Category theory and diagrammatic reasoning

27th February 2019
Last updated: 10th March 2019

5 The Curry-Howard-Lambek correspondence

In the previous lectures, we have considered two main examples of monoidal categories: the monoidal category $\mathbf{Set}_\times$ of sets and functions, and meet-semilattices. We have seen that, in fact, both are *cartesian* monoidal categories. There is more: $\mathbf{Set}_\times$ has right Kan extensions of every 1-cell along any 1-cell.

Definition 1. A *symmetric monoidal closed category* is a symmetric monoidal category X together with a right Kan extension $\mathrm{ev}_{a,b}$ of every 1-cell b along every 1-cell a, called an *evaluator*.

If X is cartesian, we say that it is a *cartesian closed category*.

Hence $\mathbf{Set}_\times$ is cartesian closed with evaluation functions as evaluators; so is every meet-semilattice P *with implications*. Now, consider P as an model of propositional logic: elements of P model propositions $A, B, \ldots$ and the partial order models an *entailment* relation, written $A \vdash B$ and read "B follows from A". The same structure of a cartesian closed category is expressed in quite different ways:

cartesian closed category	$\mathbf{Set}_\times$	P
1-cells	sets	propositions
2-cells	functions	entailment
1-cell composition	cartesian product $S \times T$	conjunction $A \wedge B$
identity 1-cell	one-element set $\{*\}$	true proposition $\top$
right Kan extension	function set T^S	implication $A \to B$
braiding	$(x, y) \mapsto (y, x)$	$A \wedge B \vdash B \wedge A$
copy	$x \mapsto (x, x)$	$A \vdash A \wedge A$
discard	$x \mapsto *$	$A \vdash \top$

This is a classical instance of a three-way correspondence between *categories, processes* and *logic*, called the *Curry-Howard-Lambek correspondence.* Since its discovery in the 1970s, it has illuminated aspects of each of its branches, and is considered today a fundamental concept in theoretical computer science.

To understand this correspondence, let us spell out what it means for P to be a "model of propositional logic". Logical reasoning is commonly formalised as a *sequent calculus*: this is given by proof rules of the form

$$\frac{\{\Gamma_i \vdash \Delta_i\}_{i \in I}}{\Gamma \vdash \Delta}\,\mathrm{R} \tag{1}$$

where the $\Gamma_i, \Gamma, \Delta_i, \Delta$ are sequences of formulas, and the rule states that when Δ_i follows from Γ_i for each $i \in I$, then Δ follows from Γ. The top line of a rule is called the *premise*, and the bottom line the *consequence* of a rule.

Each expression of the form $\Gamma_i \vdash \Delta_i$ is called a *sequent*. The sequent calculus is said to be *intuitionistic* if the right-hand of each sequent contains a single formula.

The rules of a sequent calculus are usually divided into *structural rules*, not related to any particular logical connective, and *logical rules*. Logical rules often come in pairs, one pair for each connective:

- in a *natural deduction-style* sequent calculus, each connective has one *introduction* rule and one *elimination* rule;
- in a *Gentzen-style* sequent calculus, each connective has two introduction rules, one for each side of a sequent.

We will only use Gentzen-style sequent calculus. We assume to have fixed a set of propositional atoms, and use $A, B, \ldots$ as variables for generic formulas.

Definition 2. The sequent calculus $\mathsf{LJ}_{\wedge\rightarrow}$ for the fragment of intuitionistic logic with connectives $(\wedge, \rightarrow)$ has a set of formulas defined as follows:

- propositional atoms are formulas;
- if A and B are formulas, then $A \wedge B$ and $A \rightarrow B$ are formulas.

It has the following structural rules:

$$\frac{}{A \vdash A}\ \text{AX} \qquad \frac{\Gamma_1, A, B, \Gamma_2 \vdash C}{\Gamma_1, B, A, \Gamma_2 \vdash C}\ \text{EXC}$$

$$\frac{\Gamma, A, A \vdash B}{\Gamma, A \vdash B}\ \text{CTR} \qquad \frac{\Gamma \vdash B}{\Gamma, A \vdash B}\ \text{WKN}$$

called *axiom*, *exchange*, *contraction*, and *weakening*, respectively, together with the following logical rules:

$$\frac{\Gamma, A, B \vdash C}{\Gamma, A \wedge B \vdash C}\ \wedge_L \qquad \frac{\Gamma_1 \vdash A \qquad \Gamma_2 \vdash B}{\Gamma_1, \Gamma_2 \vdash A \wedge B}\ \wedge_R$$

$$\frac{A, \Gamma \vdash B}{\Gamma \vdash A \rightarrow B}\ \rightarrow_R \qquad \frac{\Gamma_1 \vdash A \qquad B, \Gamma_2 \vdash C}{\Gamma_1, A \rightarrow B, \Gamma_2 \vdash C}\ \rightarrow_L$$

Example 3. The following is a proof of the sequent $(A \to B) \wedge (B \to C) \vdash A \to C$, that is, transitivity of implication, in $\mathsf{LJ}_{\wedge\to}$:

$$\dfrac{\dfrac{\dfrac{}{A \vdash A}\,\text{AX} \quad \dfrac{}{B \vdash B}\,\text{AX}}{A, A \to B \vdash B}\,{\to_L} \quad \dfrac{\dfrac{}{B \vdash B}\,\text{AX} \quad \dfrac{}{C \vdash C}\,\text{AX}}{B, B \to C \vdash C}\,{\to_L}}{\dfrac{A, A \to B, B \to C \vdash C}{\dfrac{A \to B, B \to C \vdash A \to C}{(A \to B) \wedge (B \to C) \vdash (A \to C)}\,\wedge_L}\,{\to_R}}\,{\to_L}$$

In addition to the basic rules, one can state several *admissible* rules: a rule is admissible in a sequent calculus if from any proof of its premise, one can construct a proof of its consequence.

Exercise 4. Show that the following rules are all admissible for $\mathsf{LJ}_{\wedge\to}$:

$$\dfrac{}{\Gamma, A \vdash A}\,\text{AX}' \qquad \dfrac{\Gamma \vdash A}{\sigma(\Gamma) \vdash A}\,\text{EXC}'$$

where σ is an arbitrary permutation of the formulas in Γ.

Consider the following rule:

$$\dfrac{\Gamma_1 \vdash A \qquad A, \Gamma_2 \vdash B}{\Gamma_1, \Gamma_2 \vdash B}\,\text{CUT}$$

This can be seen as a sequent-version of *modus ponens*, and is the logical rule that most closely mirrors our practical reasoning: instead of proving the "theorem" $\Gamma \vdash B$ directly, we decompose the "hypotheses" Γ into two subsets Γ_1, Γ_2, and then decompose the theorem into a "lemma" $\Gamma_1 \vdash A$ with fewer hypothesis, and another lemma $A, \Gamma_2 \vdash B$. A celebrated result by Gentzen, which he proved in greater generality, says the following when restricted to $\mathsf{LJ}_{\wedge\to}$.

Theorem 5 (Cut-elimination)**.** *The rule* CUT *is admissible in* $\mathsf{LJ}_{\wedge\to}$.

The proof proceeds by a case analysis on the last steps of the derivation of the premise of a CUT rule: at each step, the CUT rule is "moved up" in the derivation tree, until finally we reach cuts where one premise is obtained by an axiom rule, which can safely be eliminated. Cut-elimination has a number of remarkable consequences: most importantly, that provability of a sequent in $\mathsf{LJ}_{\wedge\to}$ is *decidable.*

Let us return to the semantics. A meet-semilattice with implications P is a model for $\mathsf{LJ}_{\wedge\to}$ in the following sense. Given any interpretation $[\![A]\!] := x$ of propositional atoms as elements of P, we can

1. define interpretations $[\![A \wedge B]\!] := [\![A]\!] \wedge [\![B]\!]$ and $[\![A \to B]\!] := [\![A]\!] \to [\![B]\!]$ of composite formulas,

2. so that, for each rule R of $\mathsf{LJ}_{\wedge\to}$ written as in (1), if $[\![\Gamma_i]\!] \leq [\![\Delta_i]\!]$ for all $i \in I$, then $[\![\Gamma]\!] \leq [\![\Delta]\!]$ in P.

Rephrasing in categorical language: given any interpretation $[\![A]\!]$ of propositional atoms as 1-cells, we can

1. define 1-cells $[\![A \wedge B]\!]$ and $[\![A \to B]\!]$;

2. for each rule R, given any interpretation $[\![\Gamma_i \vdash \Delta_i]\!]$ of the premises of R as 2-cells $([\![\Gamma_i]\!]) \Rightarrow ([\![\Delta_i]\!])$, we can define a 2-cell $([\![\Gamma]\!]) \Rightarrow ([\![\Delta]\!])$.

We can now see that any cartesian closed category is a model of $\mathsf{LJ}_{\wedge\to}$ in this generalised sense.

Construction 6. Let X be a cartesian closed category, $c_{a,b} : (a, b) \Rightarrow (a \times b)$ its compositors, and $\mathrm{ev}_{a,b} : (a, a \to b) \Rightarrow (b)$ its evaluators.

We fix an interpretation $[\![A]\!]$ of propositional atoms as 1-cells of X. We extend this to all formulas of $\mathsf{LJ}_{\wedge\to}$ by

$$[\![A \wedge B]\!] := [\![A]\!] \times [\![B]\!], \qquad [\![A \to B]\!] := [\![A]\!] \to [\![B]\!].$$

Next, we define the interpretation of each rule of $\mathsf{LJ}_{\wedge\to}$. To avoid being pedantic, we will omit brackets in the source and target of 2-cells when there is no ambiguity.

Structural rules.

1. *Axiom rule.* There are no premises; the consequence is interpreted as $\mathrm{id}_{[\![A]\!]} : [\![A]\!] \Rightarrow [\![A]\!]$.

2. *Exchange rule.* For any interpretation $f : [\![\Gamma_1]\!], [\![A]\!], [\![B]\!], [\![\Gamma_2]\!] \Rightarrow [\![C]\!]$ of the premise, the consequence is interpreted as

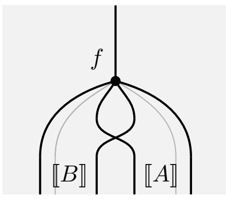

.

3. *Contraction rule.* For any interpretation $f : [\![\Gamma]\!], [\![A]\!], [\![A]\!] \Rightarrow [\![B]\!]$ of the premise, the consequence is interpreted as

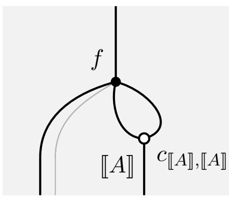

.

4. *Weakening rule.* For any interpretation $f : [\![\Gamma]\!] \Rightarrow [\![B]\!]$ of the premise, the consequence is interpreted as

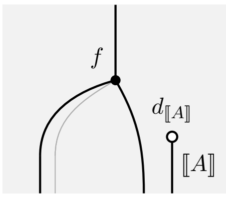

.

Logical rules.

1. *Left $\wedge$-introduction.* For any interpretation $f : [\![\Gamma]\!], [\![A]\!], [\![B]\!] \Rightarrow [\![C]\!]$ of the premise, the consequence is interpreted as

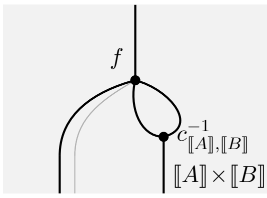

. (2)

2. *Right $\wedge$-introduction.* For any interpretation $f : [\![\Gamma_1]\!] \Rightarrow [\![A]\!]$ and $g : [\![\Gamma_2]\!] \Rightarrow [\![B]\!]$ of the premises, the consequence is interpreted as

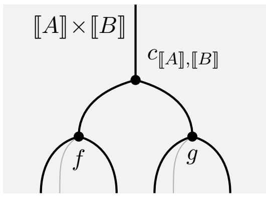

.

3. *Right $\to$-introduction.* For any interpretation $f : [\![A]\!], [\![\Gamma]\!] \Rightarrow [\![B]\!]$ of the premise, the consequence is interpreted as the unique 2-cell $\lambda_1 f : [\![\Gamma]\!] \Rightarrow [\![A]\!] \to [\![B]\!]$ satisfying

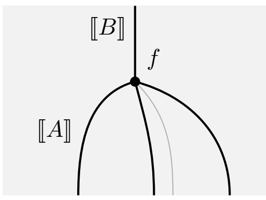

=

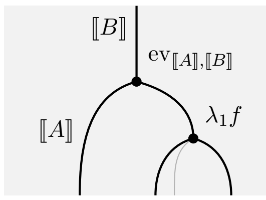

.

4. *Left $\to$-introduction.* For any interpretation $f : [\![\Gamma_1]\!] \Rightarrow [\![A]\!]$ and $g : [\![B]\!], [\![\Gamma_2]\!] \Rightarrow [\![C]\!]$ of the premises, the consequence is interpreted as

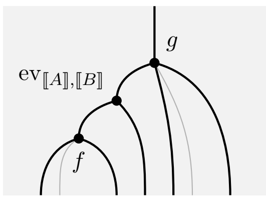

.

Remark 7. Because CUT is admissible, we expect to be able to interpret it. Indeed, for any interpretation $f : [\![\Gamma_1]\!] \Rightarrow [\![A]\!]$ and $g : [\![A]\!], [\![\Gamma_2]\!] \Rightarrow [\![B]\!]$ of its premises, the consequence is interpreted as

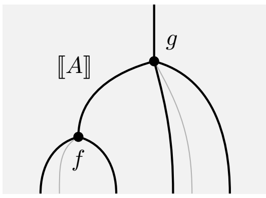

.

Remark 8. The interpretations of the right $\wedge$-introduction rule and of the left $\to$-introduction rules are analogous: we "plug" the interpretations of the two premises into the universal 2-cells $c_{a,b}$ and $\mathrm{ev}_{a,b}$, in the only possible way.

The interpretation of right $\to$-introduction is the solution of a division problem. It is possible to rephrase the interpretation of left $\wedge$-introduction in a way that it mirrors it, using the characterisation of isomorphisms as universal morphisms: (2) is, equivalently, the unique $\tilde{f}$ satisfying

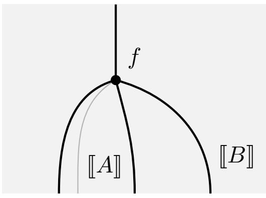

=

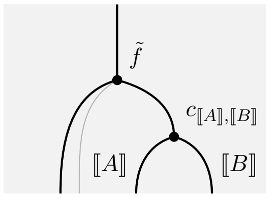

.

The other way around, we can say that each of the 2-cells $c_{a,b}$ and $\mathrm{ev}_{a,b}$ gives rise both to a "composition problem" (what can we compose along the wires labelled a and b?) and to a "division problem" by their universal property, and their solutions are the interpretation of the left and the right introduction rule associated to a connective.

The same analysis can be applied to other logical connectives, which we will not treat here.

Exercise 9. Exhibit a cartesian closed category with 1-cells a, b such that there are no 2-cells $((a \to b) \to a) \Rightarrow (a)$. Deduce that the sequent $(A \to B) \to A \vdash A$ is not provable in $\mathsf{LJ}_{\wedge\to}$.

Does A follow from $(A \to B) \to A$ in classical logic?

Through the interpretation of sequent calculus in a cartesian closed category X, we turn the process of *proving*, or *deriving a sequent*, into the process of *constructing a 2-cell* of X. This 2-cell is the interpretation of a derived sequent, but importantly, if X is rich enough, it can contain information on the particular way that the sequent has been proven.

Indeed, in the traditional algebraic models, such as a meet-semilattice P, because there is at most one 2-cell with any given source and target, any proof of a sequent will give rise to the same interpretation. Thus, the interpretation of a sequent $\Gamma \vdash \Delta$ in P only exhibits the fact that $\Gamma \vdash \Delta$ is *provable*.

On the other hand, in a richer model — such as $\mathbf{Set}_\times$ — we can actually distinguish two proofs of the same sequent via their interpretation. For example, consider the following two proofs of $A \wedge A \vdash A$:

$$\dfrac{\dfrac{\dfrac{}{A \vdash A}\,\mathrm{AX}}{A, A \vdash A}\,\mathrm{WKN}}{A \wedge A \vdash A}\,\wedge_L \qquad \text{and} \qquad \dfrac{\dfrac{\dfrac{\dfrac{}{A \vdash A}\,\mathrm{AX}}{A, A \vdash A}\,\mathrm{WKN}}{A, A \vdash A}\,\mathrm{EXC}}{A \wedge A \vdash A}\,\wedge_L$$

Their interpretations in an arbitrary cartesian closed category are

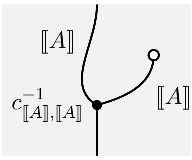

,
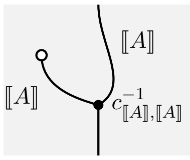

,

which are the two different projections of $[\![A]\!] \times [\![A]\!]$ onto $[\![A]\!]$. In $\mathbf{Set}_\times$, if $[\![A]\!]$ has at least two elements, these are different functions: $(x, y) \mapsto x$ and $(x, y) \mapsto y$.

Therefore, the categorical models of logic imply a shift: from a semantics of *provability* ("if something is provable, then it has a model") to a semantics of *proofs*.

Remark 10. Decades before categorical semantics were discovered, a number of logicians proposed what is known as the *Brouwer-Heyting-Kolmogorov* (BHK) interpretation of logic. According to BHK, the meaning of a formula is given by its proofs: for an atomic formula A, we are free to choose what constitutes a proof of A; for a composite formula, the standard of proof is fixed by that of its components. For example:

- a proof of $A \wedge B$ is a *pair* of a proof of A and a proof of B;
- a proof of $A \to B$ is a *process* turning a proof of A into a proof of B.

If we read "process" as "function", this is compatible with the semantics of $\mathsf{LJ}_{\wedge\to}$ in $\mathbf{Set}_\times$, under the interpretation that $[\![A]\!]$ is the "set of proofs of A". Then $[\![A \wedge B]\!]$ is exactly the "set of pairs of a proof of A and a proof of B", and $[\![A \to B]\!]$ the "set of functions sending proofs of A to proofs of B".

This is a useful picture to have in mind, but by itself it does no justice either to the BHK interpretation, or to categorical semantics:

1. on the one hand, both the notion of proof and the notion of process in BHK are deliberately vague, so they are compatible with a number of more interesting and dynamic interpretations than "elements of a set" and "functions";

2. on the other hand, there are categorical semantics of logic in which any “functional” or “process-theoretic” interpretation is a stretch; for example topological interpretations where a proof is interpreted as a *tangle*.

Since in deriving a sequent in $\mathsf{LJ}_{\wedge\to}$, we are also constructing a 2-cell in a cartesian closed category, we may think of labelling each sequent with a term which denotes that 2-cell. Fortunately, the right language for the job predates even the definition of category by a few decades: it is the λ-*calculus*.

Given the ample literature available on the subject, we will be brief on the definition of the λ-calculus, of which we consider a version with pairings and projections.

Definition 11. The set of *terms* of the λ-calculus (or λ-terms), together with a set $\mathsf{free}(t)$ of free variables of each term t, are defined inductively as follows:

- each variable x is a term with $\mathsf{free}(x) := \{x\}$;
- if t and u are terms, then $\langle t, u\rangle$ and tu are terms, with $\mathsf{free}(\langle t, u\rangle) = \mathsf{free}(tu) := \mathsf{free}(t) \cup \mathsf{free}(u)$;
- if t is a term, then $\pi_1 t$ and $\pi_2 t$ are terms, with $\mathsf{free}(\pi_1 t) = \mathsf{free}(\pi_2 t) := \mathsf{free}(t)$;
- if t is a term and x a variable, then $\lambda x.t$ is a term, with $\mathsf{free}(\lambda x.t) := \mathsf{free}(t) \setminus \{x\}$.

Terms are considered modulo α-equivalence, that is, the renaming of *bound* (non-free) variables and their binders: for example, $\lambda x.x$ is α-equivalent to $\lambda y.y$.

The *substitution* $u[t/x]$ of a term t for a variable x inside a term u is defined by:

- $x[t/x] := t$, and $y[t/x] := y$ if $y \neq x$;
- $\langle u_1, u_2\rangle[t/x] := \langle u_1[t/x], u_2[t/x]\rangle$;
- $(\pi_i u)[t/x] := \pi_i(u[t/x])$, for $i = 1, 2$;
- $u_1 u_2[t/x] := (u_1[t/x])(u_2[t/x])$;
- $(\lambda y.u)[t/x] := \lambda y.u[t/x]$ assuming $y \notin \mathsf{free}(x) \cup \mathsf{free}(t)$.

The condition in the last clause can always be realised working modulo α-equivalence. We will also write $u[t/x, t'/y]$ for the simultaneous substitution of t for x and t' for y.

The relation $\rightsquigarrow_\beta$ of β-*reduction* on terms is defined by

- $(\lambda x.t)u \rightsquigarrow_\beta t[u/x]$,
- $\pi_1\langle t, u\rangle \rightsquigarrow_\beta t$ and $\pi_2\langle t, u\rangle \rightsquigarrow_\beta u$,
- if $t \rightsquigarrow_\beta t'$, then $tu \rightsquigarrow_\beta t'u$, $ut \rightsquigarrow_\beta ut'$, $\pi_i t \rightsquigarrow_\beta \pi_i t'$, $\langle t, u\rangle \rightsquigarrow_\beta \langle t', u\rangle$, $\langle u, t\rangle \rightsquigarrow_\beta \langle u, t'\rangle$, and $\lambda x.t \rightsquigarrow_\beta \lambda x.t'$.

We write $\rightsquigarrow_\beta^*$ for the transitive, reflexive closure of $\rightsquigarrow_\beta$. A term t is said to be in *normal form* if there is no β-reduction from t to another term.

Construction 12. Now, we will use terms of the λ-calculus to label sequents of $\mathsf{LJ}_{\wedge\to}$ in a way that denotes their interpretation in a cartesian closed category. The labelling of a sequent $A_1, \ldots, A_n \vdash B$ by a term t is done in type-theoretic style, as a *typing judgment* for t:

$$x_1 : A_1, \ldots, x_n : A_n \vdash t : B,$$

where $\mathsf{free}(t) \subseteq \{x_1, \ldots, x_n\}$ and the x_i are all distinct. The left-hand side of the sequent is called a *context* for the typing judgment.

We will still use $\Gamma, \Delta, \ldots$ for sequences of *typed variables* $x_1 : A_1, \ldots, x_n : A_n$. We call an entire sequent $\Gamma \vdash t : B$ a *typed term*, and say that t has type B in the context Γ.

Structural rules.

$$\frac{}{x : A \vdash x : A}\ \textsc{ax} \qquad \frac{\Gamma_1, x : A, y : B, \Gamma_2 \vdash t : C}{\Gamma_1, y : B, x : A, \Gamma_2 \vdash t : C}\ \textsc{exc}$$

$$\frac{\Gamma, x : A, y : A \vdash t : B}{\Gamma, z : A \vdash t[z/y, z/x] : B}\ \textsc{ctr} \qquad \frac{\Gamma \vdash t : B}{\Gamma, z : A \vdash t : B}\ \textsc{wkn}$$

where in CTR and WKN the variable z is fresh (does not appear elsewhere in the context).

Logical rules.

$$\frac{\Gamma, x : A, y : B \vdash t : C}{\Gamma, z : A \wedge B \vdash t[\pi_1 z/x, \pi_2 z/y] : C}\ \wedge_L \qquad \frac{\Gamma_1 \vdash t : A \qquad \Gamma_2 \vdash u : B}{\Gamma_1, \Gamma_2 \vdash \langle t, u\rangle : A \wedge B}\ \wedge_R$$

$$\frac{x : A, \Gamma \vdash t : B}{\Gamma \vdash \lambda x.t : A \to B}\ \to_R \qquad \frac{\Gamma_1 \vdash u : A \qquad x : B, \Gamma_2 \vdash t : C}{\Gamma_1, z : A \to B, \Gamma_2 \vdash t[zu/x] : C}\ \to_L$$

where in $\wedge_L$ and $\to_L$ the variable z is fresh.

It is clear that the term so constructed is only unique up to α-equivalence.

Example 13. The following is the derivation of the typed λ-term corresponding to the proof of Example 3:

$$\dfrac{\dfrac{\dfrac{\dfrac{}{x : A \vdash x : A}\ \textsc{ax} \qquad \dfrac{}{x' : B \vdash x' : B}\ \textsc{ax}}{x : A, y : A \to B \vdash yx : B}\ \to_L \qquad \dfrac{\dfrac{}{w : B \vdash w : B}\ \textsc{ax} \qquad \dfrac{}{w' : C \vdash w' : C}\ \textsc{ax}}{w : B, z : B \to C \vdash zw : C}\ \to_L}{\dfrac{x : A, y : A \to B, z : B \to C \vdash z(yx) : C}{y : A \to B, z : B \to C \vdash \lambda x.z(yx) : A \to C}\ \to_R}\ \to_L}{w : (A \to B) \wedge (B \to C) \vdash \lambda x.\pi_2 w((\pi_1 w)x) : A \to C}\ \wedge_L$$

Exercise 14. Give a proof in $\mathsf{LJ}_{\wedge\to}$ of the following sequents:

$$A \to (A \to B) \vdash A \to B, \qquad A \to (B \to C) \vdash (A \to B) \to (A \to C).$$

In both cases, construct the corresponding typed λ-term.

The typed λ-terms that are derivable in $\mathsf{LJ}_{\wedge\to}$ have a quite restrictive property: they are all in normal form. We can obtain non-normalised typed λ-terms by adding the following, labelled CUT rule:

$$\frac{\Gamma_1 \vdash t : A \qquad x : A, \Gamma_2 \vdash u : B}{\Gamma_1, \Gamma_2 \vdash u[t/x] : B}\ \textsc{cut}$$

For example, we can obtain a β-reducible term of the form $(\lambda x.t)u$ as follows:

$$\dfrac{\dfrac{\dfrac{\vdots\,(a)}{x : A, \Gamma_1 \vdash t : B}}{\Gamma_1 \vdash \lambda x.t : A \to B}\to_R \qquad \dfrac{\dfrac{\vdots\,(b)}{\Gamma_2 \vdash u : A} \qquad \dfrac{}{z : B \vdash z : B}\ \textsc{ax}}{y : A \to B, \Gamma_2 \vdash yu : B}\to_L, \textsc{exc}}{\Gamma \vdash (\lambda x.t)u : B}\ \textsc{cut}$$

The typed terms that we can construct in $\mathsf{LJ}_{\wedge\to} + \textsc{cut}$ are called *simply-typed*.

A step of Gentzen's cut-elimination proof for $\mathsf{LJ}_{\wedge\to}$ turns this into the following:

$$\dfrac{\dfrac{\vdots\,(b)}{\Gamma_2 \vdash u : A} \qquad \dfrac{\vdots\,(a)}{x : A, \Gamma_1 \vdash t : B}}{\Gamma \vdash t[u/x] : B}\ \textsc{cut}, \textsc{exc}$$

which is a derivation of the immediate β-reduction of the original term.

One can do a similar analysis of the entire cut-elimination proof, and obtain a correspondence

$$\text{cut-elimination steps} \Leftrightarrow \beta\text{-reduction steps}$$

and ultimately

$$\text{cut-elimination} \Leftrightarrow \text{reduction to normal form.}$$

In this sense, proofs of $\mathsf{LJ}_{\wedge\to} + \textsc{cut}$ with cut-elimination as a dynamics and simply-typed λ-terms with β-reduction as a dynamics are equivalent models of computation.

What do the categorical models make of this? Given a derivation of $\Gamma \vdash t : B$, we write $[\![t]\!] : [\![\Gamma]\!] \Rightarrow [\![B]\!]$ for the corresponding 2-cell in a cartesian closed category X.

Proposition 15. *Suppose that $\Gamma \vdash t : B$ is derivable in $\mathsf{LJ}_{\wedge\to} + \textsc{cut}$, and $t \rightsquigarrow_\beta t'$. Then in any cartesian closed category X, under any interpretation of propositional atoms, $[\![t]\!] = [\![t']\!]$.*

What this result intuitively says is that the dynamics of sequent calculus and λ-calculus are invisible to the categorical models: if a proof or typed λ-term is the definition of a *process* which computes a function, then the categorical interpretation can tell at most *what function is being computed*, and not *how* it is computed.

This is often expressed by saying that categorical semantics are a *denotational semantics*, giving the "denotation" of a program (what it computes), rather than an *operational semantics*, modelling how the program is executed.

Remark 16. Any categorical interpretation also identifies terms related by *η-expansion*: this is a relation restricted to simply-typed λ-terms, given by

- $t \rightsquigarrow_\eta \lambda x.tx$ if t has a type of the form $A \to B$ and $x \notin \mathsf{free}(t)$, and
- $t \rightsquigarrow_\eta \langle \pi_1 t, \pi_2 t\rangle$ if t has a type of the form $A \wedge B$.

The equivalence relation obtained as the union of the reflexive, symmetric, transitive closures of $\rightsquigarrow_\beta$ and $\rightsquigarrow_\eta$, called *λ-conversion*, is the minimal equivalence relation respected by all interpretations in cartesian closed categories. In fact, one can define a "syntactic model" whose 1-cells are formulas of $\mathsf{LJ}_{\wedge\to}$, and 2-cells are simply-typed λ-terms modulo λ-conversion, and prove that it is a cartesian closed category.

Remark 17. As cut-elimination always reduces a proof in $\mathsf{LJ}_{\wedge\to} + \textsc{cut}$ to a proof in $\mathsf{LJ}_{\wedge\to}$, so β-reduction always reduces a simply-typed λ-term to normal form.

However, the latter is not true of general λ-terms: λ-terms can encode recursion, in such a way that β-reduction diverges. Famously, the term

$$\mathsf{Y} := \lambda f.(\lambda x.f(xx))(\lambda x.f(xx))$$

satisfies $\mathsf{Y}t \rightsquigarrow_\beta^* t(\mathsf{Y}t)$ for all terms t, and iterating we have

$$\mathsf{Y}t \rightsquigarrow_\beta^* t(\mathsf{Y}t) \rightsquigarrow_\beta^* t(t(\mathsf{Y}t)) \rightsquigarrow_\beta^* \ldots,$$

a divergent computation. In particular, Y is not simply-typed.

However, we can still interpret a generic λ-term in a cartesian closed category with a 1-cell a such that $a \times a$ and $a \to a$ are both isomorphic to a, so that we can choose a compositor $c_{a,a} : (a, a) \Rightarrow (a)$ and an evaluator $\mathrm{ev}_{a,a} : (a, a) \Rightarrow (a)$.

If we interpret all propositional atoms as a, we will have $[\![A]\!] = a$ for all formulas A of $\mathsf{LJ}_{\wedge\to}$. We can then "forget the typing" in $\mathsf{LJ}_{\wedge\to} + \textsc{cut}$, and obtain a pure calculus of λ-terms: for example, the $\textsc{cut}$ rule becomes

$$\frac{\Gamma_1 \vdash t \qquad x, \Gamma_2 \vdash u}{\Gamma_1, \Gamma_2 \vdash u[t/x]}\,\textsc{cut}$$

In this system, we can actually construct any λ-term. For example, the generic application tu of a term t to a term u is obtained as

$$\dfrac{\dfrac{\vdots}{\Gamma_1 \vdash t} \qquad \dfrac{\dfrac{\vdots}{\Gamma_2 \vdash u} \qquad \dfrac{}{z \vdash z}\,\mathrm{AX}}{y, \Gamma_2 \vdash yu}\,\to_L, \mathrm{EXC}}{\Gamma_1, \Gamma_2 \vdash tu}\,\mathrm{CUT}$$

Still, we can interpret any such term as a 2-cell $(a, \ldots, a) \Rightarrow (a)$, following the steps of Construction 6.

Exercise 18. Complete Remark 17, by constructing derivations of all λ-terms in the "type-free" sequent calculus.

Example 19. There is a cartesian closed category whose only non-identity 1-cell is the set $\mathbb{N}$ of natural numbers, and 2-cells $(\mathbb{N}, \overset{n}{\ldots}, \mathbb{N}) \Rightarrow (\mathbb{N}, \overset{m}{\ldots}, \mathbb{N})$ are *computable numerical partial functions* of n variables returning m natural numbers (or more precisely, partial functions whose projections are partial recursive functions of n variables).

There is a computable bijective *pairing* function $\langle -, - \rangle$ encoding pairs of natural numbers as natural numbers; this will be our compositor $c_{\mathbb{N},\mathbb{N}} : (\mathbb{N}, \mathbb{N}) \Rightarrow (\mathbb{N})$. Moreover, we can fix a *Gödel numbering* of partial recursive functions, in such a way that there is a computable function sending a pair of numbers (n, e) to $f_e(n)$ if and only if e is the Gödel number of f_e and f_e is defined on n: this will be our evaluator $\mathrm{ev}_{\mathbb{N},\mathbb{N}} : (\mathbb{N}, \mathbb{N}) \Rightarrow (\mathbb{N})$.

Proceeding as in Remark 17, we obtain a model of λ-terms as partial recursive functions.

Remark 20. In the previous lecture, we saw how we can interpret algebraic theories in cartesian monoidal categories. Then, we noticed that some theories do not need the entire structure of a cartesian monoidal category: a symmetric monoidal category, or even a monoidal category can suffice.

The same reasoning can be applied to the semantics of logic in cartesian closed categories. The cartesian structure is only used in the interpretation of *contraction* and of *weakening*: a proof which does not make use of these rules may be interpreted in a *symmetric monoidal closed category*, a symmetric monoidal category with all right Kan extensions.

The sequent calculus that one obtains from $\mathsf{LJ}_{\wedge\to}$ by removing contraction and weakening is a fragment of *linear logic*, introduced 30 years ago by Jean-Yves Girard. This is much more than just a "subtraction process": for example, certain rules which are usually mutually derivable become independent without contraction and weakening. In particular, one can decompose $\wedge$ into a pair of independent connectives, $\otimes$ (*tensor*, or *multiplicative* conjunction) and & (*with*, or *additive* conjunction), and similarly for disjunction.

We can also consider removing the exchange rule, which leads to a *non-commutative logic*, with semantics in a (non-symmetric) monoidal category. Non-commutative logic is compatible with the existence of a "right implication" (modelled by right Kan extensions) and a "left implication" (modelled by right Kan *lifts*, duals of extensions via $(-)^{\mathrm{op}}$). While somewhat less developed than linear logics, non-commutative logics have been studied for the modelling of the syntax of natural language.

In general, logics that reject certain structural rules are called *substructural*; they form a rich field of research on their own, one that has been particularly illuminated by categorical semantics.