

Supplementary Notes on Objects

15-312: Foundations of Programming Languages
Jonathan Aldrich

Lecture 20-21
November 4-6, 2003

In this lecture we examine the design goals and technical properties of object-oriented languages. It is possible to study some features of objects by encoding them in terms of constructs we already know; however, realistic encodings of objects are so complex that it is often simpler to study objects as first-class constructs. Rather than study objects in the context of a completely different language (such as Java) we will look at how objects can be added naturally to already powerful functional languages like ML.

A Brief History of Objects.

The first object-oriented language was Simula 67, developed by Ole-Johan Dahl and Kristen Nygaard (who won the 2002 Turing award for this work). Simula 67 had many of the same concepts as modern object-oriented programming languages, including objects, classes, virtual methods, and inheritance. The motivation for the Simula languages came due to Kristen Nygaard's work on operational research. Nygaard needed a tool for the precise description and simulation of complex systems involving both people and machines. Simula 67 was designed to support both a description of complex systems and a simulation or execution of these systems.

The design goal of modeling real-world systems influenced the features of Simula 67, as well as more recent object-oriented languages. Real-world systems are built of entities that have state and behave in certain ways; these entities are modeled as objects. We often classify real-world entities into hierarchical categories. Each category provides a partial description of the state and behavior of entities in that category: for example, animals

typically move, while plants typically photosynthesize. These classification structures are modeled with classes that describe the state and behavior of object instances. Our classification structures are often hierarchical, so that mammals are a particular kind of animal that is warm-blooded, furry, gives birth to live young, etc. Similarly, object-oriented programming languages allow one class to inherit from another, so that the state and behavior specifications in one class can be reused in its subclasses. Virtual methods are provided so that subclasses can refine the behavior defined in their superclasses.

Most object-oriented languages are imperative. This is not strictly necessary—a number of purely functional object-oriented languages have been defined. However, imperative objects form a natural model of the world, because we usually think of the world as having state that changes from moment to moment. As in the real world, the behavior of an object-oriented system are derived from the interactions between objects, as described by their methods.

The design goals of object-oriented languages also provide insight into the software engineering tradeoffs in using object-oriented languages. For example, in order to build a useful program, you need to both verify it (build the system right) and validate it (build the right system). Functional languages are based closely on the mathematical theory of computation; thus, they facilitate verification, or building the system right. Object-oriented languages are “messier” in a sense: functions are split into methods, references are all over the place, and object-oriented dispatch makes it hard to tell what code is running. Therefore, it is often harder to verify that an object-oriented system is built right. However, object-oriented languages are very good at modeling the world using abstractions that are close to real-world concepts. Thus, it is often easier to validate object-oriented systems—ensuring that the programmers wrote the right system, one that fulfills the needs of users. In practice, of course, we need both verification and validation, and the relative benefits and drawbacks of different languages may depend on the domain in which they are used, as well as many other factors.

Object-Oriented and Functional Programming.

In many ways, objects are similar to features provided by other languages. For example, a class is a lot like a constructor in a datatype, and a method in class C is a lot like the case for datatype constructor C in a ML function definition. However, object-oriented languages provide several properties

not found in ML or other functional languages.

First, classes and methods are extensible, unlike datatype and function declarations. In object-oriented languages, you can add a new class at any time. One important benefit of this property is that you can extend or build on top of a system by adding new classes even when you cannot modify the system's source code. Each class can include its own implementation for methods inherited from superclasses. In contrast, when you declare a datatype in ML, you have to give all of the constructors in one place, and for each function declared over the datatype, a case must be given for each constructor at the same location in the source code. Thus, in order to add a new datatype case in ML you have to modify the datatype declaration and all functions over the datatype, which can be a major problem if the source code is not available.

Second, class extension is hierarchical. If class C extends another class B, then C inherits all of the fields and methods from its superclass. Any code written for class B can be customized and reused by class C—and this is true whether or not the author of B intended for the code to be reused. Functional languages like ML also provide support for customized reuse through first-class functions and functors. Compared to these mechanisms, inheritance is more structured than passing around first-class functions, but more lightweight and flexible than writing a functor. Inheritance also typically requires less advanced planning for reuse, since extensibility is “built in” to the language construct.

The hierarchical nature of class extension is also useful in that subclasses can define methods that are not present in their superclasses. In functional languages, this is equivalent to writing a function that is defined on some constructors in a datatype but not others. As we will see, a primary benefit of object-oriented languages is that the type system can ensure that such partial functions are never called on classes on which the function is not defined. Current CMU research on refinement types is focused in part on providing this benefit in the context of functional languages.

EML: Objects in ML.

Unfortunately many object-oriented languages do not take advantage of the benefits of advanced functional programming languages. For example, Java makes it easy to add new classes to a program externally by subclassing from existing classes. However, there is no way to externally add a new function with cases for each class—instead the function must be split up and added as one method to each relevant class in the system, something that is

impossible if the system's source code is not available. In comparison, this task is easy in ML—you simply define a new function with cases for each relevant datatype constructor.

A better language design combines the benefits of both functional and object-oriented languages. Here we follow the design of Extensible ML (EML), a recently-developed language that provides object-oriented features in ML by making datatypes and functions hierarchically extensible. The grammar of EML's class and function declarations is given as follows:

$$\begin{aligned} \text{decl} ::= & \text{[abstract] class } C \text{ [extends } C'] \\ & \text{of } \{\bar{l} : \bar{\tau}\} \\ | & \text{fun } f : \bar{C} \rightarrow \tau \\ | & \text{extend fun } f (\bar{x} \text{ as } \bar{C}) = e \end{aligned}$$

In this example, we use the overbar notation \bar{l} to denote a tuple of terms that have the form t . We assume that expressions e have the same form as in MinML programs we have already studied. The EML language is declaration-oriented, and there are three forms of declarations: `class` declarations, function signature declarations, and function cases.

A class declaration in EML is an extended form of an ML datatype declaration, where the body of the datatype is an extensible record. Class declarations can be declared `abstract`, meaning that the class cannot be instantiated, only extended. Inheritance is represented by an `extends` clause. Finally, the body of the class is a record mapping labels \bar{l} to types $\bar{\tau}$.

Function declarations are split into two parts. A `fun` declaration declares the type of a function from a tuple of classes \bar{C} to a type τ . In the full EML language, functions can have any argument and result type, but we are simplifying the system for the purposes of modeling it formally.

Function cases are declared with the `extend fun` declaration. The `extend fun` keywords are followed with a pattern expression describing the conditions under which the function case is applicable. In our formal treatment, we assume that the pattern is a tuple of bindings of variables to classes; the full EML language permits an arbitrary pattern here. The body of the function e is evaluated in a scope where the variables \bar{x} from the pattern match are bound to the actual parameters of the function.

A multi-method language is an object-oriented language where function dispatch can depend on all of the arguments to the function, not just the receiver of the message ("this" in Java or C++). EML is a multi-method language, because function cases can specify dispatch on every argument in the tuple that is passed to the function.

We also need new syntax for creating objects: this syntax the same one we would use to create a ML datatype that is implemented as a record:

$$e ::= C\{\bar{l} = \bar{e}\}$$

Here we create an object of class C , initializing the fields with the expressions \bar{e} .

Operational Semantics for EML

EML's operational semantics follows the same rules as in MinML, except for calls to functions. Function calls must be treated differently because the system must use the actual classes of the arguments to determine which function case to execute. We define function case lookup using an auxiliary function, *lookup*, which chooses the most specific applicable function case for a given tuple of classes.

$$\frac{\bar{v} = \overline{C\{\dots\}} \quad \text{lookup}(f, \bar{C}) = \bar{x}.e}{f(\bar{v}) \mapsto \{\bar{v}/\bar{x}\}e} \text{FnApp}$$

$$\frac{(\text{extend fun } f(\bar{x} \text{ as } \bar{C}') = e) \in \text{Decls} \quad \bar{C} <: \bar{C}' \quad (\text{extend fun } f(\bar{x}' \text{ as } \bar{C}'') = e') \in \text{Decls} \Rightarrow (\bar{C}' <: \bar{C}'' \wedge \bar{C}'' \not<: \bar{C}')}{\text{lookup}(f, \bar{C}) = \bar{x}.e} \text{Lookup}$$

Client-Side Typechecking in EML

The typing rules for EML are divided into client-side and implementation-side typechecking. Client-side typechecking is the same kind of typechecking we have been studying up to this point in the class: we check that expressions e have types τ , and that class declarations and function declarations are well formed, written *decl OK*. Implementation-side typechecking, discussed in a later section, verifies that each generic function is implemented completely and consistently.

EML's rules for client-side typechecking build on the rules we have been studying for MinML. The rules that are new are given below. Several of the rules refer to *Decls*, the (global) set of declarations visible in the program.

$$\frac{(\text{fun } f : \bar{C} \rightarrow \tau) \in \text{Decls}}{\Gamma \vdash f : \bar{C} \rightarrow \tau} \text{FnTyp}$$

$$\frac{C \text{ class } \quad C \text{ concrete} \quad \Gamma \vdash \bar{e} : \bar{\tau} \quad \text{fields}(C) = \{\bar{l} : \bar{\tau}\}}{\Gamma \vdash C\{\bar{l} = \bar{e}\} : C} \text{ClsTyp}$$

$$\frac{\bar{C} \text{ class} \quad \tau \text{ type}}{\text{fun } f : \bar{C} \rightarrow \tau \text{ OK}} \text{FnOK}$$

$$\frac{(\text{fun } f : \bar{C}' \rightarrow \tau) \in \text{Decls} \quad \bar{C} <: \bar{C}' \quad \bar{x} : \bar{C} \vdash e : \tau}{\text{extend fun } f (\bar{x} \text{ as } \bar{C}) = e \text{ OK}} \text{ExtOK}$$

$$\frac{C' \text{ class} \quad \bar{\tau} \text{ type} \quad \{\bar{l}\} \cap \text{fields}(C') = \emptyset}{[\text{abstract}] \text{ class } C [\text{extends } C'] \text{ of } \{\bar{l} : \bar{\tau}\} \text{ OK}} \text{ClsOK}$$

The first rule gives a type to a variable in client code that refers to a function that has been declared. The second rule states that an object of class C is well-typed and has C if C is a class and C is concrete. It also requires that the expressions assigned to the fields of C have types that match the class declaration.

The rule for function declarations ensures that the declared classes and types exist. The rule for function cases checks that the classes in the pattern expression are subtypes of the classes in the function signature. The rule also checks that the body is well-formed, assuming the arguments have appropriate types. The rule for classes just checks that the types and classes referred-to actually exist, and also verifies that the new fields are distinct from the ones being inherited. In a real language, we would allow one field to shadow another, disambiguating them using the class in which the fields were defined.

To be precise, we should also define auxiliary judgments of the form C class and $\text{fields}(C)$.

$$\frac{([\text{abstract}] \text{ class } C \dots) \in \text{Decls}}{C \text{ class}} \text{Class}$$

$$\frac{C \text{ class}}{C \text{ type}} \text{ClassType}$$

$$\frac{([\text{abstract}] \text{ class } C \text{ of } \{\bar{l} : \bar{\tau}\}) \in \text{Decls}}{\text{fields}(C) = \{\bar{l} : \bar{\tau}\}} \text{FieldsBase}$$

$$\frac{([\text{abstract}] \text{ class } C \text{ extends } C' \text{ of } \{\bar{l} : \bar{\tau}\}) \in \text{Decls}}{\text{fields}(C) = \{\bar{l} : \bar{\tau}\} \cup \text{fields}(C')} \text{FieldsInherit}$$

The first two rules state that a class exists if it is declared in *Decls*, and that each class defines a type with the same name. The fields of a class are computed based on the union of the inherited and the newly-declared fields. We can define the subtyping relation as the reflexive, transitive closure of the declared inheritance relation:

$$\frac{([\text{abstract}] \text{ class } C \text{ extends } C' \dots) \in \text{Decls}}{C <: C'} \text{SubBase}$$

$$\frac{}{C <: C} \text{SubReflex}$$

$$\frac{C <: C' \quad C' <: C''}{C <: C''} \text{SubTrans}$$

Implementation-Side Typechecking in EML

To complete the typing rules for EML, we must give rules for implementation-side typechecking as well. Implementation-side typechecking verifies that each function in the system is completely and consistently implemented. For example, consider the following code:

```
abstract class Shape of ...
class Rectangle extends Shape of ...
class Circle extends Shape of ...
```

```
fun draw : Shape -> unit
extend fun draw (c as Circle) = ...
```

```
fun intersect : Shape*Shape -> boolean
extend fun intersect (c as Circle, s as Shape) = ...
extend fun intersect (s as Shape, r as Rectangle) = ...
```

In the code above, the draw function is incomplete, because there is no case for Rectangle. It's actually fine that there is no case for Shape, because since Shape is an abstract class, and so we can never create a Shape at run time. However, if we create a Rectangle at run time and pass it to draw, we will get a message not understood error. This terminology stems from the object-oriented practice of calling functions "messages" that are sent to objects like Rectangle; in this case, there is no function case for Rectangle in the message, so we say that Rectangle did not understand the message.

Similarly, in the code above, the intersect function is ambiguous, because if the function is invoked with the tuple (c,r), where c is a Circle and r is a Rectangle, then either case of the intersect function applies to the tuple, but neither is more specific than the other. This is known as a message ambiguous error.

The goal of implementation-side typechecking is to statically guarantee that message not understood and message ambiguous errors cannot occur. We will first define a global typechecking algorithm formally, and then discuss informally how to derive a modular typechecking algorithm from the global one. The global algorithm is given by the rule below:

$$\frac{\begin{array}{l} \forall(\text{fun } f : \overline{C} \rightarrow \tau) \in \text{Decls} \\ \forall \overline{C}' \text{ where } \overline{C}' \text{ concrete and } \overline{C}' <: \overline{C} \\ \text{lookup}(f, \overline{C}') = \overline{x}.e \end{array}}{\text{Decls OK}} \text{Global-ITC}$$

The rule above checks, for each function f and for each tuple of concrete classes \overline{C}' that subtype the classes declared in the function signature, whether there is a unique function case for that tuple of classes. The same *lookup* function that is used for dispatching at run time is used to check for a unique function case.

Modular Implementation Side Typechecking.

The implementation side typechecking rule defined above is global, which is undesirable since it means that we can't ensure that two modules that typecheck in isolation will also typecheck when they are combined. It is preferable to have a modular typechecking algorithm. We say a typechecking algorithm is modular if it can typecheck a module by looking at only the implementation of that module and the interfaces of the modules it statically depends on. In EML, we broaden the definition of interface slightly to include the list of cases defined for each function in a module.

The following table summarizes the rules that Java, ML, and EML use to modularly check for the absence of message not understood and message ambiguous errors. The table shows that EML's strategy is a combination of the strategies of ML and of Java.

Implementation Side Typechecking	Message Not Understood	Message Ambiguous
Java	Concrete classes must define inherited abstract methods, and all argument types except the receiver must match the signature	Each function cases must be defined in the receiver class body
ML	All cases must be covered in a single local declaration (compiler warning if cases are not exhaustive)	All cases are declared together, and the order resolves any ambiguity
EML	External functions must provide a default case where function is declared. For internal functions, a local default case must be defined whenever a new concrete subclass of an abstract class is defined.	All function cases must be declared either in the same module where the function is declared, or in the module where the receiver is declared.

With the restrictions above applied to the EML language, the global typechecking algorithm can be run one module at a time, taking the *Decls* set to be just the locally-declared declarations as well as the declarations

in modules that the current module statically depends on. The authors of EML have proved that these restrictions are sufficient to guarantee that when the typechecking algorithm is run on several modules in isolation, global typechecking will succeed when these modules are linked. The proof is beyond the scope of this course, however.

Using Inheritance Correctly.

Consider the code below.

```
class Engine of { horsepower:int }
class Automobile extends Engine of { mileage:int }
class Airplane extends Engine of { prop:Propeller }
```

This code represents a broken use of inheritance. Essentially, it mixes up the subtyping relationship with the containment relationship. An Automobile has an Engine, and an Airplane has an Engine as well. However, “has a” relationships should be expressed with containment. For example, an Airplane has a Propeller, and this is properly expressed by giving the Airplane a prop field of type Propeller. Inheritance imposes a subtyping relationship, meaning that the declarations above state (incorrectly) that an Automobile is a Engine. This breaks down in an obvious way when you consider that not all operations which apply to engines also apply to Airplanes. For example, it makes sense to change the oil of an engine, but changing the oil of an airplane doesn’t make sense in the same way. Using inheritance for a “has a” relationship also breaks down when we consider an Airplane that has more than one engine.

The general rule is: only use inheritance to represent an “is a” relationship between entities, never a “has a” relationship.