

Lecture 2: Big-Step Semantics

1 Representing Abstract Syntax

These are examples of *arithmetic expressions*:

- $2 * 4$
- $1 + 2 + 3$
- $5 * 4 * 2$
- $1 + 2 * 3$

We all know how to evaluate these expressions in our heads. But, when we do, we resolve several ambiguities. For example, should we evaluate $1 + 2 * 3$ like this:

$$\begin{aligned} & 1 + 2 * 3 \\ = & (1 + 2) * 3 \\ = & 3 * 3 \\ = & 9 \end{aligned}$$

or like this:

$$\begin{aligned} & 1 + 2 * 3 \\ = & 1 + (2 * 3) \\ = & 1 * 6 \\ = & 6 \end{aligned}$$

The latter is the convention in mathematics, but it is an arbitrary choice. A programming language could adopt either convention or adopt a completely different notation. For example, the following three programs, written in three different languages, are trying to express the same thing:

- $1 + 2$ infix syntax, from the C language
- $(+ 1 2)$ parenthesized prefix syntax, from the Scheme language
- $1 2 +$ postfix syntax, from the Forth language

These three *concrete syntaxes* are very different, but all mean “the sum of the number three and the number four”.

Concrete syntax is important, because it is the human-computer interface to a programming language. It is easy to find acrimonious debates on the Web about the virtues of Python’s indentation-sensitive syntax, whether semicolons should be optional in JavaScript,

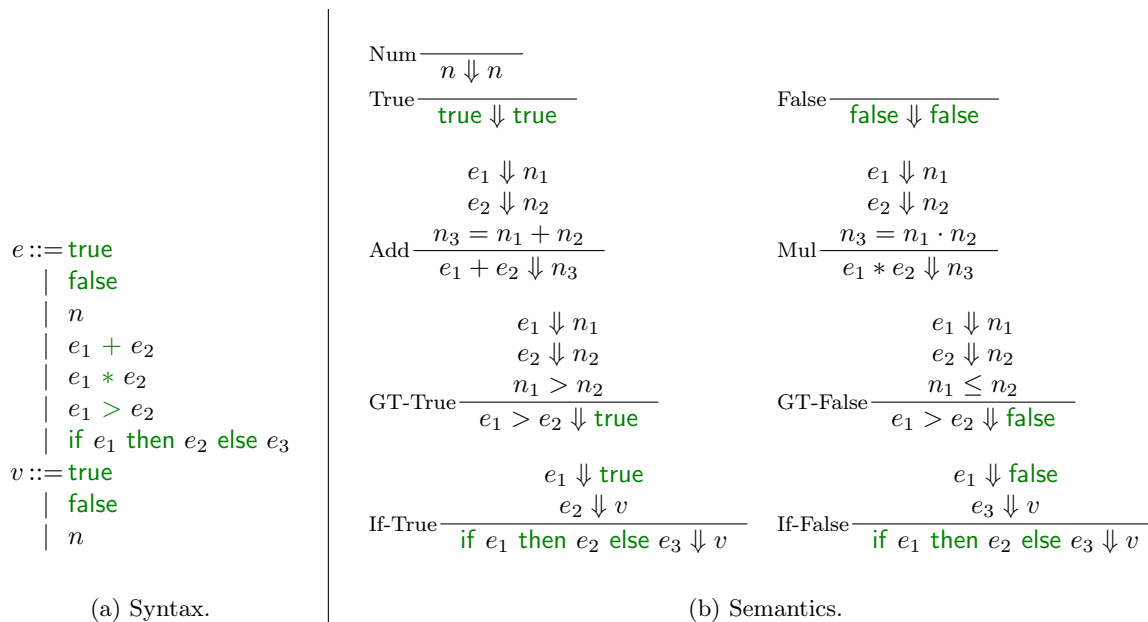


Figure 3.1: Syntax and semantics of a language with arithmetic and boolean expressions.

how C code should be indented, and so on. But, **this course will almost completely ignore concrete syntax** because it is irrelevant to the semantics of programming languages. We will instead work with *abstract syntax*, which is an abstract “tree-shaped” representation of syntax that suffers none of the ambiguities of concrete syntax. In compilers, the *parser* consumes a concrete-syntax string and produces an equivalent *abstract syntax tree*. This course will largely ignore parsing and instead work directly with the abstract syntax tree.

OCaml makes it easy to define a type that represents the abstract syntax of arithmetic expressions:

```

type exp =
| Num of int
| Add of exp * exp
| Mul of exp * exp
| Div of exp * exp

```

Here are some examples of abstract arithmetic expressions and their concrete representations (written in normal, mathematical notation):

Concrete Syntax	Abstract Syntax
1 + 2 + 3	Add (Add (Num 1, Num 2), Num 3)
1 + 2 * 3	Add (Num 1, Mul (Num 2, Num 3))
(1 + 2) / 3	Div (Add (Num 1, Num 2), Num 3)

2 Syntax as Sets

We start with a very rigorous mathematical definition of the abstract syntax of a small language of arithmetic and boolean expressions.

Definition 1 (Syntax of arithmetic expressions). *Let \mathcal{E} denote the set of arithmetic expressions. We define \mathcal{E} to be the smallest set that is generated by the following rules:*

- $\text{true} \in \mathcal{E}$.
- $\text{false} \in \mathcal{E}$.
- If $n \in \mathbb{Z}$ then $n \in \mathcal{E}$.
- If $e_1 \in \mathcal{E}$ and $e_2 \in \mathcal{E}$ then $e_1 + e_2 \in \mathcal{E}$.
- If $e_1 \in \mathcal{E}$ and $e_2 \in \mathcal{E}$ then $e_1 * e_2 \in \mathcal{E}$.
- If $e_1 \in \mathcal{E}$ and $e_2 \in \mathcal{E}$ then $e_1 > e_2 \in \mathcal{E}$.
- If $e_1 \in \mathcal{E}$, $e_2 \in \mathcal{E}$, and $e_3 \in \mathcal{E}$ then $\text{if } e_1 \text{ then } e_2 \text{ else } e_3 \in \mathcal{E}$.

Hopefully, you'll agree that defining syntax in this way is extremely tedious. We will never do this again and instead define syntax using the notation in fig. 3.1a.

This notation is much terser and is what you'll find when you read the programming languages literature. But, you should be aware that it is just shorthand for the more verbose definition given above. Also note that we are "abusing notation" and using the metavariable e to denote the set of expressions and elements of the set (and similarly for n). Again, this is standard practice.

3 Semantics as Relations

Inference rules In this section, we use *inference rules* to define the semantics of our language. These are some examples of inference rules:

$$\frac{\frac{X}{Y} \quad \frac{A \quad B}{C}}{W}$$

You should read these rules as:

- If X holds, then Y holds.
- If A and B hold, then C holds.
- W holds (i.e., W is an axiom).

Semantics Programs evaluate expressions until they become *values*. Intuitively, a value is an expression that cannot be further simplified. For the language defined above, the values (v) are just integers and booleans. Note that $v \subset e$.

Given our definition of expressions and values, we can define the semantics of the language as an *evaluation relation*. The evaluation relation is a binary relation that relates expressions e to equivalent values v . We write the evaluation relation using the following notation:

$$e \Downarrow v$$

You should pronounce this as " e evaluates to v ".

The $e \Downarrow v$ relation is defined by a set of inductively defined rules, similar to the definition of expressions. However, instead of first soldiering through a long, definition in prose, we'll

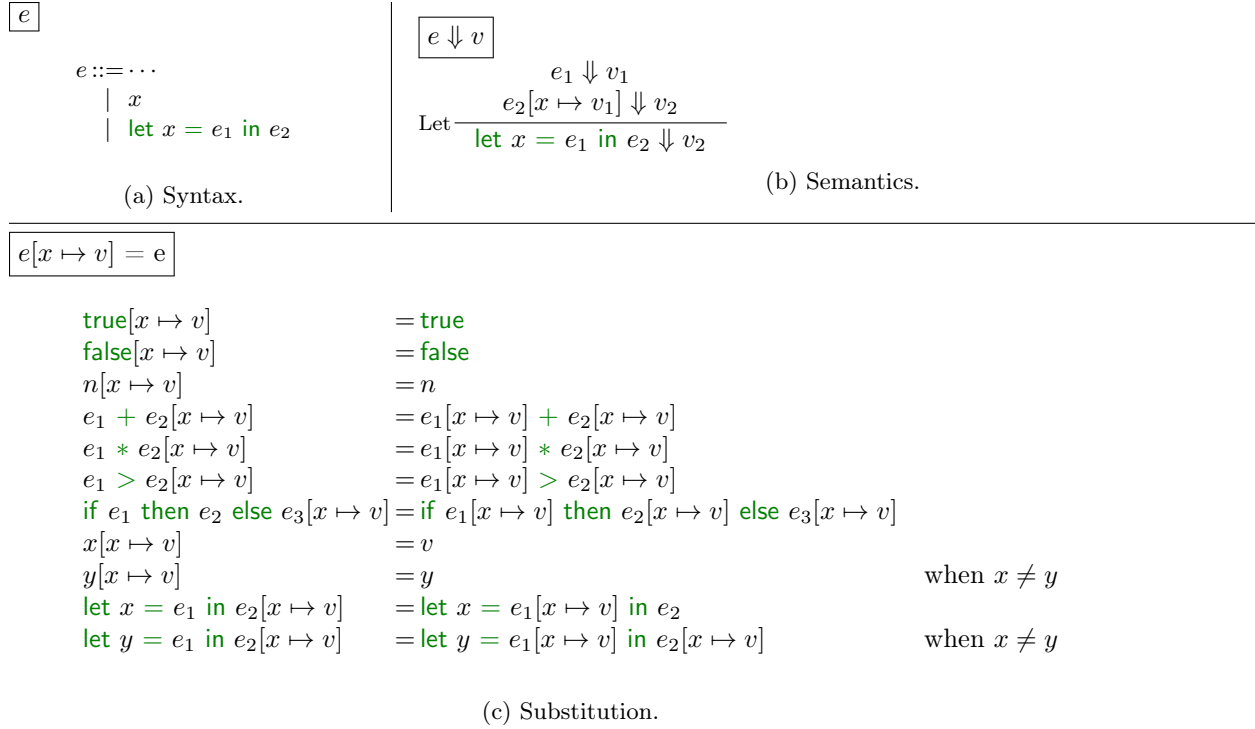


Figure 3.2: Identifiers and let-expressions. Extends fig. 3.1.

use the inference rules in fig. 3.1b to define $e \Downarrow v$. Notice that $e \Downarrow v$ is not a function. In this case, $e \Downarrow v$ is not defined for all e and v . E.g., $\text{true} + 5$ is not defined. In this simple case, $e \Downarrow v$ is a partial function. But, in a richer language with non-deterministic features, such as threads or random number generators, $e \Downarrow v$ will not even be a partial function.

4 Let-Binding and Substitution

Even when we write simple arithmetic expressions, we often need to reuse the same expression in several places. If we try to repeat the same expression, we are likely to propagate mistakes. More fundamentally, repeated expressions consume more storage space and require more steps to evaluate. Therefore, almost all programming languages provide a means to name expressions.

Figure 3.2a extends the syntax of our language with *let-expressions* and *identifiers*. A let-expression, such as $\text{let } x = 1 + 2 \text{ in } x * x$ has three parts:

- A *binding identifier* x ,
- A *named expression* $1 + 2$, and
- The *body* $x * x$.

In this example, the body has two bound identifiers that refer to the binding identifier.

The semantics of this language extension are given in fig. 3.2b. Notice that there is no inference rule for evaluating identifiers. Instead, the inference rule for let expressions

substitutes identifiers with the value of the named expression. The substitution function, $e[x \mapsto v]$, is typically read as “substitute x with v in e ”.

Here is an example of a derivation that uses substitution:

$$\text{LET} \frac{\text{ADD} \frac{\text{NUM} \frac{1 \Downarrow 1}{1} \quad \text{NUM} \frac{2 \Downarrow 2}{2}}{1 + 2 \Downarrow 3} \quad (x * x)[x \mapsto 3] = 3 * 3 \quad \text{MUL} \frac{\text{NUM} \frac{3 \Downarrow 3}{3} \quad \text{NUM} \frac{3 \Downarrow 3}{3}}{3 * 3 \Downarrow 9}}{\text{let } x = 1 + 2 \text{ in } x * x \Downarrow 9}$$

This is a very straightforward example since it only has one let-expression. In general, a program may have several let-expressions and even reuse the same identifier in several places. The semantics has been carefully designed to handle these cases in a natural way.

This example defines two names by nesting let-expressions:

$$\text{LET} \frac{\text{LET} \frac{\vdots}{\text{let } y = 20 + 10 \text{ in } y * x \Downarrow 300}}{\text{let } x = 10 \text{ in let } y = 20 + x \text{ in } y * x \Downarrow 300}$$

This example has two let-expressions that use the same name:

$$\text{LET} \frac{\text{NUM} \frac{10 \Downarrow 10}{10} \quad (\text{let } x = 20 \text{ in } 1 + x)[x \mapsto 10] = \text{let } x = 20 \text{ in } 1 + x \quad \text{LET} \frac{\vdots}{\text{let } x = 20 \text{ in } 1 + x \Downarrow 21}}{\text{let } x = 10 \text{ in let } x = 20 \text{ in } 1 + x \Downarrow 21}$$

As the proof tree shows, the identifier in the expression $1 + x$ is bound to the inner definition of x . We say that the inner x *shadows* the enclosing definition.

4.1 Substitution and Evaluation Order

The inference rule for let-expressions first evaluates the named expression to a value and then substitutes that value into the body. Consider this alternate definition, which applies substitution first:

$$\text{LET}' \frac{e'_2 = e_2[x \mapsto e_1] \quad e'_2 \Downarrow v}{\text{let } x = e_1 \text{ in } e_2 \Downarrow v}$$

For example, the following proof that $\text{let } x = 10 + 10 \text{ in } x * x \Downarrow 400$ uses this alternate rule for let. Notice that we use the Add rule twice and the Mul rule once:

$$\text{LET}' \frac{\text{MUL} \frac{\text{ADD} \frac{\text{NUM} \frac{10 \Downarrow 10}{10} \quad \text{NUM} \frac{10 \Downarrow 10}{10}}{10 + 10 \Downarrow 20} \quad \text{ADD} \frac{\text{NUM} \frac{10 \Downarrow 10}{10} \quad \text{NUM} \frac{10 \Downarrow 10}{10}}{10 + 10 \Downarrow 20}}{(10 + 10) * (10 + 10) \Downarrow 400}}{\text{let } x = 10 + 10 \text{ in } x * x \Downarrow 400}$$

However, using the original Let rule, we use the the Add and Mul rules once each:

$$\text{LET} \frac{\text{ADD} \frac{\text{NUM} \frac{10 \Downarrow 10}{10} \quad \text{NUM} \frac{10 \Downarrow 10}{10}}{10 + 10 \Downarrow 20} \quad x * x[x \mapsto 20] = 20 * 20 \quad \text{MUL} \frac{\text{NUM} \frac{20 \Downarrow 20}{20} \quad \text{NUM} \frac{20 \Downarrow 20}{20}}{20 * 20 \Downarrow 400}}{\text{let } x = 10 + 10 \text{ in } x * x \Downarrow 400}$$

Loosely speaking, the larger proof tree corresponds to an evaluation that “takes more time”. In fact, proofs with Let’ will be larger whenever there are multiple occurrences of the identifier in the body of the let expression.

However, a deeper problem arises when we have the named expression is faulty. For example, let e be the expression $\text{let } x = \text{true} + 10 \text{ in } 200$. Using the Let rule, there does not exist a v , such that $e \Downarrow v$. However, using the Let’ rule, $e \Downarrow 200$, since x is unused in the body.

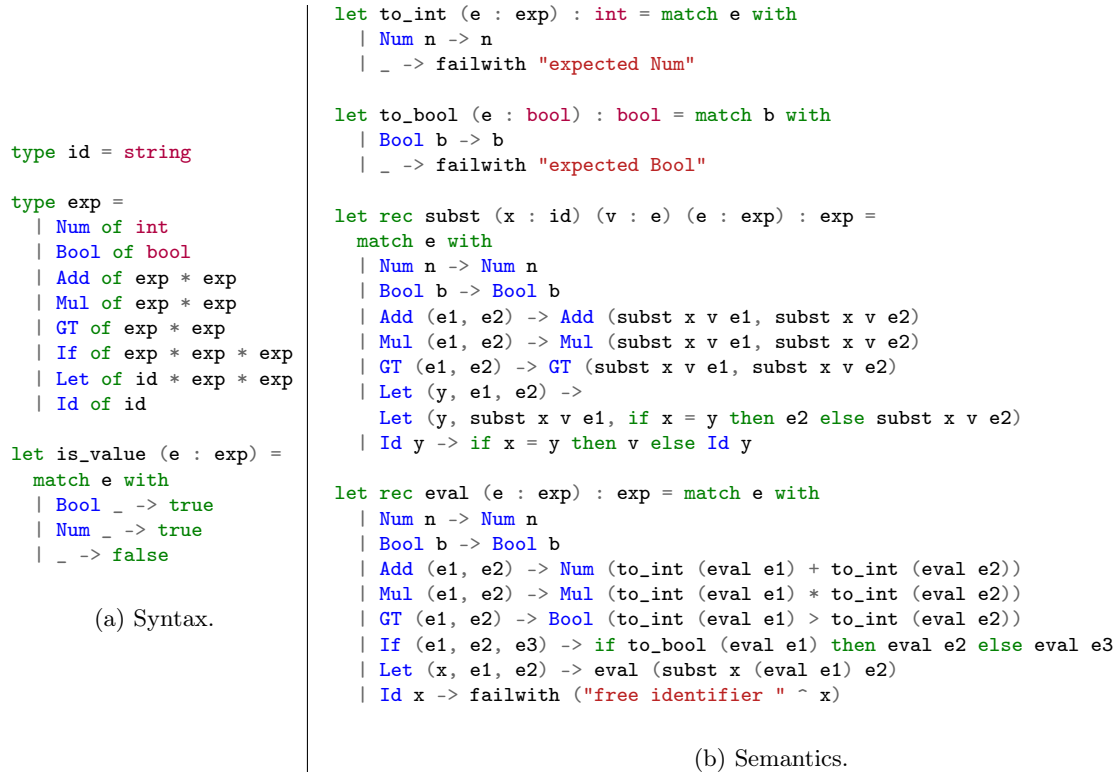


Figure 3.3: OCaml implementation of fig. 3.1

We'll investigate this in more depth in a few weeks.

4.2 Implementation

Figure 3.3 is an OCaml implementation of our language. You should study this definition and the inference rules in fig. 3.1 to ensure that they are in close correspondence.

5 Functions

In this section, we extend our language with *functions*, similar to functions in OCaml. The syntax and semantics of this extension are given in fig. 3.4. In this extension, functions

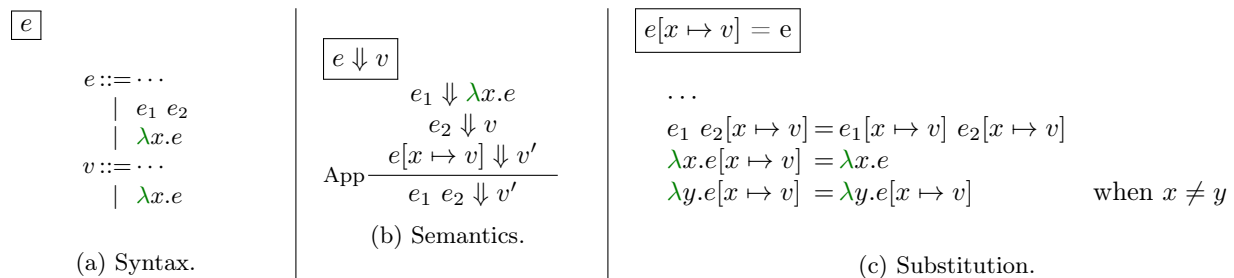


Figure 3.4: First-class functions. Extends fig. 3.2.

are values, which means that they are values just like numbers and booleans. This allows functions to be passed to other functions, returned from functions, bound to identifiers, and so on. When a language gives functions all the rights and privileges of simpler values, they are known as *first-class functions*.

An apparent shortcoming of this extension is that all functions take exactly one argument. However, we can encode multi-argument functions by leveraging first-class functions. For example, suppose we want to write a function that takes two arguments x and y and calculates $10 * x * y$. Although we cannot write a two-argument function, we can do this:

$$\lambda x. \lambda y. 10 * x * y$$

This is known as *currying*.

