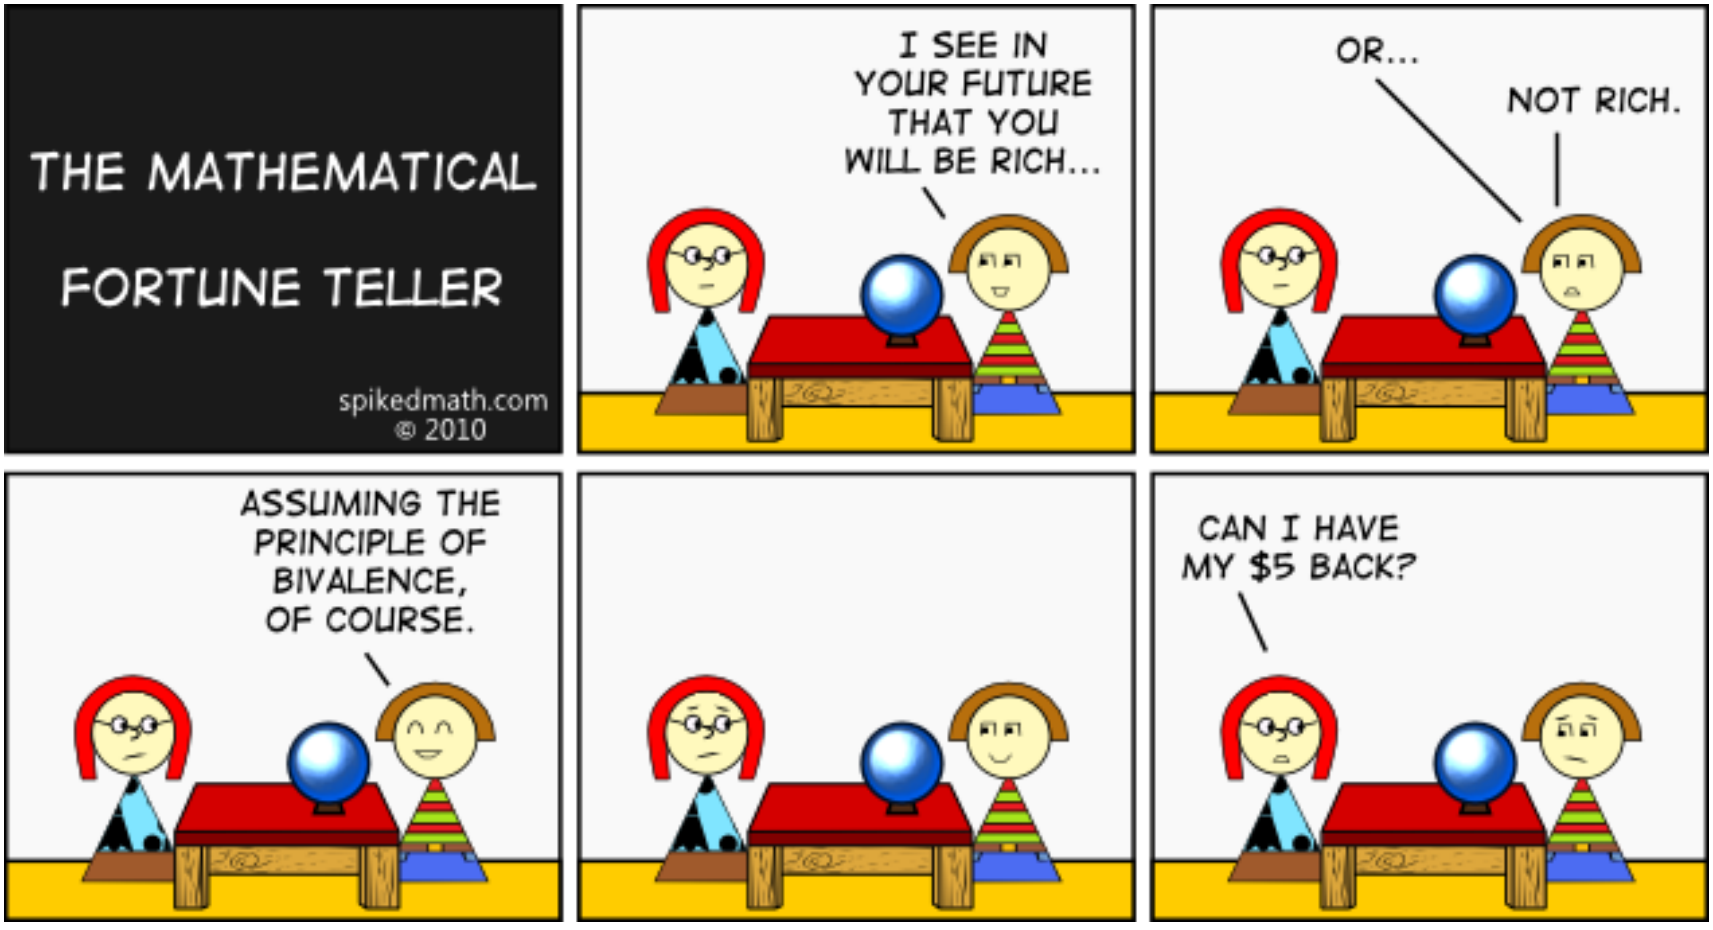
THE MATHEMATICAL
FORTUNE TELLER
spikedmath.com
® 2010
I SEE IN YOUR FUTURE THAT YOU WILL BE RICH...
OR...
NOT RICH.
ASSUMING THE PRINCIPLE OF BIVALENCE, OF COURSE.
CAN I HAVE MY $5 BACK?

Double-negation translation and CPS transformation

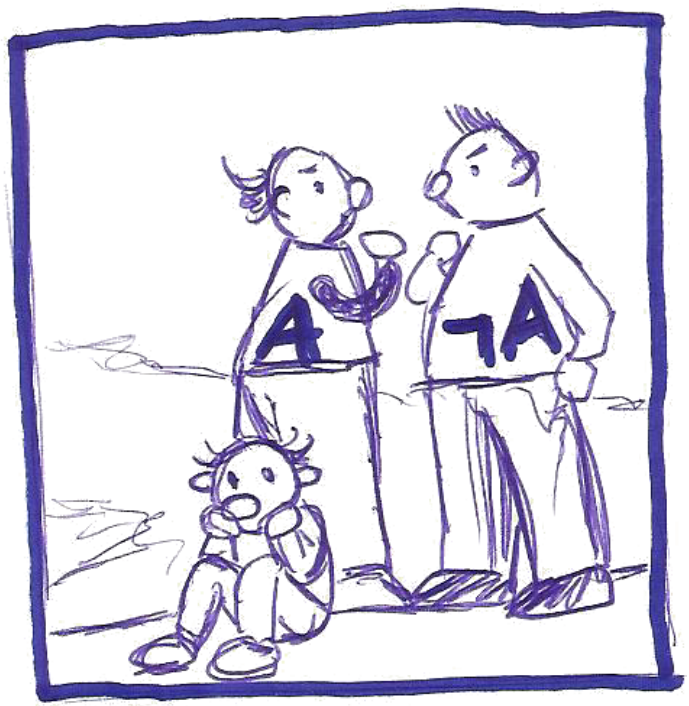

Ingo Blechschmidt

June 3rd, 2015 at KU Leuven

Abstract

Constructive mathematicians don't use the law of excluded middle, which approximately says that for any proposition P, either P is true or $\neg P$ is true. Several advantages emerge from this rejection, for instance one can mechanically extract algorithms from constructive proofs of existence statements and rigorously work with non-standard *dream axioms* which are plainly false in classical mathematics, such as *any function is smooth*.

For communicating with classicial mathematicians, constructive mathematicians can employ the *double-negation translation*. This device associates to any formula a translated formula in such a way that a given formula holds classically if and only if its translation holds constructively.

The talk will give an introduction to these topics and discuss the intriguing relationship of the double-negation translation with the well-known continuation-passing style transformation: In some sense, they are the same. This is a beautiful facet of *computational trinitarianism*.

For the first part of the talk, no background in formal logic or constructive mathematics is required. For the second part of the talk, one should be vaguely familiar with the continuation-passing style transformation.

Teaser

By the Curry–Howard correspondence intuitionistic proofs have computational content. By the double-negation translation, we see that *classical proofs too have computational content.*

As long as we stay in the continuation monad, the required bluffing and cheating will not be apparent.

Care must be taken when leaving the continuation monad (for instance by supplying the identity continuation), since then we might obtain *incorrect results.*

Outline

1. Constructive mathematics
   - The law of excluded middle
   - Interpretation of intuitionistic logic
   - Applications
2. The double-negation translation
   - The doubly-negated law of excluded middle
   - The fundamental result
   - Game-theoretical interpretation
3. Continuations
   - The Curry–Howard correspondence
   - Computational content of classical proofs
4. Outlook

Non-constructive proofs

Theorem. There exist **irrational** numbers x, y such that x^y is rational.

Proof. Either $\sqrt{2}^{\sqrt{2}}$ is rational or not.

In the first case we are done.

In the second case take $x := \sqrt{2}^{\sqrt{2}}$ and $y := \sqrt{2}$.
Then $x^y = 2$ is rational.

The proof is nice and short. However, after having seen the proof, we are still not able to give an example of irrational numbers x, y such that x^y is rational! The proof was *non-constructive*. If we want to extract explicit witnesses from the proof, the proof has to be constructive, such as this one:

> *Set* $x := \sqrt{2}$ *and* $y := \log_{\sqrt{2}} 3$. *Then* $x^y = 3$ *is rational. The proof that y is irrational is even easier than the proof that* $\sqrt{2}$ *is irrational.*

It turns out that from all the axioms of classical logic, exactly one is responsible for non-constructivity: the law of excluded middle.

The law of excluded middle

"For any formula A, we may deduce $A \vee \neg A$."

Classical logic =
intuitionistic logic + law of excluded middle.

Classical interpretation	
$\bot$	There is a contradiction.
$A \wedge B$	A and B are true.
$A \vee B$	A is true or B is true.
$A \Rightarrow B$	If A holds, then also B.
$\forall x{:}X.\, A(x)$	For all $x : X$ it holds that $A(x)$.
$\exists x{:}X.\, A(x)$	There is an $x : X$ such that $A(x)$.

The law of excluded middle

"For any formula A, we may deduce $A \vee \neg A$."

Classical logic =
intuitionistic logic + law of excluded middle.

Constructive interpretation	
$\bot$	There is a contradiction.
$A \wedge B$	We have evidence for A and for B.
$A \vee B$	We have evidence for A or for B.
$A \Rightarrow B$	We can transform evidence for A into one for B.
$\forall x{:}X.\, A(x)$	Given $x : X$, we can construct evidence for $A(x)$.
$\exists x{:}X.\, A(x)$	We have an $x : X$ together with evidence for $A(x)$.

More precisely, one should say: Classical mathematics = intuitionistic logic + law of excluded middle + a set theory including the axiom of choice.

The constructive interpretation of the axiom of excluded middle is: For any formula A, we have evidence for A or for $\neg A$. This is an absurd statement – recall that there exist undecidable statements ("Gödel sentences").

By "we" in the definition of the constructive intepretation, we don't literally refer to some group of people. It should be read in a generic mathematical way. Look up *Brouwer–Heyting–Kolmogorov interpretation* and *Realizability Theory* (see for instance Andrej Bauer's notes) for a formal treatment.

Several years ago a video showing Kate Moss consuming drugs surfaced. From the video it was clear that the drugs were either of some type A or of some type B, but there was no direct evidence for either type. Kate Moss was not prosecuted; in this sense, Great Britain's judicial system operated intuitionistically. Check out Dan Piponi's blog post about this topic.

Note that constructive mathematicians do *not* claim that the law of excluded middle is false (that is, that its negation holds). In fact, some instances of the law of excluded middle are true intuitionistically: For example one can show by induction that any natural number is zero or is not zero. Constructive mathematicians simply don't suppose that the law of excluded holds generally.

Negated statements

"$\neg A$" is syntactic sugar for $(A \Rightarrow \bot)$
and means: There can't be any evidence for A.

Constructive interpretation	
$\bot$	There is a contradiction.
$A \wedge B$	We have evidence for A and for B.
$A \vee B$	We have evidence for A or for B.
$A \Rightarrow B$	We can transform evidence for A into one for B.
$\forall x{:}X.\, A(x)$	Given $x : X$, we can construct evidence for $A(x)$.
$\exists x{:}X.\, A(x)$	We have an $x : X$ together with evidence for $A(x)$.

Note that the word "contradiction" is not generally forbidden in intuitionistic logic. For instance, the usual proof that $\sqrt{2}$ is not rational, deducing $\bot$ from the assumption that $\sqrt{2}$ were rational, is perfectly fine intuitionistically.

Colloquially, those proofs are called "proof by contradiction", but this labeling is deceptive. A true proof by contradiction runs like this:

> *We want to show A. Assume* $\neg A$*. Then . . ., so* $\bot$*. Therefore* $\neg\neg A$*. Thus A.*

The last step needs the axiom of double negation elimination, $\neg\neg A \Rightarrow A$, which is not available in intuitionistic logic. (In fact, the statement that double negation elimination holds for all A is equivalent to the statement that the law of excluded middle holds for all A.)

Doubly-negated statements

"$\neg\neg A$" means: There can't be any evidence for $\neg A$.

Trivially, we have $A \Longrightarrow \neg\neg A$.
We can't deduce $\neg\neg A \Longrightarrow A$.

Constructive interpretation

$\bot$	There is a contradiction.
$A \wedge B$	We have evidence for A and for B.
$A \vee B$	We have evidence for A or for B.
$A \Rightarrow B$	We can transform evidence for A into one for B.
$\forall x{:}X.\, A(x)$	Given $x : X$, we can construct evidence for $A(x)$.
$\exists x{:}X.\, A(x)$	We have an $x : X$ together with evidence for $A(x)$.

Doubly-negated statements

"$\neg\neg A$" means: There can't be any evidence for $\neg A$.

Trivially, we have $A \implies \neg\neg A$.
We can't deduce $\neg\neg A \implies A$.

Where is the key?

$\neg\neg(\exists x.$ the key is at position $x)$

versus

$\exists x.$ the key is at position x

If we know that the key to our apartment has to be somewhere in the apartment (since we used it to enter last night) but we can't find it, we can constructively only justify

$$\neg\neg(\exists x.\ \text{the key is at position } x),$$

not the stronger statement

$$\exists x.\ \text{the key is at position } x.$$

(Of course, this example does not quite work, since "we" now really has to refer to us key-seekers.)

Applications

Intuitionistic logic …

- can guide to more elegant proofs,
- is good for the mental hygiene, and
- allows to make finer distictions.

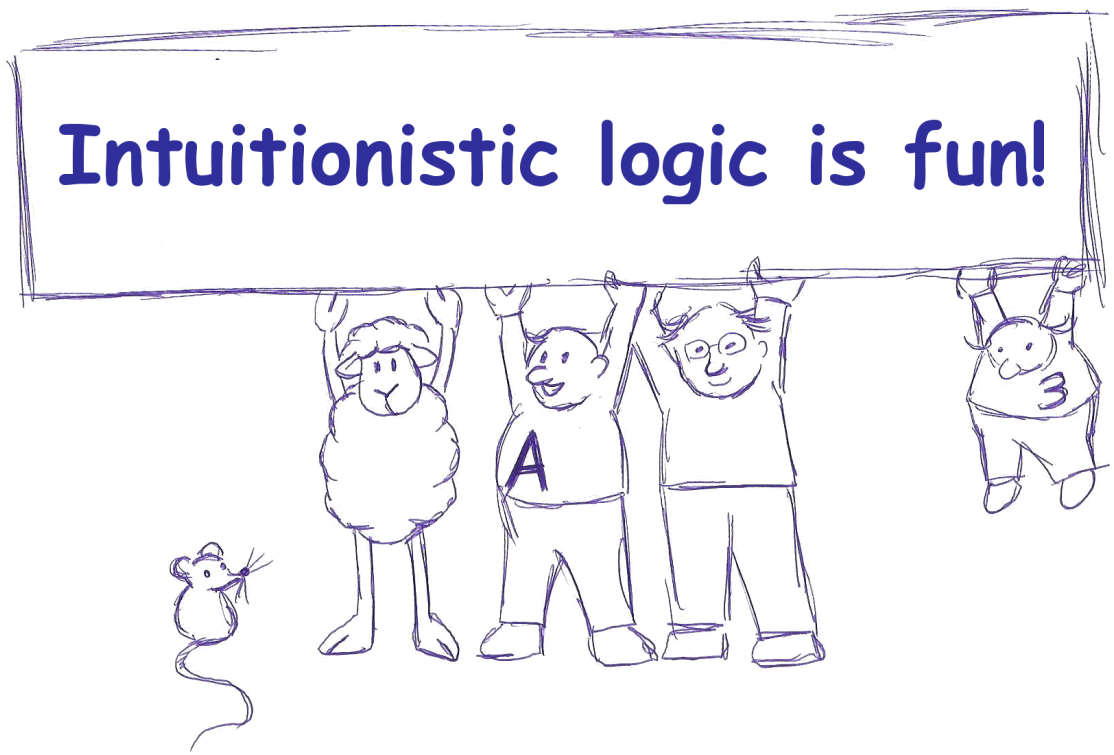

Applications

- We can **mechanically extract algorithms** from intuitionistic proofs of existence statements.
- The **internal language of toposes** is intuitionistic.
- **Dream mathematics** only works intuitionistically.

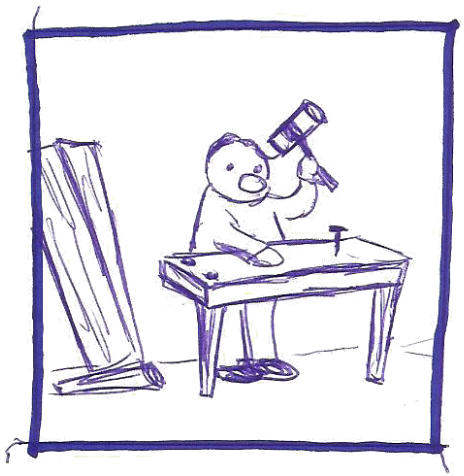

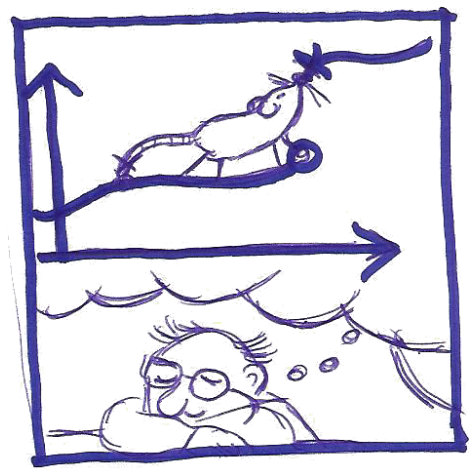

Regarding intuitionistic logic as an elegance assisting device: In classical logic, we might have the habit of beginning proofs like this:

> *If the set X is empty, the claim holds trivially. So suppose that X is not empty; then consider ...*

Often it's the case that the second part of the proof applies just as fine to the trivial case – if we don't fear the empty set. In these cases, the proof can be streamlined by skipping the case distiction.

Intuitionistically, the case distinction is not possible without further hypotheses on X. Therefore, by trying to make a proof intuitionistically acceptable, we are nudged to remove unnecessary case distictions and related issues.

Here is a concrete example. Let $f\colon X \to Y$ be a map and let

$$\begin{aligned} \varphi : \quad \mathcal{P}(Y) &\longrightarrow \mathcal{P}(X), \\ U &\longmapsto f^{-1}[U] := \{x \in X \mid f(x) \in U\} \end{aligned}$$

be the inverse image operation. Then it's a standard exercise to show that f is surjective if and only if φ is injective.

Can you do it, especially the "$\Rightarrow$" direction? There is a very short and elegant proof for it! But the proof which might first come to your mind is non-constructive and unnecessarily cumbersome:

> *Assume that f fails to be surjective. Then there exists $y \in Y$ such that y is not an element of the image of f. Therefore $\varphi(\{y\}) = \emptyset = \varphi(\emptyset)$. This is a contradiction to the injectivity of φ.*

Spoiler. We have $\varphi(\operatorname{im} f) = X = \varphi(Y)$, thus $\operatorname{im} f = Y$.

Here is a basic example for extracting algorithms from proofs. Consider the statement

> *"There are infinitely many prime numbers."* or somewhat more explicitly, *"For any finite list $p_1, \ldots, p_n$ of prime numbers, there exists an additional prime number q not on that list."*

The standard proof, attributed to Euclid, goes like this:

> *Consider the number $N := p_1 \cdots p_n + 1$. Since $N \geq 2$, there exists some prime factor q of N. (If N is itself prime, we can take $q := N$.) This prime is not equal to any p_i, since the numbers p_i don't divide N whereas q does.*

The algorithm for constructing q can be directly read off from the proof. Different proofs result in different algorithms; in particular, there exist (more complex) proofs whose algorithms produce better (smaller) witnesses.

See the wonderful book *Applied Proof Theory: Proof Interpretations and their Use in Mathematics* by Kohlenbach for details. Already its introduction is a very worthwhile reading.

Tangentially, observe that the stated constructive version of Euclid's proof is less prone to misunderstandings than its well-known counterpart which uses proof by contradiction:

> *Assume that there are only a finite number of primes, $p_1, \ldots, p_n$. Then consider $N := p_1 \cdots p_n + 1$. This number is either prime or composite. Since no prime number divides N (by assumption the only primes are the p_i and these don't divide N), it cannot be composite. Therefore N is prime. Since N doesn't equal any of the p_i, this is a contradiction.*

From this proof one might think that for primes $p_1, \ldots, p_n$ the number $N := p_1 \cdots p_n + 1$ is always prime. But this only holds in a counterfactual world where there are only finitely many primes. In fact, the number

$$N := 2 \cdot 3 \cdot 5 \cdot 7 \cdot 11 \cdot 13 + 1 = 59 \cdot 509$$

is composite. A smaller example is

$$N := 2 \cdot 7 + 1 = 3 \cdot 5.$$

Topos power

Any finitely generated vector space does *not not* possess a basis.

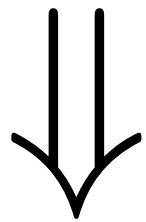

Any sheaf of modules of finite type on a reduced scheme is locally free on a dense open subset.

Toposes are certain kinds of categories, thought of as *mathematical universes*. The usual topos in which we do mathematics in is the category of sets and maps between sets, but there are many others:

- In the *effective topos*, any map is computable.
- In the *sheaf topos* of a topological space X the objects and morphisms depend on our position in X.

A metatheorem states that *intuitionistically provable statements hold in any topos*. This greatly expands the scope of an intuitionistic theorem.

A side project of mine is to recognize the basic concepts and statements of algebraic geometry as topos-theoretic interpretations of simple concepts and statements of ordinary first-year linear algebra. See `https://github.com/iblech/internal-methods` for expository notes on this topic (directed at geometers). The example of the slide is taken from these notes. The simple statement about vector spaces *automatically implies* the more complicated statement about sheaves.

Dream mathematics

Synthetic differential geometry

Any map $\mathbb{R} \to \mathbb{R}$ is smooth. There are infinitesimal numbers ε such that $\varepsilon^2 = 0$ and $\varepsilon \neq 0$.

Synthetic domain theory

For any set X there exists a map

$$\mathsf{fix} : (X \to X) \to X$$

such that $f(\mathsf{fix}(f)) = \mathsf{fix}(f)$ for any $f : X \to X$.

Synthetic computability theory

There are only countably many subsets of $\mathbb{N}$.

Dream mathematics is working with dream axioms – axioms which are classically false, but very convenient:

- If you adopt synthetic differential geometry, you can do calculus like 300 hundred years ago, by manipulating infinitesimals. See a blog post by Andrej B
- If you adopt synthetic domain theory, you can use the ordinary mathematical notions of sets and maps to give semantics to programming languages. See slides of a talk by Alex Simpson and a paper by Hyland.
- If you adopt synthetic computability theory, you can use the ordinary notions of sets and maps to secretly talk about enumerable sets and computable maps. You can drop any of the usual adjectives like "effectively" or "enumerable". See a paper by Andrej Bauer.

All of these dream axioms can be made to work: By dropping the law of excluded middle. More precisely, there are alternate toposes in which the law of excluded middle does not hold but the given dream axiom does. Also there are metatheorems which guarantee that results obtained in the dream universe also hold in the the usual universe.

Warning. For space reasons, the axioms are not presented faithfully on the previous slide. Check the references for precise formulations.

The doubly-negated LEM

Even intuitionistically "$\neg\neg(A \vee \neg A)$" holds.

Proof. Assume $\neg(A \vee \neg A)$, we want to show $\bot$.

If A, then $A \vee \neg A$, thus $\bot$.

Therefore $\neg A$.

Since $\neg A$, we have $A \vee \neg A$, thus $\bot$.

The $\neg\neg$-translation

$$
\begin{aligned}
A^{\square} &:\equiv \neg\neg A \text{ for atomic formulas } A \\
(A \wedge B)^{\square} &:\equiv \neg\neg(A^{\square} \wedge B^{\square}) \\
(A \vee B)^{\square} &:\equiv \neg\neg(A^{\square} \vee B^{\square}) \\
(A \Rightarrow B)^{\square} &:\equiv \neg\neg(A^{\square} \Rightarrow B^{\square}) \\
(\forall x{:}X.\, A(x))^{\square} &:\equiv \neg\neg(\forall x{:}X.\, A^{\square}(x)) \\
(\exists x{:}X.\, A(x))^{\square} &:\equiv \neg\neg(\exists x{:}X.\, A^{\square}(x))
\end{aligned}
$$

Theorem. A classically $\iff$ $A^{\square}$ intuitionistically.

The gray $\neg\neg$'s can be omitted: One can prove by structural induction that translating with those double negations yields logically equivalent formulas as translating without those.

The blue $\neg\neg$'s in contrast are crucial.

One could say that the only difference between intuitionistic logic and classical logic is in the meaning of disjunction and existential quantification.

If A does not contain "$\Rightarrow$" and "$\forall$", then $A^{\square} \Leftrightarrow \neg\neg A$. Such formulas are also called *geometric*, because their truth value is preserved by so-called *geometric functors* in topos theory. In general, there is no implicational relation between $A^{\square}$ and $\neg\neg A$.

If you can read German, then see these pizza seminar notes for a detailed treatment. Else see *Constructivism in Mathematics: An Introduction* by Troelstra and van Dalen. The proof is routine; for each of the logical rules of classical logic you show that its double-negation translation is valid intuitionistically.

Note that $\bot^{\Box} \equiv \neg\neg\bot \Leftrightarrow \bot$. Also recall that Heyting arithmetic is the same as Peano arithmetic, only with intuitionistic instead of classical logic.

Corollary. Peano arithmetic and Heyting arithmetic are equiconsistent.

Proof. It is clear that inconsistency of Heyting arithmetic implies inconsistency of Peano arithmetic.

For the converse direction, write Ax for the axioms of Peano arithmetic, thought of as a single formula by conjunction. If Peano arithmetic proves $\bot$, that is if $\mathrm{Ax} \Rightarrow \bot$ classically, then by the theorem $\mathrm{Ax}^{\Box} \Rightarrow \bot^{\Box}$ intuitionistically. By inspection $\mathrm{Ax} \Rightarrow \mathrm{Ax}^{\Box}$ intuitionistically. Therefore $\mathrm{Ax} \Rightarrow \bot$ intuitionistically.

A classical logic fairy tale

Narrator. Once upon a time, in a kingdom far, far away, the queen of the land and of all Möbius strips called for her royal philosopher.

Queen. Philosopher! I ask you to carry out the following order. Get me the Philosopher's Stone, or alternatively find out how one could produce arbitrary amounts of gold with it!

Philosopher. But my queen! I haven't studied anything useful! How could I fulfill this order?

Queen. That is not my concern. I'll see you again tomorrow. Should you not accomplish the task, I will take your head off.

Narrator. After a long and wakeful night the philosopher was called to the queen again.

Queen. Tell me! What do you have to report?

Philosopher. It was not easy and I needed to consult lots of books, but finally I actually found out how to use the Philosopher's Stone to produce arbitrary amounts of gold. But only I can conduct this procedure, your royal highness.

Queen. Alright. So be it.

Narrator. And so years passed by, during which the philosopher imagined himself to be safe. The queen searched for the stone on her own, but as long as she hadn't found it, the philosopher didn't need to worry. Yet one day the impossible happened: The queen has found the stone! And prompty called for her philosopher.

Queen. Philosopher, look! I have found the Philosopher's Stone! Now live up to your promise! *[She hands over the stone.]*

Philosopher. Thank you. *[He inspects the stone.]* This is indeed the Philosopher's Stone. Many years ago you asked me to either acquire the Philosopher's Stone or find out how to produce arbitrary amounts of gold using it. Now it's my pleasure to present to you the Philosopher's Stone. *[He returns the stone.]*

A classical logic fairy tale

A classical logic fairy tale

A intuitionistically $\iff$ we can defend A in any dialog.

A classically $\iff$ we can defend $A^{\square}$ in any dialog.

A classical logic fairy tale

A intuitionistically $\iff$ we can defend A in any dialog.

A classically $\iff$ we can defend $A^{\square}$ in any dialog.

$\iff$ we can defend A in any dialog
with jumps back in time allowed.

Recall the usual dialog metaphor for proofs:

> ***Alice***. *I claim that* $\forall x : X.\, A(x) \Rightarrow B(x)$.
> ***Eve***. *Really? Take this particular* $x_0 : X$*!*
> ***Alice***. *Sure. Then I claim that* $A(x_0) \Rightarrow B(x_0)$.
> ***Eve***. *My* x_0 *satisfies* $A(x_0)$*, see? …*
> ***Alice***. *Right. And from this I'm able to show* $B(x_0)$*: …*

This metaphor can be formalized. Then it's a theorem that a statement A holds intuitionistically if and only if the proponent has a winning strategy for dialogs about A. See the survey article by Helger Rückert in *Essays on Non-Classical Logic* for details.

Reading "A" for "Philosopher's Stone" and "$A \Rightarrow \bot$" for "using the Philosopher's Stone to produce arbitrary amounts of gold", the classical logic fairy tale gives a proof of the law of excluded middle – in dialog form, with time jumps. See this blog post by Edward Yang for a different rendition of the tale. It was popularized by Philip Wadler in his CbV is dual to CbN-paper and might be due to Peter Selinger.

It is not a coincidence that the tale feels "continuation-y".

Curry–Howard correspondence

logic	programming
formula A	type A
intuitionistic proof $p : A$	term $p : A$
conjunction $A \wedge B$	product type (A, B)
disjunction $A \vee B$	sum type Either A B
implication $A \Rightarrow B$	function type $A \to B$

Curry–Howard correspondence

logic	programming
formula A	type A
intuitionistic proof $p : A$	term $p : A$
conjunction $A \wedge B$	product type (A, B)
disjunction $A \vee B$	sum type Either A B
implication $A \Rightarrow B$	function type $A \to B$
$\neg\neg$-translation	**CPS transformation**

Curry–Howard correspondence

logic	programming
formula A	type A
intuitionistic proof $p : A$	term $p : A$
conjunction $A \wedge B$	product type (A, B)
disjunction $A \vee B$	sum type Either A B
implication $A \Rightarrow B$	function type $A \to B$
$\neg\neg$-translation	**CPS transformation**
$\neg\neg A$	??

Curry–Howard correspondence

logic	programming
formula A	type A
intuitionistic proof $p : A$	term $p : A$
conjunction $A \wedge B$	product type (A, B)
disjunction $A \vee B$	sum type Either A B
implication $A \Rightarrow B$	function type $A \to B$
$\neg\neg$-translation	**CPS transformation**
$(A \Rightarrow \bot) \Rightarrow \bot$	??

Curry–Howard correspondence

logic	programming
formula A	type A
intuitionistic proof $p : A$	term $p : A$
conjunction $A \wedge B$	product type (A, B)
disjunction $A \vee B$	sum type Either A B
implication $A \Rightarrow B$	function type $A \to B$
$\neg\neg$-translation	**CPS transformation**
$(A \Rightarrow \bot) \Rightarrow \bot$	$(A \to r) \to r$

On the previous slide r should denote the void type (having no inhabitants). But this is a bit boring; we can also use any type for r and employ a variant of the double-negation translation: One where we don't use $\neg\neg$, but $\neg_r\neg_r$, where $\neg_r$ is defined as $\neg_r A :\equiv (A \Rightarrow r)$. For this variant, it still holds that we can transform classical proofs of A into intuitionistic proofs of $A^{\square_r}$.

To see that the double-negation translation corresponds to the CPS transformation, simply note that the type $(a \to r) \to r$ is also known under the name Cont r a.

Different but logically equivalent versions of the double-negation translation yield different variants of the CPS transformation (call by name, call by value, ...).

Computational content of classical proofs

```
type Cont r a = ((a -> r) -> r)

-- Decide an arbitrary statement a.
lem :: Cont r (Either a (a -> Cont r b))
lem k = k $ Right $ \x -> (\k' -> k (Left x))

-- Calculate the minimum of an infinite list
-- of natural numbers.
min :: [Nat] -> Cont r (Int, Int -> Cont r ())
min xs = ...
```

To decide an arbitrary statement A, we proceed as follows. When asked whether A or $\neg A$ holds, we bluff and immediately claim that $\neg A$ holds. Since $\neg A$ is defined as $(A \Rightarrow \bot)$, *our opponent has to work in order to rebut our claim.* As soon as our opponent presents evidence for A, we rewind time and make it look as if we always claimed that A holds from the start.

Similarly, to find the minimum of an infinite list $[x_0, x_1, \ldots]$ of natural numbers, we simply claim that the first number x_0 of the list is the minimum. This might actually be true. Should our opponent later present a smaller element x_i, we rewind time (take a previously-stored continuation) and claim that x_i is the minimum. This process of refining our initial guess will terminate after at most x_0 many time jumps.

The logical analog to this algorithm is the following. Let $[x_0, x_1, \ldots]$ be an infinite list of natural numbers. Then we can't intuitionistically verify that this list has a minimal element. But we can verify that there is *not not* a minimal element.

Food for thought

By the Curry–Howard correspondence intuitionistic proofs have computational content. By the double-negation translation, we see that *classical proofs too have computational content.*

As long as we stay in the continuation monad, the required bluffing and cheating will not be apparent.

Care must be taken when leaving the continuation monad (for instance by supplying the identity continuation), since then we might obtain *incorrect results.*

See the article Computational content of classical logic by Coquand for details.

Also note that the double-negation translation of the axiom of choice is not valid intuitionistically. Therefore one has to work harder to extract computational content from classical proofs which use this particular axiom.

Here is a practical example from algorithmic number theory for the use of the "computational law of excluded middle". Let x be an algebraic number, that is a complex number which is a zero of a normed polynomial $f(X) \in \mathbb{Q}[X]$. Consider the field extension $\mathbb{Q}(x)$ of $\mathbb{Q}$ generated by x; this field contains the rational numbers, the number x, and every number which can be obtained from these by addition, subtraction, multiplication, and division.

We want to describe an algorithm for computing the inverse of a nonzero element in $\mathbb{Q}(x)$ and expressing this inverse as a polynomial in x. In principle, this can be done as follows.

Factor $f(X)$ into irreducible polynomials. One of the factors, say $g(X)$, will be zero at x. This factor is called the *minimal polynomial* of x and general theory tells us that $\mathbb{Q}(x) \cong \mathbb{Q}[X]/(g(X))$; so working in $\mathbb{Q}(x)$ is the same as working in the polynomial ring $\mathbb{Q}[X]$ modulo $g(X)$. Since the extended Euclidean algorithm provides an efficient method for finding modular inverses, it looks like we are done.

However, factoring $f(X)$ into irreducible polynomials is computationally expensive. Is there a way to avoid that?

Yes! Let $[h(X)] \in \mathbb{Q}[X]/(f(X))$. We want to compute a modular inverse to $h(x)$ in $\mathbb{Q}(x)$. Note that since $f(X)$ might not be irreducible, the ring $\mathbb{Q}[X]/(f(X))$ might not be a field (any factor of $f(X)$ is a zero divisor in this ring). Nevertheless, we can use the extended Euclidean algorithm to compute the (monic) greatest common divisor $d(X)$ of $f(X)$ and $h(X)$ and a *Bézout representation*

$$d(X) = a(X)f(X) + b(X)h(X).$$

Three cases can occur:

1. $d(X) = 1$. Then the equation shows that $b(X)$ is inverse to $h(X)$ modulo $f(X)$ (and therefore, *a fortiori*, modulo the inaccessible $g(X)$). In particular, $b(x)$ is inverse to $h(x)$ in $\mathbb{Q}(x)$.

2. $d(X) = f(X)$. Then the equation shows that $h(X)$ is a multiple of $f(X)$ and therefore $h(x)$ is zero. In this case we don't need (and can't) compute an inverse.

3. Else $d(X)$ is a nontrivial factor of $f(X)$. At least one of the polynomials $d(X)$ and $f(X)/d(X)$ still has x as a zero. Call this polynomial $\tilde{f}(X)$. Then restart the calculation using $\tilde{f}(X)$ instead of $f(X)$.

In this way we can work in $\mathbb{Q}[X]/(g(X))$ without having to explicitly compute $g(X)$. "Leaving the continuation monad" is not a problem, since inverses modulo $f(X)$ or $\tilde{f}(X)$ or $\tilde{\tilde{f}}(X)$ are also inverses modulo the proper $g(X)$.

Outlook

- CPS transformation = Yoneda embedding
- What about delimited continuations?
- Geometrical interpretation:

$$\mathrm{Sh}(X) \models A^{\Box} \quad \Longleftrightarrow \quad \mathrm{Sh}(X_{\neg\neg}) \models A$$

- Generalize from $\neg\neg$ to arbitrary **modal operators** (monads): Relevant axioms are
  1. $A \Rightarrow \Box A$
  2. $\Box\Box A \Rightarrow \Box A$
  3. $\Box(A \wedge B) \Leftrightarrow \Box A \wedge \Box B$

Outlook

- CPS transformation = Yoneda embedding
- What about delimited continuations?
- Geometrical interpretation:

$$\mathrm{Sh}(X) \models A^{\square} \quad \Longleftrightarrow \quad \mathrm{Sh}(X_{\neg\neg}) \models A$$

- Generalize from $\neg\neg$ to arbitrary **modal operators** (monads): Relevant axioms are
  1. $A \Rightarrow \square A$
  2. $\square\square A \Rightarrow \square A$
  3. $\square(A \wedge B) \Leftrightarrow \square A \wedge \square B$

/iblech/talk-constructive-mathematics

Recommended reading:

- Oleg Kiselyov on the law of excluded middle.
- Chetan Murthy's PhD thesis
  Extracting Constructive Content from Classical Proofs.
- The blog of Andrej Bauer, especially the posts tagged `constructive-math`.

Not directly related to the topic of these slides, but if you haven't read yet Andrej Bauer's short book contribution *Intuitionistic mathematics and realizability in the physical world* or don't follow Dan Piponi's blog, you totally should. Actually I envy you, because you can look forward to some great reading!

If you want to know more about the geometrical interpretation, check out `https://github.com/iblech/internal-methods` for expository notes (prerequisites for the relevant sections: general topology, but not algebraic geometry).