a < b$, then for each real number x , either $a < x$ or $x < b$.

2) If $(x > 0)$ is impossible, then $x \leq 0$.

Note, though, that the statement

If $(x \geq 0)$ is impossible, then $x < 0$

implies (actually, is equivalent to) another constructively dubious principle,
Markov's principle:

MP: If a is a binary sequence and it is impossible that $a_n = 0$ for all n , then there exists N such that $a_N = 1$.

The real line

Starting with the set \mathbf{N} of natural numbers, we can build the sets \mathbf{Z} (of integers) and \mathbf{Q} (of rationals) by elementary algebraic means.

By a *real number* we mean a subset \mathbf{x} of $\mathbf{Q} \times \mathbf{Q}$ such that

- ▷ for all (q, q') in \mathbf{x} , $q \leq q'$;
- ▷ for all (q, q') and (r, r') in \mathbf{x} , the closed intervals $[q, q']$ and $[r, r']$ in \mathbf{Q} intersect in points of \mathbf{Q} ;
- ▷ for each positive rational ε there exists (q, q') in \mathbf{x} such that $q' - q < \varepsilon$.

The last of these properties ensures that the set x is inhabited—that is, we can construct elements of \mathbf{x} .

Any rational number q gives rise to a canonical real number

$$\mathbf{q} = \{(q, q)\}$$

with which the original rational q is identified.

Two real numbers \mathbf{x} and \mathbf{y} are

- *equal*, written $\mathbf{x} = \mathbf{y}$, if for all $(q, q') \in \mathbf{x}$ and all $(r, r') \in \mathbf{y}$, the intervals $[q, q']$ and $[r, r']$ in \mathbf{Q} have a rational point in common;
- *unequal* (or *distinct*), written $\mathbf{x} \neq \mathbf{y}$, if there exist $(q, q') \in \mathbf{x}$ and $(r, r') \in \mathbf{y}$ such that the intervals $[q, q']$ and $[r, r']$ in \mathbf{Q} are disjoint.

Taken with the equality and inequality we have defined above, the collection of real numbers forms a set: the real line \mathbf{R} .

Let \mathbf{x}, \mathbf{y} be real numbers. We say that

▷ $\mathbf{x} > \mathbf{y}$, and that $\mathbf{y} < \mathbf{x}$, if there exist $(q, q') \in \mathbf{x}$ and $(r, r') \in \mathbf{y}$ such that $r' < q$;

▷ $\mathbf{x} \geq \mathbf{y}$, and that $\mathbf{y} \leq \mathbf{x}$, if for all $(q, q') \in \mathbf{x}$ and all $(r, r') \in \mathbf{y}$ we have $q' \geq r$.

We pass over the (complicated) definitions of the algebraic operations on real numbers.

The set \mathbf{R} is uncountable: if $(\mathbf{a}_n)_{n \geq 1}$ is a sequence of real numbers, then there exists $\mathbf{x} \in [0, 1]$ such that $\mathbf{x} \neq \mathbf{a}_n$ for each n .

The set \mathbf{R} is complete: every Cauchy sequence of real numbers converges to a limit in \mathbf{R} . (The proof requires the principle of dependent choice.)

What about the order-completeness of \mathbf{R} ?

Let S be a subset of \mathbf{R} .

An *upper bound* of/for S is a real number \mathbf{b} such that $\mathbf{x} \leq \mathbf{b}$ for each $\mathbf{x} \in S$.

We say that \mathbf{b} is the *supremum*, $\sup S$, of S if (i) it is an upper bound for S and (ii) for each $\mathbf{x} < \mathbf{b}$ there exists $s \in S$ such that $\mathbf{x} < s$.

Let S be a subset of \mathbf{R} .

An *upper bound* of/for S is a real number \mathbf{b} such that $\mathbf{x} \leq \mathbf{b}$ for each $\mathbf{x} \in S$.

We say that \mathbf{b} is the *supremum*, $\sup S$, of S if (i) it is an upper bound for S and (ii) for each $\mathbf{x} < \mathbf{b}$ there exists $s \in S$ such that $\mathbf{x} < s$.

We say that S is *upper order located* if for all rational numbers a, b with $a < b$, either $\mathbf{x} \leq b$ for all $\mathbf{x} \in S$ or else there exists $\mathbf{x} \in S$ such that $\mathbf{x} > a$.

Let S be a subset of \mathbf{R} .

An *upper bound* of/for S is a real number \mathbf{b} such that $\mathbf{x} \leq \mathbf{b}$ for each $\mathbf{x} \in S$. We say that \mathbf{b} is the *supremum*, $\sup S$, of S if (i) it is an upper bound for S and (ii) for each $\mathbf{x} < \mathbf{b}$ there exists $s \in S$ such that $\mathbf{x} < s$.

We say that S is *upper order located* if for all rational numbers a, b with $a < b$, either $\mathbf{x} \leq b$ for all $\mathbf{x} \in S$ or else there exists $\mathbf{x} \in S$ such that $\mathbf{x} > a$.

The *constructive least-upper-bound principle*:

Let S be an inhabited set of real numbers that is bounded above. Then $\sup S$ exists if and only if S is upper order located.

Analogous definitions and results hold for the infimum, $\inf S$, of S .

The upper order locatedness cannot be dropped from the hypotheses of the constructive least-upper-bound principle.

Consider any statement P . The set

$$S \equiv \{0\} \cup \{\mathbf{x} \in \mathbf{R} : x = 1 \wedge (P \vee \neg P)\}$$

is inhabited by 0 and bounded above by 1. Suppose that $\sigma \equiv \sup S$ exists. Then $\sigma \leq 1$. If $\sigma < 1$, then $\neg(P \vee \neg P)$, which is absurd. Hence $\sigma = 1$ and there exists $s \in S$ with $s > 1/2$. It follows that

$$s \in \{\mathbf{x} \in \mathbf{R} : x = 1 \wedge (P \vee \neg P)\},$$

so $P \vee \neg P$.

The upper order locatedness cannot be dropped from the hypotheses of the constructive least-upper-bound principle.

Consider any statement P . The set

$$S \equiv \{0\} \cup \{\mathbf{x} \in \mathbf{R} : x = \mathbf{1} \wedge (P \vee \neg P)\}$$

is inhabited by 0 and bounded above by 1. Suppose that $\sigma \equiv \sup S$ exists. Then $\sigma \leq 1$. If $\sigma < 1$, then $\neg(P \vee \neg P)$, which is absurd. Hence $\sigma = 1$ and there exists $s \in S$ with $s > 1/2$. It follows that

$$s \in \{\mathbf{x} \in \mathbf{R} : x = \mathbf{1} \wedge (P \vee \neg P)\},$$

so $P \vee \neg P$.

From now on, we drop boldface notation for real numbers.

Normed linear spaces

We'll skip over the theory of metric spaces.

Normed linear spaces

We'll skip over the theory of metric spaces.

Let X be a linear space over the field \mathbb{K} (either \mathbf{R} or \mathbf{C}). An inequality relation \neq on X is said to be *compatible with the linear structure* on X if, for all $x, y \in X$ and $t \in \mathbb{K}$,

$$\begin{aligned}x \neq y &\Leftrightarrow x - y \neq 0, \\x + y \neq 0 &\Rightarrow x \neq 0 \vee y \neq 0, \\tx \neq 0 &\Rightarrow t \neq 0 \wedge x \neq 0.\end{aligned}$$

Then

$$x \neq y \Rightarrow \forall z \in X (x + z \neq y + z).$$

From now on, “linear space” means “linear space with a compatible inequality”.

A *seminorm* on a linear space X is mapping $x \rightsquigarrow \|x\|$ of X into the nonnegative real line \mathbf{R}^{0+} such that for all x, y in X and all t in \mathbb{K} ,

- $\|x\| > 0 \Rightarrow x \neq 0$,
- $\|tx\| = |t| \|x\|$, and
- $\|x + y\| \leq \|x\| + \|y\|$.

Then $(X, \| \cdot \|)$ —or just X itself—is a *seminormed (linear) space* over \mathbb{K} . If the inequality satisfies

$$x \neq 0 \Leftrightarrow \|x\| > 0,$$

then $\| \cdot \|$ is called a *norm* on X .

Let X be a normed space. Then the mapping $(x, y) \rightsquigarrow \|x - y\|$ of $X \times X$ into \mathbf{R} provides the associated metric ρ on X .

The *unit ball* of X is the closed ball with centre 0 and radius 1 ,

$$B_X = \overline{B}_X(0, 1) = \overline{B}(0, 1) = \{x \in X : \|x\| \leq 1\},$$

relative to that metric. This ball, like any open or closed ball in a normed space, is located.

We pass over most of the standard examples, notions, and elementary properties familiar from the classical theory of normed spaces.

A mapping u between vector spaces X, Y is *linear* if

$$u(x + y) = u(x) + u(y) \quad \text{and} \quad u(tx) = tu(x)$$

whenever $x, y \in X$ and $t \in \mathbb{K}$.

If $X = Y$, then u is called an *operator* on X .

If $Y = \mathbb{K}$, then u is called a *linear functional* on X .

A linear mapping $u : X \rightarrow Y$ between normed spaces is continuous on X if and only if it is *bounded*, in the sense that there exists $c > 0$ such that $\|u(x)\| \leq c \|x\|$ for each $x \in X$.

This is not enough to ensure that u is *normed/normable*, in the sense that

$$\|u\| \equiv \sup \{ \|u(x)\| : x \in X, \|x\| \leq 1 \}$$

exists.

There is a criterion for the normability of nonzero bounded linear functionals.

Proposition: *A nonzero linear functional u on a normed space X is normed if and only if*

$$\ker u \equiv \{x \in X : u(x) = 0\}$$

is located in X .

Basic idea of the proof: for each $x \in X$,

$$\rho(x, \ker u) = \frac{|u(x)|}{\|u\|},$$

provided either $\rho(x, \ker u)$ or $\|u\|$ exists.

The classically redundant notion of normability plays a key role in the Riesz representation theorem:

Theorem: *A bounded linear functional u on a Hilbert space is normed if and only if there exists a unique vector $a \in H$ such that $u(x) = \langle x, a \rangle$ for each $x \in H$.*

Proving “only if” is the harder part, in which the classical argument goes through if $\|u\| > 0$.

For the general case, we use a little trick.

We consider the direct sum $H \oplus \mathbb{K}$, a Hilbert space with the inner product

$$\langle (x, \zeta), (x', \zeta') \rangle \equiv \langle x, x' \rangle + \zeta \zeta',$$

on which we define a nonzero bounded linear functional v by

$$v(x, \zeta) = u(x) + \zeta.$$

A little work shows that v is normed. By the first part of the proof, there exists $a \in X$ such that

$$v(x, \zeta) = \langle (x, \zeta), (a, 1) \rangle$$

for each $(x, \zeta) \in H \oplus \mathbb{K}$. Then $u(x) = \langle x, a \rangle$ for each $x \in H$.

For any not-necessarily-bounded operator T on H , we define the *adjoint* T^* , if it exists, by the equation

$$\langle Tx, y \rangle = \langle x, T^*y \rangle \quad (x, y \in H), \quad (1)$$

in which case we refer to T as *jointed*.

Classically, the Riesz representation theorem enables us to prove the existence of T^* for any bounded operator T on H .

For any not-necessarily-bounded operator T on H , we define the *adjoint* T^* , if it exists, by the equation

$$\langle Tx, y \rangle = \langle x, T^*y \rangle \quad (x, y \in H), \quad (2)$$

in which case we refer to T as *jointed*.

Classically, the Riesz representation theorem enables us to prove the existence of T^* for any bounded operator T on H .

Constructively, the universal existence of adjoints for bounded operators on l_2 implies **LPO**.

Can we characterise those bounded operators for which the adjoint exists?

Can we characterise those bounded operators for which the adjoint exists?

Yes, by the following result of Ishihara and Richman.

Proposition. *A bounded operator T on a Hilbert space H is jointed if and only if it maps the unit ball of H to a located set—that is, if and only if*

$$\rho(x, T(B)) \equiv \inf \{ \rho(x, Ty) : \|y\| \leq 1 \}$$

exists for each $x \in H$.

Best approximation theory

Let X be a linear space.

Vectors e_1, \dots, e_n in X are *linearly independent* if for all scalars $\lambda_1, \dots, \lambda_n$ such that $\sum_{i=1}^n |\lambda_i| > 0$ we have $\sum_{i=1}^n \lambda_i e_i \neq 0$.

We say that X is *finite-dimensional* if either $X = \{0\}$ or else it contains finitely many linearly independent vectors e_1, \dots, e_n such that for each $x \in X$ there exist scalars $\lambda_1, \dots, \lambda_n$ for which $x = \sum_{i=1}^n \lambda_i e_i$.

In the first case, X is *0-dimensional*.

In the second, X is *n-dimensional* and $\{e_1, \dots, e_n\}$ is a *basis* of X . The *coordinates* $u_i(x) \equiv \lambda_i$ are uniquely determined by x , and the *coordinate functionals* $u_i : X \rightarrow \mathbb{K}$ are linear mappings.

Inducting on the dimension, we can prove, in turn, that

- (i) the coordinate functionals on a finite-dimensional normed space are bounded,
and

- (ii) every linear mapping of a finite-dimensional normed space into a normed space is bounded and normed.

Inducting on the dimension, we can prove, in turn, that

- (i) the coordinate functionals on a finite-dimensional normed space are bounded, and
- (ii) every linear mapping of a finite-dimensional normed space into a normed space is bounded and normed.

We also have this familiar

Proposition: *A normed space is finite-dimensional if and only if its closed unit ball is compact.*

A subspace Y of a metric space (X, ρ) is called *proximal* if each element of X has a *best approximation* in Y : that is, if for each $a \in X$ there exists $b \in Y$ such that $\rho(a, b) \leq \rho(a, y)$ for all $y \in Y$. In that case, Y is *located* in X :

$$\rho(x, Y) \equiv \inf \{ \rho(x, y) : y \in Y \}$$

exists for each $x \in X$.

Classical fundamental theorem of approximation theory: a finite-dimensional subspace V of a real normed space X is proximal.

A subspace Y of a metric space (X, ρ) is called *proximal* if each element of X has a *best approximation* in Y : that is, if for each $a \in X$ there exists $b \in Y$ such that $\rho(a, b) \leq \rho(a, y)$ for all $y \in Y$. In that case, Y is *located* in X :

$$\rho(x, Y) \equiv \inf \{ \rho(x, y) : y \in Y \}$$

exists for each $x \in X$.

Classical fundamental theorem of approximation theory: a finite-dimensional subspace V of a real normed space X is proximal.

The constructive content of the classical proof of this result is simply that the finite-dimensional subspace V is located in X . The existence of a best approximation in the case of general X and V implies **LLPO**.

For a constructive counterpart to the theorem, we introduce new notions.

We say that a has *at most one best approximation* in Y if for all distinct points y, y' in Y , there exists $z \in Y$ such that

$$\max \{ \rho(a, y), \rho(a, y') \} > \rho(a, z).$$

We call Y *quasiproximinal* if each point of X with at most one best approximation in Y actually has a (perforce unique) best approximation in Y .

Proximinal implies quasiproximinal.

We say that a has *at most one best approximation* in Y if for all distinct points y, y' in Y , there exists $z \in Y$ such that

$$\max \{ \rho(a, y), \rho(a, y') \} > \rho(a, z).$$

We call Y *quasiproximinal* if each point of X with at most one best approximation in Y actually has a (perforce unique) best approximation in Y .

Proximinal implies quasiproximinal.

The converse cannot be proved in **BISH** but can be proved using **LEM**:

Suppose that $a \in X$ has no best approximation in a quasiproximinal subspace Y of X . Then (classically!) a has at most one best approximation in Y ; so, by quasiproximality, a has a best approximation in Y , which is a contradiction.

The next lemma is crucial for the proof of our approximation theorem. In it,

$$\mathbf{R}e \equiv \{\lambda e : \lambda \in \mathbf{R}\}.$$

Lemma: *Let x, e be elements of a real normed space X with $e \neq 0$, and let $d \geq 0$. Suppose that*

$$\max \{ \|x - te\|, \|x - t'e\| \} > d$$

whenever t, t' are distinct real numbers. Then there exists $\tau \in \mathbf{R}$ such that if $\|x - \tau e\| > d$, then $\rho(x, \mathbf{R}e) > d$.

The next lemma is crucial for the proof of our approximation theorem. In it,

$$\mathbf{R}e \equiv \{\lambda e : \lambda \in \mathbf{R}\}.$$

Lemma: *Let x, e be elements of a real normed space X with $e \neq 0$, and let $d \geq 0$. Suppose that*

$$\max \left\{ \|x - te\|, \|x - t'e\| \right\} > d$$

whenever t, t' are distinct real numbers. Then there exists $\tau \in \mathbf{R}$ such that if $\|x - \tau e\| > d$, then $\rho(x, \mathbf{R}e) > 0$.

Constructive fundamental theorem of approximation theory: *Every finite-dimensional subspace of a real normed space is quasiproximinal.*

Proved by induction on the dimension n of the subspace. The case $n = 0$ is trivial; the case $n = 1$ follows easily from the lemma. The lemma is also used in the induction step.

Our fundamental theorem can be used to give an algorithmic proof of the existence of best Chebyshev approximations—i.e. in the case $X = C[0, 1]$, and the finite-dimensional space is generated by the monomials $1, x, x^2, \dots, x^n$.

Our fundamental theorem can be used to give an algorithmic proof of the existence of best Chebyshev approximations—i.e. in the case $X = C[0, 1]$, and the finite-dimensional space is generated by the monomials $1, x, x^2, \dots, x^n$.

What about the well-known *Remez algorithm* for computing best Chebyshev approximations?

Our fundamental theorem can be used to give an algorithmic proof of the existence of best Chebyshev approximations—i.e. in the case $X = C[0, 1]$, and the finite-dimensional space is generated by the monomials $1, x, x^2, \dots, x^n$.

What about the well-known *Remez algorithm* for computing best Chebyshev approximations?

The classical proof of its convergence has an essentially nonconstructive step and therefore does not provide rates of convergence!

Our fundamental theorem can be used to give an algorithmic proof of the existence of best Chebyshev approximations—i.e. in the case $X = C[0, 1]$, and the finite-dimensional space is generated by the monomials $1, x, x^2, \dots, x^n$.

What about the well-known *Remez algorithm* for computing best Chebyshev approximations?

The classical proof of its convergence has an essentially nonconstructive step and therefore does not provide rates of convergence!

A constructive version of the Remez algorithm, with a fully constructive proof of convergence, was given by dsb in 1978.

The scope of constructive mathematics

Analysis

Complex analysis, including the Picard theorems and the Riemann mapping theorem.

Abstract integration and measure theory.

Haar measure on locally compact groups.

Spectral theory for normal operators.

Banach algebras.

Operator algebras: characterisation of ultraweakly continuous linear functionals; double commutant theorem for commutative von Neumann algebras.

Algebra

Constructive counterparts of large tracts of classical algebraic theories, including Galois theory and the Hilbert basis theorem (!).

Topology

Formal (pointfree) topology (Sambin, 2013).

Apartness and uniform spaces (Bridges and Vîţă, 2011).