Lecture 9: Adjunctions

Summary.

(1) Motivation: ‘Free’ and ‘forgetful’ transformations.

(2) Adjunctions: Definition and examples. Relation to universal properties. Properties.

(3) Adjunctions on ordered structures: Galois connections.

(4) Exponentials: The $- \times C \dashv -^C$ case study. Cartesian closed categories.

Luís Soares Barbosa,

UNIV. MINHO *(Informatics Dep.)* & INL *(Quantum Software Engineering Group)*

Motivation.

If categories can be thought of as particular mathematical spaces and functors as structure-preserving translations between them, an *adjunction* between, say, two functors $F : C \longrightarrow D$ and $G : D \longrightarrow C$, can be regarded as a source of *universals* in C and D. In fact, products and coproducts, final and initial objects and, in general, any universal construction arise in such a context. The notion of an adjunction pervades category theory and, in a sense, Mathematics as a whole.

As a motivation, recall the free monoid construction discussed in Lecture 4 (exercise 3), captured by the free functor $F : \mathsf{Set} \longrightarrow \mathsf{Mon}$ which builds a ‘syntactic’ monoid of ‘words’ over a set S. The forgetful functor $U : \mathsf{Mon} \longrightarrow \mathsf{Set}$ ‘undoes’ this construction returning the set of words over S. It is not difficult to verify that there exists a sort of symmetry between arrows involving these two functors. In detail, giving a set S and a monoid M, for each function $f : S \longrightarrow U(M)$ there is a unique monoid homomorphism $f^\bullet : F(S) \longrightarrow M$ making the diagram below left to commute. Or, starting from the other end, for each monoid homomorphism $h : F(S) \longrightarrow M$, there is a unique function $h_\bullet : S \longrightarrow U(M)$ so that the diagram in the right commutes, where η and ϵ are the natural transformations defined in the exercise mentioned above.

$$\begin{tikzcd} UF(S) \arrow[r, "U(f^\bullet)"] & U(M) \\ S \arrow[u, "\eta_S"] \arrow[ur, "f"'] & \end{tikzcd} \qquad \begin{tikzcd} F(S) \arrow[r, "U(h_\bullet)"] \arrow[dr, "h"'] & FU(M) \arrow[d, "\epsilon_M"] \\ & M \end{tikzcd}$$

$$F(S) \xrightarrow{\quad f^\bullet \quad} M \qquad\qquad S \xrightarrow{\quad h_\bullet \quad} U(M)$$

We’ve just captured a universal property. Recall that, by an entity being universal among a collection of similar ones, it is understood that there exists a unique way in which every other entity in the collection can be reduced to (or factored through). What we’ve just observed is that each component η_S of natural transformation η is universal among the arrows $f : S \longrightarrow U(M)$ in the sense that, for each such arrow, there exists a unique arrow which factors uniquely through η_S. And similarily for ϵ.

The notion of an *adjunction* captures a sort of symmetry and is therefore a source of universality, as discussed in the exercises below. We write $F \dashv U$, calling F the *left* and U the *right* adjoint functor. Natural transformations $\eta : \mathrm{Id} \Longrightarrow GF$ and $\epsilon : FG \Longrightarrow \mathrm{Id}$ are called the *unit* and *counit* of the adjunction, respectively.

Universality entails the existence of a natural isomorphism

$$
\begin{tikzcd}
\mathrm{Hom}_{\mathrm{Mon}}(F(S), M) \arrow[rr, bend right, "i"'] \arrow[rr, phantom, "\cong"] & & \mathrm{Hom}_{\mathrm{Set}}(S, U(M)) \arrow[ll, bend right, "j"']
\end{tikzcd}
$$

Exercise 1

Let $F \dashv G$. Compute $\eta : \mathrm{Id} \Longrightarrow GF$ and $\epsilon : FG \Longrightarrow \mathrm{Id}$ from the underlying natural isomorphism between homsets.

Exercise 2

Consider functors $! : \mathbf{1} \longrightarrow C$ and $\triangle : C \longrightarrow C \times C$, where $\mathbf{1}$ is the final object in Cat and $\triangle(A) = (A, A)$. Derive, for each of them, a right and a left adjoint. Comment the following statement: *all limits come from the right adjoints; all colimits from the left ones*.

Exercise 3

Suppose functors T and S compose and both have a left adjoint. Show that their composition TS has a left adjoint as well.

Exercise 4

Show that the unit and counit of an adjunction $F \dashv G$ satisfy the following conditions, known as the *triangle equalities*:

$$
\begin{aligned}
\epsilon F \cdot F\eta &= \mathrm{id}_F \\
G\epsilon \cdot \eta G &= \mathrm{id}_G
\end{aligned}
$$

Draw the relevant diagrams.

Exercise 5

An adjunction $f \dashv g$ between posets regarded as categories, say $P = (P, \leq)$ and $Q = (Q, \sqsubseteq)$ is a Galois connection:

$$\underbrace{f}_{\text{left adjoint}}(b) \leq a \quad \Leftrightarrow \quad b \sqsubseteq \underbrace{g}_{\text{right adjoint}}(a)$$

Draw the corresponding diagrams and explain why the adjunction unit and counit boil down to inequalities

$$f(g\ a) \leq a \quad \text{and} \quad b \sqsubseteq g(f\ b)$$

Exercise 6

In a Galois connection $f(b) \leq a \Leftrightarrow b \sqsubseteq g(a)$ the adjuncts determine each other uniquely: for example $f(b)$ is the greatest lower bound of all elements a such that $b \sqsubseteq g(a)$. Thus,

$$f\ b = \bigwedge\{a \mid b \sqsubseteq g\ a\} \quad \text{and} \quad g\ a = \bigsqcup\{b \mid f\ b \leq a\}$$

Using this fact, show that $f(b \sqcup b') = (f\ b) \vee (f\ b')$ and $g(a' \wedge a) = (g\ a') \sqcap (g\ a)$. Relate this result to the general fact that left adjoints preserve colimits and right adjoints preserve limits.

Exercise 7

Ler Rel be the poset of binary relations ordered by set inclusion, and consider the *converse* operation which computes the converse of a given relation. The usual relational laws

$$(R^\circ)^\circ = R$$
$$(R \cap S)^\circ = R^\circ \cap S^\circ$$
$$(R \cup S)^\circ = R^\circ \cup S^\circ$$

correspond to a particular adjunction over Rel. Can you identify it?

Exercise 8

Consider the following Galois connections in Rel where f and g are functions (thus, special cases of relations): $(f\cdot) \dashv (f^\circ\cdot)$ and $(\cdot f^\circ) \dashv (\cdot f)$. Write down the corresponding isomorphisms, known in the relational calculus as the *shunting* laws, and use them to conclude that

$$f \subseteq g \;\Leftrightarrow\; f = g \;\Leftrightarrow\; f \supseteq g$$

$\boxed{\textbf{Exercise } 9}$

Several laws in the calculus of binary relations are consequences of specific Galois connections. Consider the following operators, called the *right* and *left division*, respectively, and often useful to compute with relational data:

$$
\begin{aligned}
a(R \setminus S)c &\;\Leftrightarrow\; \forall_b\,.\,(bRa) \Rightarrow (bSc)\\
c(S \,/\, R)a &\;\Leftrightarrow\; \forall_b\,.\,(aRb) \Rightarrow (cSb)
\end{aligned}
$$

To quickly grasp the meaning of the right division, observe that if R relates flights with passengers, and S flights to the air-companies in charge of them, assertion $a\ (R \setminus S)\ c$ states that passenger a only flies with company c. Give a similar explanation for the meaning of left division.

Both divisions can be actually defined through Galois connections $(R\cdot) \dashv (R\setminus)$ and $(\cdot R) \dashv (/R)$, i.e.

$$
\begin{aligned}
R \cdot X \subseteq S &\;\Leftrightarrow\; X \subseteq R \setminus S\\
X \cdot R \subseteq S &\;\Leftrightarrow\; X \subseteq S \,/\, R
\end{aligned}
$$

and related to each other by still another adjunction: $(R/) \dashv (\setminus R)$. Show that the following laws are immediate consequences of these facts:

$$
\begin{tikzcd}
A \arrow[dr, "R"'] & & C \arrow[ll, "X\subseteq R\setminus S"'] \arrow[dl, "S"] \\
& B &
\end{tikzcd}
\qquad\qquad
\begin{tikzcd}
A \arrow[rr, "X\subseteq S/R"] & & C \\
& B \arrow[ul, "R"] \arrow[ur, "S"'] &
\end{tikzcd}
$$

$$
\begin{aligned}
R \cdot (S \cup T) &= (R \cdot S) \cup (R \cdot T)\\
(S \cup T) \cdot R &= (S \cdot R) \cup (T \cdot R)\\
R \setminus (S \cap T) &= (R \setminus S) \cap (R \setminus T)\\
(S \cap T) \,/\, R &= (S \,/\, R) \cap (T / R)\\
R \,/\, (S \cup T) &= (R \,/\, S) \cap (R \,/\, T)\\
(S \cup T) \setminus R &= (S \setminus R) \cap (T \setminus R)\\
R \setminus (S \setminus T) &= (S \cdot R) \setminus T
\end{aligned}
$$

$\boxed{\textbf{Exercise } 10}$

Every binary relation $R : A \longrightarrow B$ induces a function $Im_R : \mathcal{P}(A) \longrightarrow \mathcal{P}(B)$ mapping each $S \subseteq A$ to $\{b \in B \mid \exists_{a \in A}.\ (a, b) \in R\}$. This relation has a right adjoint: $[R] : \mathcal{P}(B) \longrightarrow \mathcal{P}(A)$. Draw the diagram corresponding to $Im_R \dashv [R]$ and show that a possible definition for $[R]$ is

$$[R](S') \;\widehat{=}\; \{a \in A \mid \forall_{b \in B}.\ (a, b) \in R \Rightarrow b \in S'\}$$

Observe that, in the context of transition systems, where R is an accessibility relation over a set of states, $[R]$ gives the semantics of the modal logic operator $\Box$ discussed in the first module of this course.

Case study on adjunctions: internalising 'arrow spaces' — $- \times C \dashv -^C$.

We will further illustrate the concept of an adjunction through a case study on an adjunction defining a fundamental universal construction which turns out not to be neither a limit nor a colimit. Actually, the categorical version of the usual notion of a function space in $\mathcal{S}et$ arises, as one could expect, from an adjunction. Let us briefly detail this construction.

Let C be an object of $\mathcal{C}$ and suppose that functor $- \times C$ has a right adjoint which we shall denote by $-^C$. This means that for all $f : X \times C \longrightarrow Y$, there exists a unique $f_\bullet : X \longrightarrow Y^C$ such that $f = \epsilon_Y \cdot (f_\bullet \times C)$, both the object Y^C and the universal ϵ_Y being uniquely determined up to isomorphism. Diagrammatically,

$$\begin{tikzcd} X \times C \arrow[r, "f_\bullet \times id"] \arrow[dr, "f"'] & Y^C \times C \arrow[d, "\epsilon_Y"] \\ & Y \end{tikzcd}$$

$$X \xrightarrow{\quad f_\bullet \quad} Y^C$$

Construction $-^C$ extends to a functor, the covariant *exponential* functor, by defining

$$h^C : A^C \longrightarrow B^C \;=\; (h \cdot \epsilon_A)_\bullet$$

for $h : A \longrightarrow B$. Note that Y^C has exactly the characteristic properties of the set of functions from C to Y in $\mathcal{S}et$. Bijection $f \leftrightsquigarrow f_\bullet$ corresponds, in this particular context, to *currying*: the well-known isomorphism between (binary) functions from $X \times C$ to Y and (unary) functions from X to the set of functions from C to Y. Being so popular, this terminology is also adopted in an arbitrary category: $f_\bullet$ is called the *curry* of f and written $\overline{f}$.

The family $\epsilon_X : X^C \times C \longrightarrow X$ is, of course, the counit of the adjunction

$$- \times C \;\dashv\; -^C$$

On the other hand, its unit has $\eta_X : X \longrightarrow (X \times C)^C$ as components. In $\mathcal{S}et$, ϵ corresponds to function evaluation and η to a function constructor:

$$\begin{aligned} \epsilon_Y\,(g, c) &= g\,(c) \qquad \text{(for } g : X \longrightarrow Y \text{)} \\ \eta_X\,(x)(c) &= (x, c) \end{aligned}$$

Counit ϵ is more commonly named ev, after *evaluation*. We shall also refer to η as sp, after *stamping*, and, again, such designations will carry over to general case.

The universal property captured by the $- \times C \;\dashv\; -^C$ adjunction diagram above can be written as the following equivalence (the concrete component of ev being of course determined by the type of f):

$$k = \overline{f} \;\;\Leftrightarrow\;\; f = ev \cdot (k \times id)$$

Note that the left to right implication expresses existence, while the converse one entails uniqueness (why?).

In an arbitrary category with exponentials C, X^C represents, as a C-object, the arrows from C to X. Consequently, the action of $-^C$ on each arrow $f : X \longrightarrow Y$ should *internalise* composition with f. In Set it is easy to verify that this is indeed the case. For $g : C \longrightarrow X$ and $c \in C$, a simple calculation yields,

$$
\begin{aligned}
& f^C(g)(c) \\
= & \quad \{\ f^C = \overline{(f \cdot ev)}, \text{ as discussed above} \} \\
& \overline{(f \cdot ev)}(g)(c) \\
= & \quad \{\ \text{uncurrying} \ \} \\
& f \cdot ev\ (g, c) \\
= & \quad \{\ \text{function composition} \ \} \\
& f(ev(g, c)) \\
= & \quad \{\ ev \text{ definition} \} \\
& f(g(c)) \\
= & \quad \{\ \text{function composition} \ \} \\
& (f \cdot g)\ (c)
\end{aligned}
$$

In an arbitrary category, however, we cannot talk about 'applying' a morphism to an 'element' of an object. We have, then, to state the intended result in the language of generalised elements (see Lecture 1, exercise 21). A generalised element of an exponential X^C is an arrow $\overline{g} : T \longrightarrow X^C$, which corresponds uniquely, under the adjunction, to $g : T \times C \longrightarrow X$. Keeping in mind that, in the generalised elements notation, $f^C\ (\overline{g})$ corresponds to $f^C \cdot \overline{g}$, the 'internalisation' result takes the form of an 'absorption' property for exponentials:

$$\overline{f \cdot g} \ = f^C \cdot \overline{g}$$

Taking g as a *point*, i.e. $\overline{g} : \mathbf{1} \longrightarrow X^C$, $f^C\ (\overline{g})$ equals $\overline{f \cdot g}$ as proved above, but now $\overline{f \cdot g}$ is itself a point of B^C, which corresponds to morphism $f \cdot g$. In other words,

$$f^C \ = \ f \cdot \_$$

Furthermore, the exponential functor above can be made into a *bifunctor* by defining, for each $h : C \longrightarrow D$, an arrow $X^h : X^D \longrightarrow X^C$ as follows:

$$X^h \ \widehat{=} \ \ X^D \xrightarrow{sp} (X^D \times C)^C \xrightarrow{(id_{X^D} \times h)^C} (X^D \times D)^C \xrightarrow{\quad ev^C \quad} X^C$$

Note that the exponential bifunctor is *contravariant* in its second argument. Moreover, X^h can be proved equal to post-composition with h, i. e. $X^h = \_ \cdot h$.

A category with finite products is called Cartesian and provides the right setting for discussing the existence of exponentials. When they exist, the category is called Cartesian closed.

Exercise 11

Using the universal property entailed by the adjunction $- \times C \dashv -^C$, show that

$$\overline{ev} = id_{X^C} \quad \text{and} \quad sp = \overline{id_{X \times C}}$$

Exercise 12

In the context of the previous exercise, derive the following results, known in the Bird and Moor algebra of programs [2] as the exponential *cancellation* and *fusion* laws, respectively.

$$f = ev \cdot (\overline{f} \times id) \quad \text{and} \quad \overline{g} \cdot f = \overline{g \cdot (f \times id)}$$

Exercise 13

Consider the diagram below. Why do the left triangle and right square commute?

$$\begin{tikzcd}
T \times C \arrow[r, "\overline{g} \times C"] \arrow[dr, "g"'] & A^C \times C \arrow[r, "f^C \times C"] \arrow[d, "ev_A"] & B^C \times C \arrow[d, "ev_B"] \\
& A \arrow[r, "f"'] & B
\end{tikzcd}$$

Fill in the explanations in the following conclusion of the proof that $\overline{f \cdot g} = f^C \cdot \overline{g}$:

$$\begin{aligned}
& f \cdot g = ev_B \cdot (f^C \times C) \cdot (\overline{g} \times C) \\
\equiv \quad & \quad \{ \cdots \} \\
& f \cdot g = ev_B \cdot (f^C \cdot \overline{g} \times C) \\
\equiv \quad & \quad \{ \cdots \} \\
& \overline{f \cdot g} = f^C \cdot \overline{g}
\end{aligned}$$

Exercise 14

Exponentials can be defined in any category with products such that, for every object X, the functor $(- \times X)$ is a left adjoint. Consider the category Graph of finite graphs. An object T in Graph is a pair of parallel functions $s_T, t_T : T_e \longrightarrow T_v$ from the set of edges T_e to the set of vertices T_v specifying the source and target of each edge, respectively. An arrow $h : T \longrightarrow R$ is a homomorphism of graphs defined as a pair of functions (h_v, h_e) such that the following diagram commutes:

$$\begin{tikzcd}
T_e \arrow[r, "h_e"] \arrow[d, shift right, "T_e"'] \arrow[d, shift left, "T_v"] & R_e \arrow[d, shift right, "R_e"'] \arrow[d, shift left, "R_v"] \\
T_v \arrow[r, "h_v"] & R_v
\end{tikzcd}$$

The category has products defined pointwise: in particular, an object $T \times R$ of $\mathsf{Graph} \times \mathsf{Graph}$ is given by $s_T \times s_R, t_T \times t_R : T_e \times R_e \longrightarrow T_v \times R_v$. The exponential object T^R is defined in [1] as a graph whose vertices are maps $\phi : T_v \longrightarrow R_v$. An edge θ connecting vertices ϕ to ψ is a map $\theta : T_e \longrightarrow R_e$ making the following diagram commute:

$$\begin{tikzcd}
T_v \arrow[d, "\phi"'] & T_e \arrow[l, "s_T"'] \arrow[r, "t_T"] \arrow[d, "\theta"] & T_e \arrow[d, "\psi"] \\
R_v & R_e \arrow[l, "s_R"'] \arrow[r, "t_R"] & R_e
\end{tikzcd}$$

i.e. a family $(\theta_e)_{e \in T_v}$ such that $s_R(\theta_e) = \phi(s_T(e))$ and $t_R(\theta_e) = \psi(t_T(e))$.

Thinking about maps ϕ and ψ as two different images of the vertices of graph T in graph R, θ is a family of edges in T, labeled by the edges of R, each connecting the source vertex in ϕ to the corresponding target one in ψ.

The evaluation arrow $ev : T^R \times R \longrightarrow T$ maps a vertex (ϕ, r) to the vertex $\phi(r)$, and an edge (θ, e) to the edge θ_e. On the other hand, the curry $\overline{h} : S \longrightarrow T^R$ of a graph homomorphism $h : S \times R \longrightarrow T$ takes a vertex $a \in T_v$ to the map $h(a, -) : R_v \longrightarrow T_v$, and an edge $c : a \longrightarrow b \in T_e$ to the map $h(c, -) : R_e \longrightarrow T_e$.

Verify that these data defines exponentials in Graph. Draw all necessary diagrams.

Exercise 15

Products are defined pointwise in the category Pos of partially ordered sets, i.e. given $(P, \leq)$ and $Q = (Q, \sqsubseteq)$,

$$(P, \leq) \times (Q, \sqsubseteq) \;=\; (P \times Q, \leq \times \sqsubseteq)$$

The exponential Q^P is defined as

$$(\{h : P \longrightarrow Q \mid f \text{ is monotone}\}, \dot{\leq})$$

where $h \mathrel{\dot{\leq}} h' \mathrel{\widehat{=}} \forall_{x \in P}.\ h(x) \leq h'(x)$. The ev natural transformation and the curry $\overline{f} : X \longrightarrow Q^P$ of $f : X \times P \longrightarrow Q$ are defined as in Set.

Complete the exponential construction in Pos showing that all functions involved are indeed monotone.

References

[1] S. Awodey. *Category Theory*. Oxford Logic Guides. Oxford University Press, 2006.

[2] R. Bird and O. Moor. *The Algebra of Programming*. Series in Computer Science. Prentice-Hall International, 1997.