Intuitionistic Mathematics and Realizability in the Physical World

Andrej Bauer
University of Ljubljana

January 18, 2012

Abstract

Intuitionistic mathematics perceives subtle variations in meaning where classical mathematics asserts equivalence, and permits geometrically and computationally motivated axioms that classical mathematics prohibits. It is therefore well-suited as a logical foundation on which questions about computability in the real world are studied.

The realizability interpretation explains the computational content of intuitionistic mathematics, and relates it to classical models of computation, as well as to more speculative ones that push the laws of physics to their limits. Through the realizability interpretation Brouwerian continuity principles and Markovian computability axioms become statements about the computational nature of the physical world.

1 Intuitionistic understanding of truth

Constructive mathematics, whose main proponent was Erret Bishop,[1] lives at the fringe of mainstream mathematics. It is largely misunderstood by mathematicians, and consequently by physicists as well. Contrary to the popular opinion, constructive mathematics is not poorer but *richer* in possibilities of mathematical expression than its classical counterpart. It differentiates meaning where classical mathematics asserts equivalence and thrives on geometric and computational intuitions that are banned by the classical doctrine. In this contribution I explore what constructive mathematics and the related realizability interpretation of intuitionistic logic have to offer to those who are interested in real-world computation.

If classical and constructive mathematicians just disagreed about what was true, the matter would be resolved easily. Unfortunately they use the same words to mean two different things, which is always an excellent source of confusion. The origin of the schism lies in the criteria for truth, i.e., in what makes a statement true. Speaking vaguely, intuitionistic logic demands *positive evidence* of truth, while classical logic is

[1] Bishop's constructivism is *compatible* with respect to classical mathematics, and should not be confused with the intuitionism of L.E.J. Brouwer, which assumes principles that are classically false.

happy with *lack of negative evidence*. The constructive view is closer to the criterion of truth in science, where a statement is accepted only after it has been positively confirmed by an experiment.

The Brouwer-Hetying-Kolmogorov (BHK) interpretation explains informally what counts as positive evidence:

- evidence for a conjunction $\phi \wedge \psi$ consists of evidence for ϕ together with evidence for ψ,
- evidence for a disjunction $\phi \vee \psi$ consists of evidence for ϕ, or for evidence for ψ, together with information about which disjunct the evidence is for,
- evidence for an implication $\phi \Rightarrow \psi$ is a method for converting evidence for ϕ to evidence for ψ,
- evidence for a universal quantification $\forall x \in A \,.\, \phi(x)$ is a method which takes (a representation of) any $a \in A$ can converts it to evidence for $\phi(a)$,
- evidence for an existential quantification $\exists x \in A \,.\, \phi(x)$ is (a representation of) an $a \in A$ together with evidence for $\phi(a)$,
- evidence for a negation $\neg\phi$ is lack of evidence for ϕ,
- anything is evidence for truth $\top$,
- there is no evidence for falsehood $\bot$.

The BHK interpretation leaves the meaning of "evidence" and "method" unexplained. Its formalization leads to the *realizability interpretation*, which we shall consider in detail in Section 3. Let us apply the BHK interpretation to the archetypical example, the Law of excluded middle:

> *"Every proposition is either true or false."*

The criterion by which a classical mathematician judges the law is

> *"Is it the case that each proposition is either true or false?"*

whereas an intuitionistic mathematician is more demanding:

> *"Is there a* method *for determining, given any proposition, which of the two possibilities holds?"*

Even though it might be the case that each particular proposition happens to be true or false, there might still be no method for deciding the truth and falsehood of propositions. The precise reason for there being no such method depends crucially on what counts as one. For example, if methods are required to be Turing computable, then a standard argument in computability theory shows that a method for deciding arbitrary propositions would yield a Halting oracle.

The intuitionistic mathematician does not accept the Law of excluded middle. When asked to produce counterexamples, he states the there are none, and makes things worse by claiming that there is in fact no proposition which is neither true nor false. Such a position looks thoroughly nonsensical to the classical mathematician. In terms of formal logic, the intuitionistic mathematician does not accept the Law of excluded middle

$$\forall \phi \in \mathsf{Prop} \,.\, \phi \vee \neg\phi,$$

yet he claims

$$\neg\exists \phi \in \mathsf{Prop} \,.\, \neg\phi \wedge \neg\neg\phi.$$

Are not these two propositions logically equivalent? Only if we accept the Law of excluded middle, so the intuitionist has not been caught in an inconsistency.

It is instructive to consider the difference between ϕ and $\neg\neg\phi$ in intuitionistic logic. Whereas ϕ holds if there is evidence supporting it, $\neg\neg\phi$ holds if there is no evidence that there is no evidence of ϕ. In other words $\neg\neg\phi$ holds when ϕ cannot be falsified, or when ϕ is *potentially true*. Every actually true statement is potentially true, but the converse need not hold. An example of a statement which is potentially true but not intuitionistically true would be "There is an undetectable elephant in Bertrand's room." Clearly, there can be no evidence against such a claim, but there can be no positive evidence for it either.

If a statement ϕ is intuitionistically equivalent to its double negation $\neg\neg\phi$, then there is no difference between its being actually and potentially true. In logic such statements are called $\neg\neg$-*stable*. The laws of physics are typically stated as universal statements involving equations, possibly with side conditions. Because such statements[2] are $\neg\neg$-stable the laws of physics are crisp: if they are potentially true then they are actually true.

There is a well-known translation of classical logic into intuitionistic logic, called the *double negation translation*. It transforms a given proposition by placing double negations in front of every quantifier and logical connective. In terms of our terminology, it inserts the adverb "potentially" everywhere, so the intuitionistic mathematician can interpret the utterances of his classical colleague as statements about potential truth. For example, the statement "every Turing machine either halts or runs forever" is translated to "there is no Turing machine which neither halts nor runs forever" Unfortunately, there is no such easy translation in the opposite direction. The classical mathematician must build models of intuitionistic logic to imagine how the intuitionistic friends think.

The intuitionistic and classical understanding of $\neg\neg$-stable statements coincide. This is fortunate because we can all agree on what the universal laws of physics say. However, a disagreement is reached when we discuss which things exist in our universe, for intuitionistic existence requires explicit evidence where classical existence is satisfied with lack of negative evidence. Again we see the similarity between intuitionistic and scientific thinking.

[2] A statement built from the universal quantifier $\forall$, conjunction $\wedge$, implication $\Rightarrow$, and numerical equality $=$ is $\neg\neg$-stable, as can be easily verified.

2 Synthetic differential geometry

Before focusing on the central theme, let us show how intuitionistic logic allows us to accept axioms which are useful to physicists in their everyday calculations, but are officially prohibited because of the reign of classical logic.

Nowadays we teach analysis in the style of Cauchy and Weierstrass, with ϵ-δ definitions of continuity and differentiability. We might even tell our students that the original differential calculus of Leibniz and Newton was based on a flawed concept of *infinitesimals*, which were supposed to be infinitely small non-zero quantities. Yet when the students attend a physics class they never see an ϵ-δ argument. Their professor freely differentiates everything in sight and uses the outlawed infinitesimals. Not having been told the precise rules for handling infinitesimals, students get confused. A typical calculation might go like this:

$$(x^2)' = \frac{(x+dx)^2 - x^2}{dx} = \frac{2x \cdot dx + dx^2}{dx} = 2x + dx = 2x.$$

In the last step we pretend that dx is so small in comparison to $2x$ that it can be neglected. Surely then it is also small with respect to x, and can be neglected already in the first step:

$$(x^2)' = \frac{(x+dx)^2 - x^2}{dx} = \frac{x^2 - x^2}{dx} = \frac{0}{dx} = 0 \quad ?!$$

My own experience as a student was that asking these kinds of questions about infinitesimals only lead to frustration and further confusion. No wonder the exile of infinitesimals was welcomed by mathematicians.

However, if the ϵ-δ analysis is of no use to physicists, it is mathematicians' task to provide a theory of infinitesimal calculus which does not suffer from the 17th century deficiencies. That this can be done was shown by William F. Lawvere and others, under the name *Synthetic Differential Geometry*.[3] Its fundamental axiom is the following principle, for which physicists should feel a certain degree of affinity:

> ***Principle of micro-affinity:*** *An infinitesimal change in the independent variable causes an affine (linear) change in the dependent variable.*

More precisely, if $f : \mathbb{R} \to \mathbb{R}$ is *any* function, $x \in \mathbb{R}$ and dx is an infinitesimal, then there exists a unique number $f'(x)$, called *the derivative of f at x*, such that $f(x + dx) = f(x) + f'(x)dx$ for all infinitesimals dx. A quantity dx is *infinitesimal* (of second degree) if its square dx^2 is zero.

Classical logic contradicts the Principle of micro-affinity because it shows that the only infinitesimal is 0. Indeed, if $dx^2 = 0$ then dx cannot be positive as that would imply $dx^2 > 0$, and it cannot be negative for the same reason. But if 0 is the only infinitesimal then the Principle of micro-affinity is false because $f(x) = x$ has both 0 and 1 as its derivative. We are left with no choice but to abandon classical logic.

[3] Synthetic differential geometry is not to be confused with Robinson's non-standard analysis [9], which uses classical logic and does not contain nilpotent infinitesimals.

The Principle of micro-affinity has strange consequences. Because it fails if all infinitesimals are 0, there must *potentially* exist some that are not zero. On the other hand, since infinitesimals can be neither negative nor positive, they cannot be different from zero, which is to say that they are *potentially* zero.[4] However strange this seems, it is not a contradiction. It may be helpful to think of infinitesimals as quantities so small that they cannot be experimentally distinguished from zero (they are potentially zero), but neither can they be shown to all equal zero (potentially there are some non-zero ones).[5]

The following consequence of the Principle of micro-affinity is quite useful for calculations:

Law of cancellation: *If $a \cdot dx = b \cdot dx$ for all infinitesimals dx, then $a = b$.*

To derive the law, consider the function $f(x) = ax - bx$ and compute

$$f(x + dx) - f(x) = (a - b) \cdot dx = 0 \cdot dx,$$

where we used the assumption that $(a - b) \cdot dx = 0$. Because both $a - b$ and 0 are the derivative of f, they are equal, hence $a = b$.

The Law of cancellation is important in practical calculations because it allows us to cancel infinitesimals as long as they are *arbitrary*. Without it we could do no such thing, as infinitesimals are not invertible (they are potentially zero). For illustration, let us calculate the derivative of $f(x) = x^2$, this time correctly. For an *arbitrary* infinitesimal dx, by taking into account $dx^2 = 0$ we compute

$$f'(x) \cdot dx = f(x + dx) - f(x) = (x + dx)^2 - x^2 = 2x \cdot dx,$$

and by canceling dx we get $f'(x) = 2x$. We can similarly derive all the usual rules for computing derivatives, prove the fundamental theorem of calculus in two lines, derive the wave and heat equations, etc., as long as we stick to a simple set of rules: never assume anything specific about an infinitesimal dx, other than $dx^2 = 0$; cancel infinitesimals on both sides of an equation, do not divide by them; and do not prove equations by contradiction, just calculate as physicists always do.

The Principle of micro-affinity implies that every function has derivatives of all orders, everywhere. How are we supposed to model sudden changes, such as reflection of a light-ray or a bouncing ball? Intuitionistic treatment of sudden changes is more profound that the classical one. Consider a ball which moves freely up to time t_0, bounces off a wall, and moves freely afterwards. Its position p is described as a function of time t in two parts,

$$\begin{aligned} p(t) &= p_1(t) \quad \text{if } t \leq t_0, \\ p(t) &= p_2(t) \quad \text{if } t \geq t_0, \end{aligned}$$

[4] We have used the intuitionistically acceptable principle that there is no number which is neither negative, nor zero, nor positive. The stronger law of trichotomy, which states that every number is either negative, zero, or positive, is not intuitionistically acceptable.

[5] Incidentally, we are not talking about lengths below the Planck length, as there are clearly positive reals numbers smaller than $1.6 \cdot 10^{-35}$.

where p_1 and p_2 are smooth functions and $p_1(t_0) = p_2(t_0)$. Because p is defined separately for $t \leq t_0$ and for $t \geq t_0$, its domain of definition is the union of two half-lines $(-\infty, t_0] \cup [t_0, \infty)$. Obviously this is a subset of the reals $\mathbb{R}$, but it may come as a bit of surprise that it is *not* the whole $\mathbb{R}$ because that would amount to the statement

> *"For every real number t, either $t \leq t_0$ or $t \geq t_0$."*

which is inconsistent with the Principle of micro-affinity, and is generally not acceptable intuitionistically.[6] Therefore, the domain of p is a proper subset of $\mathbb{R}$ and so p is *not* defined everywhere on $\mathbb{R}$. But neither is there an instant of time at which p is undefined! The domain of p should not be imagined as "$\mathbb{R}$ with missing points" but rather as "$\mathbb{R}$ with extra information". In the smooth world, any sudden changes in physical quantities are recorded at the level of logic as wrinkles in space-time.

Consistency of Synthetic differential geometry is secured by sheaf-theoretic models based on a geometric interpretation of intuitionistic logic that measures the truth of a statement as a region of space in which the statement holds locally. Let us explain this by way of example. If T denotes temperature, then in some places we have $T < 273$ and $T \geq 273$ in others. We say that $T < 273$ holds *locally* at a given place if it holds everywhere in a small neighborhood of the place, and likewise for $T \geq 273$. The statement

> *"Either $T < 273$ or $T \geq 273$."*

is then said to hold at a given place when either $T < 273$ holds there locally, or $T \geq 273$ does. Even if we assume that T varies smoothly, there might be a place where the temperature is exactly 273 so that $T \geq 273$ holds, but an arbitrarily small displacement leads to $T < 273$. The above statement, which happens to be an instance of the Law of excluded middle, therefore need not hold universally. The interested readers are invited to consult Bell's booklet [1] and the more advanced texts [6, 8] for further reading on this topic.

3 The Realizability Interpretation

Formalization of the BHK interpretation leads to the *realizability interpretation* of intuitionistic mathematics. The interpretation acts as a bridge between constructive mathematics and computational models of various kinds.

The essential idea of realizability theory is expressed mathematically with the *realizability relation*, written as

$$r \Vdash \varphi$$

and read as "*r realizes* φ". Depending on the context, we also say that r represents, implements, or witnesses φ. We usually think of r as something concrete (a program,

[6] To see that there is no realistic method for determining whether $t \leq t_0$ or $t \geq t_0$, consider a real-valued measured quantity t. Whatever experiment we perform, there will always be a small measurement uncertainty $t \pm \Delta t$. If by luck $t + \Delta t < t_0$ or $t - \Delta t > t_0$, then we can decide which of $t \leq t_0$ and $t \geq t_0$ holds. Otherwise, we might perform a more precise measurement, but there is no guarantee that we shall succeed in finite time.

a number, a sequence of bits), and of φ as something abstract (an element of a set, a function, a logical statement), even though they are both mathematical objects. Later on we shall speculate on r being a real-world entity, but for now we focus on the mathematical aspects of realizability.

For a sound interpretation of intuitionistic mathematics, the realizers should support certain operations, such as:[7] a *pairing* which combines realizers r and s into a single one $\langle r, s\rangle$, together with projections that recover the components of a pair; an *application* operation $r \cdot s$ which applies the function encoded by the realizer r to the realizer s; and a suitable encoding of natural numbers which assigns to each $n \in \mathbb{N}$ a *numeral* $\overline{n}$. The BHK interpretation is then formalized as follows:

$$
\begin{array}{rll}
\langle r, s\rangle \Vdash \phi \wedge \psi & \text{iff} & r \Vdash \phi \text{ and } s \Vdash \psi \\
\langle r, s\rangle \Vdash \phi \vee \psi & \text{iff} & r = \overline{0} \text{ and } s \Vdash \phi, \text{ or } r = \overline{1} \text{ and } s \Vdash \psi \\
r \Vdash \phi \Rightarrow \psi & \text{iff} & \text{for all } s, \text{ if } s \Vdash \phi \text{ then } r \cdot s \Vdash \psi \\
r \Vdash \forall x \in A \,.\, \phi(x) & \text{iff} & \text{for all } s \text{ and } a, \text{ if } s \Vdash (a \in A) \text{ then } r \cdot s \Vdash \phi(a) \\
\langle r, s\rangle \Vdash \exists x \in A \,.\, \phi(x) & \text{iff} & \text{there is } a \text{ such that } r \Vdash (a \in A) \text{ and } s \Vdash \phi(a) \\
r \Vdash \neg\phi & \text{iff} & \text{there is no } s \text{ such that } s \Vdash \phi \\
r \Vdash \top & \text{always} & \\
r \Vdash \bot & \text{never} &
\end{array}
$$

In realizability a set must always be introduced together with realizers for its membership relation. In other words, we have to explain how the elements of the set are represented by the realizers. For instance, a natural number $n \in \mathbb{N}$ is represented by the corresponding numeral $\overline{n}$. As always we need to keep in mind that n is an abstract mathematical entity whereas $\overline{n}$ is something concrete, such as a sequence of bits in computer memory, or a string of symbols on paper.

We can use the realizability relation to compute the realizers of any logical statement. For example, if we unravel the realizability interpretation of the principle of mathematical induction, which is formally expressed by the formula

$$\phi(0) \wedge (\forall k \in \mathbb{N} \,.\, \phi(k) \Rightarrow \phi(k+1)) \Rightarrow \forall n \in \mathbb{N} \,.\, \phi(n),$$

we discover that its realizers correspond precisely to a known concept in programming, namely *primitive* recursion. There are many examples where the realizability interpretation of a well-known statement in mathematics turns out to be a well-known programming concept. At a more abstract level the connection between logic and computation is expressed by the following soundness theorem:

> ***Soundness Theorem:*** *from an intuitionistic proof of a statement we can extract a realizer for it.*

The theorem has at least three uses. First, because the extraction process is completely mechanical, we can actually compute realizers from formal proofs and constructions, a

[7] The technical requirement is that the realizers should form a *partial combinatory algebra*.

task best left to computer proof systems.[8] Second, it is often easier to prove a statement intuitionistically than it is to construct a realizer for it. Third, we can conclude that a statement is not realized if we show that it implies another statement which is known not to be realized.

4 Realizability in the real world

We have so far not specified a particular model of computation on top of which realizability is built. Kleene's original number realizability [3] was based on partial recursive functions, encoded by natural numbers. Later Kleene gave a realizability model based on functions [5], and there are still other realizability models based on programming languages, topological spaces, and even games, see [13] for a general theory of realizability.

We would like to use "real-world computation" as the underlying model for realizability. We do not know precisely what such a model amounts to. Is it possible to send a computer into a black hole so it performs infinitely many steps of computation? Can we compute in parallel universes? Is it possible to perform a computation which whose inner workings are hidden from the rest of the universe? Is the space-time made of tiny cellular automata? Since realizability can accommodate all kinds of computational models, we need not make a particular commitment about the true nature of the real world. Instead, we leave open all possibilities and ask what it would take to realize various logical and mathematical principles, such as the Axiom of choice and Brouwer's Continuity principle. We will be naturally lead to questions about the limits of real-world computation.

Law of excluded middle. What would it take to realize the Law of excluded middle? Some instances are realized, for example decidability of equality of natural numbers

> *"Every two natural numbers are either equal or distinct."*

This is so because we can prove the statement intuitionistically by induction. However, decidability of equality of real numbers

> *"Every two real numbers are either equal or distinct."*

is more problematic because its realizer would exceed the power of Turing machines. For if the statement were realized then we could tell whether any given Turing machine t halts by comparing zero and the computable real whose k-th digit is 1 if t has halted at step k, and 0 otherwise. In fact, decidability of equality of real numbers and the Halting problem are equivalent in terms of computational power:

> ***Theorem:*** *Decidability of reals is real-world realized if, and only if, we can build the Halting oracle for Turing machines.*

[8] A famous proof system which uses the realizability interpretation for extraction of programs from proofs is Coq. [7]

Proposals have been given on just how we might solve the Halting problem by sending a machine into a black hole in such a way that it would perform infinitely many computational steps in time that appears finite to us. Until such machines appear in computer stores, it seems safer not to rely on them. In any case, even if there are super-machines that perform amazing computational feats, some instances of the Law of excluded middle will still fail to be realized. Indeed, the (formalization of) the statement

> *"Every real-world machine halts or runs forever."*

fails to be realized, for its realizer would be a Halting oracle for real-world machines, which does not exist by the usual argument.[9]

Axiom of Choice. The Axiom of choice is perhaps the most controversial of the standard axioms of set theory. Some mathematicians accept it grudgingly only because it is necessary for certain desirable theorems in algebra and analysis. There are several formulations of the Axiom of choice, but the one whose realizability interpretation is most easily analyzed is

> *"Every total relation contains a function."*

In this generality the Axiom of choice is not realized because a theorem of Diaconescu's tells us that the Axiom of choice implies the Law of excluded middle. Since the latter is not realized neither is the former.

Nevertheless, some instances of the Axiom of choice are realized, namely those for which the domain of the total relation has *canonical realizers*. In general, an element may have many realizers, for example, a function has as many realizers as there are programs for computing it. We say that a set A has canonical realizers if there is a realizer r which computes canonical realizers: if s_1 and s_2 both realize $x \in A$ then so does $r \cdot s_1$ and moreover $r \cdot s_1 = r \cdot s_2$.

The natural numbers have canonical realizers. Recall that the only realizer for a natural number n is the corresponding numeral $\overline{n}$. Thus we are never even in a position to contemplate two different realizers for n. Consequently, the axiom of *Number Choice* is realized, which is lucky as large parts of constructive analysis depend on it.

A more complicated instance of the Axiom of Choice is *Function Choice*. For it to be realized we need to compute canonical realizers for number-theoretic functions $\mathbb{N} \to \mathbb{N}$. The realizers of a function $h : \mathbb{N} \to \mathbb{N}$ are programs that compute the function. How could we select one of them, in a realized way? In Kleene's function realizability this is accomplished by realizing h with an (infinite tape containing) the sequence $h(0), h(1), h(2), \ldots$ We could try doing the same in the real world, as follows. We build a machine which accepts a realizer r for h, and constructs a tape containing the values $h(0), h(1), h(2), \ldots$. We can then use the tape to get values of h by simply looking them up. The tape itself depends only on h and not on the particular realizer r from which it was generated. However, since the tape is infinite its construction is an unfinished process, so in principle we could always observe the machine building it,

[9] The usual argument assumes that there is a universal machine, which may or may not exist in the real-world. But even if there isn't one, there will be still other instances of the Law of excluded middle that fail to be realized. A realizability model is always properly intuitionistic, as long as there are at least two realizers.

and within it the original realizer r. What we actually need is to place r inside a *black box*, i.e., isolate it in such a way that its internal workings cannot be observed, and the only way to interact with it is to feed it inputs and observe outputs. To summarize:

> ***Theorem:*** Function choice *is real-world realized if there are black boxes in which computations can be hidden.*

In a science fiction story a black box might be an impenetrable "force field" or a "quantum cage". Should physicists declare that there are no such things, a bureaucratic approach might help: if we change the definition of real-world realizers so that any realizer labeled as "black box" never gets inspected, then black boxes are easily obtained by means of a sticker and a pen.

Continuity principles. In Section 2 we considered an intuitionistic system in which all real functions were smooth. Realizability does not validate such a setting because we can realize the absolute value map $x \mapsto |x|$, which is not smooth. In fact, we can realize some fairly complicated functions, such as Weierstrass's continuous but nowhere differentiable functions. However, realizing the jump function

$$\begin{aligned} j(x) &= 0 \quad \text{if } x \leq 0, \\ j(x) &= 1 \quad \text{if } x > 0, \end{aligned}$$

amounts to decidability of equality on $\mathbb{R}$. Indeed, by computing $j(|x - y|)$ with precision $1/3$ suffices to tell whether its value is 0 or 1, and consequently whether $x = y$ or not. The same sort of reasoning can be applied to any function with a sudden jump, and so we wonder whether the continuity principle

> *"All functions are continuous."*

is real-world realized. Let us consider just the specific instance

> *"Every real function is continuous."*

This is essentially realized by a machine C which accepts as input a machine F for computing a real function $f : \mathbb{R} \to \mathbb{R}$, a realizer t for a real number $x \in \mathbb{R}$, and a numeral $\overline{n}$. The machine C is supposed to figure out how many digits of x suffice to compute the first n digit of $f(x)$. One way to accomplish this as follows. The machine C feeds into F a specially crafted realizer t' for the real number x which acts exactly like t, except that it communicates back to C what digits have been accessed by F. Now C waits for F to output the first n digits of $f(x)$ and records the maximum value m reported by t'. Clearly, m digits of the input suffice to compute the first n digits of $f(x)$, and therefore C may safely output $\overline{m}$.

There is a potential problem with our construction. The realizer F could analyze the inner workings of t' and discover that t' is secretly communicating with the outside world, namely with C. In such a case F could simply cut the communication between t' and C, or modify it in some way. We may dispense with such problems by assuming that t' can be hidden in a black box *with a private communication channel*, i.e., t'

can be constructed in such a way that F cannot analyze it, and neither can it detect or disrupt its communication with C.

There is a second problem with C. If by a sleight of physics F manages to inspect infinitely many digits of its input in finite time, the maximum value reported to C will be infinity and f need not actually be continuous. So we need one more assumption to have C function properly:

> ***Theorem:*** *If black boxes and private communication channels exist, and only finitely many computation steps can be performed in finite time, then the Continuity principle is real-world realized.*

The Church-Turing thesis. An obvious question to ask is how real-world machines compare to Turing machines. One has the feeling that every Turing machine can be built in the real world, at least if we ignore certain questions about finiteness of space-time. That the converse holds is expressed by

> ***Real-world Church-Turing thesis:*** *a function is realizable in the real-world if, and only if, it is computed by a Turing machine.*

The principle states that we cannot exceed the computational power of Turing machines by using black holes, quantum mechanics, and other engineering tricks. Since I am not an engineer I do not wish to pass a judgment about the issue, but I can still relate the Church-Turing thesis with realizability of mathematical statements.

Kleene constructed an unbounded Turing computable binary tree which has no Turing computable infinite paths [4]. To understand how strange such a tree is, contemplate the following situation. A master machine has built a system of underground tunnels, beginning with a hole at the surface. The tunnels always go down, they may split into two, or lead to dead ends. We are told that there is no limit to the depth of the tunnels. Yet, even if the master machine is known to us we cannot build a machine that would enter the hole and always go down without ever hitting a dead end.

Students of mathematics learn certain basic geometric facts, such as:

> ***Theorem:*** *A continuous real function defined on a closed interval is bounded.*

and

> ***Theorem:*** *A closed interval is compact.*

These theorems take some getting used to, but are indispensable tools in mathematical analysis. They express basic geometric intuitions about the continuum. You may wonder what analysis and the Kleene tree have to do with the real-world Church-Turing thesis. They are all related! For if the real-world Church-Turing thesis holds, then the following statement is realized by the Kleene tree:

> *"There is an unbounded binary tree without an infinite path."*

From such a tree we can construct counterexamples to real-world realizability of the above two theorems [12], namely a real-world realized continuous unbounded map defined on a closed interval, and a real-world realized infinite cover of a closed interval such that every finite subfamily fails to be a cover. Unfortunately, reviewing the constructions here would exceed the scope of the contribution.

Should we reject the real-world Church-Turing thesis, or do we give up basic geometric intuitions about space? There is a profound move that lets us eat the cake and have it too. In the usual conception of geometry a space is a set of points equipped with extra structure, such as metric or topology. But we can switch to a different view in which the extra structure is primary and points are derived ideal objects. For example, a topological space is not viewed as a set of points with a topology anymore, but rather just the topology, given as an abstract lattice with suitable properties, known as a *locale*.[10] In constructive mathematics such treatment of the notion of space is much preferred to the usual one. The desirable geometric intuitions are restored [11, 14], even under the assumption that everything in sight is Turing computable. Locale theory has certain advantages over traditional topology even in classical mathematics, where it allows for an appealing construction of the space of random sequences and resolves the Banach-Tarski paradox [10]. These are strong indicators that the last word on the nature of computation and geometry has not been said yet.

References

[1] J. L. Bell. *A primer of infinitesimal analysis*. Cambridge University Press, 1998.

[2] P. T. Johnstone. *Stone spaces*. Cambridge University Press, 1982.

[3] S. C. Kleene. On the interpretation of intuitionistic number theory. *Journal of Symbolic Logic*, 10:109–124, 1945.

[4] S. C. Kleene. Recursive functions and intuitionistic mathematics. In L.M. Graves, E. Hille, P.A. Smith, and O. Zariski, editors, *Proceedings of the International Congress of Mathemaiticans, August 1950. Cambridge, Mass.*, pages 679–685, 1952.

[5] S. C. Kleene and R. E. Vesley. *The Foundations of Intuitionistic Mathematics, especially in relation to recursive functions*. North-Holland Publishing Company, 1965.

[6] A. Kock. *Synthetic differential geometry*. Cambridge University Press, 1981.

[7] The Coq development team. *The Coq proof assistant reference manual*. LogiCal Project, 2004. Version 8.0.

[8] I. Moerdijk and G. E. Reyes. *Models for smooth infinitesimal analysis*. Springer-Verlag, New York, 1991.

[9] A. Robinson. *Non-standard analysis*. North-Holland Publishing Co., 1966.

[10] The suitable properties ask for a complete lattice in which finite infima distribute over arbitrary suprema [2].

[10] A. Simpson. Measure, randomness and sublocales. *Annals of Pure and Applied Logic*, 2012.

[11] P. Taylor. Tychonovs theorem in Abstract Stone Duality. In *Domains VII*, 2004. `http://www.paultaylor.eu/ASD/tyctas/`.

[12] A.S. Troelstra and D. van Dalen. *Constructivism in Mathematics, An Introduction, Vol. 1*. Number 121 in Studies in Logic and the Foundations of Mathematics. North-Holland, 1988.

[13] J. van Oosten. *Realizability: An Introduction to its Categorical Side*, volume 152 of *Studies in Logic and the Foundations of Mathematics*. Elsevier, 2008.

[14] S. Vickers. Some constructive roads to Tychonoff. In L. Crosilla and P. Schuster, editors, *From Sets and Types to Topology and Analysis: Practicable Foundations for Constructive Mathematics*, number 48 in Oxford Logic Guides, pages 223–238. Oxford University Press, 2005.