

Lecture 8: Type Inference

Type annotations help you document your code and catch programming errors. However, it can be quite annoying to have to write down every single type in a program. If you've used Java Generics or something similar, you've probably encountered the annoyance of type annotations in real code. The type-checker you wrote earlier also has the same property.

The type inference problem is to take a program with no type annotations, such as the following program:

```
fix f . fun x . if n = 0 then 1 else n * f(n - 1)
```

and produce an equivalent program with type annotations:

```
fix (f : int → int) . fun (x : int) . if n = 0 then 1 else n * f(n - 1)
```

We can type-check the latter program in the usual way.

One way to infer types is to come up with ad hoc heuristics. However, we see an approach known as the *Hindley-Milner algorithm*, which is the foundation for type inference in Standard ML, OCaml, Haskell, and many other ML-family programming languages. The Hindley-Milner algorithm has two key properties:

- It is *sound*, so it always produces correct types. Therefore, you can trust the output of the algorithm and there is no need to type-check after inference. (Assuming your implementation is correct!)
- It is *complete*, so it always produces a type, if there is a type to be found. Therefore, if the algorithm fails to find a type, then the program is truly un-typable.

The algorithm has one other important property, which we'll look at later.

An Informal Example Suppose we are given the untyped factorial function and were asked to infer its type in our heads. We'll run through a detailed description of how we may do so in the table below. For brevity we use the following notation.

Notation

- We write $\llbracket e \rrbracket$ to mean “the type of the expression e ”.
- We use Greek letters to denote unknown types. These Greek letters are called “metavariables”.

Here is how we may reason about the type of factorial:

	Claim	Justification
1	$\llbracket \text{fix } f . \text{ fun } n . \dots \rrbracket = \llbracket \text{fun } n . \dots \rrbracket$	fix evaluates to the body.
2	$\llbracket f \rrbracket = \llbracket \text{fix } f . \dots \rrbracket$	f is bound to the whole expression during evaluation.
3	$\llbracket \text{fun } n . \dots \rrbracket = \alpha \rightarrow \beta$	The expression is a function, so it must have an argument and result (α and β).
4	$\llbracket n \rrbracket = \alpha$	n is the argument of the function that takes argument of type α .
5	$\llbracket \text{if } n = 0 \text{ then } 1 \text{ else } n * f(n - 1) \rrbracket = \beta$	The expression is the body of a function that produces a value of type β .
6	$\llbracket n = 0 \rrbracket = \text{bool}$	The expression is used as the condition in an if-expression.
7	$\llbracket n \rrbracket = \text{int}$	n is used in a comparison. (We can only compare integers.)
8	$\llbracket 1 \rrbracket = \llbracket n * f(n - 1) \rrbracket$	Both branches must have the same type.
9	$\llbracket 1 \rrbracket = \llbracket \text{if } n = 0 \text{ then } 1 \text{ else } n * f(n - 1) \rrbracket$	Type of a conditional is the type of either branch.
10	$\llbracket f \rrbracket = \gamma \rightarrow \delta$	f appears in function position in an application.
11	$\llbracket \gamma \rrbracket = \llbracket (n - 1) \rrbracket$	The argument to f , which has type γ .
12	$\llbracket n - 1 \rrbracket = \text{int}$	Subtraction produces integers.
13	$\llbracket f(n - 1) \rrbracket = \delta$	The return type of f is the type of the application.
14	$\llbracket \delta = \text{int} \rrbracket$	The expression of type δ is used as an argument of $*$.

At this point, we can solve the a few constraints to figure out the types of the metavariables α , β , γ , and δ :

Binary Operators

$op_2 ::= + \mid > \mid \dots$

Constants

$c ::= true$	True
$ false$	False
$ n$	Integers

Expressions

$e ::= c$	Constants
$ x$	Identifiers
$ op_2(e_1, e_2)$	Bin. Ops.
$ if \ e_1 \ then \ e_2 \ else \ e_3$	Conditionals
$ e_1 \ e_2$	Applications
$ fun \ x \ . \ e$	Functions
$ fix \ f \ . \ e$	Fixpoints

(a) Implicitly-Typed Syntax.

Types

$\tau ::= int$	Integer type
$ bool$	Boolean type
$ \tau_1 \rightarrow \tau_2$	Function types
$ \alpha$	Type metavariables

Binary Operators

$op_2 ::= + \mid > \mid \dots$

Constants

$c ::= true$	True
$ false$	False
$ n$	Integers

Expressions

$e ::= c$	Constants
$ x$	Identifiers
$ op_2(e_1, e_2)$	Bin. Ops.
$ if \ e_1 \ then \ e_2 \ else \ e_3$	Conditionals
$ e_1 \ e_2$	Applications
$ fun \ (x:\tau) \ . \ e$	Functions
$ fix \ (f:\tau) \ . \ e$	Fixpoints

(b) Explicitly-Typed Syntax.

Figure 12.1: Syntaxes for type inference.

- $\alpha = int$ by (4) and (7),
- $\beta = int$ by (5), (9) (given that $\llbracket 1 \rrbracket = int$),
- $\gamma = int$ by (11) and (12),
- $\delta = int$ by (14).

We should also ensure that no constraints are contradictory, since contradictory constraints indicate that the program has a type error. Once that is done, we can reconstruct the type annotations in the program.

1 Introduction

Type inference works with two syntaxes and two syntaxes that we will use are shown in fig. 12.1. The *implicitly-typed syntax*, on the left-hand side, is the syntax in which the user writes the program. The *explicitly-typed syntax*, on the right-hand side, is the result of type-inference. The difference between these two syntaxes is that the explicitly-typed syntax has type annotations that can be used to type-check the program. In contrast, the implicitly-typed syntax has no type annotations. However, it is important to note that **the implicitly-typed language is typed**. The types aren't written down, but we could still write typing rules for this language.¹ The explicitly-typed syntax just makes types manifest, so that a simple type-checker can type-check the program.

Types in the explicitly-typed language include metavariables, α , β , etc. Metavariables are not types themselves, but are placeholders that stand for types. We need metavariables to describe type inference. However, the result of type inference will produce a program that has no metavariables.

Type inference has several steps:

1. We transform a program in the implicitly-typed syntax to an identical program in the explicitly-typed syntax, using unique metavariables for each type annotation. For example, we transform $fun \ x \ . \ fun \ y \ . \ x + y$ to $fun \ (x:\alpha) \ . \ fun \ (y:\beta) \ . \ x + y$. These metavariables will allow us to refer to the type of an identifier without getting scope mixed up.
2. We generate a set of constraints by recursively processing the program. Each constraint equates two types to each other. To handle variables and scoping correctly, we use an environment that maps identifiers to metavariables. For example, the expression we have produces the constraints $\alpha = int$ and $\beta = int$.

¹We would have to “guess” the type of binding identifiers when building typing derivations.

<p>Constraints $C \in \tau = \tau$ Constraint Sets $\mathcal{C} \in \{C_1, \dots, C_n\}$ Type Environments $\Gamma \in x \mapsto \alpha$</p> <hr style="border: 0.5px solid black;"/> <div style="border: 1px solid black; padding: 2px; display: inline-block; margin-bottom: 10px;">$\Gamma \vdash e \Rightarrow \tau \times \mathcal{C}$</div> <p>T-INT $\Gamma \vdash n \Rightarrow (int, \emptyset)$ T-BOOL $\Gamma \vdash b \Rightarrow (bool, \emptyset)$ T-ID $\Gamma \vdash x \Rightarrow (\Gamma(x), \emptyset)$</p>	<p>T-ADD $\frac{\Gamma \vdash e_1 \Rightarrow (\tau_1, \mathcal{C}_1) \quad \Gamma \vdash e_2 \Rightarrow (\tau_2, \mathcal{C}_2)}{\Gamma \vdash e_1 + e_2 \Rightarrow (int, \{\tau_1 = int, \tau_2 = int\} \cup \mathcal{C}_1 \cup \mathcal{C}_2)}$</p> <p>T-IF $\frac{\Gamma \vdash e_1 \Rightarrow (\tau_1, \mathcal{C}_1) \quad \Gamma \vdash e_2 \Rightarrow (\tau_2, \mathcal{C}_2) \quad \Gamma \vdash e_3 \Rightarrow (\tau_3, \mathcal{C}_3)}{\Gamma \vdash \text{if } e_1 \text{ then } e_2 \text{ else } e_3 : (\tau_2, \{\tau_1 = bool, \tau_2 = \tau_3\} \cup \mathcal{C}_1 \cup \mathcal{C}_2 \cup \mathcal{C}_3)}$</p> <p>T-FUN $\frac{\Gamma[x \mapsto \alpha] \vdash e \Rightarrow (\tau, \mathcal{C})}{\Gamma \vdash \text{fun } (x:\alpha) . e \Rightarrow (\alpha \rightarrow \tau, \mathcal{C})}$</p> <p>T-FIX $\frac{\Gamma[x \mapsto \alpha] \vdash e \Rightarrow (\tau, \mathcal{C})}{\Gamma \vdash \text{fix } (x:\alpha) . e \Rightarrow (\alpha, \{\alpha = \tau\} \mathcal{C})}$</p> <p>T-APP $\frac{\Gamma \vdash e_1 \Rightarrow (\tau, \mathcal{C}_1) \quad \Gamma \vdash e_2 \Rightarrow (\tau', \mathcal{C}_2) \quad \alpha \text{ is a fresh metavariable}}{\Gamma \vdash e_1 e_2 \Rightarrow (\alpha, \{\tau' \rightarrow \alpha = \tau\} \cup \mathcal{C}_1 \cup \mathcal{C}_2)}$</p> <p style="text-align: center;">(a) Type Checking.</p>
---	--

Figure 12.2: Constraint Generation

3. We solve the constraints, using a classic algorithm called *unification*. Constraint-solving produces a *substitution* from metavariables to types that have no metavariables within them. For example, solving the constraints above produces the substitution $[\alpha \mapsto int, \beta \mapsto int]$. Several things can go wrong during constraint-solving. For example, we may derive an unsatisfiable constraint, such as $int = bool$, which indicates that the program has a type error.
4. We annotate the program by applying the substitution to the type variables we generated in the first step. In this example, we get the program $\text{fun } (x:int) . \text{fun } (y:int) . x + y$.
5. Finally, we can type-check the resulting program to verify the result of type-inference. However, if the steps above are perfectly correct, there is no need to verify the result.

2 Constraint Generation

Figure 12.2 gives the constraint generation rules for our language. We can read the rule $\Gamma \vdash e \Rightarrow T \times \mathcal{C}$ as *in the environment Γ , e has type T and produces the set of constraints \mathcal{C}* . There are two key differences between the type-inference rules and the type-checking rules:

- Unlike the type-checking rules, the type T may not be a complete type and may include metavariables (α , β , etc.).
- Constraint-generation does not catch type-errors. For example, the expression `true + 10` will fail to type-check. However, the constraint generation rules will produce the type int and the set of constraints $\{bool = int, int = int\}$.

However, constraint-generation does catch unbound-identifier errors. If a program has a free variable, then it won't be bound to a metavariable in the environment.

Notice that the types produced by constraint-generation are used to further generate additional constraints. For example, in the expression $\text{fun } (x:\alpha) . x + 3$, the bound identifier x produces the type α , which is used to produce the constraint $\alpha = int$ when generating constraints for $x+3$.

3 Solving Type Constraints

After constraint-generation completes, we can solve constraints by *unification*. The UNIFY algorithm, shown in fig. 12.3, takes a single constraint and produces as output a *substitution*, which is a finite map from metavariables to types. A substitution can be applied to a type to replace metavariables with concrete types.

For example, $\text{UNIFY}(\alpha \rightarrow int = bool \rightarrow \beta)$ produces the substitution $[\alpha \mapsto int, \beta \mapsto bool]$. If we apply this substitution to $\alpha \rightarrow int$ or to $bool \rightarrow \beta$, we get the type $bool \rightarrow int$. Notice that we get the same type on either

Substitutions

$$\begin{aligned} \pi &\in \alpha \rightarrow T \\ (\pi_1 \cdot \pi_2)(\alpha) &= \pi_2(\pi_1(\alpha)) \end{aligned}$$

Unification

$$\begin{aligned} \text{UNIFY} &\in C \rightarrow \pi \\ \text{UNIFY}(T = T) &= \cdot \\ \text{UNIFY}(\alpha = T) &= [\alpha \mapsto T] \text{ if } \alpha \text{ does not occur in } T \\ \text{UNIFY}(T = \alpha) &= [\alpha \mapsto T] \text{ if } \alpha \text{ does not occur in } T \\ \text{UNIFY}(S_1 \rightarrow S_2 = T_1 \rightarrow T_2) &= \text{UNIFY}(\pi(S_2), \pi(T_2)) \cdot \pi \\ &\quad \text{where } \pi = \text{UNIFY}(S_1, T_1) \end{aligned}$$

Figure 12.3: Unification

side, which is a key property of the algorithm. However, $\text{UNIFY}(\gamma \rightarrow \text{int} = \text{bool} \rightarrow \gamma)$ should fail, since there is no substitution that can be applied to either side to produce the same type.

UNIFY is a simple recursive algorithm with three cases:

- $\text{UNIFY}(\text{int} = \text{int})$ should produce the empty substitution, whereas $\text{UNIFY}(\text{int} = \text{bool})$ should fail. In general, unifying two base types should succeed with the empty substitution if the types are the same or fail if the types are different.
- $\text{UNIFY}(\alpha = T)$ and $\text{UNIFY}(T = \alpha)$ should produce the substitution $[\alpha \mapsto T]$. However, there is an important caveat discussed below, called the *occurs check*.
- $\text{UNIFY}(S_1 \rightarrow S_2 = T_1 \rightarrow T_2)$ needs to recursively unify the argument and the result types. However, the two substitutions have to be *composed* together.

The Occurs Check Consider the constraint $\alpha = \text{int} \rightarrow \alpha$. If unification is done naively, we will produce the substitution $[\alpha \mapsto \text{int} \rightarrow \alpha]$. If we apply this substitution to either side, we get the constraint $\text{int} \rightarrow \alpha = \text{int} \rightarrow (\text{int} \rightarrow \alpha)$. We can unify this new constraint to get a larger type on either side. In fact, we can repeatedly apply unification to get larger and larger types. The real problem is that the original constraint is circular and contradictory. There is no substitution π such that $\pi(\alpha) = \pi(\text{int} \rightarrow \alpha)$. To avoid this infinite regress, unification needs an *occurs check* to detect cyclic constraints. It is enough to run the occurs check when unifying metavariables with types, as shown in fig. 12.3.

Unifying a set of constraints Given UNIFY, it is straightforward to unify a set of constraints. After unifying each constraint in the set, we need to apply the intermediate substitution to the remaining constraints.