

1 Evaluation Contexts

The rules for structural operational semantics can be classified into two types:

- *reduction rules*, which describe the actual computation steps; and
- *evaluation order rules*, which constrain the choice of reductions that can be performed next.

For example, the CBV reduction strategy for the λ -calculus was captured in the following rules:

$$\overline{(\lambda x. e) v \rightarrow e\{v/x\}} \quad (1)$$

$$\frac{e_1 \rightarrow e'_1}{e_1 e_2 \rightarrow e'_1 e_2} \quad \frac{e_2 \rightarrow e'_2}{v e_2 \rightarrow v e'_2} \quad (2)$$

Rule (1), β -reduction, is a reduction rule, whereas rules (2) are evaluation order rules. The rules (2) say essentially that a reduction may be applied to a redex on the left-hand side of an application anytime, and may be applied to a redex on the right-hand side of an application provided the left-hand side is already fully reduced.

Although there are only two evaluation order rules in the CBV λ -calculus, there are typically many more in real-world programming languages. This motivates the desire to find a more compact representation for such rules.

Evaluation contexts provide a mechanism to do just that. An evaluation context E , sometimes written $E[\bullet]$, is a λ -term or a metaexpression representing a family of λ -terms with a special variable $[\bullet]$ called the *hole*. If $E[\bullet]$ is an evaluation context, then $E[e]$ represents E with the term e substituted for the hole.

Every evaluation context $E[\bullet]$ represents a *context rule*

$$\frac{e \rightarrow e'}{E[e] \rightarrow E[e']},$$

which says that we may apply the reduction $e \rightarrow e'$ in the context $E[e]$.

For the case of the CBV λ -calculus, the two evaluation order rules (2) are specified by the two evaluation context schemes $[\bullet] \cdot e$ and $v \cdot [\bullet]$. These are just a compact way of representing the rules (2). Thus we could specify the CBV λ -calculus simply by writing

$$(\lambda x. e) v \rightarrow e\{v/x\} \quad [\bullet] \cdot e \quad v \cdot [\bullet].$$

The CBN λ -calculus has an equally compact specification:

$$(\lambda x. e) e' \rightarrow e\{e'/x\} \quad [\bullet] \cdot e.$$

2 Nested Contexts

Note that in CBV, the evaluation contexts $[\bullet] \cdot e$ and $v \cdot [\bullet]$ do not specify *all* contexts in which the reduction rule (1) may be applied. There are also compound contexts obtained from nested applications of the rules (2). For example, the context

$$(v \cdot [\bullet]) \cdot e \quad (3)$$

is also a valid evaluation context for CBV, since it can be derived from two applications of the rules (2):

$$\frac{\frac{e_1 \rightarrow e_2}{v e_1 \rightarrow v e_2}}{(v e_1) e \rightarrow (v e_2) e}. \quad (4)$$

Here we have applied the right-hand rule of (2) in the first step and the left-hand rule of (2) in the second. The evaluation context (3) represents the abbreviated rule

$$\frac{e_1 \rightarrow e_2}{(v e_1) e \rightarrow (v e_2) e}$$

obtained by collapsing the two steps of (4).

The set of *all* valid evaluation contexts for the CBV λ -calculus is represented by the grammar

$$E ::= [\bullet] \mid E \cdot e \mid v \cdot E.$$

3 Annotated Proof Trees

We can also use evaluation contexts to indicate exactly where a reduction is applied in each step of a proof tree. For example, consider the annotated proof tree

$$\frac{\frac{\frac{(\lambda x. x) 0 \rightarrow 0}{(\lambda x. x) ((\lambda x. x) 0) \rightarrow (\lambda x. x) 0} \quad ((\lambda x. x) [\bullet])}{(\lambda x. x) ((\lambda x. x) 0) \lambda z. zz \rightarrow (\lambda x. x) 0 \lambda z. zz} \quad ([\bullet] \lambda z. zz)}$$

We have labeled each step to indicate the context in which the β -reduction was applied.

As above, we can simplify the tree by collapsing the two steps and annotating the resulting abbreviated tree with the corresponding nested context:

$$\frac{(\lambda x. x) 0 \rightarrow 0}{(\lambda x. x) ((\lambda x. x) 0) \lambda z. zz \rightarrow (\lambda x. x) 0 \lambda z. zz} \quad ((\lambda x. x) [\bullet] \lambda z. zz)$$

4 Error Propagation

Evaluation contexts can be used to define the semantics of error exceptions. If we have a special error value `error`, we can very easily propagate it using the evaluation order rule

$$E[\text{error}] \rightarrow \text{error}.$$

This obviates the need to show in painstaking detail how error propagates up through a series of applications of rewrite rules. We will revisit this idea later on when we talk about exception handling mechanisms.

The benefits of evaluation contexts will become exceedingly clear in the future as we add more features to the language.