

Object Inheritance Without Classes

Timothy Jones¹, Michael Homer², James Noble³, and Kim Bruce⁴

- 1 Victoria University of Wellington, Wellington, New Zealand
tim@ecs.vuw.ac.nz
- 2 Victoria University of Wellington, Wellington, New Zealand
mwh@ecs.vuw.ac.nz
- 3 Victoria University of Wellington, Wellington, New Zealand
kix@ecs.vuw.ac.nz
- 4 Pomona College, Claremont, California, USA
kim@cs.pomona.edu

Abstract

Which comes first: the object or the class? Language designers enjoy the conceptual simplicity of object-based languages (such as Emerald or Self) while many programmers prefer the pragmatic utility of classical inheritance (as in Simula and Java). Programmers in object-based languages have a tendency to build libraries to support traditional inheritance, and language implementations are often contorted to the same end. In this paper, we revisit the relationship between classes and objects. We model various kinds of inheritance in the context of an object-oriented language whose objects are not defined by classes, and explain why class inheritance and initialisation cannot be easily modelled purely by delegation.

1998 ACM Subject Classification D.3.3 Classes and objects

Keywords and phrases Inheritance, Objects, Classes, Operational semantics

Digital Object Identifier 10.4230/LIPIcs.ECOOP.2016.13

Supplementary Material ECOOP Artifact Evaluation approved artifact available at <http://dx.doi.org/10.4230/DARTS.2.1.6>

1 Introduction

Class-based object-oriented languages have a simple story about the relationships between objects and the classes that create them: an object is an instance of a class [2]. A specialised, ‘one-off’ object is just an instance of a specialised, one-off, anonymous class [27]. Inheritance is between classes, and new objects are constructed and initialised by their classes.

This simple story comes at the expense of a more complicated story about classes — especially so if classes are themselves objects. Thirty years ago, Alan Borning identified eight separate roles that classes can play in most class-based object-oriented languages [4], each of these roles adding to the complexity of the whole language, which typically leads inexorably to various kinds of infinite regress in meta-object systems [21, 29, 44].

To address this problem, prototype-based object-oriented languages, beginning with Lieberman’s work inspired by LOGO [32] and popularised by Self [51], adopted a conceptually simpler model in which objects were the primary concept, defined individually, without any classes. Inheritance-like sharing of state and behaviour was handled by delegation between objects, rather than inheritance between their defining classes. Special-purpose objects could be defined directly, while new objects could be created in programs by cloning existing objects.



© Timothy Jones, Michael Homer, James Noble, and Kim Bruce;
licensed under Creative Commons License CC-BY

30th European Conference on Object-Oriented Programming (ECOOP 2016).

Editors: Shriram Krishnamurthi and Benjamin S. Lerner; Article No. 13; pp. 13:1–13:26

Leibniz International Proceedings in Informatics



LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany



```

method graphic(canvas) {
  object {
    method image { abstract }
    method draw { canvas.render(image) }
    var name
    canvas.register(self)
    draw // Local method request
  }
}

def amelia = object {
  inherit graphic(canvas)
  // Override using a field getter
  def image = images.amelia
  // Assign to inherited field
  name := "Amelia"
}

```

■ **Figure 1** An example object constructor method and an inheriting object.

Emerald [3] went one step further and aimed to eschew all implicit sharing mechanisms, supporting neither inheritance nor delegation.

Programmers using object-based languages have found the need to reintroduce classes — several times over in many of these languages. The Emerald compiler, and later the Self IDE, added explicit support for class-style inheritance. Languages with object inheritance such as Lua [34], JavaScript [46], and Tcl [52] have a variety of libraries implementing classes in terms of objects. Most recently classes have been added explicitly to the recent ECMAScript standard [53], to bring some order to the profusion of libraries already offering classes.

The problem we address in this paper is precisely the tension between the conceptual simplicity of objects and the practical utility of classes: can a language be based conceptually on ‘objects first’ and include a relatively familiar notion of inheritance? In this paper we present models of eight inheritance mechanisms for a language without an inherent class construct: forwarding, delegation, concatenation, merged identity, uniform identity, multiple uniform identity, method transformations, and positional inheritance. The first three correspond to the foundational object-based models, while merged identity and uniform identity introduce more classical behaviour that parallels C++ and Java, and the remaining models introduce multiple object inheritance with different conflict resolution techniques.

We evaluate the tradeoffs between power and complexity of these models, particularly in their object initialisation semantics, and compare them to the behaviour of other languages, demonstrating that the typical class initialisation semantics are fundamentally at odds with prototypical object inheritance. We have also implemented all of the models in PLT Redex [17], making the models executable and allowing direct comparison of the differences in execution for any program.

2 Inheritance Without Classes

The term ‘inheritance’ typically refers to reuse relationships between classes, but there are also a number of ‘objects-first’ languages that eschew classes and permit reuse relationships directly between objects. Consider the example `graphic` method in Figure 1: the method constructs a fresh object with some methods and a mutable (**var**) field, and runs some initialisation code including a local call to the `draw` method before returning. If we can inherit from objects created by this method, what is the resulting program behaviour? Results vary for different interpretations of inheritance. We can inherit from the objects this method creates, or assign special semantics to methods which directly construct fresh objects like this and inherit from the method itself. Because of the presence of initialisation code in the class, these interpretations have different — and potentially unexpected — behaviours.

In its most basic form, inheritance permits reuse by providing a mechanism for an object to defer part of its implementation to an already implemented part of the program, but the reality is that there is much more to consider: the value of ‘self’ in method bodies and during initialisation, intended and accidental method overriding, whether objects are composed of several object identities in an inheritance chain or a single unified identity, the meaning of method requests which are not qualified with a receiver, and so on. We can draw different conclusions about the state of our program after `amelia` in Figure 1 is constructed, depending on which inheritance semantics is in play. We group these into our relevant concerns:

- **Registration.** Is the identity of a super-object stored during initialisation, either explicitly or through lexical capture, the same as the final object? This is clearly the intention of the call to `register` in `graphic`’s initialisation.
- **Down-calls.** Can a method in a super-object call down into a method in a lower object? Can it do so during initialisation? The implementation of the `draw` method relies on a down-call to the `image` method.
- **Action at a Distance.** Can operations on an object implicitly affect another object? If the registered `graphic` object is different to `amelia`, what is the value of its `name` field after `amelia` is initialised?
- **Stability.** Is the implementation of methods in an object the same throughout its lifetime? Which `image` method will be invoked by the request to `draw` at the end of `graphic`? Can the structural type of an object change after it has been constructed?
- **Preëxistence.** Can an object inherit from any other object it has a reference to? Does `amelia` have to inherit a call to the `graphic` method, or will a preëxisting object suffice?
- **Multiplicity.** Can an object inherit from multiple other objects? If `amelia` also wished to have the behavior of another object, can a second inherit clause be added? If so, how are multiple methods with the same name resolved, and where are fields located?

We are also interested in the general semanticses of the inheritance systems, such as what order the parts of initialisation execute in, what visibility and security concerns arise, and how local method requests are resolved. We also point out particular curiosities, such as the absence of overloaded methods from super references and inheriting from definitions already inherited from some other parent, but these features are specific to individual models.

These are the concerns that we will judge in the following object inheritance models. While our `graphic` example clearly assumes some of these features in its implementation, these features are not universally desirable, and each provides certain abilities or guarantees at some cost. Our intention is to use these concerns to provide an accurate description of the tradeoffs involved with each inheritance model. Some of our models of inheritance attempt to interpret `graphic` as a class, but only in the sense that it is a factory that guarantees fresh `graphic` objects. We also compare the models to existing languages that use each form of inheritance, particularly JavaScript, which is capable of implementing all of the models.

3 Graceless

In order to provide a formal semantics for the various object inheritance systems, we first present a base model of the Grace programming language without inheritance, and then proceed to extend this base model in different ways to construct the different behaviours. Grace is a useful language to model object inheritance features in, as it is not class-based, and it does not permit mutation of an existing object’s structure, so we can consider pure

Syntax

$e ::= o \mid r.m(\bar{e}) \mid m(\bar{e}) \mid \mathbf{self} \mid v$	$\boxed{\mid e; e \mid f \xrightarrow{x} \mid f \xleftarrow{x} e}$	<i>(Expression)</i>
$S ::= \mathbf{def} \ x = e \mid \mathbf{var} \ x \mid \mathbf{var} \ x := e \mid e$	<i>(Statement)</i>	$v ::= \mathbf{done} \ \boxed{\mid \ell}$ <i>(Value)</i>
$M ::= \mathbf{method} \ m(\bar{x}) \{ \bar{e}; e \}$	<i>(Method)</i>	$r ::= e$ <i>(Method receiver)</i>
$m ::= x \mid x :=$	<i>(Method name)</i>	$f ::= \mathbf{self} \ \boxed{\mid \ell}$ <i>(Field receiver)</i>
$o ::= \mathbf{object} \ \{ \overline{M} \ \overline{S} \}$	<i>(Object expression)</i>	$F ::= \boxed{x \mapsto v}$ <i>(Field)</i>
$\sigma ::= \boxed{\cdot \mid \ell \mapsto \langle \overline{F}, \overline{M} \rangle, \sigma}$	<i>(Object store)</i>	$s ::= \boxed{v/x \mid \mathbf{self}.m/m}$ <i>(Substitution)</i>

Evaluation context

$$E ::= [] \mid E.m(\bar{e}) \mid v.m(\bar{v}, E, \bar{e}) \mid m(\bar{v}, E, \bar{e}) \mid E; e \mid \ell \xleftarrow{x} E$$

■ **Figure 2** The grammar of the Graceless base model, with runtime-only components boxed.

object concerns without being overcome by ‘open objects’ with trivially mutable structure, as found in JavaScript.

Our core model resembles the existing Tinygrace language [28], but with more features, such that ‘Tiny’ is not an appropriate moniker. As the language is still not a complete implementation of Grace, we have opted to name it *Graceless*. We have modelled object references, mutable and immutable fields, and method requests unqualified by a receiver, in order to demonstrate the wide-ranging effects of changes to the behaviour of inheritance. Explicit types have been removed, in order to focus on the dynamic semantics.

The grammar for Graceless is provided in Figure 2. The boxed areas in the grammar represent forms which only exist at runtime, not in the user-side syntax of programs. As in Grace, we omit empty argument and parameter parentheses, making some method requests indistinguishable from local variables and field lookups per Meyer’s uniform access principle [37]. Within an object are method definitions M and statements S . A statement is either a field declaration or an expression, the latter allowing method requests in object initialisation code. Field declarations are either constant **definitions** or **variable**, and **var** fields may be declared without an initial value.

Expressions e in the user-side syntax are object expressions o , method requests either qualified or unqualified by a receiver, or the sentinel value **done**. At runtime, expressions can also include references ℓ to locations in the heap, sequences $e; e$, field fetch $\mathbf{self} \xrightarrow{x}$, and field assign $\mathbf{self} \xleftarrow{x} e$, with \mathbf{self} eventually substituted for some ℓ . The field operations are distinguished from method calls $e.x$ and $e.x := e$.

Objects can be nested inside of each other (as method bodies) and unqualified requests can be made on the resulting local scope. We assume Barendregt’s rule [1] for \mathbf{self} references, such that each \mathbf{self} variable introduced by nested object expressions is unique. Because method names appear in the local scope as well as the public interface of their directly surrounding object, we cannot assume Barendregt’s rule that their names are unique — an object may need to shadow an outer object’s method in order to conform to a given interface — so substitution cannot continue past a shadowing definition.

Graceless has two forms of substitution. The typical substitution $[v/x]e$ replaces the term x with the value v in the term e . A *qualifying substitution* $[\mathbf{self}.m/m]e$ replaces any local request

$$\boxed{\langle \sigma, e \rangle \rightsquigarrow \langle \sigma', e' \rangle}$$

$ \frac{}{\langle \sigma, v; e \rangle \rightsquigarrow \langle \sigma, e \rangle} $ <p>(E-NEXT)</p>	$ \frac{\langle \sigma, e \rangle \rightsquigarrow \langle \sigma', e' \rangle}{\langle \sigma, E[e] \rangle \rightsquigarrow \langle \sigma', E[e'] \rangle} $ <p>(E-FETCH)</p>	$ \frac{\text{method } m(\bar{x}) \{ e \} \in \sigma(\ell)}{\langle \sigma, \ell.m(\bar{v}) \rangle \rightsquigarrow \langle \sigma, [\ell/\text{self}][\bar{v}/\bar{x}]e \rangle} $ <p>(E-REQUEST)</p>
$ \frac{}{\langle \sigma, \ell \xrightarrow{x} \rangle \rightsquigarrow \langle \sigma, \sigma(\ell)(x) \rangle} $ <p>(E-OBJECT)</p>	$ \frac{\ell \text{ fresh} \quad \bar{m} = \text{names}(\overline{M}, \overline{S}) \quad \langle \overline{M}_f, \bar{e} \rangle = \text{body}(\overline{S}) \quad \bar{m} \text{ unique}}{\langle \sigma, \mathbf{object} \{ \overline{M} \overline{S} \} \rangle \rightsquigarrow \langle \sigma(\ell \mapsto \langle \emptyset, [\text{self}.m/m]\overline{M} \overline{M}_f \rangle), [\ell/\text{self}][\text{self}.m/m]\bar{e}; \ell \rangle} $	
$ \frac{}{\langle \sigma, \ell \xleftarrow{x} v \rangle \rightsquigarrow \langle \sigma(\ell)(x \mapsto v), \mathbf{done} \rangle} $ <p>(E-ASSIGN)</p>		

Auxiliary Definitions

$$\text{names}(\overline{\mathbf{method}} m(\bar{x}) \{ e \}, \overline{S}) = \bar{m} \cup \bar{m}_f \quad \text{where } \langle \overline{\mathbf{method}} m_f(\bar{y}) \{ e_f \}, e \rangle = \text{body}(\overline{S})$$

$$\text{body}(\emptyset) = \langle \emptyset, \emptyset \rangle$$

$$\text{body}(\mathbf{def} x = e, \overline{S}) = \langle \text{accessors}(\mathbf{def}, x, y) \overline{M}, \text{self} \xleftarrow{y} e \bar{e} \rangle \quad \text{where } y \text{ fresh and } \langle \overline{M}, \bar{e} \rangle' = \text{body}(\overline{S})$$

$$\text{body}(\mathbf{var} x, \overline{S}) = \langle \text{accessors}(\mathbf{var}, x, y) \overline{M}, \bar{e} \rangle \quad \text{where } y \text{ fresh and } \langle \overline{M}, \bar{e} \rangle' = \text{body}(\overline{S})$$

$$\text{body}(\mathbf{var} x := e, \overline{S}) = \langle \text{accessors}(\mathbf{var}, x, y) \overline{M}, \text{self} \xleftarrow{y} e \bar{e} \rangle \quad \text{where } y \text{ fresh and } \langle \overline{M}, \bar{e} \rangle' = \text{body}(\overline{S})$$

$$\text{body}(e, \overline{S}) = \langle \overline{M}, e \bar{e} \rangle \quad \text{where } \langle \overline{M}, \bar{e} \rangle' = \text{body}(\overline{S})$$

$$\text{accessors}(\mathbf{def}, x, y) = \mathbf{method} x \{ \text{self} \xrightarrow{y} \}$$

$$\text{accessors}(\mathbf{var}, x, y) = \mathbf{method} x \{ \text{self} \xrightarrow{y} \} \mathbf{method} x := (z) \{ \text{self} \xleftarrow{y} z \}$$

■ **Figure 3** Term reduction.

$m(\bar{e})$ with the qualified form $\text{self}.m(\bar{e})$ in the term e . Both forms of substitution are ended by a shadowing definition, which can be either an adjacent method name or a surrounding parameter definition: substitutions into an object expression $[v/x]o$ and $[\text{self}.m/m]o$ do not modify o if it contains a method x or m respectively, and the substitutions into a method $[v/x]M$ and $[\text{self}.x/x]M$ do not modify the method body if M has a parameter x . Because the local variable x and the local method request x are indistinguishable, the ordering of the substitution matters: $[\text{self}.x/x][v/x]x$ produces v , but $[v/x][\text{self}.x/x]x$ produces $\text{self}.x$.

The reduction judgment $\langle \sigma, e \rangle \rightsquigarrow \langle \sigma', e' \rangle$ is defined in Figure 3, indicating an expression e with store σ is reduced to an expression e' with a potentially modified store σ' . Rule E-CONTEXT uses the evaluation context to perform congruence reduction, and Rule E-NEXT evaluates to the next expression in a sequence when the current one has finished evaluating. Rules E-FETCH and E-ASSIGN handle operations on a field store. The store access $\sigma(\ell)(x)$ looks up the field x in the object at location ℓ , getting stuck if the field has not been initialised yet. The store update $\sigma(\ell)(x \mapsto v)$ sets the field x to the value v in the object at location ℓ , introducing a new field if one was not already present.

Rule E-REQUEST processes requests by looking up the corresponding method in the receiver. In the method body, it substitutes both the arguments for parameter names, and the receiver for the name self . Rule E-OBJECT takes an object expression and builds a corresponding object in the store, with no fields, the methods in the object expression, and

13:6 Object Inheritance Without Classes

Extended Syntax

$$I ::= \mathbf{inherit} \ e \ \boxed{\bar{s}} \quad (\text{Inherit clause}) \quad v ::= \dots \ \boxed{|\ (\ell \ \mathbf{as} \ v)|} \quad (\text{Value})$$

$$o ::= \dots \ | \ \mathbf{object} \ \{ I \ \bar{M} \ \bar{S} \} \quad (\text{Object expression}) \quad r ::= \dots \ | \ \mathbf{super} \ \boxed{|\ (\ell \ \mathbf{as} \ e)|} \quad (\text{Receiver})$$

$$s ::= \dots \ \boxed{|\ (\ell \ \mathbf{as} \ \mathbf{self})/\mathbf{super}|} \quad (\text{Substitution})$$

$$\frac{\langle \sigma, e \rangle \rightsquigarrow \langle \sigma, e \rangle \quad \frac{\mathbf{method} \ m(\bar{x}) \ \{ e \} \in \sigma(\ell_{\uparrow})}{\langle \sigma, (\ell_{\uparrow} \ \mathbf{as} \ \ell_{\downarrow}).m(\bar{v}) \rangle \rightsquigarrow \langle \sigma, [\ell_{\downarrow}/\mathbf{self}][v/x]e \rangle}}{\langle \sigma, \mathbf{object} \ \{ \mathbf{inherit} \ e \ \bar{s} \ \bar{M} \ \bar{S} \} \rangle \rightsquigarrow \langle \sigma', \mathbf{object} \ \{ \mathbf{inherit} \ e' \ \bar{s} \ \bar{M} \ \bar{S} \} \rangle} \quad (\text{E-INHERIT/CONTEXT})$$

$$\frac{\langle \sigma, e \rangle \rightsquigarrow \langle \sigma', e' \rangle}{\langle \sigma, \mathbf{object} \ \{ \mathbf{inherit} \ e \ \bar{s} \ \bar{M} \ \bar{S} \} \rangle \rightsquigarrow \langle \sigma', \mathbf{object} \ \{ \mathbf{inherit} \ e' \ \bar{s} \ \bar{M} \ \bar{S} \} \rangle} \quad (\text{E-INHERIT})$$

$$\frac{\langle \bar{F}, \bar{M}_{\uparrow} \rangle = \sigma(\ell) \quad \bar{M}_{\uparrow}' = \text{override}(\bar{M}_{\uparrow}, \text{names}(\bar{M}, \bar{S}))}{\langle \sigma, \mathbf{object} \ \{ \mathbf{inherit} \ \ell \ \bar{s} \ \bar{M} \ \bar{S} \} \rangle \rightsquigarrow \langle \sigma, \mathbf{object} \ \{ \bar{M}_{\uparrow}' \ \bar{s} \} [(\ell \ \mathbf{as} \ \mathbf{self})/\mathbf{super}](\bar{M} \ \bar{S}) \rangle} \quad (\text{E-INHERIT})$$

Auxiliary Definitions

$$\text{override}(\emptyset, \bar{m}) = \emptyset$$

$$\text{override}(\mathbf{method} \ m(\bar{x}) \ \{ e \} \ \bar{M}, \bar{m}) = \begin{cases} m \in \bar{m} & \text{override}(\bar{M}, \bar{m}) \\ m \notin \bar{m} & \mathbf{method} \ m(\bar{x}) \ \{ e \} \ \text{override}(\bar{M}, \bar{m}) \end{cases}$$

■ **Figure 4** Object inheritance extension.

the generated getter and setter methods for the fields. The rule also converts the fields in the object expression into a series of assignments, which ultimately result in the new reference. The internal field names are (globally) fresh to avoid having to worry about overridden names under inheritance. Both the methods and the assignments have the relevant qualifying substitutions applied, and the body of the object has **self** bound to the new reference.

Note that we could bind **self** in the bodies of an object's methods either in the Rule E-OBJECT (early binding, when the object is allocated) or in Rule E-REQUEST (late binding, when a method is requested). In this model, either choice produces the same behaviour. As we add inheritance to Graceless, our choice to use late binding in the base model will matter: it is trivial to overwrite by also using early binding, but not the other way around.

4 Object Inheritance

We now extend Graceless with various implementations of object inheritance, and consider the complexity and impacts of the changes. The extensions are presented in a rough ordering of implementation complexity. These represent the three foundational strands of object inheritance: forwarding, as used in E; delegation, as found in JavaScript [53], Lua [25], and Self [51, 10]; and concatenation, as in Kevo [48, 49] and numerous libraries and idioms for languages with open objects.

The extended syntax for object inheritance is given in Figure 4. The extension introduces two new components of user-facing syntax: the bodies of object expressions may now begin with an **inherit** e clause, and requests can now be qualified by the special variable **super**.

These two components each result in a runtime complication. This intermediate extension of Graceless already implements a form of object inheritance similar to concatenation with all fields reset to uninitialised, but we do not know of any language with this behaviour, nor can we see why it would be desirable.

Inheritance introduces the methods from the super-object into the local scope of the inheriting object using an ‘up, then out’ rule: inherited definitions take precedence over those introduced in a surrounding scope. Because the inherit clause contains an arbitrary expression, we might not know what the names are that the clause will introduce. To counter this, substitutions are delayed by an inherit clause in an object expression: while the substitution will transform the expression in the clause itself, it *will not* proceed into the body of the object expression, and gets ‘stuck’ on the clause instead. Once the expression in an inherit clause is resolved to an object reference, the substitution can proceed into the body of the surrounding object expression, where it may be removed by shadowing.

Although there is an explicit super-object in this model of inheritance, making a request to `super` is not the same as making a direct request to the super-object, as the value of `self` will be bound to the inheriting object. At runtime, the variable `super` is substituted for a special ‘up-call’ receiver (`ℓ as self`), which indicates the method to call should be sourced from the object at location `ℓ`, but `self` in the body of that method should be bound to the eventual value of `self` at the site of the request.

The evaluation context has not been extended as expected, to avoid issues in the future with constraints on what can be evaluated in an inherit clause. Evaluation in an inherit clause is instead handled by Rule E-INHERIT/CONTEXT, which is modified by future models. Rule E-INHERIT transforms an object expression with an inherit clause into one without, by copying the methods from the super-object which are not overridden by a method with the same name into the body of the inheriting object (directly copying the methods has the same behaviour as creating new methods which delegate to the super-object with the same arguments, or searching through a chain of objects for the method, which is representative of typical object inheritance implementations). It also applies the delayed substitutions to the body, after substituting `super` for an up-call receiver to the inherited object. Rule E-REQUEST/SUPER simply applies an up-call as described.

By making subtle modifications to the existing rules in this extended form of Graceless, we can produce models for various implementations of object inheritance.

4.1 Forwarding

Under forwarding, inherited methods are simply redirected to the super-object. The super-method receives the same arguments and `self` binding. In the example from earlier, if `amelia` receives a request for the `draw` method, which is not implemented directly inside of `amelia`, the request is passed on to the `graphic` super-object instead. The value of `self` in the resulting invocation is the identity of the super-object: in this example, the `draw` method crashes, complaining the `graphic` has not implemented its `image` method, because the local request to `image` has been resolved to the `graphic` object and not passed back down to `amelia`.

The modification to the existing Graceless dynamic semantics to implement forwarding is provided in Figure 5. The modification makes one subtle change (highlighted) to the Rule E-OBJECT, by early-binding the value of `self` when an object is created in both its methods and field accessors. The result of this change is that the late-binding of `self` in requests (both normal and to `super`) no longer achieves anything, because `self` has already been bound to the object that the method or field originally appeared in. Any forwarded request behaves as though the original object had received the request directly.

$$\begin{array}{c}
\text{(E-OBJECT/FORWARDING)} \\
\hline
\ell \text{ fresh} \quad \bar{m} = \text{names}(\bar{M}, \bar{S}) \quad \langle \bar{M}_f, \bar{e} \rangle = \text{body}(\bar{S}) \\
\langle \sigma, \text{object} \{ \bar{M} \bar{S} \} \rangle \rightsquigarrow \langle \sigma(\ell \mapsto \langle \emptyset, [\ell/\text{self}]([\text{self}.m/m]\bar{M} \bar{M}_f)), [\ell/\text{self}][\text{self}.m/m]\bar{e}; \ell \rangle
\end{array}
\quad \bar{m} \text{ unique}$$

■ **Figure 5** Forwarding modifications.

The value of **self** is always the object a method was defined in, so down-calls are not possible; similarly, the value of **self** during initialisation is the distinct identity of the super-object, so registration cannot occur then. In this model, an object can only inherit from another before it has run any of its own initialisation, so every object has a stable structure, but one object may be inherited many times. While this model does not permit multiple super-objects, it would only require updating Rule E-INHERIT to include methods from multiple inherited objects, with some arbitrary mechanism to resolve multiply-defined methods (such as placing significance on the order of the inherit clauses, or requiring that a multiply-defined method be overridden in the inheriting object).

There is no concept of confidential access of methods between implementations, as a forwarded request is sent as a regular request, and so will only be handled by the public interface of the super-object. Forwarding also does not permit downcalls: an inherited object cannot invoke an inheriting object’s method. On the other hand, an object can forward messages to a preëxisting object, and many forwarding objects can share a single target. Inherited fields are shared between all inheriting objects, and the mutation of an inherited field will implicitly affect the super-object and all of its heirs.

In the E language, there is no explicit self reference, as all object definitions are explicitly named. The authors of E refer to the language’s inheritance mechanism as “delegation”, but in the absence of self references the behaviour aligns with what we have called forwarding.

```

def graphic { to draw() { canvas.render(graphic.image()) } }
def amelia extends graphic { to image() { images.cat() } }

```

Even though `amelia` defines a method `image`, the call in `draw` clearly looks for the method in the `graphic` object. Methods cannot be requested on the local scope, so the receiver must always be explicit. In order to achieve down-calls, the inheriting object must explicitly be passed to the super-object, which is a standard pattern for simulating class behaviour in E.

4.2 Delegation

Delegation is an implementation of object inheritance that aimed to be at least as powerful as inheritance, if not more so [32, 51, 33]. Delegation has a subtle distinction from forwarding: a **self** request in a method called under delegation goes *back to the original object*, while under forwarding a **self** request to a delegatee will be handled only by that delegatee. This allows delegation to support down-calls, where forwarding cannot.

In the previous example, it was not safe to request the `draw` method on `amelia` under forwarding, as the implementation of `draw` expects to be able to see an overridden implementation of the `image` method. Under delegation, this now works as expected: `draw` executes with **self** as `amelia`, so the local request to `image` calls down into `amelia` and successfully retrieves the image of Amelia.

The modification to the existing Graceless dynamic semantics to implement delegation is provided in Figure 6. Like forwarding, the modification early-binds the value of `self` in the

$$\begin{array}{c}
\text{(E-OBJECT/DELEGATION)} \\
\frac{\ell \text{ fresh} \quad \bar{m} = \text{names}(\bar{M}, \bar{S}) \quad \langle \bar{M}_f, \bar{e} \rangle = \text{body}(\bar{S})}{\langle \sigma, \mathbf{object} \{ \bar{M} \ \bar{S} \} \rangle \rightsquigarrow \langle \sigma(\ell \mapsto \langle \emptyset, [\mathbf{self}.m/m] \bar{M} \ [\ell/\mathbf{self}] \bar{M}_f \rangle), [\ell/\mathbf{self}] [\mathbf{self}.m/m] \bar{e}; \ell \rangle} \bar{m} \text{ unique}
\end{array}$$

■ **Figure 6** Delegation modifications.

Rule E-OBJECT, but this time *only in the field accessors*. The value of `self` remains late-bound in the bodies of methods to the receiver of a method request, allowing super-objects to perform down-calls to methods implemented in a sub-object. Fields are shared between the original object and any inheriting objects, and, as under forwarding, mutation of any field which originated in a super-object is reflected in all of its heirs.

Delegation makes no requirement of freshness, which combined with down-calls produces a further complication to information hiding in a language with confidential methods that can only be called on `self`. These methods may provide access to secret data or capabilities, and the ability to access them from arbitrary code could be a security concern. Delegation to pre-existing objects opens the way for the ‘vampire’ problem: any object to which a reference has been obtained can be taken over and fully controlled from the outside, merely by defining a new child object. If access to confidential methods is instead *not* provided to delegators, simulating common classical patterns becomes difficult or impossible, and an odd asymmetry is introduced: the delegator can override an inherited method, changing its behaviour, but has no access to use the overridden method in its own implementation.

Unlike forwarding, delegation permits down-calls after the object has been constructed, but *not during initialisation*. Objects under delegation cannot perform registration during initialisation, as a captured `self` reference in a super-object refers to that super-object, which may have no other references and will not have access to any overriding definitions from the child. These two limitations of object initialisation are the major barriers to simulating the typical behaviour of class-based inheritance under delegation.

Object structure and behaviour is not stable during construction, as new methods may appear on `self` and existing methods may have different implementations depending on the stage of inheritance. Like forwarding, delegation permits inheritance from a pre-existing object; if this is allowed, stability does not exist after construction either. Delegation, like forwarding, can be easily extended to multiple inheritance by introducing multiple inherit clauses and some resolution of multiply-defined methods.

Including the `inherit` clause in the object constructor ensures that objects can delegate either to a pre-existing prototype object or construct a new object with custom arguments and delegate to that, as an equivalent to calling a super-constructor in a class system. This distinguishes it from the prototypical inheritance in JavaScript, which requires that the `prototype` property of the constructor be set before any inheriting object is constructed. In Self, object expressions can set their parents as a fresh object constructor call:

```
( | parent* = factory new. | )
```

This has the same semantics in terms of `self` binding in the super-constructor as presented in our model of delegation.

$$\begin{array}{c}
\text{(E-INHERIT/CONCATENATION)} \\
\frac{\langle \bar{x} \mapsto \bar{v}, \bar{M}_\uparrow \rangle = \sigma(\ell) \quad \bar{M}'_\uparrow = \text{override}(\bar{M}_\uparrow, \text{names}(\bar{M}, \bar{S}))}{\langle \sigma, \text{object } \{ \text{inherit } \ell \bar{s} \bar{M} \bar{S} \} \rangle \rightsquigarrow \langle \sigma, \text{object } \{ \bar{M}'_\uparrow [\bar{s}] [(\ell \text{ as self})/\text{super}] (\bar{M} \text{ self } \xleftarrow{\bar{x}} \bar{v} \bar{S}) \} \rangle}
\end{array}$$

■ **Figure 7** Concatenation modification.

4.3 Concatenation

Concatenation is an alternative approach to both forwarding and delegation that aimed to have the power of inheritance without the drawbacks of delegation [48, 49]. Under concatenation, one object inherits from another by (conceptually) taking a shallow copy of the methods and fields of its parent into itself, and then appending local overriding definitions. Concatenation supports down-calls, but unlike delegation does not allow subsequent changes in either the parent or the child to affect each other.

The modification to the existing Graceless dynamic semantics to implement concatenation is provided in Figure 7. Unlike the previous two models, this modification instead changes the Rule E-INHERIT, to copy the fields from the super-object into the inheriting object as assignments. The existing late-binding of `self` in methods is sufficient to provide the desired behaviour: any inherited method executes in the context of the inheriting object, and any request to an inherited field accessor will access the copied field in the inheriting object as well. The only relationship the inheriting object has with its super-object is explicit up-calls, but it is impossible to access or modify the state of the super-object without explicitly referring to it through an existing reference. The resulting behaviour is equivalent to delegating to a clone of the super-object.

Concatenation also makes no requirement of freshness. Concatenation with pre-existing objects does not quite permit the ‘vampirism’ of delegation, but does allow ‘mind reading’: any confidential state in an object can be read simply by inheriting from it, but the existing object cannot be manipulated by the child. Unlike the two previous models, mutations to inherited fields do not cause action at a distance, as the mutation will always affect a field in the actual receiver of the request (even for super-calls). With lexical scoping, the two objects are also not as independent as they seem: methods exist in the same scope in both objects, and any lexically captured state is shared between the two.

Like delegation, concatenation permits down-calls after the object has been constructed, but still not during construction. Concatenation does not allow registration, as a captured `self` reference in a super-object refers to the parent. Object structure and behaviour is not stable during construction; as for delegation, if inheritance from pre-existing objects is allowed then stability does not exist afterwards either. Concatenation can be straightforwardly extended to multiple inheritance by inserting the contents of each parent into the child, with some resolution of multiply-defined methods.

Objects with mutable structure can trivially implement concatenation, by manually copying the structure of the inherited object into the inheriting one. JavaScript complicates this story with unenumerable properties, implicit field accessors, and existing delegation relationships, but for the most part this is a valid implementation of concatenation:

```
for (var name in inherited) { inheriting[name] = inherited[name]; }
```

It is also possible to use mutable object structure to implement either forwarding or delegation, by assigning methods (or field accessors) to the inheriting object that directly forward or delegate to the inherited object, as in our models. In a JavaScript constructor:

$$\begin{array}{c}
\text{(E-INHERIT/CONTEXT)} \\
\frac{\langle \sigma, e \rangle \rightsquigarrow \langle \sigma', e' \rangle \quad e = v.m(\bar{v}) \implies e' = \bar{e}; o}{\langle \sigma, \mathbf{object} \{ \mathbf{inherit} \ e \ \bar{s} \ \bar{M} \ \bar{S} \} \rangle \rightsquigarrow \langle \sigma', \mathbf{object} \{ \mathbf{inherit} \ e' \ \bar{s} \ \bar{M} \ \bar{S} \} \rangle}
\end{array}$$

■ **Figure 8** Fresh inheritance modification.

```

var self = this;
this.foo = function () { self.bar(); };

```

If the `foo` method is called on a sub-object, the call to `bar` is guaranteed to not perform a down-call, because `self` is bound directly to the original object.

JavaScript’s built-in object inheritance otherwise works the same way as Graceless delegation, but field assignments are directly available in the language instead of only through accessor methods. The result is that an object shares each field of its inherited object, but when a field is assigned to it, the field is unique to that object, shadowing the inherited one. In any language where the objects have mutable structure it is necessary to retain a parent reference in order to implement delegation, in order to accurately defer to the current implementation of a parent object. Implementing more complicated forms of inheritance in JavaScript, such as with traits or mixins, tends to involve combining the built-in delegation alongside manual concatenation.

5 Emulating Classes

The inheritance models in the previous section represent the three foundational strands of purely object-based inheritance. Class-based languages tend to provide different semantics, and programmers and language designers may wish to use those behaviours, or may unconsciously expect object-based languages to behave similarly.

It is possible to construct object-based models that parallel many of these classical behaviours. In some languages with very flexible semantics, such as JavaScript and Lua, libraries exist to provide “classes” as a second-class construct by mutating objects or leveraging specially-constructed objects with the existing inheritance systems [46, 34]. The two models in this section approximate the inheritance behaviour of C++ (Section 5.1) and Java (Section 5.2) in an object-based system. They remain purely object-based, but trade off some of the flexibilities of the object inheritance models for the classical functionalities they provide, such as registration, down-calls, and stability.

The new models use as their base a further extension to Graceless inheritance, provided in Figure 8, which enforces that the inherited object is a ‘fresh’ reference. This extension modifies Rule E-INHERIT/CONTEXT so that method requests ready to be computed directly in an `inherit` clause may only be reduced if they syntactically return an object expression. This modification to the rule still permits the object constructor to then be resolved to a reference before handling the `inherit` clause.

The requirement of strict syntactic freshness is not required for merged identity, but it prevents potentially dangerous manipulation of object identity from outside of an object’s creator. Merged identity builds directly on the new extension, while uniform identity further modifies the context rule to also prevent an object expression from being computed in an `inherit` clause.

$$\begin{array}{c}
\text{(E-INHERIT/MERGED)} \\
\frac{\langle \overline{F}, \overline{M}_\uparrow \rangle = \sigma(\ell) \quad \ell_\uparrow \text{ fresh} \quad \overline{m}_\downarrow = \text{names}(\overline{M}, \overline{S})}{\overline{M}'_\uparrow = \text{override}(\overline{M}_\uparrow, \overline{m}_\downarrow) \quad \overline{m} = \overline{m}_\downarrow \cup \text{names}(\overline{M}'_\uparrow, \emptyset) \quad \langle \overline{M}_f, \overline{e} \rangle = \text{body}(\overline{S}) \quad \overline{m} \text{ unique}}{\langle \sigma, \text{object} \{ \text{inherit } \ell \ \overline{M} \ \overline{S} \} \rangle \rightsquigarrow} \\
\langle \sigma(\ell_\uparrow \mapsto \langle \emptyset, \overline{M}_\uparrow \rangle)(\ell \mapsto \langle \overline{F}, \overline{M}'_\uparrow [s][\text{self.m/m}][(\ell_\uparrow \text{ as self})/\text{super}]\overline{M} \ \overline{M}_f \rangle), \\
[s][\ell/\text{self}][\text{self.m/m}][(\ell_\uparrow \text{ as self})/\text{super}]\overline{e}; \ell \rangle
\end{array}$$

■ **Figure 9** Merged identity modification.

5.1 Merged Identity

In this model, an inheriting object takes over the identity of its parent, ‘becoming’ that object but putting in place all of its own method definitions, in effect the reverse of concatenation. The parent object is constructed and initialised before the child, and then mutated at the point of inheritance. Rather than mutating preëxisting objects, inheritance must now occur from a fresh object, one newly created and immediately returned from a method call. Without the requirement of freshness, objects can directly steal the identity of any existing object: the resulting ‘body-snatchers’ problem is substantially worse than delegation’s vampirism, which can only control objects internally. Overridden methods from the parent remain accessible through **super**, and will always execute with the final identity of the object.

Merged identity provides essentially the C++ inheritance behaviour: the apparent type of the object changes during construction, as each layer of inheritance is processed. After object construction is complete, down-calls resolve to their final overridden method, but until then they obtain the most recent definition from ((great-)grand)parent to child.

The use of **self** for registration in **graphic** will now correctly store the value of **amelia** when initialising the super-object. Although the object is *not amelia* when it is registered, **amelia becomes** the object that was registered once it has finished initialising, merging the two identities together. The initialisation of **amelia** will still fail at the local request of **draw**, because **amelia**’s overriding of the abstract method has not yet been merged into the identity of the object.

The modification to **Graceless** with fresh inheritance to implement merged identity is provided in Figure 9. The Rule E-INHERIT/MERGED overwrites the existing Rule E-INHERIT. Where the old rule resulted in a new object expression with some of the methods from the super-object, this new rule skips straight to returning the resulting reference. This is necessary because the resulting reference is not fresh: it is the reference of the inherited object. The inherited object is updated with the new methods in the store, including keeping all of the old field values, but removing overridden methods.

The new rule is a combination of the behaviour of the old E-INHERIT and Rule E-OBJECT. We have **highlighted** the changes to their combined behaviour. There is no longer an object which corresponds to **super** to perform an up-call substitution on the body of the object as in Rule E-INHERIT, so rather than creating a fresh location for the resulting object, we create a fresh location for the super-object instead. The super-object has all of the methods of the inherited object before it was modified, while the merged object retains its own fields instead of beginning with an empty field store. All of the models which emulate classes with super-references create these ‘part-objects’, which exist purely to store the super-methods for super calls, never have any fields of their own, and are only ever referenced under an alias for the ‘real’ object that is bottom-most in the hierarchy.

Like delegation, merged identity enables down-calls after, but not during, initialisation. Unlike delegation, a single identity is preserved throughout the construction process, so registration is possible. It does not provide classical stability during initialisation, but once objects are complete they cannot change again. The meta-mutation of merging objects means that it is unsafe to permit inheritance from arbitrary objects, so only freshly-constructed objects are valid parents. Merged identity is really the only model presented in this paper that does not lend itself well to multiple inheritance, because only a single parent can have its identity preserved.

As a reverse form of concatenation, merged identity is just as easy to implement for objects with mutable structure. A constructor can call into the super-constructor, then concatenate its own definitions into the resulting object. Copying the original object beforehand ensures that a super-reference is still available.

5.2 Uniform Identity

The uniform identity design is a direct attempt to match as closely as possible the behaviour of a typical class-based inheritance system, but based on objects rather than on classes. In this design, the *first* observable action is to create a new object identity in the bottom-most ‘child’ object constructor; this exactly mirrors the merged identity design, which creates a single identity of the *topmost* parent. All inherited declarations are assembled in a single object associated with this identity, but no fields are initialised and no inline code run until the direct inheritance completes. Inheritance occurs with the identity “passed along”: declarations are attached to the original object, without initialisation, until the topmost object with no parent is reached. Finally, the initialisation code runs from top to bottom.

All initialisation occurs in the context of the final object, including its behaviour. No new methods or overrides are added visibly during construction. While initialisation code always sees objects as a consistent type, it may observe uninitialised fields, including constant fields. We model uninitialised field access as a fatal error, and leave static detection to a higher layer in the language. This model essentially aligns with the semantics of Java-like languages.

The behaviour of *amelia* is the same as in the merged identity semantics, except that the local request to *draw* in the initialisation of *graphic* now successfully down-calls into *amelia*’s implementation, as the object *is amelia* during all initialisation in the hierarchy. The result is still an error, as *amelia*’s initialisation of the *image* field has not yet been executed. Super-constructors may still encounter uninitialised fields that will be assigned to in some sub-constructor, particularly when methods can be overridden by fields.

The modification to Graceless with fresh inheritance to implement uniform identity is provided in Figure 10. As with merged identity, the Rule E-INHERIT/UNIFORM overwrites the existing Rule E-INHERIT. The Rule E-INHERIT/CONTEXT has been further refined to also prevent the evaluation of object expressions directly inside an inherit clause.

Where Rule E-INHERIT/MERGED was a modified application of E-INHERIT and then E-OBJECT, the uniform identity modification applies this in reverse. As the body of the inherit clause is an object expression, we have to apply many of the same processes for reducing a regular object expression, but then move the sequence of expressions resulting from the body of the super-object down into the body of the inheriting object. The evaluation *does* create a new object reference for the super-object, but only for up-calls to *super*. All fields will end up in the ultimate inheriting object.

Because the inherited object expression has not been processed into a runtime object, the rule needs to manually build its field methods and body. The substitutions to local definitions \bar{m} need to be applied to the inherited methods (both those in the store, and the

$$\begin{array}{c}
\text{(E-INHERIT/CONTEXT)} \\
\frac{\langle \sigma, e \rangle \rightsquigarrow \langle \sigma', e' \rangle \quad e = v.m(\bar{v}) \implies e' = \bar{e}; o \quad e \neq o}{\langle \sigma, \mathbf{object} \{ \mathbf{inherit} \ e \ \bar{s} \ \bar{M} \ \bar{S} \} \rangle \rightsquigarrow \langle \sigma', \mathbf{object} \{ \mathbf{inherit} \ e' \ \bar{s} \ \bar{M} \ \bar{S} \} \rangle} \\
\text{(E-INHERIT/UNIFORM)} \\
\frac{\ell \text{ fresh} \quad \bar{m} = \text{names}(\bar{M}, \bar{S}) \quad \bar{M}_{\uparrow} = [\mathbf{self}.m/m]\bar{M} \quad \langle \bar{M}_f, \bar{e} \rangle = \text{body}(\bar{S}) \quad \bar{m}_{\downarrow} = \text{names}(\bar{M}_{\downarrow}, \bar{S}_{\downarrow}) \quad \bar{M}'_{\uparrow} = \text{override}(\bar{M}_{\uparrow} \ \bar{M}_f, \bar{m}_{\downarrow}) \quad \bar{m} \text{ unique}}{\langle \sigma, \mathbf{object} \{ \mathbf{inherit} \ \mathbf{object} \{ \bar{M} \ \bar{S} \} \ \bar{s} \ \bar{M}_{\downarrow} \ \bar{S}_{\downarrow} \} \rangle \rightsquigarrow \langle \sigma(\ell \mapsto \langle \emptyset, \bar{M}_{\uparrow} \ \bar{M}_f \rangle), \mathbf{object} \{ \bar{M}'_{\uparrow} \ [s][(\ell \ \mathbf{as} \ \mathbf{self})/\mathbf{super}]\bar{M}_{\downarrow} \ [\mathbf{self}.m/m]\bar{e} \ [s][(\ell \ \mathbf{as} \ \mathbf{self})/\mathbf{super}]\bar{S}_{\downarrow} \} \rangle}
\end{array}$$

■ **Figure 10** Uniform identity inheritance modification.

non-overridden methods in the inheriting object). The same substitution into the body of the inheriting object occurs as before, but now `self` is not bound in the body: `self` will be bound by the ultimate application of Rule E-OBJECT. The substitutions applied simultaneously to M_{\downarrow} and S_{\downarrow} in the inheriting object by Rule E-INHERIT now have to be split between the two, as the super-body e appears in between, but it amounts to the same behaviour.

Uniform identity permits down-calls both after and during construction, as well as registration, as the identity and structure of the object are both constant, also guaranteeing stability. It does not allow inheriting from preexisting objects, instead requiring that parents be fresh. It does not support multiple inheritance, but there is a logical extension to do so.

Constructors have a special role in JavaScript, and must have their `prototype` property set before they are used to construct a new object. As a result, inheriting from a preexisting prototype object is simple, but inheriting from another constructor, particularly one that requires arguments specific to a particular inheriting object, is more difficult. `Object.create` allows the creation of an object from a prototype without actually invoking the constructor:

```
function Bar(arg) { Foo.call(this, arg); }
Bar.prototype = Object.create(Foo.prototype);
```

The ability to bind the value of `this` in a function call allows JavaScript to implement uniform identity, ensuring that the initialisation code in a super-constructor can be executed in the context of the inheriting object instead of creating a fresh object and inheriting from that. JavaScript's class syntax, introduced by ECMAScript 2015, is just sugar for this behaviour.

The most natural forms of inheritance in JavaScript are the built-in single delegation and, by extension, single uniform inheritance, now codified directly in the language with the class syntax. Our primary concern is that without the particular dynamism of the necessary JavaScript features, simulating class initialisation is impossible without reinterpreting factory methods as constructors with special semantics, at which point the language has arguably just implemented classes.

6 Multiple Inheritance

Reusing behaviour from multiple parents is a widespread desire, but is less commonly supported in language designs in practice. In this section we show three logical extensions enabling multiple object inheritance. The first extends the uniform identity model with multiple `inherit` statements. The second is a separate model following the tradition of trait systems, where method names are unique in any object. The third extends any of the

base models with the ability to include multiple **inherit** statements in an object, processed imperatively. All of these extensions process **inherit** clauses as under uniform identity, but they could all be simplified to perform inheritance from preexisting objects as per forwarding, delegation, or concatenation. We have demonstrated this by constructing all possible combinations in our Redex implementation.

All systems enable code reuse from arbitrarily-many parents. All parents are treated symmetrically in the first two models, but further restrictions or privileges could be accorded to some parents without substantially affecting the models. In order to handle conflicts in symmetric inheritance, methods can be **abstract** in their body, which causes an error if that method is called. Alternatively, the implementations could ban the construction of an object with abstract methods, but this would only be valid for uniform identity, as the overriding of a concrete implementation would not take effect until the object was inherited under any of the other interpretations.

6.1 Multiple Uniform

The multiple uniform model allows a sequence of **inherit** statements to appear at the start of an object body. Each statement includes an **as name** clause, which binds **name** locally to have **super** semantics with regard to that inheritance tree. The single-inheritance uniform identity model is recovered with **inherit parent as super**. When the same method name is inherited from multiple parents, none has priority and an abstract method of that name is inserted instead. The programmer must provide a local override calling the version from a particular named parent if desired. All methods are collected and installed before any initialisation code from the object bodies executes, so a consistent set of method implementations is seen throughout the initialisation.

The implementation of multiple uniform identity is presented in Figure 11, as an extension to Graceless inheritance with the caveat that **super** is no longer a special form of receiver. The Rule E-INHERIT/CONTEXT implements fresh inheritance in objects with potentially multiple **inherit** clauses, ensuring each clause is evaluated in order.

Rule E-INHERIT/MULTIPLE is essentially the same as Rule E-INHERIT/UNIFORM, but processing multiple **inherit** clauses at once (hence the extra multiplicities for many of the bindings). Once all of the **super**-methods are collected, conflicting methods are resolved with the join auxiliary function that, for each unique method name in the collection of methods, accepts exactly one concrete implementation of a method with that name and removes all of the abstract implementations, or provides a single abstract method with that name and removes all other implementations.

Multiple uniform supports registration and downcalls exactly as in uniform identity. It supports multiple inheritance from freshly-created parents, and is stable because methods are collected first. These are all properties of uniform inheritance: applying this modification to the simpler object inheritance models each retains their own particular properties as well.

6.2 Method Transformations

Multiple inheritance under method transformations resembles trait-like composition of objects, representing object values as a single mapping of method names to methods, with no equivalent to **super** in the previous designs. Instead, an **inherit** statement can have any number of **alias** or **exclude** clauses associated, which respectively create an alternative name for an inherited method and exclude its implementation. If a method is overridden locally and still needs to be accessed, the inherited method can be aliased to a different name and

13:16 Object Inheritance Without Classes

Extended Syntax

$I ::= \text{inherit } e \text{ as } x$ (Inherit clause) $o ::= \text{object } \{\bar{I} \boxed{\bar{s}} \bar{M} \bar{S}\}$ (Object expression)

$e ::= \dots \mid \text{abstract}$ (Expression) $s ::= \dots \mid (\ell \text{ as self})/x$ (Substitution)

$I^\circ ::= \text{inherit object } \{\bar{M} \bar{S}\} \text{ as } x$ (Evaluated inherit clause)

$$\begin{array}{c}
 \text{(E-INHERIT/CONTEXT)} \\
 \frac{\langle \sigma, e \rangle \rightsquigarrow \langle \sigma', e' \rangle \quad e = v.m(\bar{v}) \implies e' = \bar{e}; o \quad e \neq o}{\langle \sigma, \text{object } \{\bar{I}^\circ \text{ inherit } e \bar{I} \bar{M} \bar{S}\} \rangle \rightsquigarrow \langle \sigma', \text{object } \{\bar{I}^\circ \text{ inherit } e' \bar{I} \bar{M} \bar{S}\} \rangle} \\
 \\
 \text{(E-INHERIT/MULTIPLE)} \\
 \frac{\ell \text{ fresh} \quad \bar{m} = \text{names}(\bar{M}, \bar{S}) \quad \bar{M}' = \overline{[\text{self}.m/m]\bar{M}} \quad \langle \bar{M}_f, \bar{e} \rangle = \text{body}(\bar{S})}{\frac{\bar{M}_\uparrow = \text{join}(\bar{M}', \bar{M}_f) \quad \bar{m}_\downarrow = \text{names}(\bar{M}_\downarrow, \bar{S}_\downarrow) \quad \bar{M}'_\uparrow = \text{override}(\bar{M}_\uparrow, \bar{m}_\downarrow)}{\bar{m} \text{ unique}}}{\langle \sigma, \text{object } \{\text{inherit object } \{\bar{M} \bar{S}\} \text{ as } x \bar{s} \bar{M}_\downarrow \bar{S}_\downarrow\} \rangle \rightsquigarrow \langle \sigma(\ell \mapsto \langle \emptyset, \bar{M}' \bar{M}_f \rangle), \\ \text{object } \{\bar{M}'_\uparrow [s][(\ell \text{ as self})/x]\bar{M}_\downarrow \overline{[\text{self}.m/m]\bar{e}} [s][(\ell \text{ as self})/x]\bar{S}_\downarrow\} \rangle}
 \end{array}$$

Auxiliary Definitions

$$\begin{array}{l}
 \text{join}(\emptyset) = \emptyset \\
 \text{join}(\text{method } m(\bar{x}) \{e\} \bar{M}) = \begin{cases} m \notin \text{names}(\bar{M}, \emptyset) & \text{method } m(\bar{x}) \{e\} \text{join}(\bar{M}) \\ e \equiv \text{abstract} & \text{join}(\bar{M}) \\ \text{method } m(\bar{y}) \{\text{abstract}\} \in \bar{M} & \text{join}(\bar{M} \text{ method } m(\bar{x}) \{e\}) \\ \text{otherwise} & \text{method } m \{\text{abstract}\} \text{join}(\text{override}(\bar{M}, m)) \end{cases}
 \end{array}$$

■ **Figure 11** Multiple uniform identity modification.

accessed through that name within the object. An object contains at most one method with any given name, and there are no part-objects.

Multiple **inherit** statements can appear in an object, treated symmetrically. If the same name is inherited from multiple parents, all but one must be **excluded**, or a local overriding method declared, if a concrete implementation of that method is to appear in the object. All inheritance expressions are evaluated and the final method set of the object assembled before any initialisation code executes. Initialisation occurs from top to bottom (depth-first search) within each branch of the inheritance hierarchy. Methods are decoupled from their names because aliases may provide multiple equivalent names all reaching the method, while exclusion means that local definitions may not be implemented in the final object.

If **amelia** wished to simultaneously be a **graphic** and a **gunslinger**, then two inherit clauses can be included under the transformation model. If the **gunslinger** class also contained a **draw** method for drawing her gun, **amelia** would be required to choose a single implementation by excluding one of the two. The excluded method can still be accessed if it is aliased:

```

def amelia = object {
  inherit gunslinger
  inherit graphic alias draw as render exclude draw
  def image = images.amelia
  var name := "Amelia"
}

```

Even if the **gunslinger** class also has a **name** method, **amelia** has successfully combined the two by overriding both. Note that method resolution is consistent throughout the entire object:

Extended Syntax

$$I ::= \mathbf{inherit} \ e \ \mathbf{alias} \ \overline{m} \ \mathbf{as} \ \overline{m} \ \mathbf{exclude} \ \overline{m} \quad (\text{Inherit clause})$$

$$\begin{array}{c}
 \text{(E-INHERIT/TRANSFORM)} \\
 \frac{\ell \text{ fresh} \quad \overline{m} = \text{names}(\overline{M}, \overline{S}) \quad \overline{M}' = [\text{self.m/m}]\overline{M} \quad \langle \overline{M}_f, \overline{e} \rangle = \text{body}(\overline{S})}{\overline{M}_a = \text{aliases}(\overline{m}_a \ \mathbf{as} \ \overline{m}'_a, \overline{M}' \ \overline{M}_f) \quad \overline{M}_e = \text{excludes}(\overline{m}_e, \overline{M}_a) \quad \overline{M}_\uparrow = \text{join}(\overline{M}_e)}{\overline{m}_\downarrow = \text{names}(\overline{M}_\downarrow, \overline{S}_\downarrow) \quad \overline{M}'_\uparrow = \text{override}(\overline{M}_\uparrow, \overline{m}_\downarrow) \quad \overline{m} \text{ unique}} \\
 \frac{\langle \sigma, \mathbf{object} \{ \mathbf{inherit} \ \mathbf{object} \{ \overline{M} \ \overline{S} \} \ \mathbf{alias} \ \overline{m}_a \ \mathbf{as} \ \overline{m}'_a \ \mathbf{exclude} \ \overline{m}_e \ \overline{s} \ \overline{M}_\downarrow \ \overline{S}_\downarrow \} \rangle \rightsquigarrow}{\langle \sigma, \mathbf{object} \{ \overline{M}'_\uparrow \ [\overline{s}]\overline{M}_\downarrow \ [\text{self.m/m}]\overline{e} \ [\overline{s}]\overline{S}_\downarrow \} \rangle}
 \end{array}$$

Auxiliary Definitions

$$\begin{aligned}
 \text{aliases}(\emptyset, \overline{M}) &= \overline{M} & \text{aliases}(m \ \mathbf{as} \ m' \ \overline{m} \ \mathbf{as} \ m', \overline{M}) &= \text{aliases}(\overline{m} \ \mathbf{as} \ m', \text{alias}(\overline{M}, m \ \mathbf{as} \ m')) \\
 \text{excludes}(\emptyset, \overline{M}) &= \overline{M} & \text{excludes}(m \ \overline{m}, \overline{M}) &= \text{excludes}(\overline{m}, \text{exclude}(\overline{M}, m)) \\
 \text{alias}(\emptyset, m \ \mathbf{as} \ m') &= \emptyset \\
 \text{alias}(\mathbf{method} \ m(\overline{x}) \{ e \} \ \overline{M}, m \ \mathbf{as} \ m') &= \mathbf{method} \ m(\overline{x}) \{ e \} \ \mathbf{method} \ m'(\overline{x}) \{ e \} \ \text{alias}(\overline{M}, m \ \mathbf{as} \ m') \\
 \text{alias}(M \ \overline{M}, m \ \mathbf{as} \ m') &= M \ \text{alias}(\overline{M}, m \ \mathbf{as} \ m') \\
 \text{exclude}(\emptyset, m) &= \emptyset \\
 \text{exclude}(\mathbf{method} \ m(\overline{x}) \{ e \} \ \overline{M}, m) &= \mathbf{method} \ m \ \{ \mathbf{abstract} \} \ \text{exclude}(\overline{M}, m) \\
 \text{exclude}(M \ \overline{M}, m) &= M \ \text{exclude}(\overline{M}, m)
 \end{aligned}$$

■ **Figure 12** Method transformation modification, with context rule omitted.

the local request to draw in the graphic initialisation will unholster instead of rendering.

The implementation of method transformation multiple uniform inheritance is given in Figure 12 as a modification of the previous uniform multiple inheritance extension. Inherit clauses no longer have super-names, using the method aliasing and excluding syntax instead. The obvious modifications to the context rule to account for method transformations instead of super-names are omitted.

Rule E-INHERIT/TRANSFORM pre-processes the methods of each inherited object before passing them to join, with the aliases and excludes auxiliary functions. Both of these functions proceeds as expected: fold the transformations over the methods, applying the alias or exclude rules in order. Because these rules are ordered, it is possible to create an alias of an existing alias that occurs earlier in the list, and it is possible to exclude aliases. Exclusion is always processed after aliases, so a directly excluded method cannot be aliased.

Uniform method transformation permits both down-calls and registration. It provides stability through time, but not through local analysis of visible declarations. It imposes the same freshness requirements as the other models, and supports multiple inheritance. Applying method transformation to the simpler object inheritance models continues to maintain their own particular properties, as they were never stable to begin with.

6.3 Positional

The previous multiple inheritance models treated the inherit clauses as symmetric, such that no definition in any one was preferred in the case of conflicts, requiring resolution either by

overriding or method transformations. Moreover, they did not permit any initialisation until all of the direct inheritance was completed. This can be a problem if, for instance, fields need to be initialised for down-calls in super-initialisation code, as with `amelia`'s `image` field. Positional inheritance addresses these concerns, at the cost of visibly mutating the object during construction, similarly to what can be achieved with mutable object structure.

Under positional inheritance, multiple **inherit** clauses can appear in an object, amongst the typical statements. These inheritances are processed imperatively where they appear, in order, using the same semantics as the selected base model of inheritance. As under multiple uniform, each **inherit** statement can have a name associated. Positional could also be used for single inheritance, allowing some initialisation before the super-constructor is called, but the visible mutation is still present, and the distinct ordering of the **inherit** clauses in an object body implies a natural method conflict resolution.

In the most general version, an **inherit** clause can appear anywhere in the object body, and have other code before, after, and in between, with the semantics of the inheritance taking effect at the point of appearance and later parents having precedence over earlier. This allows some interesting programmer choices with some of the base models. Altering the order of parents affects which versions of same-named methods are accessed, and interleaving other code in between exposes both at different times. The availability and safety of upcalls and down-calls are affected by the placement of the inheritance, field initialisation, and other code. A more restrained approach could limit inheritance to appearing all at the top (or at the bottom) of the object body.

Positional delegation with named **supers** is essentially the behaviour of `Self` [10], where multiple parent pointers may exist in a single object; `Self` does not allow initialisation code to execute in the context of the object under construction or have any priority between parents, but does allow parent pointers to be mutable. The nature of concatenation fundamentally supports positional multiple inheritance, simply copying in the contents of the inherited object in place of the **inherit** statement, and the limitation in the base model was purely syntactic. Named **supers** (or a next-method functionality) are also necessary to access overridden methods. Multiple forwarding is quite straightforward, and strictly named **supers** are not required: because forwarding only accesses the public interface, an ordinary reference to each parent suffices. One of the primary use-cases of positional inheritance — initialising fields before invoking a super-constructor — is irrelevant without uniform identity, as the super-constructor cannot make a down-call into the inheriting object anyway.

Merged identity does not lend itself to the positional extension because it relies on taking over the identity of the parent object, which with multiple parents would result in repeated identity changes, some of which may even lose methods. A different extension could merge multiple identities together, or resolve the resulting issues in some other way, but we do not address this combination here and simply exclude it from consideration as confusing at best.

Positional inheritance is the only one of our multiple-inheritance models that permits inheriting from something obtained from a parent. A parent could define a number of specialised “inner classes”, with the intention that its child would in turn inherit from one of those specialisations as well. It is not obvious to us that such an ability is useful, but nor is it obvious that it is not. We note this unique ability, but do not pursue it further. The most interesting version of positional inheritance is based on uniform identity, so we will focus on that extension for the remainder of this section. Each **inherit** is processed in order of appearance, using the same structure-then-initialisation process as under uniform identity. Again, a single identity exists for all initialisation, and again down-calls are valid throughout.

Extended Syntax

$$\begin{aligned}
e & ::= \dots \boxed{\text{super inherit } e \text{ as } x \bar{s} \mid \bar{i} \text{ inherit } e \text{ as } x \bar{s}} && \text{(Expression)} \\
S & ::= \dots \mid \text{inherit } e \text{ as } x && \text{(Statement)} \quad r ::= \dots \boxed{\mid (\ell \text{ as } \ell)} && \text{(Receiver)} \\
o & ::= \text{object } \{ \boxed{\bar{s}} \overline{M} \overline{S} \} && \text{(Object expression)} \quad i ::= \langle \ell, \overline{M}, \bar{s} \rangle && \text{(Inherit context)} \\
s & ::= \dots \mid (\ell \text{ as } \ell)/x \mid \bar{i}/\text{super} && \text{(Substitution)}
\end{aligned}$$

$$\begin{array}{c}
\text{(E-OBJECT/POSITIONAL)} \\
\hline
\ell \text{ fresh} \quad \overline{m} = \text{names}(\overline{M}, \overline{S}) \quad \langle \overline{M}_f, \overline{e} \rangle = \text{body}(\overline{S}) \quad \overline{m} \text{ unique} \\
\hline
\langle \sigma, \text{object } \{ \boxed{\bar{s}} \overline{M} \overline{S} \} \rangle \rightsquigarrow \\
\langle \sigma(\ell \mapsto \langle \emptyset, \boxed{\bar{s}}[\text{self.m/m}]\overline{M} \overline{M}_f \rangle), \boxed{\bar{s}}[\langle \ell, \overline{M} \overline{M}_f, \bar{s} \rangle/\text{super}][\ell/\text{self}][\text{self.m/m}]\overline{e}; \ell \rangle
\end{array}$$

$$\begin{array}{c}
\text{(E-INHERIT/POSITIONAL)} \\
\hline
\ell \text{ fresh} \quad \overline{m} = \text{names}(\overline{M}, \overline{S}) \quad \overline{M}_\uparrow = \boxed{\bar{s}_\uparrow}[\text{self.m/m}]\overline{M} \\
\langle \overline{M}_f, \overline{e}_\uparrow \rangle = \text{body}(\overline{S}) \quad \ell_\downarrow = \text{first}(\text{last}(i)) \quad i' = \text{add-subst}((\ell \text{ as } \ell_\downarrow)/x, \bar{i}) \quad \overline{m} \text{ unique} \\
\hline
\langle \sigma, \bar{i} \text{ inherit object } \{ \boxed{\bar{s}_\uparrow} \overline{M} \overline{S} \} \text{ as } x \bar{s}; e \rangle \rightsquigarrow \langle \text{update}(\sigma(\ell \mapsto \langle \emptyset, \overline{M}_\uparrow \overline{M}_f \rangle), \overline{M}_\uparrow \overline{M}_f, \bar{i}'), \\
\boxed{\bar{s}_\uparrow}[\langle \ell, \overline{M} \overline{M}_f, \bar{s}_\uparrow \rangle \bar{i}'/\text{super}][\ell_\downarrow/\text{self}][\text{self.m/m}]\overline{e}_\uparrow; \boxed{\bar{s}}[\text{self.m/m}][i'/\text{super}][(\ell \text{ as } \ell_\downarrow)/x]e \rangle
\end{array}$$

Extended Auxiliary Definitions

$$\text{body}(\text{inherit } e \text{ as } x \overline{S}) = \langle \overline{M}, \text{super inherit } e \text{ as } x \overline{e} \rangle \quad \text{where } \langle \overline{M}, \overline{e} \rangle = \text{body}(\overline{S})$$

$$\text{add-subst}(s, \langle \ell, \overline{M}, \bar{s} \rangle \bar{i}) = \langle \ell, \overline{M}, \bar{s} s \rangle \bar{i}$$

$$\begin{aligned}
& \text{update}(\sigma, \overline{M}_\uparrow, \emptyset) = \sigma \\
& \text{update}(\sigma, \overline{M}_\uparrow, \langle \ell, \overline{M}, \bar{s} \rangle \bar{i}) = \text{update}(\sigma(\ell \mapsto \langle \overline{F}, \overline{M}'_\uparrow \overline{M}' \overline{M}'_\downarrow \rangle), \overline{M}'_\uparrow \overline{M}' \overline{M}'_\downarrow, \bar{i}) \\
& \text{where } \overline{m} = \text{names}(\overline{M}, \emptyset) \text{ and } \overline{m}_\uparrow = \text{names}(\overline{M}_\uparrow, \emptyset) \text{ and } \overline{M}'_\uparrow = \text{override}(\overline{M}_\uparrow, \overline{m}) \text{ and } \langle \overline{F}, \overline{M}'_\downarrow \rangle = \sigma(\ell) \\
& \text{and } \overline{m}_\downarrow = \text{names}(\overline{M}'_\downarrow, \emptyset) \text{ and } \overline{M}' = \boxed{\bar{s}}[\text{self.m}_\uparrow/\text{m}_\uparrow][\text{self.m}_\downarrow/\text{m}_\downarrow]\overline{M} \text{ and } \overline{M}'_\downarrow = \text{override}(\overline{M}'_\downarrow, \overline{m} \overline{m}_\downarrow)
\end{aligned}$$

■ **Figure 13** Positional uniform identity modification, with context rule omitted.

Positional inheritance reintroduces mutation during construction to uniform identity, because each **inherit** adds new methods to the object. When multiple methods are inherited by the same name, the last-inherited method wins out. An unusual aspect is that while down-calls are always available, during construction ‘side-calls’ to co-parents of a common child can be made only to parents whose **inherit** preceded this one. An object can even define a fresh constructor directly inside of itself, and then inherit from it.

Each line of initialisation occurs after preceding inheritance statements and before subsequent inheritance statements. If **inherit** precedes a field initialisation or other expression, upcalls to that parent are available from that expression; if **inherit** follows a field initialisation, down-calls from that parent accessing that field are safe.

The implementation of positional uniform identity inheritance is provided in Figure 13 as an extension to the core Graceless language; the rule for fresh inheritance is omitted again. The primary difficulty with implementing positional inheritance is that the meaning of local definitions can change imperatively: a local request might refer to some definition in the surrounding scope, but after processing an inherit clause during the initialisation

phase that request might now refer to an inherited definition instead. This presents a particular difficulty for substitution, which irreversibly binds an unqualified name to a particular definition. Substitutions are now delayed by all object expressions, as the inherit clauses are now nested in the object statements. Rule E-OBJECT/POSITIONAL replaces the Rule E-OBJECT, handling the new statement form with the updated body translation and applying the delayed substitution.

In order to implement the dynamic scoping, we introduce an *inherit context* i , which records the reference ℓ of the object that an inherit clause appeared inside, the source of the methods \overline{M} that appeared directly in that object, and the delayed substitutions \overline{s} that were on that object. Any processed inherit clause has an ordered stack of these contexts on it, from the actual object the inherit clause appeared in, down to the bottom-most inheriting object. The `super` prefix is used as a placeholder on newly created inherit expressions, so that Rule E-OBJECT/POSITIONAL can substitute it out for the initial inherit context.

By retaining the source of the methods and the substitution scope they appeared in, the methods can be re-substituted in any new scope that appears. The update auxiliary function applies this for any newly inherited methods \overline{M}_\uparrow to every object in the current stack of contexts. Rule E-INHERIT/POSITIONAL handles any inherit expression, constructing a new part-object ℓ , and using update to include the new methods in the original object and every intervening part-object, after applying overrides from the existing methods. The value of `self` is bound in the inherited object body by the last object reference in the inherit context, as that is the location of the original object. The inherited body is processed in the same way as in Rule E-OBJECT/POSITIONAL. One of the key distinctions here is that the value of `self` to substitute in the inherited body already exists, whereas under single uniform identity the body is concatenated into an object expression and `self` is substituted once that is evaluated.

Note that inherit expressions still delay substitutions, preventing them from applying to expressions later in any sequence. After each inherit expression is evaluated, the substitutions are then applied to the following expressions, after being shadowed by inherited definitions and the `super`-name defined by the `as` clause. Positional uniform inheritance preserves the other traits of uniform identity from Section 5.2, with the exception of stability: during construction, an object's apparent structure and behaviour can change. Applied to the simpler object inheritance models, it preserves each of their unique properties.

7 Discussion

Table 1 compares the models according to the criteria established in Section 2. Each model provides a different mix of the criteria, which may be appropriate for different circumstances or languages. The uniform identity design provides the closest match to Java semantics (given at the bottom of the table) while the other multiple-inheritance models following trade off one or more of these properties. No model provides every property; indeed, stability, downcalls, and inheriting from existing objects are fundamentally in conflict. As well, the complexity of each design and its implementation roughly increases down the table, which is a further trade-off for language designers to consider.

While delegation, forwarding, and concatenation can fundamentally support inheriting from arbitrary objects, the other models lean towards supporting planned reuse rather than ad-hoc reuse — that is, inheriting from objects that have been designed to be inherited from, rather than from any arbitrary object. Both planned and unplanned reuse have solid software-engineering motivations; indeed, language features exist both specifically to prevent inheritance (final or sealed classes) and to enable ad-hoc reuse (structural types).

■ **Table 1** Comparison of models of object-first inheritance. * indicates answer holds during construction, but is reversed after. Overload indicates multiple definitions of the same method name in an object, accessible through super-references. Parent indicates ability to inherit from something obtained through another parent. All other columns relate to criteria from Section 2.

	Reg.	Down.	Dist.	Stable	Exist.	Mult.	Overl.	Par.
Forwarding	no	no	yes	yes	yes	no	yes	no
Delegation	no	no*	yes	no	yes	no	yes	no
Concatenation	no	no*	no	no	yes	no	yes	no
Merged	yes	no*	no	no*	fresh	no	yes	no
Uniform	yes	yes	no	yes	fresh	no	yes	no
Mult. Uniform	yes	yes	no	yes	fresh	yes	yes	no
Transform U.	yes	yes	no	no	fresh	yes	no	no
Positional U.	yes	yes	no	no	fresh	yes	yes	yes
Java	yes	yes	no	yes	class	no	yes	no

We do not wish to present one or another choice as better, but to draw attention to a potentially-unintended side effect of various points in the solution space. Nonetheless, it is possible for any of the fresh-object-based systems to support delegation or forwarding semantics simply by exposing a method, accepting any object as an argument, that returns a fresh object whose methods provide the behaviour in question. Concatenation semantics can similarly be supported by inheriting from a standard clone.

Diamond inheritance (repeatedly inheriting from the same class or trait two or more times) has long been recognised as a problem in object-oriented language design. Eiffel and C++ both offer the same essential solution to the problem: arranging that some classes can be replicated each time they are inherited, while other classes will be inherited only once. Malayeri and Aldrich [36] present a good discussion of the problems diamonds cause for inheritance, and then argue that diamond inheritance can be prevented in languages, partly by supporting a *requires* clause inspired by Scala which indicates that a trait depends upon the eventual final **self** object providing a set of methods, but not actually implementing those methods itself. The two multiple-inheritance systems we describe are open to this sort of collaboration, but do not require it. Because they are object-based, rather than class-based, some issues of diamond inheritance do not arise, as each instance of a parent is (unavoidably) separately obtained and constructed. In an object-based system, coalescing similar ancestors is a dubious activity, as side effects may occur on the path to construction and be semantically meaningful, which cannot happen in a static, declarative class system.

In our formal model we have made a conscious effort to handle as many errors as possible in the operational semantics (i.e., at run time) rather than by defining erroneous programs as ill-formed (i.e., at compile time). There are several reasons for this, but most important is that we see further layers on top — such as type systems or checks for diamond inheritance — as an important, but separable part of the design process. Omitting such definitions highlights the inheritance designs, and enables the core language of the model to be smaller and more general. In fact, a language could provide the ability to impose additional families of static restrictions at the surface level [23].

Similarly, we take a pragmatic approach to object initialisation: access to uninitialised variables raises a runtime error. This choice has helped simplify our inheritance designs, as we have not needed to contort the models to ensure that e.g. all variables are definitively initialised. Just as a range of static type systems can be layered over the language’s semantics, so we expect that a safe initialisation scheme, such as Delayed Types [16], Masked Types [42],

Hard Hats [20, 54], Freedom Before Commitment [47], or the Billion Dollar Fix [43] should be able to be layered on top of the initialisation semantics in the model. Highlighting where such runtime errors are liable to occur in the different designs helps language designers to choose where to trade off for or against additional safety.

8 Related Work

Class-based object-oriented languages begin with SIMULA [2], and much of the conceptual framework of object-orientation descends from the Scandinavian school founded by Dahl and Nygaard, viewing programming as simulation, and consequently programs as models of phenomena in the real world [35]. Taivalaari argues the class-based understanding of programming is also ‘classical’ in the sense of descending from the classical philosophy of Plato and Aristotle [49].

Simula’s model of objects as instances of classes was greatly expanded on by Smalltalk [4]: unlike SIMULA, Smalltalk classes are also instances of other (meta-) classes themselves, importing the power (and also the complexity) of Lisp-style computational reflection into object-oriented languages [45]. Lisp then returned the favour with a series of object-oriented extensions culminating in the CLOS Meta-Object Protocol [29]. We also owe the notion of ‘static’ (per class) declarations, distinct from per-instance declarations, to Smalltalk.

The complexity of Smalltalk’s meta-model clearly inspired Lieberman to propose languages based purely on objects, with delegation as the sharing mechanism [32], which then led to a general interest in ‘prototype-based’ programming languages, i.e. languages, following Self [51] that create new instances by copying existing instances, rather than by an apostrophe to a class. Emerald [3] also lacked classes, but objects were created by literal expressions.

The fates of Emerald and Self are instructive. By 1991 Emerald had a “*syntactic construct called a class that provides the functionality normally expected of classes*” [24]. At the same time, Self’s programming style was based on separate objects corresponding to instance prototypes (defining all the instance variables in one object) and method suits (called “traits”) with each conceptual class giving rise to both a trait, and a prototype delegating to that trait; traits were also linked by delegation corresponding to the conceptual class inheritance [50]. By 1995, Self objects were effectively given a single “copy down parent” attribute, and slots from that object would be copied into the object whenever it was edited. This is how e.g. if an extra ‘z’ slot was added into the graphic object, it would also be added into the gRectangle object. Eventually, a “Subclass me” button was added to the IDE, which copied both instance object and trait objects, recreating the class-like structure with delegation and copy-down parent links configured correctly.

Dony et al. modelled a range of different designs for class-, object-, and prototype- based languages, although by writing a software product line of definitional interpreters in Smalltalk, rather than modelling language features formally [13, 12]. Taivalaari et al. [40] surveyed other contemporaneous research along these lines.

Relatively early on, Eiffel incorporated a sophisticated class-based multiple inheritance with deep renaming (rather than shallow aliasing) and exclusion (‘undefine’) and repeated inheritance [38]. Eiffel’s design has many advantages, notably that an object implementing several different types can have multiple implementations of methods with the same name, depending on the type in which those definitions originate. Eiffel’s development environment can generate the ‘flat form’ of the class: unfortunately because of Eiffel’s type-aware semantics, the flattened form cannot represent all the possible behaviours of the class. Over the years, other languages incorporated many of the features of Eiffel’s design, notably C++.

The other main stream of work is based on mixins, rather than classes [7]. Unlike CLOS or Eiffel multiple superclasses, mixins are applied one at a time, in a linear order specified by the programmer. Bracha’s Jigsaw formalised mixin composition in a class-based style, along with a rich trait algebra including merge, restrict, select, project, overriding, and rename operators [8, 5]. Flatt et al. [18, 19] develop a semantics for classes and class-like mixins (without composition operators) in a core language, and have incorporated mixins and traits into Scheme (now Racket) based on classes and macros. Lagorio et al. [30] modeled Jigsaw in a class-based formalism based on Featherweight Java [26], and then Corradi et al. [11] extended the formalism to handle family polymorphism [15]. More recently, class mixins have been incorporated into Newspeak (alongside family polymorphism) [9] and Dart [6].

Traits were revived for multiple inheritance in Smalltalk [14]: in this design, a class can inherit from one single superclass, and then incorporate multiple traits based with a trait algebra comprising sum, aliasing, exclusion operators. A key property of Smalltalk traits is that their conceptual model was based on flattening, rather than dynamic dispatch, although both semantics have been proven equivalent [39]. Scala [41] and Java 8 [22] incorporate traits, although relying on ‘super’ calls rather than aliasing or renaming.

A discussion of object-based inheritance systems would be incomplete without referring to OCaml [31], which is not dissimilar to the language of our base model. Both languages have object constructors, classes as sugar for methods returning fresh objects, symmetric multiple inheritance, and are structurally typed. We work on the theory of traditional object-style polymorphism rather than OCaml’s row polymorphism.

More significantly, even if OCaml classes are described as syntactic sugar, objects (or other classes) can only inherit from classes, whereas in all of our models objects can inherit from (at least) any manifestly fresh objects whether or not defined by classes. OCaml also has a more complex, but in some ways less powerful, initialisation model than what we specify here. In OCaml, field initialisers are evaluated in the enclosing lexical scope, and initialisation code must be sequestered into initialisation blocks which are run late — neither of which reflect a straightforward reading of program’s source code — while in our models initialisation is always in a straight line, within the context of the object being constructed.

9 Conclusion

Object-based inheritance is unexpectedly complicated, especially when commonplace desires for functionality available in classical models are involved, and programmers have resorted to increasingly complex workarounds in existing object-based languages. In this paper we showed that object inheritance without classes is both viable and desirable, avoiding the conceptual complexity of an additional kind of entity in an object-oriented language without losing functionality through careful feature selection, and set out a range of options with their various trade-offs made explicit.

We presented and discussed seven models of object inheritance, including the well-known approaches of delegation, forwarding, and concatenation. We presented a novel extended operational semantics for a base language incorporating advanced but standard features affected by inheritance, and formalised the models as extensions to that single base language, formally demonstrating the subtle behavioural differences of each model. In particular, we addressed the complex questions of downcalls, object registration, stability, inheriting pre-existing objects, action at a distance, and multiple inheritance, and their interactions, illustrating that object-based inheritance has the full range of possibilities of classical inheritance, and showed that many of these models can be used as effectively as purely declarative classes, but particular combinations — especially class initialisation semantics

combined with inheritance from preexisting objects — require a specific set of features usually reserved for very dynamic and reflective languages.

References

- 1 Henk Barendregt. *The Lambda Calculus*. North-Holland, revised edition, 1984.
- 2 Graham M. Birtwistle, Kristen Nygaard, Bjørn Myhrhaug, and Ole-Johan Dahl. *Simula Begin*. Studentlitteratur, 1979.
- 3 Andrew P. Black, Eric Jul, Norman Hutchinson, and Henry M. Levy. The development of the Emerald programming language. In *History of Programming Languages III*, 2007.
- 4 Alan Borning. Classes versus prototypes in object-oriented languages. In *Proc. ACM Fall Joint Computer Conference*, pages 36–40, 1986.
- 5 Gilad Bracha. *The Programming Language Jigsaw: Mixins, Modules, and Multiple Inheritance*. PhD thesis, University of Utah, 1992.
- 6 Gilad Bracha. Mixins in Dart, 2015. [Online; accessed 30-November-2015]. URL: <https://www.dartlang.org/articles/mixins>.
- 7 Gilad Bracha and William Cook. Mixin-based inheritance. In *OOPSLA*, 1990.
- 8 Gilad Bracha and Gary Lindstrom. Modularity meets inheritance. In *Proc. International Conference on Computer Languages*, pages 282–290, 1992.
- 9 Gilad Bracha, Peter von der Ahé, Vassili Bykov, Yaron Kashai, William Maddox, and Eliot Miranda. Modules as objects in Newspeak. In *ECOOP*, pages 405–428, 2010.
- 10 Craig Chambers, David Ungar, Bay-Wei Chang, and Urs Hölzle. Parents are shared parts of objects: inheritance and encapsulation in Self. *Lisp and Symbolic Computation*, 1991.
- 11 Andrea Corradi, Marco Servetto, and Elena Zucca. DeepFJig: modular composition of nested classes. In *PPPJ*, pages 101–110, 2011.
- 12 Christophe Dony, Jacques Malenfant, and Pierre Cointe. Prototype-based languages: From a new taxonomy to constructive proposals and their validation. In *OOPSLA*, 1992.
- 13 Christophe Dony, Jacques Malenfant, and Pierre Cointe. Classifying prototype-based programming languages. In James Noble, Antero Taivalsaari, and Ivan Moore, editors, *Prototype-Based Programming: Concepts, Languages and Applications*, chapter 2. 1999.
- 14 Stéphane Ducasse, Oscar Nierstrasz, Nathanael Schärli, Roel Wuyts, and Andrew P. Black. Traits: A mechanism for fine-grained reuse. *ACM Transactions on Programming Languages and Systems*, 2005.
- 15 Erik Ernst. Family polymorphism. In *ECOOP*, 2001.
- 16 Manuel Fähndrich and Songtao Xia. Establishing object invariants with delayed types. In *OOPSLA*, pages 337–350, 2007.
- 17 Matthias Felleisen, Robert Bruce Findler, and Matthew Flatt. *Semantics Engineering with PLT Redex*. MIT Press, 2009.
- 18 Matthew Flatt, Shriram Krishnamurthi, and Matthias Felleisen. Classes and mixins. In *POPL*, 1998.
- 19 Matthew Flatt, Shriram Krishnamurthi, and Matthias Felleisen. A programmer’s reduction semantics for classes and mixins. In *Formal Syntax and Semantics of Java*, 1998.
- 20 Joseph (Yossi) Gil and Tali Shragai. Are we ready for a safer construction environment? In *ECOOP*, pages 495–519, 2009.
- 21 Adele Goldberg and David Robson. *Smalltalk-80: The Language and its Implementation*. Addison-Wesley, 1983.
- 22 James Gosling, Bill Joy, Guy Steele, Gilad Bracha, and Alex Buckley. *The Java Language Specification*. Oracle, 2015.
- 23 Michael Homer, Timothy Jones, James Noble, Kim B. Bruce, and Andrew P. Black. Graceful dialects. In *ECOOP*, pages 131–156, 2014.

- 24 Norman C. Hutchinson, Rajendra K. Raj, Andrew P. Black, Henry M. Levy, and Eric Jul. The Emerald programming language report. Computer Science, UBC, October 1991.
- 25 Roberto Ierusalimsky, Luiz Henrique de Figueiredo, and Waldemar Celes. The evolution of Lua. In *History of Programming Languages III*, 2007.
- 26 Atsushi Igarashi, Benjamin C. Pierce, and Philip Wadler. Featherweight Java: A minimal core calculus for Java and GJ. *ACM Transactions on Programming Languages and Systems*, 23(3):396–450, 2001.
- 27 Carlton Egremont III. *Mr. Bunny's Big Cup o'Java*. Addison-Wesley, 1999.
- 28 Timothy Jones and James Noble. Tinygrace: A simple, safe and structurally typed language. In *Formal Techniques for Java-like Programs*. 2014.
- 29 Gregor Kiczales, Jim des Rivières, and Daniel G. Bobrow. *The Art of the Metaobject Protocol*. MIT Press, 1991.
- 30 Giovanni Lagorio, Marco Servetto, and Elena Zucca. Featherweight Jigsaw: Replacing inheritance by composition in Java-like languages. *Inf. Comput.*, 214:86–111, 2012.
- 31 Xavier Leroy, Damien Doligez, Alain Frisch, Jacques Garrigue, Didier Rémy, and Jérôme Vouillon. The OCaml system release 4.00 documentation and user's manual, 2012.
- 32 Henry Lieberman. Using prototypical objects to implement shared behavior in object oriented systems. In *OOPSLA*, November 1986.
- 33 Henry Lieberman, Lynn Andrea Stein, and David Ungar. Treaty of Orlando. In *Addendum to OOPSLA Proceedings*, May 1988.
- 34 Lua-Users. Object oriented programming, 2014. [Online; accessed 30-November-2015]. URL: <http://lua-users.org/wiki/ObjectOrientedProgramming>.
- 35 Ole Lehrmann Madsen, Birger Møller-Pedersen, and Kristen Nygaard. *Object-Oriented Programming in the BETA Programming Language*. Addison-Wesley, 1993.
- 36 Donna Malayeri and Jonathan Aldrich. Cz: Multiple inheritance without diamonds. In *OOPSLA*, 2009.
- 37 Bertrand Meyer. *Object-oriented Software Construction*. Prentice Hall, 1988.
- 38 Bertrand Meyer. *Eiffel: The Language*. Prentice Hall, 1992.
- 39 Oscar Nierstrasz, Stéphane Ducasse, and Nathanael Schärli. Flattening traits. *Journal of Object Technology*, 5:66–90, 2006.
- 40 James Noble, Antero Taivalsaari, and Ivan Moore, editors. *Prototype-Based Programming: Concepts, Languages, Applications*. Springer-Verlag, 1997.
- 41 Martin Odersky, Lex Spoon, and Bill Venners. *Programming In Scala*. Artima, 2011.
- 42 Xin Qi and Andrew C. Myers. Masked types for sound object initialization. In *POPL*, pages 53–65, 2009.
- 43 Marco Servetto, Julian Mackay, Alex Potanin, and James Noble. The billion-dollar fix—safe modular circular initialisation with placeholders and placeholder types. In *ECOOP*, pages 205–229, 2013.
- 44 Pat Shaughnessy. *Ruby Under A Microscope*. No Starch Press, 2013.
- 45 Brian C. Smith. Reflection and semantics in Lisp. In *Proceedings of the ACM Conference on Principles of Programming Languages*, 1984.
- 46 Stack Overflow. Which JavaScript library has the most comprehensive class inheritance support?, 2015. [Online; accessed 30-November-2015]. URL: <https://stackoverflow.com/questions/711209>.
- 47 Alex J. Summers and Peter Müller. Freedom before commitment - a lightweight type system for object initialisation. In *OOPSLA*, pages 1013–1032, 2011.
- 48 Antero Taivalsaari. Delegation versus concatenation or cloning is inheritance too. *Oops Messenger*, 6(3), 1995.

13:26 Object Inheritance Without Classes

- 49 Antero Taivalsaari. Classes vs. prototypes: Some philosophical and historical observations. In James Noble, Antero Taivalsaari, and Ivan Moore, editors, *Prototype-Based Programming: Concepts, Languages and Applications*, chapter 1. Springer-Verlag, 1999.
- 50 David Ungar, Craig Chambers, Bay-Wei Chang, and Urs Hölzle. Organizing programs without classes. *Lisp and Symbolic Computation*, 4(3), June 1991.
- 51 David Ungar and Randall B. Smith. SELF: The power of simplicity. *Lisp and Symbolic Computation*, 4(3), June 1991.
- 52 Wikipedia. Snit, 2015. [Online; accessed 30-November-2015]. URL: <https://en.wikipedia.org/w/index.php?title=OTcl&oldid=663399454>.
- 53 Allen Wirfs-Brock, editor. *ECMAScript 2015 Language Specification*. Ecma International, 6th edition, 2015.
- 54 Yoav Zibin, David Cunningham, Igor Peshansky, and Vijay Saraswat. Object initialization in X10. In *ECOOP*, pages 207–231, 2012.