

# Object-Oriented Programming Languages: Application and Interpretation

Éric Tanter

June 22, 2023

last updated: Thursday, June 22nd, 2023

Copyright © 2010-2023 by Éric Tanter

The electronic version of this work is licensed under the Creative Commons Attribution Non-Commercial No  
Derivatives License

This booklet exposes fundamental concepts of object-oriented programming languages in a constructive and progressive manner. It follows the general approach of the PLAI book by Shriram Krishnamurthi (or at least I'd like to think it does). The document assumes familiarity with the following Parts of PLAI: I-V (especially first-class functions, lexical scoping, recursion, and state), as well as XII (macros).

OOPLAI is also available in PDF version. Note however that OOPLAI is subject to change at any time, so accessing it through the web is the best guarantee to be viewing the latest version.

I warmly welcome comments, suggestions and fixes; just send me a mail!

As of June 2018, OOPLAI has been translated to Chinese by ChongKai Zhu.

**Acknowledgments:** I am thankful to the members of the PLEIAD lab and the students of the Programming Languages course at University of Chile for detecting bugs and suggesting enhancements.

# 1 From Functions to Simple Objects

This exploration of object-oriented programming languages starts from what we know already from PLAI, as well as our intuition about what objects are.

## 1.1 Stateful Functions and the Object Pattern

An object is meant to encapsulate in a coherent whole a piece of state (possibly, but not necessarily, mutable) together with some behavior that relies on that state. The state is usually called *fields* (or *instance variables*), and the behavior is provided as a set of *methods*. Calling a method is often considered as *message passing*: we send a message to an object, and if it understands it, it executes the associated method.

In higher-order procedural languages like Scheme, we have seen similar creatures:

```
(define add
  (λ (n)
    (λ (m)
      (+ m n))))

> (define add2 (add 2))
> (add2 5)
7
```

The function `add2` encapsulates some hidden state (`n = 2`) and its behavior effectively depends on that state. So, in a sense, a closure is an object whose fields are its *free variables*. What about behavior? well, it has a unique behavior, triggered when the function is *applied* (from a message passing viewpoint, `apply` is the only message understood by a function).

If our language supports mutation (`set!`), we can effectively have a stateful function with changing state:

```
(define counter
  (let ([count 0])
    (λ ()
      (begin
        (set! count (add1 count))
        count))))
```

We can now effectively observe that the state of `counter` changes:

```
> (counter)
1
```

```
> (counter)
2
```

Now, what if we want a bi-directional counter? The function must be able to do either +1 or -1 on its state depending on... well, an argument!

```
(define counter
  (let ([count 0])
    (λ (msg)
      (match msg
        ['dec (begin
                 (set! count (sub1 count))
                 count)]
        ['inc (begin
                 (set! count (add1 count))
                 count)]))))
```

Note how `counter` uses `cmd` to discriminate what action to perform.

```
> (counter 'inc)
1
> (counter 'dec)
0
```

This looks quite like an object with two methods and one instance variable, doesn't it?

Let us generalize a bit, in order to support multiple arguments and organize methods as a list of functions indexed by their name. We will also introduce another instance variable `step` for controlling the amount by which the counter is increased or decreased. We add a few more methods as well:

```
(define counter
  (let ([count 0]
        [step 1])
    (let ([methods
           (list
            (cons 'inc (λ () (set! count (+ count step))
                       count))
            (cons 'dec (λ () (set! count (- count step))
                       count))
            (cons 'reset (λ () (set! count 0)))
            (cons 'step! (λ (v) (set! step v)))))]
      (λ (msg . args)
        (let ([found (assoc msg methods)])
          (if found
```

```
(apply (cdr found) args)
(error "message not understood:" msg))))))
```

We use the dot notation for the arguments of the dispatch function: this enables the function to receive one mandatory argument (the `msg`) as well as zero or more optional arguments (available in the body as a list bound to `args`).

To define the dispatch function in a generic fashion, we first put all methods in an association list (ie. a list of pairs) called `methods`, which associates a symbol (aka. message) to the corresponding method (ie. a lambda). We use `assoc` to lookup an existing method for handling the `msg`. We get the result as `found`, which will be a pair symbol-lambda if found. We then use `apply` to apply the function to the (possibly empty) list of arguments. If no method was found, `found` will be `#f`; in that case, we generate an informative error message.

Let's try that:

```
> (counter 'inc)
1
> (counter 'step! 2)
> (counter 'inc)
3
> (counter 'dec)
1
> (counter 'reset)
> (counter 'dec)
-2
> (counter 'hello)
message not understood: 'hello
```

## 1.2 A (First) Simple Object System in Scheme

We can now use macros to embed a simple object system that follows the pattern identified above.

```
(defmac (OBJECT ([field fname fval] ...)
               ([method mname mparams mbody ...] ...))
  #:keywords field method
  (let ([fname fval] ...)
    (let ([methods (list (cons 'mname (lambda (mparams mbody ...) ...))]
                          (lambda (msg . args)
                            (let ([found (assoc msg methods)])
                              (if found
                                  (apply (cdr found) args)
                                  (error "message not understood:" msg))))))
```

Note that in this booklet, we use `defmac` to define macros. `defmac` is like `define-syntax-rule`, but it also supports the specification of keywords and captures of identifiers (using the `#:keywords` and `#:captures` optional parameters). It is provided by the `play` package.

We additionally define a specific notation for sending a message to an object, using an arrow  $\rightarrow$ , such as `( $\rightarrow$  counter step! 3)`. This little macro hides from the user the fact that an object is a function and that the message is treated as a symbol.

```
(defmac ( $\rightarrow$  o m arg ...)
  (o 'm arg ...))
```

We can now use our embedded object system to define our counter in a nice way:

```
(define counter
  (OBJECT
    ([field count 0]
     [field step 1])
    ([method inc () (set! count (+ count step)) count]
     [method dec () (set! count (- count step)) count]
     [method reset () (set! count 0)]
     [method step! (v) (set! step v)])))
```

and use it:

```
> ( $\rightarrow$  counter inc)
1
> ( $\rightarrow$  counter step! 2)
> ( $\rightarrow$  counter inc)
3
> ( $\rightarrow$  counter dec)
1
> ( $\rightarrow$  counter reset)
> ( $\rightarrow$  counter dec)
-2
> ( $\rightarrow$  counter hello)
message not understood: 'hello'
```

Note that object fields (like `count` and `step`) are *strongly encapsulated* here. There is no way to access them directly without going through a method.

### 1.3 Constructing Objects

Up to now, our objects have been created as unique specimen. What if we want more than one point object, possibly with different initial coordinates?

In the context of functional programming, we have already seen how to craft various similar functions in a proper way: we can use a higher-order function, parameterized accordingly, whose role is to produce the specific instances we want. For instance, from the `add` function defined previously, we can obtain various single-argument adder functions:

```

> (define add4 (add 4))
> (define add5 (add 5))
> (add4 1)
5
> (add5 1)
6

```

Because our simple object system is embedded in Scheme, we can simply reuse the power of higher-order functions to define *object factories*:

```

(define (make-counter [init-count 0] [init-step 1])
  (OBJECT
    ([field count init-count]
     [field step init-step])
    ([method inc () (set! count (+ count step)) count]
     [method dec () (set! count (- count step)) count]
     [method reset () (set! count 0)]
     [method step! (v) (set! step v)])))

```

This approach can also be used in object-based languages such as JavaScript (JS has various ways to create similar objects).

The `make-counter` function takes the initial count and step as *optional* parameters, and returns a freshly created object, properly initialized.

```

> (let ([c1 (make-counter)]
        [c2 (make-counter 10 5)])
    (+ (→ c1 inc) (→ c2 inc)))
16

```

## 1.4 Dynamic Dispatch and Polymorphism

Our simple object system is sufficient to show the fundamental aspect of object-oriented programming: dynamic dispatch, which gives rise to the kind of polymorphism that makes object-oriented programs extensible.

See how the `inc-all` function below sends the `inc` message to each object in the list `lst`, without knowing how it will handle it (assuming it can!):

```

(define (inc-all lst)
  (map (λ (x) (→ x inc)) lst))

> (inc-all (list (make-counter)
                 (make-counter 10 5)
                 (OBJECT () ((method inc () "hola")))))
'(1 15 "hola")

```

In particular, the last object is not at all a counter, it just happens to be a weird object that handles `inc` in its own distinctive way. `inc-all` is completely oblivious to the kinds of objects it talks to.

Likewise, notice how, in the following object-oriented implementation of a tree, a node object sends the `sum` message to each of its children without knowing whether it is a leaf or a node:

```
(define (make-node l r)
  (OBJECT
    ([field left l]
     [field right r])
    ([method sum () (+ (→ left sum) (→ right sum))]))))

(define (make-leaf v)
  (OBJECT
    ([field value v])
    ([method sum () value])))

> (let ([tree (make-node
                (make-node (make-leaf 3)
                           (make-node (make-leaf 10)
                                       (make-leaf 4)))
                (make-leaf 1))])
    (→ tree sum))
18
```

We have successfully embedded a simple object system in Scheme that shows the connection between lexically-scoped first-class functions and objects. However, we are far from done, because the object system we have is still incomplete and primitive. In particular, objects lack their "self" so far, and we have no advanced reuse mechanism yet. However, as simple as it may seem, this object system is entirely enough to illustrate the fundamental data abstraction mechanism that objects really are.

See the chapter on the benefits and limits of objects for more on this matter.

## 2 Looking for the Self

In the previous section, we have built a simple object system, whose main limitation is that objects are not "aware" of themselves. Object-oriented programming languages all give this self-awareness to objects, which allows methods to reuse others, and to be (mutually) recursive.

Consider a simple extension of our counters from the previous chapter with an `inc-by!` method, which does two things: it changes the `step` by its argument (like `step!`) and then actually increments the counter (like `inc`). Of course we'd want to implement this "compound" method by invoking the existing methods on the currently-executing object:

```
(method inc-by! (v)
  (-> self step! v)
  (-> self inc))
```

Here we used `self` to denote the currently-executing object; in other languages, like Java and JavaScript, it is called `this`. Clearly, our account of OOP so far does not inform us as to what `self` is.

### 2.1 What is Self?

Let us go back to our first definition of an object (without macros). We see that an object is a function; and so we want, from within that function, to be able to refer to itself. How do we do that? We already know the answer from our study on recursion! We just have to use a recursive binding (with `letrec`) to give a name to the function-object and then we can use it in method definitions:

```
(define counter
  (letrec
    ([self
      (let ([count 0]
            [step 1])
        (let ([methods
              (list
                (cons 'inc (lambda () (set! count (+ count step))
                       count))
                (cons 'dec (lambda () (set! count (- count step))
                       count))
                (cons 'reset (lambda () (set! count 0)))
                (cons 'step! (lambda (v) (set! step v)))
                (cons 'inc-by! (lambda (n) (self 'step! n)
                               (self 'inc))))))]
          self))
```

```

        (λ (msg . args)
          (let ([found (assoc msg methods)])
            (if found
                (apply (cdr found) args)
                (error "message not understood:" msg))))))
  self))

```

Note that the body of the `letrec` simply returns `self`, which is bound to the recursive procedure we have defined.

```

> (counter 'inc)
1
> (counter 'inc-by! 20)
21
> (counter 'inc)
41

```

## 2.2 Self with Macros

Let us take the pattern above and use it in our `OBJECT` macro:

```

(defmacro (OBJECT ([field fname fval] ...)
                 ([method mname mparams mbody ...] ...))
  #:keywords field method
  (letrec
    ([self
      (let ([fname fval] ...)
        (let ([methods (list (cons 'mname (λ mparams mbody ...)) ...))]
          (λ (msg . args)
            (let ([found (assoc msg methods)])
              (if found
                  (apply (cdr found) args)
                  (error "message not understood:" msg))))))
        self))
      (defmac (→ o m arg ...)
        (o 'm arg ...))

```

Now let us try it out with some points:

```

(define (make-counter [init-count 0] [init-step 1])
  (OBJECT
    ([field count init-count]
     [field step init-step])

```

```

      ([method inc () (set! count (+ count step)) count]
       [method dec () (set! count (- count step)) count]
       [method reset () (set! count 0)]
       [method step! (v) (set! step v)]
       [method inc-by! (v) (→ self step! v) (→ self inc)])))

> (let ([c (make-counter)])
    (→ c inc-by! 20))
self: undefined;
cannot reference an identifier before its definition
in module: 'program

```

What?? But we did introduce `self` with `letrec`, so why isn't it defined? The reason is... because of *hygiene*! Remember that Scheme's `syntax-rules` is hygienic, and for that reason, it transparently renames all identifiers introduced by macros such that they don't accidentally capture/get captured where the macro is expanded. It is possible to visualize this very precisely using the macro stepper of DrRacket. You will see that the identifier `self` in the `inc-by!` method is not the same color as the same identifier in the `letrec` expression.

Luckily for us, `defmac` supports a way to specify identifiers that can be used by the macro user code even though they are introduced by the macro itself. The only thing we need to do is therefore to specify that `self` is such an identifier:

```

(defmac (OBJECT ([field fname fval] ...)
              ([method mname mparams mbody ...] ...))
 #:keywords field method
 #:captures self
 (letrec
  ([self
   (let ([fname fval] ...)
     (let ([methods (list (cons 'mname (λ mparams mbody ...)) ...)])
       (λ (msg . args)
         (let ([found (assoc msg methods)])
           (if found
              (apply (cdr found) args)
              (error "message not understood:" msg)))))))]
    self))

> (let ([c (make-counter)])
    (→ c inc-by! 20))
20

```

## 2.3 Mutually-Recursive Methods

The previous section already shows that methods can use other methods by sending messages to `self`. This other example shows mutually-recursive methods.

```
(define odd-even
  (OBJECT ()
    ([method even (n) (match n
                        [0 #t]
                        [1 #f]
                        [else (→ self odd (- n 1))])]
     [method odd (n) (match n
                        [0 #f]
                        [1 #t]
                        [else (→ self even (- n 1))])]))))

> (→ odd-even odd 15)
#t
> (→ odd-even odd 14)
#f
> (→ odd-even even 1423842)
#t
```

Try the same definition in Java, and compare the results for "large" numbers. Yes, our small object system does enjoy the benefits of tail-call optimization! Can you explain why?

We now have an object system that supports `self`, including returning `self`, and sending messages to `self`. Notice how `self` is bound in methods at object creation time: when the methods are defined, they capture the binding of `self` and this binding is fixed from then on. We will see in the following chapters that this eventually does not work if we want to support delegation or if we want to support classes.

## 2.4 Nested Objects

Because objects and methods are compiled into lambdas in Scheme, our objects inherit interesting properties. First, as we have seen, they are first-class values (otherwise what would be the point?). Also, as we have just seen above, method invocations in tail position are treated as tail calls, and therefore space efficient. We now look at another benefit: we can use higher-order programming patterns, such as objects producing objects (usually called *factories*). That is, we can define *nested objects*, with proper lexical scoping.

Consider the following example:

```
(define counter-factory
  (OBJECT
    ([field default-count 0]
     [field default-step 1]))
```

```

([method df-count! (v) (set! default-count v)]
 [method df-step! (v) (set! default-step v)]
 [method make ()
  (OBJECT
   ([field count default-count]
    [field step default-step])
   ([method inc () (set! count (+ count step)) count]
    [method dec () (set! count (- count step)) count]
    [method reset () (set! count 0)]
    [method step! (v) (set! step v)]
    [method inc-by! (v) (→ self step! v) (→ self inc)])))]))

> (define c1 (→ counter-factory make))
> (→ c1 inc)
1
> (→ counter-factory df-count! 10)
> (→ counter-factory df-step! 5)
> (→ c1 inc)
2
> (define c2 (→ counter-factory make))
> (→ c2 inc)
15
> (→ c1 inc)
3
> (→ c2 inc)
20

```

Convince yourself that these results make sense.

## 3 Benefits and Limits of Objects

In the language course, we have been programming by defining a data type and its variants and then defining all the "services" over these structures using procedures that work by case analysis. This style of programming is sometimes called "procedural paradigm" or "functional design" (note that "functional" here does not refer to "side-effect free").

In PLAI, we have done this with `define-type` to introduce the data type and its variants, and using `type-case` in case-analyzing procedures. This procedural approach is common in other languages like C (unions), Pascal (variants), ML and Haskell's algebraic data types, or plain Scheme's tagged data.

So, what does object-oriented programming really bring us? What are its weaknesses? As it turns out, using an object-oriented language does not mean that programs are "object oriented". Many Java programs are not, or at least sacrifice some of the fundamental benefits of objects.

The aim of this intermediary chapter is to step back from our step-by-step construction of OOP to contrast objects with the procedural approach, with the aim of clarifying the pros and cons of each approach. Interestingly, the simple object system we have built so far is entirely sufficient to study the essential benefits and limits of objects—delegation, classes, inheritance, etc. are all interesting features, but are not *essential* to objects.

This chapter is based on the article *On Understanding Data Abstraction*, Revisited by William R. Cook (2009).

### 3.1 Abstract Data Types

Let us first look at abstract data types (ADTs). An ADT is a data type that hides its representation and only supplies operations to manipulate its values.

For instance, an *integer set* ADT can be defined as follows:

```
adt Set is
  empty : Set
  insert : Set x Int -> Set
  isEmpty? : Set -> Bool
  contains? : Set x Int -> Bool
```

There are many possible *representations* for such an integer set ADT. For instance, one could implement it with Scheme's lists:

```
(define empty '())

(define (insert set val)
  (if (not (contains? set val))
      (cons val set)))
```

```

    set))

(define (isEmpty? set) (null? set))

(define (contains? set val)
  (if (null? set) #f
      (if (eq? (car set) val)
          #t
          (contains? (cdr set) val))))

```

The following client program can then use ADT values, without being aware of the underlying representation:

```

> (define x empty)
> (define y (insert x 3))
> (define z (insert y 5))
> (contains? z 2)
#f
> (contains? z 5)
#t

```

We could as well implement the set ADT with another representation, such as using PLAI's `define-type` mechanism to create a variant type to encode a set as a linked list.

```

(define-type Set
  [mtSet]
  [aSet (val number?) (next Set?)])

(define empty (mtSet))

(define (insert set val)
  (if (not (contains? set val))
      (aSet val set)
      set))

(define (isEmpty? set) (equal? set empty))

(define (contains? set val)
  (type-case Set set
    [mtSet () #f]
    [aSet (v next)
      (if (eq? v val)
          #t
          (contains? next val))]))

```

The sample client program above runs exactly the same, even though the underlying representation is now changed:

```
> (define x empty)
> (define y (insert x 3))
> (define z (insert y 5))
> (contains? z 2)
#f
> (contains? z 5)
#t
```

## 3.2 Procedural Representations

We can as well consider sets as being defined by their *characteristic function*: a function that, given a number, tells us whether or not this number is part of the set. In that case, a set is simply a function `Int -> Bool`. (In PLAI, we saw that in Chapter 11, when studying the *procedural representation* of environments.)

What is the characteristic function of the empty set? a function that always returns false. And the set obtained by inserting a new element?

```
(define empty (λ (n) #f))

(define (insert set val)
  (λ (n)
    (or (eq? n val)
        (contains? set n))))

(define (contains? set val)
  (set val))
```

Because a set is represented by its characteristic function, `contains?` simply applies the function to the element. Note that the client program is again undisturbed:

```
> (define x empty)
> (define y (insert x 3))
> (define z (insert y 5))
> (contains? z 2)
#f
> (contains? z 5)
#t
```

What do we gain with the procedural representation of sets? flexibility! For instance, we can define the set of all even numbers:

```
(define even
  (λ (n) (even? n)))
```

It is not possible to fully represent this set in any of the ADT representations we considered above. (Why?) We can even define non-deterministic sets:

```
(define random
  (λ (n) (> (random) 0.5)))
```

With the procedural representation, we have much more freedom to define sets, and in addition, they can interoperate with existing set operations!

```
> (define a (insert even 3))
> (define b (insert a 5))
> (contains? b 12)
#t
> (contains? b 5)
#t
```

In contrast, with the ADT representations we have seen above, different representations *cannot* interoperate. A set-as-list value cannot be used by a set-as-struct operation, and vice versa. ADTs abstract the representation, but they only allow *a single representation* at a time.

### 3.3 Objects

In essence, *sets as functions are objects!* Note that objects do *not* abstract type: the type of a set-as-function is very concrete: it is a function `Int -> Bool`. Of course, as we have seen in the first chapters, an object is a generalization of a function in that it can have multiple methods.

#### 3.3.1 Object Interfaces

We can define a notion of *object interface* that gathers the signature of the methods of an object:

```
interface Set is
  contains? : Int -> Bool
  isEmpty? : Bool
```

Let us use our simple object system to implement sets as objects:

```

(define empty
  (OBJECT ()
    ([method contains? (n) #f]
     [method isEmpty? () #t])))

(define (insert s val)
  (OBJECT ()
    ([method contains? (n)
      (or (eq? val n)
          (-> s contains? n))]
     [method isEmpty? () #f])))

```

Note that `empty` is an object, and `insert` is a factory function that returns objects. A set object implements the `Set` interface. The `empty` object does not contain any value, and `isEmpty?` returns `#t`. `insert` returns a new object whose `contains?` method is similar to the set characteristic function we have seen before, and `isEmpty?` returns `#f`.

A client program is unchanged for the set construction part, and then has to use message sending to interact with set objects:

```

> (define x empty)
> (define y (insert x 3))
> (define z (insert y 5))
> (-> z contains? 2)
#f
> (-> z contains? 5)
#t

```

Note that object interfaces are essentially higher-order types: methods are functions, so passing objects around means passing groups of functions around. This is a generalization of higher-order functional programming. Object-oriented programs are inherently higher-order.

### 3.3.2 Principles of Object-Oriented Programming

**Principle:** *An object can only access other objects through their public interfaces*

Once we create an object, like the one bound to `z` above, the *only* thing we can do with it is to interact by sending messages. We cannot "open it". No attribute of the object is visible, only its interface. In other words:

**Principle:** *An object can only have detailed knowledge about itself*

This is fundamentally different from the way we program with ADT values: in a `type-case` analysis (recall the implementation of `contains?` in the ADT implementation with

`define-type`), one is opening the value and gaining direct access to its attributes. ADTs provide encapsulation, but for the clients of the ADT; not for its implementation. Objects go further in this regard. Even the implementation of the methods of an object cannot access attributes of objects other than itself.

From this we can derive another fundamental principle:

**Principle:** *An object is the set of observations that can be made upon it*, as defined by its object interface.

This is a strong principle, that says that if two objects behave the same for a certain experiment (ie., a number of observations), then they should be undistinguishable otherwise. This means that the use of identity-related operations (like pointer equality) are violating this principle of OOP. With `==` in Java, we can distinguish two objects that *are* different even though they *behave* in the same way.

### 3.3.3 Extensibility

The above principles can be considered the characteristic feature of OOP. As Cook puts it: *"Any programming model that allows inspection of the representation of more than one abstraction at a time is NOT object oriented"*

The Component Object Model (COM) is one of the purest OO programming model in practice. COM enforces all these principles: there is no built-in equality, there is no way to determine if an object is an instance of a given class. COM programs are therefore highly extensible.

Note that the extensibility of objects is in fact completely independent from inheritance! (We don't even have classes in our language) It instead comes from the use of interfaces.

### 3.3.4 What about Java?

Java is not a pure object-oriented language, not importantly because it has primitive types, but because it supports many operations that violate the principles we have described above. Java has primitive equality `==`, `instanceof`, casts to class types, that make it possible to distinguish two objects even though they behave the same. Java makes it possible to declare a method that accepts objects based on their classes, not their interfaces (in Java, a class name is also a type). And of course, Java allows objects to access the internals of other objects (public fields, of course, but even private fields are accessible by objects of the same class!).

This means that Java also supports ADT-style programming. There is no nothing wrong with that! But it is important to understand the design tradeoffs involved, to make an informed choice. For instance, in the JDK, certain classes respect OO principles on the surface (al-

lowing extensibility), but in their implementation using ADT techniques (not extensible, but more efficient). If you're interested, look at the [List](#) interface, and the [LinkedList](#) implementation.

Programming in "pure OO" in Java basically means not using class names as types (ie. use class names only after `new`), and never use primitive equality (`==`).

### 3.4 The Extensibility Problem

Object-oriented programming is often presented as the panacea in terms of extensible software. But what exactly is meant with "extensible"?

The extensibility problem is concerned with defining a data type (structure + operations) in a way that two kinds of extension are properly supported: adding new representational variants, or adding new operations.

As it turns out, ADTs and objects each nicely support one dimension of extensibility, but fail in the other. Let us study this with a well-known example: an interpreter of simple expressions.

#### 3.4.1 ADT

We first consider the ADT approach. We define a data type for expressions with three variants:

```
(define-type Expr
  [num (n number?)]
  [bool (b boolean?)]
  [add (l Expr?) (r Expr?)])
```

Now we can define the interpreter as a function that type-cases on the abstract syntax tree:

```
(define (interp expr)
  (type-case Expr expr
    [num (n) n]
    [bool (b) b]
    [add (l r) (+ (interp l) (interp r))]))
```

This is good-old PLAI practice. With a little example:

```
> (define prog (add (num 1)
                    (add (num 2) (num 3))))
> (interp prog)
6
```

Here, we use the term ADT in line with Cook's usage. It is important to clarify however that the discussion of the extensibility problem here actually contrasts objects with variant types (aka. algebraic data types). We are concerned with how to have an extensible *implementation*. The interface abstraction is not relevant here.

### Extension: New Operation

Let us now consider that we want to add a new operation on expressions. In addition to interpret an expression, we want to typecheck an expression, that is, to determine the type of value it will produce (here, either `number` or `boolean`). This is fairly trivial in our case, but still makes it possible to detect without interpretation that a program is bound to fail because it adds things that are not both numbers:

```
(define (typeof expr)
  (type-case Expr expr
    [num (n) 'number]
    [bool (b) 'boolean]
    [add (l r) (if (and (equal? 'number (typeof l))
                        (equal? 'number (typeof r)))
                    'number
                    (error "Type error: not a number")))]))
```

We can determine the type of the program we defined previously:

```
> (typeof prog)
'number
```

And see that our typechecker reject non-sensical programs:

```
> (typeof (add (num 1) (bool #f)))
Type error: not a number
```

If we reflect on this extension case, we see that it all went smoothly. We wanted a new operation, and just had to define a new function. This extension is modular, because it is defined in a single place.

### Extension: New Data

We now turn to the other dimension of extensibility: adding new data variants. Suppose we want to extend our simple language with a new expression: `ifc`. We extend the datatype definition:

```
(define-type Expr
  [num (n number?)]
  [bool (b boolean?)]
  [add (l Expr?) (r Expr?)]
  [ifc (c Expr?) (t Expr?) (f Expr?)])
```

Changing the definition of `Expr` to add this new variant breaks all existing function definitions! `interp` and `typeof` are now invalid, because they type case on expressions, but do not include any case for `ifc`. We need to modify them all to include the behavior associated to the `ifc` case:

```

(define (interp expr)
  (type-case Expr expr
    [num (n) n]
    [bool (b) b]
    [add (l r) (+ (interp l) (interp r))]
    [ifc (c t f)
     (if (interp c)
         (interp t)
         (interp f))]))

(define (typeof expr)
  (type-case Expr expr
    [num (n) 'number]
    [bool (b) 'boolean]
    [add (l r) (if (and (equal? 'number (typeof l))
                        (equal? 'number (typeof r)))
                   'number
                   (error "Type error: not a number"))]
    [ifc (c t f)
     (if (equal? 'boolean (typeof c))
         (let ((type-t (typeof t))
               (type-f (typeof f)))
           (if (equal? type-t type-f)
               type-t
               (error "Type error: both branches should have
same type"))))
         (error "Type error: not a boolean"))]))

```

This works:

```

> (define prog (ifc (bool false)
                   (add (num 1)
                        (add (num 2) (num 3)))
                   (num 5)))

> (interp prog)
5

```

This extensibility scenario was much less favorable. We had to modify the datatype definition and all the functions.

To summarize, with ADTs, adding new operations (eg. `typeof`) is easy and modular, but adding new data variants (eg. `ifc`) is cumbersome and non-modular.

### 3.4.2 OOP

How do objects perform in these extensibility scenarios?

First, we start with the object-oriented version of our interpreter:

```
(define (bool b)
  (OBJECT () ([method interp () b])))

(define (num n)
  (OBJECT () ([method interp () n])))

(define (add l r)
  (OBJECT () ([method interp () (+ (-> l interp)
                                   (-> r interp))])))
```

Note that, in line with OO design principles, each expression objects knows how to interpret itself. There is no more a centralized interpreter that deals with all expressions. Interpreting a program is done by sending the `interp` message to the program:

```
> (define prog (add (num 1)
                    (add (num 2) (num 3))))
> (-> prog interp)
6
```

#### Extension: New Data

Adding a new kind of data like a conditional `ifc` object can be done by simply defining a new object factory, with the definition of how these new objects handle the `interp` message:

```
(define (ifc c t f)
  (OBJECT () ([method interp ()
                (if (-> c interp)
                    (-> t interp)
                    (-> f interp))])))
```

We can now interpret programs with conditionals:

```
> (-> (ifc (bool #f)
           (num 1)
           (add (num 1) (num 3))) interp)
4
```

This case shows that, conversely to ADTs, adding new data variants with OOP is direct and modular: just create a new (kind of) object(s). This is a clear advantage of objects over ADTs.

## Extension: New Operation

But before concluding that OOP is the panacea for extensible software, we have to consider the other extension scenario: adding an operation. Suppose we now want to typecheck our programs, just as we did before. This means that expression objects should now also understand the "typeof" message. To do that, we actually have to modify all object definitions:

```
(define (bool b)
  (OBJECT () ([method interp () b]
              [method typeof () 'boolean])))

(define (num n)
  (OBJECT () ([method interp () n]
              [method typeof () 'number])))

(define (add l r)
  (OBJECT () ([method interp () (+ (-> l interp)
                                   (-> r interp))]
              [method typeof ()
                (if (and (equal? 'number (-> l typeof))
                          (equal? 'number (-> r typeof)))
                    'number
                    (error "Type error: not a number")))])))

(define (ifc c t f)
  (OBJECT () ([method interp ()
                (if (-> c interp)
                    (-> t interp)
                    (-> f interp))]
              [method typeof ()
                (if (equal? 'boolean (-> c typeof))
                    (let ((type-t (-> t typeof))
                          (type-f (-> f typeof)))
                      (if (equal? type-t type-f)
                          type-t
                          (error "Type error: both branches
should have same type")))
                    (error "Type error: not a boolean")))])))
```

We can if this works:

```
> (-> (ifc (bool #f) (num 1) (num 3)) typeof)
'number
> (-> (ifc (num 1) (bool #f) (num 3)) typeof)
Type error: not a boolean
```

This extensibility scenario forced us to modify all our code to add the new methods.

To summarize, with objects, adding new data variants (eg. `ifc`) is easy and modular, but adding new operations (eg. `typeof`) is cumbersome and non-modular.

Note that this is just the dual situation of ADTs!

### 3.5 Different Forms of Data Abstraction

ADTs and objects are different forms of data abstraction, each with its own advantages and drawbacks.

ADTs have a private representation type that prohibits tampering and extension. This is good for reasoning (analysis) and optimization. But it only permits one representation at a time.

Objects have behavioral interfaces, which allow definition of new implementations at any time. This is good for flexibility and extensibility. But it makes it hard to analyze code, and makes certain optimizations impossible.

Both forms of abstraction also support different forms of modular extensions. It is possible to modularly add new operations on an ADT, but supporting new data variants is cumbersome. It is possible to modularly add new representations to an object-oriented system, but adding new operations implies crosscutting modifications.

There are ways to navigate this tradeoff. For instance, one can expose certain implementation details in the interface of an object. That sacrifices some extensibility, but recovers the possibility to do some optimizations. The fundamental question is therefore a design question: what do I really need?

Now you understand why many languages support both kinds of data abstraction.

Cook's paper goes more in depth in the comparison of these forms of data abstraction. It is definitely a must-read!

## 4 Forwarding and Delegation

Using message sending, an object can always forward a message to another object in case it does not know how to handle it. With our small object system, we can do that explicitly as follows:

```
(define seller
  (OBJECT ()
    ([method price (prod)
      (* (case prod
          ((1) (→ self price1))
          ((2) (→ self price2)))
        (→ self unit))])
    [method price1 () 100]
    [method price2 () 200]
    [method unit () 1])))

(define broker
  (OBJECT
    ([field provider seller])
    ([method price (prod) (→ provider price prod)])))

> (→ broker price 2)
200
```

Object `broker` does not know how to compute the price of a product, but it can claim to do so by implementing a method to handle the `price` message, and simply forwarding it to `seller`, who does implement the desired behavior. Note how `broker` holds a reference to `seller` in its `provider` field. This is a typical example of object composition, with message forwarding.

Now, you can see that the problem with this approach is that this forwarding of messages has to be explicit: for each message that we anticipate might be sent to `broker`, we have to define a method that forwards the message to `seller`. For instance:

```
> (→ broker unit)
message not understood: 'unit'
```

### 4.1 Message Forwarding

We can do better by allowing each object to have a special "partner" object to which it automatically forwards any message it does not understand. We can define a new syntactic abstraction, `OBJECT-FWD`, for constructing such objects:

```

(defmacro (OBJECT-FWD target
           ([field fname fval] ...)
           ([method mname mparams mbody ...] ...))
  #:keywords field method
  #:captures self
  (letrec ([self
            (let ([fname fval] ...)
              (let ([methods (list (cons 'mname (λ mparams mbody ...)) ...))]
                (λ (msg . args)
                  (let ([found (assoc msg methods)])
                    (if found
                        (apply (cdr found) args)
                        (apply target msg args)))))))]
    self))

```

Note how the syntax is extended to specify a `target` object; this object is used in the dispatch process whenever a message is not found in the methods of the object. Of course, if all objects forward unknown message to some other object, there has to be a last object in the chain, that simply fails when sent a message:

```

(define root
  (λ (msg . args)
    (error "message not understood" msg)))

```

Now, `broker` can be defined simply as follows:

```

(define broker
  (OBJECT-FWD seller () ()))

```

That is, `broker` is an empty object (no fields, no methods) that forwards all messages sent to it to `seller`:

```

> (→ broker price 2)
200
> (→ broker unit)
1

```

This kind of objects is often called a *proxy*.

## 4.2 Delegation

Suppose we want to use `broker` to *refine* the behavior of `seller`; say, we want to double the price of every product, by changing the unit used in the calculation of the prices. This is easy: we just have to define a method `unit` in `broker`:

```
(define broker
  (OBJECT-FWD seller ()
    ([method unit () 2])))
```

With this definition, we should make sure that asking the price of a product to `broker` is twice the price of the same product asked to `seller`:

```
> (→ broker price 1)
100
```

Hmmm... it does not work! It seems that once we forward the `price` message to `seller`, `broker` never gets the control back; in particular, the `unit` message that `seller` sends to `self` is *not* received by `broker`.

Let us consider why this is so. To which object is `self` bound in `seller`? To `seller`! Remember, we said previously (see §2 “Looking for the Self”) that in our approach, `self` is *statically bound*: when an object is created, `self` is made to refer to the object/closure that is being defined, and will always remain bound to it. This is because `letrec`, like `let`, respects lexical scoping.

What we are looking for is another semantics, called *delegation*. Delegation requires `self` in an object to be *dynamically bound*: it should always refer to the object that originally received the message. In our example, this would ensure that when `seller` sends `unit` to `self`, then `self` is bound to `broker`, and thus the re-definition of `unit` in `broker` takes effect. In that case, we say that `seller` is the *parent* of `broker`, and that `broker` *delegates* to its parent.

How do we bind an identifier such that it refers to the value at the point of usage, rather than at the point of definition? In the absence of dynamically-scoped binding forms, the only way we can achieve this is by passing that value as parameter. So, we have to parameterize methods by the actual receiver. Therefore, instead of capturing the `self` identifier in their surrounding lexical scope, they are parameterized by `self`.

Concretely, this means that the method:

```
(λ (prod) .... (→ self unit) ....)
```

in `seller` must be kept in the methods list as:

```
(λ (self prod)....(→ self unit)....)
```

This parameterization effectively allows us to pass the current receiver after we lookup the method.

Let us now define a new syntactic form, `OBJECT-DEL`, to support the delegation semantics between objects:

Ever wondered why methods in Python must explicitly receive `self` as a first parameter? (In Java, and all other languages, they do as well, although you don't see it explicitly.)

```

(defmac (OBJECT-DEL parent
        ([field fname fval] ...)
        ([method mname (mparam ...) mbody ...] ...))
  #:keywords field method
  #:captures self
  (let ([fname fval] ...)
    (let ([methods
          (list (cons 'mname (lambda (self mparam ...) mbody ...)) ...)])
      (lambda (current)
        (lambda (msg . args)
          (let ([found (assoc msg methods)])
            (if found
                (apply (cdr found) (cons current args))
                (apply (parent current) msg args))))))))))

```

Several things have changed: first, we renamed `target` to `parent`, to make it clear that we are defining a delegation semantics. Second, all methods are now parameterized by `self`, as explained above. Note that we got rid of `letrec` altogether! This is because `letrec` was used precisely to allow objects to refer to `self`, but following lexical scoping. We have seen that for delegation, lexical scoping of `self` is not what we want.

This means that when we find a method in the method dictionary, we must first give it the actual receiver as argument. How are we going to obtain that receiver? well, the only possibility is to parameterize objects by the current receiver they have to use when applying methods. That is to say, the value returned by the object construction form is not a "`lambda (msg . vals) ....`" anymore, but a "`lambda (rcvr) ....`". This effectively parameterizes our objects by "the current receiver". Similarly, if a message is not understood by a given object, then it must send the current receiver along to its parent.

This leaves us with one final question to address: how do we send a message to an object in the first place? Remember that our definition of `→` is:

```

(defmac (→ o m arg ...)
  (o 'm arg ...))

```

But now we cannot apply `o` as a function that takes a symbol (the message) and a variable number of arguments. Indeed, an object now is a function of the form `(lambda (rcvr) (lambda (msg . args) ...))`. So before we can pass the message and the arguments, we have to specify which object is the current receiver. Well, it's easy, because at the time we are sending a message, the current receiver should be... the object we are sending the message to!

Why is the `let` binding necessary?

```

(defmac (→ o m arg ...)
  (let ([obj o])
    ((obj obj) 'm arg ...)))

```

Let us see delegation—that is, late binding of self—at work:

```
(define seller
  (OBJECT-DEL root ()
    ([method price (prod)
      (* (case prod
          [(1) (→ self price1)]
          [(2) (→ self price2)])
        (→ self unit))])
    [method price1 () 100]
    [method price2 () 200]
    [method unit () 1])))
(define broker
  (OBJECT-DEL seller ()
    ([method unit () 2])))

> (→ seller price 1)
100
> (→ broker price 1)
200
```

### 4.3 Programming with Prototypes

Object-based languages with a delegation mechanism like the one we have introduced in this chapter are called *prototype-based languages*. Examples are Self, JavaScript, and AmbientTalk, among many others. What are these languages good at? How to program with prototypes?

#### 4.3.1 Singleton and Exceptional Objects

Since objects can be created *ex-nihilo* (ie. out of an object literal expression like `OBJECT-DEL`), it is natural to create one-of-a-kind objects. As opposed to class-based languages that require a specific design pattern for this (called Singleton), object-based languages are a natural fit for this case, as well as for creating "exceptional" objects (more on this below).

Let us first consider the object-oriented representation of booleans and a simple if-then-else control structure. How many booleans are there? well, two: true, and false. So we can create two standalone objects, `true` and `false` to represent them. In pure object-oriented languages like Self and Smalltalk, control structures like if-then-else, while, etc. are not primitives in the language. Rather, they are defined as methods on appropriate objects. Let us consider the if-then-else case. We can pass two thunks to a boolean, a truth thunk and a falsity thunk; if the boolean *is* true, it applies the truth thunk; if it *is* false, it applies the falsity thunk.

```

(define true
  (OBJECT-DEL root ()
    ([method ifTrueFalse (t f) (t)])))

(define false
  (OBJECT-DEL root ()
    ([method ifTrueFalse (t f) (f)])))

```

How can we use such objects? Look at the following example:

```

(define light
  (OBJECT-DEL root
    ([field on false])
    ([method turn-on () (set! on true)]
     [method turn-off () (set! on false)]
     [method on? () on])))

> (→ (→ light on?) ifTrueFalse (λ () "light is on")
      (λ () "light is off"))
"light is off"
> (→ light turn-on)
> (→ (→ light on?) ifTrueFalse (λ () "light is on")
      (λ () "light is off"))
"light is on"

```

The objects `true` and `false` are the two unique representants of boolean values. Any conditional mechanism that relies on some expression being true or false can be similarly defined as methods of these objects. This is indeed a nice example of dynamic dispatch!

Boolean values and control structures in Smalltalk are defined similarly, but because Smalltalk is a class-based language, their definitions are more complex. Try it in your favorite class-based language.

Let us look at another example where object-based languages are practical: exceptional objects. First, recall the definition of typical point objects, which can be created using a factory function `make-point`:

```

(define (make-point x-init y-init)
  (OBJECT-DEL root
    ([field x x-init]
     [field y y-init])
    ([method x? () x]
     [method y? () y])))

```

Suppose we want to introduce an exceptional point object that has the particularity of having *random* coordinates, that change each time they are accessed. We can simply define this

`random-point` as a standalone object whose `x?` and `y?` methods perform some computation, rather than accessing stored state:

```
(define random-point
  (OBJECT-DEL root ()
    ([method x? () (* 10 (random))]
     [method y? () (→ self x?)]))))
```

Note that `random-point` does not have any fields declared. Of course, because in OOP we rely on object interfaces, both representations of points can coexist.

### 4.3.2 Sharing through Delegation

The examples discussed above highlight the advantages of object-based languages. Let us now look at delegation in practice. First, delegation can be used to factor out *shared behavior* between objects. Consider the following:

```
(define (make-point x-init y-init)
  (OBJECT-DEL root
    ([field x x-init]
     [field y y-init])
    ([method x? () x]
     [method y? () y]
     [method above (p2)
      (if (> (→ p2 y?) (→ self y?))
          p2
          self)]
     [method add (p2)
      (make-point (+ (→ self x?)
                    (→ p2 x?))
                 (+ (→ self y?)
                    (→ p2 y?)))]))
```

All created point objects have the same methods, so this behavior could be shared by moving it to a common parent of all point objects (often called a prototype). Should all behavior be moved in the prototype? well, not if we want to allow different representations of points, like the random point above (which does not have any field at all!).

Therefore, we can define a `point` prototype, which factors out the `above` and `add` methods, whose implementation is common to all points:

```
(define point
  (OBJECT-DEL root ()
    ([method above (p2)
```

```

      (if (> (→ p2 y?) (→ self y?))
          p2
          self)]
[method add (p2)
  (make-point (+ (→ self x?)
                (→ p2 x?))
              (+ (→ self y?)
                (→ p2 y?))))]]))

```

Note that as a standalone object, `point` does not make sense, because it sends messages to itself that it does not understand. But it can serve as a prototype from which different points can extend. Some are typical points, created with `make-point`, which hold two fields `x` and `y`:

```

(define (make-point x-init y-init)
  (OBJECT-DEL point
    ([field x x-init]
     [field y y-init])
    ([method x? () x]
     [method y? () y])))

```

The required accessor methods could be declared as `abstract` methods in `point`, if our language supported such a concept. In Smalltalk, one would define the methods in `point` such that they throw an exception if invoked.

While some can be exceptional points:

```

(define random-point
  (OBJECT-DEL point ()
    ([method x? () (* 10 (random))]
     [method y? () (→ self x?)])))

```

As we said, these different kinds of point can cooperate, and they all understand the messages defined in the `point` prototype:

```

> (define p1 (make-point 1 2))
> (define p2 (→ random-point add p1))
> (→ (→ p2 above p1) x?)
9.852532706532351

```

We can similarly use delegation to share *state* between objects. For instance, consider a family of points that share the same x-coordinate:

```

(define 1D-point
  (OBJECT-DEL point
    ([field x 5])
    ([method x? () x]
     [method x! (nx) (set! x nx)])))

```

```
(define (make-point-shared y-init)
  (OBJECT-DEL 1D-point
    ([field y y-init])
    ([method y? () y]
     [method y! (ny) (set! y ny)])))
```

All objects created by `make-point-shared` share the same parent, `1D-point`, which determines their x-coordinate. If a change to `1D-point` is made, it is naturally reflected in all its children:

```
> (define p1 (make-point-shared 2))
> (define p2 (make-point-shared 4))
> (→ p1 x?)
5
> (→ p2 x?)
5
> (→ 1D-point x! 10)
> (→ p1 x?)
10
> (→ p2 x?)
10
```

## 4.4 Late Binding of Self and Modularity

In the definition of the `OBJECT-DEL` syntactic abstraction, notice that we use, in the definition of message sending, the self-application pattern (`obj obj`). This is similar to the self application pattern we have seen to achieve recursive binding without mutation.

See the Why of Y.

This feature of OOP is also known as "open recursion": any sub-object can redefine the meaning of a method in one of its parents. Of course, this is a mechanism that favors *extensibility*, because it is possible to extend any aspect of an object without having to foresee that extension. On the other hand, open recursion also makes software more *fragile*, because it becomes extremely easy to extend an object in unforeseen, incorrect ways. Imagine scenarios where this can be problematic and think about possible alternative designs. To shed some more light on fragility, think about black-box composition of objects: taking two objects, developed independently, and then putting them in a delegation relation with each other. What issues can arise?

Think about what you know from mainstream languages like C++ and Java: what means do these two languages provide to address this tradeoff between extensibility and fragility?

## 4.5 Lexical Scope and Delegation

As we have seen previously, we can define nested objects in our system. It is interesting to examine the relation between lexical nesting and delegation. Consider the following exam-

ple:

```
(define parent
  (OBJECT-DEL root ()
    ([method foo () 1])))

(define outer
  (OBJECT-DEL root
    ([field foo (λ () 2)])
    ([method foo () 3]
     [method get ()
      (OBJECT-DEL parent ()
        ([method get-foo1 () (foo)]
         [method get-foo2 () (→ self foo)]))]))))

(define inner (→ outer get))

> (→ inner get-foo1)
2
> (→ inner get-foo2)
1
```

As you can see, a free identifier is looked up in the lexical environment (see `get-foo1`), while an unknown message is looked up in the chain of delegation (see `get-foo2`). This is important to clarify, because Java programmers are used to the fact that `this.foo()` is the same as `foo()`. In various languages that combine lexical nesting and some form of delegation (like inheritance), this is not the case.

So what happens in Java? Try it! You will see that the inheritance chain shadows the lexical chain: when using `foo()` if a method can be found in a superclass, it is invoked; only if there is no method found, the lexical environment (i.e., the outer object) is used. Referring to the outer object is therefore very fragile. This is why Java also supports an additional form `Outer.this` to refer to the enclosing object. Of course, then, if the method is not found directly in the enclosing object's class, it is then looked up in its superclass, rather than upper in the lexical chain.

Other languages have different take on the question. Check out Newspeak for instance.

## 4.6 Delegation Models

The delegation model we have implemented here is but one point in the design space of prototype-based languages. Study the documentation of Self, JavaScript, and AmbientTalk to understand their designs. You can even modify our object system to support a different model, such as the JavaScript model.

## 4.7 Cloning

In our language, as in JavaScript, the way to create objects is *literally*: either we create an object from scratch, or we have a function whose role is to perform these object creations for us. Historically, prototype-based languages (like Self) have provided another way to create objects: by cloning existing objects. This approach emulates the copy-paste-modify metaphor we use so often with text (including code!): start from an object that is almost as you need, clone it, and modify the clone (eg. add a method, change a field).

When cloning objects in presence of delegation, the question of course arises of whether the cloning operation should be *deep* or *shallow*. Shallow cloning returns a new object that delegates to the same parent as the original object. Deep cloning returns a new object that delegates to a clone of the original parent, and so on: the whole delegation chain is cloned.

We won't study cloning in more details in this course. You should however wonder how easy it would be to support cloning in our language. Since objects are in fact compiled into procedures (through macro expansion), the question boils down to cloning closures. Unfortunately, Scheme does not support such an operation. This is a case where the mismatch between the source and target languages shows up (recall Chapter 11 of PLAI). Nothing is perfect!

## 5 Classes

Let's go back to the factory function (see §1.3 “Constructing Objects”):

```
(define (make-counter [init-count 0] [init-step 1])
  (OBJECT
    ([field count init-count]
     [field step  init-step])
    ([method inc () (set! count (+ count step)) count]
     [method dec () (set! count (- count step)) count]
     [method reset () (set! count 0)]
     [method step! (v) (set! step v)]
     [method inc-by! (v) (→ self step! v) (→ self inc)])))

(define c1 (make-counter))
(define c2 (make-counter 10))
```

Each counter object get its own version of the methods, even though they are the same. Well, at least their signature and body are the same.

Are they the completely the same, though?

Well, methods are functions, so as values, they are closures: function definitions together with their lexical environment. In particular, each method closes over the `self` of each object: i.e., `self` in a method in `c1` refers to `c1`, while it refers to `c2` in the method of `c2`. Likewise for fields, which are also let-bound variables in scope for the method closures. In other words, methods differ by the lexical environment they capture.

### 5.1 Sharing Method Definitions

Instead of duplicating all method definitions just to be able to support different selves and field values, it makes much more sense to factor out the common part (the method bodies), and parameterize them by the variable part (the object bound to `self`).

Let us try first without macros. Recall that our definition of a plain counter object without macros is as follows (slightly simplified):

```
(define make-counter
  (λ ([init-count 0] [init-step 1])
    (letrec
      ([self
        (let ([count init-count]
              [step  init-step])
          (let ([methods
                (list
                 (cons 'inc (λ () (set! count (+ count step)))
```

```

                                count))
      (cons 'step! (λ (v) (set! step v)))
      (cons 'inc-by! (λ (n) (self 'step! n)
                (self 'inc))))])
  (λ (msg . args)
    (let ([found (assoc msg methods)])
      (if found
          (apply (cdr found) args)
          (error "message not understood:" msg))))))
  self)))

```

If we hoist the `(let ([methods...])` out of the top-level `λ` (the factory), and parametrize them by `self` (as an additional first argument), we effectively achieve the sharing of method definitions at the factory level we are looking for:

```

(define make-counter
  (let ([methods
        (list
         (cons 'inc (λ (self) (set! count (+ count step)) count))
         (cons 'step! (λ (self v) (set! step v)))
         (cons 'inc-by! (λ (self n) (self 'step! n)
                        (self 'inc))))])
    (λ ([init-count 0] [init-step 1])
      (letrec
        ([self
         (let ([count init-count]
              [step init-step])
           (λ (msg . args)
             (let ([found (assoc msg methods)])
               (if found
                   (apply (cdr found) (cons self args))
                   (error "message not understood:" msg))))))]
          self))))

```

Note how each method takes an extra first parameter. In `inc-by!`, this allows us to invoke methods on `self`. Of course, when we apply the method in the dispatcher, we need to pass it the current `self` as extra argument.

The problem is that field variables are now out of scope of method bodies. More concretely, here, it means that `count` and `step` are free in the method bodies. So we need to parameterize methods by state (field values) as well, in addition to `self`. But, fair enough, `self` can "hold" the state (it can capture field bindings in its lexical environment). We just need a way to extract (and potentially assign) field values through `self`.

## 5.2 Accessing Fields

If we only pass `self` as extra argument to methods, we need a way to access an object's fields by sending it messages (since `self` is an object, and an object is just a dispatcher function). To this end, we introduce a field access protocol consisting of two messages `-read` and `-write`. The leading dash indicates that these messages are "meta" messages, not standard messages that need to be interpreted by a corresponding method.

These meta-messages take as first argument the field name to be accessed. The issue is then how to go from a field name (as a symbol) to actually reading/assigning the variable with the same name in the lexical environment of the object. A simple solution is to use a structure to hold field values. This is similar to the way we handle method definitions already: an association between method names and method definitions. However, unlike in a method table, field bindings are (at least potentially) mutable. Racket does not allow mutation in association lists, so we will use a dictionary (more precisely, a hashtable), which is accessed with `dict-ref` and `dict-set!`.

```
(define make-counter
  (let ([methods
        (list
         (cons 'inc (lambda (self)
                     (self '-write 'count
                            (+ (self '-read 'count) (self '-read 'step)))
                     (self '-read 'count)))
         (cons 'step! (lambda (self v) (self '-write 'step v)))
         (cons 'inc-by! (lambda (self n) (self 'step! n)
                        (self 'inc))))))
    (lambda ([init-count 0] [init-step 1])
      (letrec ([self
                (let ([fields (make-hash (list (cons 'count init-
count)
                                              (cons 'step init-
step)))]
                  (lambda (msg . args)
                    (match msg
                     ['-read (dict-ref fields (first args))]
                     ['-write (dict-set! fields (first args) (second args))]
                     [_ (let ([found (assoc msg methods)])
                          (if found
                              (apply (cdr found) (cons self args))
                              (error "message not
understood:" msg)))))))]
                self))))))
```

```
> (let ([c (make-counter)])
      (c 'inc-by! 10))
10
```

Note how `make-counter` now holds the list of parametrized `methods`, and the created object captures a dictionary of `fields`, which is initialized prior to returning the object. In the method bodies, field accesses are now implemented by sending `-read` and `-write` meta-messages, interpreted accordingly in the dispatcher.

### 5.3 Classes

While we did achieve the sharing of method definitions we were after, our solution is still not very satisfactory. Why? Well, observe the definition of an object (the body of the `(λ (msg . args) ...)` above). The logic that is implemented there is, again, repeated in all objects we create with `make-counter`: each object has its own copy of what to do when it is sent a `-read` message (lookup in the `fields` dictionary), a `-write` message (assign in the `fields` dictionary), or any other message (looking in the `methods` table and then applying the method).

So, all this logic could very well be shared amongst objects. The only free variables in the object body are `fields` and `self`. In other words, we could define an object as being just its self as well as its fields, and leave all the other logic to the `make-counter` function. In that case `make-counter` starts to have more than one responsibility: it is no longer only in charge of creating new objects, it is also in charge of handling accesses to fields and message handling. That is, `make-counter` is now evolving into what is called a *class*.

How are we going to represent a class? well, for now it is just a function that we can apply (and it creates an object—an *instance*); if we need that function to have different behaviors, we can apply the same Object Pattern we saw at the beginning of this course. That is:

```
(define Point
  ....
  (λ (msg . args)
    (match msg
      ['-create create instance]
      ['-read read field]
      ['-write write field]
      ['-invoke invoke method])))
```

This pattern makes clear what the role of a class is: it produces objects, and invokes methods, reads and writes fields on its instances.

So, what is the role of an object now? Well, it is just to exist as a first-class value, know its class, and hold the values of its fields. It does not hold any behavior on its own, anymore. In other words, we can define an object as a plain data structure:

This interpretation of `-write` means that setting a non-existent field will add it to the object (because that's how `dict-set!` works). This is the semantics adopted by Python, for instance. In contrast, in Java and Scala, for instance, setting a non-existent field is an error (detected statically thanks to typing). How would you modify the code above to raise an error when setting a non-existent field?

```
(struct obj (class values))
```

As simple as this! From now on, an object will just be such a structure.

Let us see exactly how we can define the class `Counter` now:

```
(define Counter
  (let ([methods ...])
    (letrec
      ([class
        (λ (msg . args)
          (match msg
            ['-create (let ([values (make-hash '((count .
0) (step . 1)))]))
                        (obj class values))]
            ['-read (dict-ref (obj-values (first args))
                              (second args))]
            ['-write (dict-set! (obj-values (first args))
                                (second args)
                                (third args))]
            ['-invoke
              (let ([found (assoc (first args) methods)])
                (if found
                    (apply (cdr found) (rest args))
                    (error "message not
understood:" (first args))))))]
        class)))
```

We can instantiate the class `Counter` by sending it the `create` message.

```
> (Counter '-create)
#<obj>
```

Now that an object is a structure, we need a different approach to sending messages to it, as well as to access its fields. To send a message to an object `p`, we must first retrieve its class, and then send the `invoke` message to the class:

```
((obj-class c) '-invoke 'inc c)
```

And similarly for reading and accessing fields.

## 5.4 Embedding Classes in Scheme

Let us now embed classes in Scheme using macros.

It is the first time we use `struct` in these notes: it is a convenient Racket macro to define datatypes, which automatically generates a constructor (here, `obj`) and accessors for each of the field of the structure (here, `obj-class`, `obj-values`).

As we did for `self` before, the `class` identifier is also defined using `letrec`: can you see why?

Note that representing classes as dispatcher classes following the Object Pattern is clearly not the only design alternative here. We could as well push the notion of classes-as-objects further, in a uniform manner (ie. what is the class of a class?), as done in Smalltalk. We could also trim-down classes as inert structures, using plain

### 5.4.1 Macro for Classes, Take 1

We can define a `CLASS` syntactic abstraction for creating classes:

```
(defmac (CLASS ([field fname fval] ...)
              ([method mname (mparam ...) mbody ...] ...))
  #:keywords field method
  #:captures self
  (let ([methods
        (list (cons 'mname (λ (self mparam ...) mbody ...)) ...)])
    (letrec
      ([class
        (λ (msg . args)
          (match msg
            ['-create
             (obj class
              (make-hash (list (cons 'fname fval) ...)))]
            ['-read
             (dict-ref (obj-values (first args)) (second args))]
            ['-write
             (dict-set! (obj-values (first args)) (second args) (third args))]
            ['-invoke
             (let ([found (assoc (second args) methods)])
               (if found
                 (apply (cdr found) (rest args))
                 (error "message not
understood:" (second args))))))]
        class)))
```

This definition follows the schema presented above, where a class is responsible for the interpretation of all the OO mechanisms (instance creation, field accesses, method invocation).

The syntactic macro for method invocation would then be:

```
(defmac (→ o m arg ...)
  (let ([obj o])
    ((obj-class obj) '-invoke 'm obj arg ...)))
```

Why is the  
`let`-binding  
necessary?

### 5.4.2 Macro for Classes, Take 2

In the definition of `CLASS` above, the interpretation of both `-read` and `-write` have nothing to do with the class itself. Their interpretation just consists in accessing the vector of values of an object. Therefore, we could move that behavior out of classes themselves, and do it in the syntactic macros for field accesses.

Likewise, a method invocation consists of two steps: looking up the method in the `methods` dictionary of the class, and applying it. Of these steps, only the first one is specific to a given class; application per se is common to any class, and could therefore be handled by the syntactic macro for method invocation as well.

We can therefore use a more lightweight `CLASS` macro, which only handles the behavior specific to a class, and leaves the rest to the auxiliary macros.

```
(defmac (CLASS ([field fname fval] ...)
             ([method mname (mparam ...) mbody ...] ...))
  #:keywords field method
  #:captures self
  (let ([methods
        (list (cons 'mname (lambda (self mparam ...) mbody ...)) ...)])
    (letrec
      ([class
        (lambda (msg . args)
          (match msg
            ['-create
             (obj class (make-hash (list (cons 'fname fval) ...)))]
            ['-lookup
             (let ([found (assoc (first args) methods)])
               (if found
                 (cdr found)
                 (error "message not
understood:" (first args)))))))]
        class)))
```

We now define the convenient syntax to invoke methods (`→`), and introduce similar syntax for accessing the fields of the current object (`? and !`).

```
(defmac (→ o m arg ...)
  (let ([obj o])
    (((obj-class obj) '-lookup 'm) obj arg ...)))

(defmac (? f) #:captures self
  (dict-ref (obj-values self) 'f))

(defmac (! f v) #:captures self
  (dict-set! (obj-values self) 'f v))
```

We also define an auxiliary function to create new instances:

```
(define (new c)
  (c '-create))
```

This simple function is conceptually very important: it helps to hide the fact that classes are internally implemented as functions, as well as the actual symbol used to ask a class to create an instance.

Why don't we need to define `new` as a macro?

### 5.4.3 Example

Let us see classes at work:

```
(define Counter
  (CLASS
    ([field count 0]
     [field step 1])
    ([method inc () (! count (+ (? count) (? step))) (? count)]
     [method dec () (! count (- (? count) (? step))) (? count)]
     [method reset () (! count 0)]
     [method step! (v) (! step v)]
     [method inc-by! (v) (→ self step! v) (→ self inc)])))

> (define c1 (new Counter))
> (define c2 (new Counter))
> (→ c1 inc-by! 10)
10
> (→ c2 inc-by! 20)
20
```

### 5.4.4 Strong Encapsulation

We have made an important design decision with respect to field accesses: field accessors `?` and `!` only apply to `self!` i.e., it is not possible in our language to access fields of another object. This is called a language with *strongly-encapsulated* objects. Smalltalk follows this discipline (accessing a field of another object is actually a message send, which can therefore be controlled by the receiver object). Java does not: it is possible to access the field of any object (provided visibility allows it); JavaScript even less! Here, our *syntax* simply does not allow foreign field accesses.

What happens in this language if we read an undeclared field? If we assign to an undeclared field? Why? Explore variations out there (ie. Java vs. JavaScript) and in your implementation.

Another consequence of our design choice is that field accesses should only occur *within* method bodies: because the receiver object is always `self`, `self` must be defined. For instance, look at what happens if we use the field read form `?` outside of an object:

```
> (? count)
self: undefined;
cannot reference an identifier before its definition
in module: 'program'
```

It would be much better if the above could yield an error saying that `?` is undefined. In order to do this, we can simply introduce `?` and `!` as *local* syntactic forms, only defined within the confines of method bodies, instead of global ones. We do that by moving the definition of these field access forms from the top-level to a *local* scope, surrounding method definitions:

```
(defmac (CLASS ([field fname fval] ...)
             ([method mname (mparam ...) mbody ...] ...))
  #:keywords field method
  #:captures self ? !
  (let ([methods
        (local [(defmac (? f) #:captures self
                    (dict-ref (obj-values self) 'f))
                (defmac (! f v) #:captures self
                    (dict-set! (obj-values self) 'f v))]
          (list (cons 'mname (lambda (self mparam ...) mbody ...)) ...))]
        (letrec
          ([class (lambda (msg . vals) ...)])))))
```

Defining the syntactic forms `?` and `!` locally, for the scope of the definition of the list of methods only, ensures that they are available to use within method bodies, but nowhere else.

Now, field accessors are no longer defined outside of methods:

```
> (? count)
?: undefined;
  cannot reference an identifier before its definition
  in module: 'program'
```

From now on, we will use this local approach.

## 5.5 Initialization

As we have seen, the way to obtain an object from a class, i.e., to instantiate it, is to send the class the `create` message. It is generally useful to be able to pass arguments to `create` in order to specify the initial values of the fields of the object. For now, our class system only supports the specification of default field values at class-declaration time. It is not possible to pass initial field values at instantiation time.

There are many ways to do this. Here we implement a simple mechanism. First, the `new` function takes optional arguments:

```
(define (new class . init-vals)
  (apply class (cons '-create init-vals)))
```

First, the new object is created using the default values declared for each field. Then, if extra arguments are passed to `new`, the `initialize` method is invoked on the new object. This method can then set other values to fields, or do whatever initialization work is desired.

```
....
(λ (msg . args)
  (match msg
    ['-create
      (let ([o (obj class (make-hash (list (cons 'fname fval) ...)))]
            (when (not (empty? args))
              (let ([found (assoc 'initialize methods)])
                (if found
                  (apply (cdr found) (cons o args))
                  (error "initialize not implemented in:" class))))
            o])
          ....)) ....
```

Let us see if it works as expected:

```
(define Counter
  (CLASS
    ([field count 0]
     [field step 1])
    ([method initialize ([cnt 0] [stp 1]) (! count cnt) (! step stp)]
     [method inc () (! count (+ (? count) (? step))) (? count)]
     [method dec () (! count (- (? count) (? step))) (? count)]
     [method reset () (! count 0)]
     [method step! (v) (! step v)]
     [method inc-by! (v) (→ self step! v) (→ self inc)])))

> (define c1 (new Counter))
> (define c2 (new Counter 5))
> (define c3 (new Counter 5 2))
> (→ c1 inc)
1
> (→ c2 inc)
6
> (→ c3 inc)
7
```

(Note that we are able to use Racket's optional argument mechanism without any problem, for `initialize`, cute!)

This object initialization mechanism is somewhat limited. You can study the variety of mechanisms in languages out there, and implement your own take on the matter. In particular, with inheritance, initialization can become quite subtle to get right.

## 5.6 Anonymous, Local and Nested Classes

We have introduced classes in our extension of Scheme, in such a way that classes are, like objects in our earlier systems, represented as first-class functions. This means therefore that classes in our language are first-class entities, which can, for instance, be passed as parameter (see the definition of the `create` function above) and constructed dynamically. Other consequences are that our system also supports both anonymous and nested classes. Of course, all this is achieved while respecting the rules of lexical scoping.

For instance, we can introduce classes in a local scope. That is, as opposed to languages like Java where classes are first-order entities that are globally visible, we are able to define classes locally.

```
(define doubleton
  (let ([cls (CLASS ([field x 0])
                    ([method initialize (v) (! x v)]
                     [method x? () (? x)])))]
    (cons (new cls 4) (new cls 8))))

> (+ (→ (car doubleton) x?) (→ (cdr doubleton) x?))
12
```

In the above we introduce `cls` only for the sake of creating two instances and returning them in a pair. After that point, the class is not accessible anymore. In other words, it is impossible to create more instances of that class. However, of course, the two instances we created still refer to their class, so the objects can be used. Interestingly, once these objects are garbage collected, their class as well can be reclaimed.

Now try to come up with interesting examples of both *anonymous* classes and *nested* classes, and compare all these facilities with other languages out there.

## 6 Inheritance

In the presence of classes, we may want a mechanism similar to §4.2 “Delegation” in order to be able to reuse and selectively refine existing classes. We therefore extend our object system with support for *class inheritance*. As we will see, many issues have to be dealt with. As usual, we will proceed step-by-step.

### 6.1 Class Hierarchy

We introduce the capability for a class to *extend* another class (called its *superclass*). We focus on *single inheritance*, where a class only extends a single class, as in Java or C#. As you surely know, some languages support multiple inheritance, such as C++ and Python.

As a result, classes are organized in a hierarchy. The superclasses of a class (transitively) are called its *ancestors*; dually, the set of transitive *subclasses* of a class are called its *descendants*.

What are the respective benefits and drawbacks of these different approaches to inheritance? What about other reuse and composition mechanisms such as Scala’s traits?

### 6.2 Method Lookup

When we send a message to an object, we look in its class for a method that implements the message, and then we apply it. The lookup part is defined in our `CLASS` macro definition as:

```
['-lookup
 (let ([found (assoc (first args) methods)])
  (if found
    (cdr found)
    (error "message not understood:" (first args))))]
```

With inheritance, if an object is sent a message for which a method is not found in its class, we can look for a method in the superclass, instead of failing. Recursively, if the method is not found in the superclass, maybe it can be found in the super-superclass, etc. Until we reach “the end”: we introduce a “`Root`” class at the top of the tree, in order to put an end to the method lookup process:

```
(define Root
  (λ (msg . args)
    (match msg
      ['-lookup (error "message not understood:" (first args))])))
```

`Root` is implemented directly as a dispatcher function, without using the `CLASS` form, so that we don’t need to specify its superclass (it has none). In case it is sent a `-lookup` message,

it raises a message-not-understood error. Note that in this system, it is an error to send any message that is not `-lookup` to the root class.

The `CLASS` macro is extended to support the declaration of a superclass via an `extends` clause, and to delegate to the superclass the lookup of an unfound method.

Why is let-binding the superclass necessary?

```
(defmac (CLASS extends sclsexpr
        ([field fname fval] ...)
        ([method mname (mparam ...) mbody ...] ...))
  #:keywords field method extends
  #:captures self ? !
  (let ([scls sclsexpr]
        [methods
         (local [(defmac (? fd) #:captures self
                    (dict-ref (obj-values self) 'fd))
                  (defmac (! fd v) #:captures self
                    (dict-set! (obj-values self) 'fd v))]
          (list (cons 'mname (lambda (self mparam ...) mbody ...)) ...))])
    (letrec ([class
              (lambda (msg . args)
                (match msg
                  ['-create
                 (obj class (make-hash (list (cons 'fname fval) ...)))]
                  ['-lookup
                 (let ([found (assoc (first args) methods)])
                   (if found
                       (cdr found)
                       (scls '-lookup (first args))))))]
              class)))
```

Let us see an example of class inheritance at work, first with a very simple example:

```
(define A
  (CLASS extends Root ()
    ([method foo () "foo"]
     [method bar () "bar"])))

(define B
  (CLASS extends A ()
    ([method bar () "B bar"])))

> (define b (new B))
> (→ b foo)
"foo"
> (→ b bar)
"B bar"
```

This looks just fine: sending a message that is unknown in `B` works as expected, and sending `bar` also results in `B`'s refinement to be executed instead of `A`'s method. We say that method `bar` in `B` *overrides* the method of the same name defined in `A`.

Let us look at a slightly more complex example, building on our `Counter` class. In this chapter, we will progressively develop a `ReactiveCounter`, which extends `Counter` with the ability to trigger actions whenever the counter reaches certain values. But, let's first try with an empty subclass:

```
(define Counter
  (CLASS extends Root
    ([field count 0]
     [field step 1])
    ([method inc () (! count (+ (? count) (? step))) (? count)]
     [method dec () (! count (- (? count) (? step))) (? count)]
     [method reset () (! count 0)]
     [method step! (v) (! step v)]
     [method inc-by! (v) (→ self step! v) (→ self inc)])))
(define ReactiveCounter
  (CLASS extends Counter () ()))

> (define rc (new ReactiveCounter))
> (→ rc inc)
hash-ref: no value found for key
key: 'count
```

What happened? It seems that we are not able to use field `count` of a reactive counter. Fair enough, we haven't dealt at all about how fields have to be handled in the presence of inheritance.

### 6.3 Fields and Inheritance

Let us look again at how we handle object creation at the moment:

```
['-create
 (obj class (make-hash (list (cons 'fname fval) ...)))]
```

That is it: we are only initializing the `values` dictionary for the fields that are *declared* in the current class! We ought to be initializing the dictionary with values for the fields of the ancestor classes as well.

### 6.3.1 Inheriting Fields

An object should have values for all the fields declared in any of its ancestors. Therefore, when we create a new class, we should determine all the fields that its instances will have. To do that, we have to extend classes such that they keep a list of all their fields, and are able to pass that information to any subclass.

```
(defmac (CLASS extends scls-expr
        ([field fname fval] ...)
        ([method mname (mparam ...) mbody ...] ...))
  #:keywords field method extends
  #:captures self ? !
  (let* ([scls scls-expr]
         [methods ....]
         [fields (append (scls '-all-fields)
                        (list (cons 'fname fval) ...))])
    (letrec
      ([class (λ (msg . args)
               (match msg
                 ['-all-fields fields]
                 ['-create (obj class (make-hash fields))]
                 ....))]))))
```

Why are we suddenly using `let*`?

We introduce a new `fields` identifier in the lexical environment of the class. This identifier is bound to the complete list of fields that the class' instances should have. All fields of the superclass are obtained by sending it the `-all-fields` message (whose implementation simply returns the list bound to `fields`). When creating an object, we just make a fresh dictionary with all the fields.

Because we have added a new message to the vocabulary of classes, we need to wonder what happens if this message is sent to the `Root` class: what are all the fields of that class? Well, it has to be the empty list, since we are using `append` blindly:

```
(define Root
  (λ (msg . args)
    (match msg
      ['-lookup (error "message not understood:" (first args))]
      ['-all-fields '()])))
```

Let us see if this works:

```
> (define rc (new ReactiveCounter))
> (→ rc inc)
1
> (→ rc inc-by! 10)
11
```

Good! We can now define some more of the reactive counter:

```
(define ReactiveCounter
  (CLASS extends Counter
    ([field predicate (λ (n) #f)]
     [field action    (λ (n) #f)])
    ([method register (p a) (! predicate p) (! action a)]
     [method react   () (when ((? predicate) (? count))
                            ((? action) (? count)))])))

> (define rc (new ReactiveCounter))
> (→ rc register even? (λ (v) (printf "reacting to ~a~n" v)))
> (→ rc react)
reacting to 0
> (→ rc inc)
1
> (→ rc react)
> (→ rc inc)
2
> (→ rc react)
reacting to 2
```

A reactive counter holds a `predicate` and an `action`. Both are functions of the current value of the counter, `count`. The `react` method implements the reactive logic: when the predicate matches the counter value, the action is triggered. The interactions below the class definition show the expected behavior: when `react` is called, if the counter value is `even`, then a message is printed, showing the value. We will see later how to automatically invoke `react` when the counter value is modified.

### 6.3.2 Binding of Fields

Actually, there is still one more issue that we haven't considered: what happens if a subclass defines a field with a name that already exists in one of its ancestors?

```
(define A
  (CLASS extends Root
    ([field x "A"])
    ([method ax () (? x)])))

(define B
  (CLASS extends A
    ([field x "B"])
    ([method bx () (? x)])))

> (define b (new B))
```

```
> (→ b ax)
"B"
> (→ b bx)
"B"
```

In both cases, we get the value bound to the `x` field of `B`. In other words, we have late binding of fields, exactly as we do for methods.

Let us see: an object is meant to encapsulate some (possibly mutable) state behind a proper procedural interface (methods). It is clear that late binding of methods is a desirable property, because methods are what makes an object's external interface. What about fields? Fields are supposed to be hidden, internal state of the object—in other words, implementation details, not public interface. Actually, notice that in our language so far, it is not even possible to access a field of another object other than `self`! So at the very least, late binding of fields is doubtful.

Let us look at what happened with §4.2 “Delegation”. How are fields handled there? Well, fields are just free variables of a function, so they are *lexically scoped*. This is a much more reasonable semantics for fields. When methods are defined in a class, they are defined in terms of the fields that are directly defined in that class, or one of its superclass. This makes sense, because all that is information known at the time one writes a class definition. Having late binding of fields means reintroducing dynamic scoping for all free variables of a method: an interesting source of errors and headaches! (Think of examples of subclasses mixing up their superclasses by accidentally introducing fields with existing names.)

### 6.3.3 Field Shadowing

We now see how to define the semantics known as *field shadowing*, in which a field in a class shadows a field of the same name in a superclass, but a method always accesses a field as declared in the class or one of its ancestors.

Concretely, this means that an object can hold different values for fields of the same name; which one to use depends on the class in which the executing method is defined (this is known as the *host class* of the method). Because of this multiplicity, it is not possible to use a hash table anymore. Instead, we will keep in a class the list of the field names, and in the object, a *vector* of values, accessed by position. A field access will be done in two steps: first determining the position of the field according to the list of names, and then accessing the value in the vector held by the object.

For instance, for class `A` above, the list of names is `'(x y)` and the vector of values in a new instance of `A` is `#(1 0)`. For class `B`, the list of names is `'(x y x)` and the vector of values in a new instance is `#(1 0 1)`. The advantage of keeping the fields ordered this way is that, if not shadowed, a field is always at the same position within an object.

To respect the semantics of shadowing, we have (at least) two options. We can rename

Is that what you expect? is it reasonable? What happens in Java? In JavaScript? In Python?

Likewise, in languages that have visibility modifiers, are private methods late bound? Why (not)?

shadowed fields to a mangled name, for instance '(x0 y x) in B so that methods hosted in B and its descendants only see the latest definition of x, that is, the field introduced in B. An other alternative is to keep the field names unchanged, but to perform lookup starting from the end: in other words, we will want to find the *last* position of a field name in the list of names. We will go for the latter.

Defining `find-last` is left as an exercise to the reader

We update our definition of `CLASS` so as to introduce fields as a vector instead of a hashtable:

```
['-create
 (let ([values (list->vector (map cdr fields))])
   (obj class values))]
```

We obtain a vector (with the initial field values) and construct the object with it. Then, for field accesses, we must access the vector at the appropriate position, as returned by `find-last`:

```
(let* ([scls scls-expr]
      [fields (append (scls '-all-fields)
                     (list (cons 'fname fval) ...))])
  [methods
   (local [(defmac (? fd) #:captures self
              (vector-ref (obj-values self)
                          (find-last 'fd fields)))]
     (defmac (! fd v) #:captures self
              (vector-set! (obj-values self)
                           (find-last 'fd fields)
                           v))]
     ....]))]
```

For this to work, the list of fields `fields` must be in scope of the method bodies, hence we moved its definition above that of `methods`. Recall that hygiene ensures that the `fields` identifier does not accidentally interfere with the (user-provided) method bodies.

This implementation is not so satisfactory because we call `find-last` (expensive/linear) for each and every field access. Can this be avoided? How?

Let us see if all this works as expected:

```
(define A
  (CLASS extends Root
    ([field x "A"])
    ([method ax () (? x)])))

(define B
  (CLASS extends A
    ([field x "B"])
    ([method bx () (? x)])))

> (define b (new B))
```

```

> (→ b ax)
"A"
> (→ b bx)
"B"

```

Great!

## 6.4 Super Sends

When a method overrides a method in a superclass, it is sometimes useful to be able to invoke that definition. This allows many typical refinement patterns, such as adding something to do before or after a method is executed, like additional processing on its arguments and return values, among many others. This is achieved by doing what is called a *super send*.

Let us go back to our reactive counter example. We would like to actually make our counter reactive, in that it *automatically* reacts to each change in the counter value. We can do that by overriding `inc` in the `ReactiveCounter` subclass, and use a super send to reuse the original implementation.

Let us look at a first example (we use `↑` as the syntax for a super send).

```

(define ReactiveCounter
  (CLASS extends Counter
    ([field predicate (λ (n) #f)]
     [field action    (λ (n) #f)])
    ([method register (p a) (! predicate p) (! action a)]
     [method react   () (when ((? predicate) (? count))
                          ((? action) (? count)))]
     [method inc     ()
      (let ([val (↑ inc)])
        (→ self react)
        val))]))

```

```

> (define rc (new ReactiveCounter))
> (→ rc register even? (λ (v) (printf "reacting to ~a~n" v)))
> (→ rc inc)
1
> (→ rc inc)
reacting to 2
2

```

The syntax of a super send is rather different from what you might be used to. For instance, in Java one writes `super.m()`, and here, `(↑ m)`. The reason for this syntax will soon become clear.

Note how a super send allows us to reuse and extend the definition of `inc` in `Counter` in order to define the method in `ReactiveCounter`. What is the semantics of a super send?

A first thing to clarify is: what is the *receiver* of a super send? In the example above, to which object is `inc` sent using `↑? self!` That's right. In Java, the syntax `super.m()` suggests that `super` is an object, but *it is not!*

In Java, try returning `super` or passing `super` as an argument. What happens? Why?

The key lesson here is that a super send only really affects *method lookup*. So, in Java syntax, what is `super` in `super.m()` is rather the `.`, not the receiver. So this "super dot" is what we write `↑` here. The receiver is left implicit, because, as for field access in our language, the only valid receiver is `self`.

Beyond this syntactic rant, a common misunderstanding is that when doing a super send, method lookup starts from the *superclass of the receiver*, instead of its class. Let us see why this is incorrect in a small, artificial example:

```
(define A (CLASS extends Root ()
           ([method m () "A"])))
(define B (CLASS extends A ()
           ([method m () (string-append "B" (↑ m) "B")])))
(define C (CLASS extends B () ()))

(define c (new C))
(→ c m)
```

What does this program return? Let's see. `→` expands into a send of `-lookup` to `c`'s class, which is `C`. There is no methods defined for `m` in `C`, so it sends `-lookup` to its superclass, `B`. `B` finds a method for `m`, and returns it. It is then applied to the current self (`c`), with no further arguments. The evaluation of the method implies the evaluation of the three arguments of `string-append`, the second of which is a super send. With the above definition of a super send, this means that `m` is looked up not in `C` (the actual class of the receiver) but in `B` (its superclass). Is there a method for `m` in `B`? Yes, and we are actually executing it... In other words, with this understanding of super sends, the above program does NOT terminate.

What is the mistake? To consider that a super send implies looking up the method in the superclass of the receiver. In the example above, we should actually lookup `m` in `A`, not in `B`. To this end, we need to know the *superclass of the host class* of the method in which the super send occurs. Is this a value that should be bound statically in method bodies, or dynamically? Well, we've just said it: it is the superclass of the host class of the method, and that is not likely to change dynamically, at least in our language.

Some dynamic languages, like Ruby, allow the class inheritance relation to be changed at runtime. This is also common in prototype-based languages, like Self and JavaScript.

Luckily, we already have a binding in the lexical context of methods that refers to the superclass, `scls`. So we just need to introduce a new local macro for `↑`, whose expansion asks the superclass `scls` to lookup the message. `↑` can be used by user code, so it is added to the list of `#:captures` identifiers:

```
(defmac (CLASS extends superclass
        ([field f init] ...)
        ([method m params body] ...))
  #:keywords field method extends
```

```

#:captures self ? ! ↑
(let* ([scls superclass]
      [fields (append (scls '-all-fields)
                      (list (cons 'f init) ...))]
      [methods
       (local [(defmac (? fd) ...)
               (defmac (! fd v) ...)
               (defmac (↑ md . args) #:captures self
                       ((scls '-lookup 'md) self . args))]
               ....)]))

```

Note how `-lookup` is now sent to the superclass of the host class of the currently-executing method, `scls`, instead of to the actual class of the current object.

```

(define A (CLASS extends Root ()
              ([method m () "A"])))
(define B (CLASS extends A ()
              ([method m () (string-append "B" (↑ m) "B")])))
(define C (CLASS extends B () ()))

> (define c (new C))
> (→ c m)
"BAB"

```

## 6.5 Inheritance and Initialization

We have previously seen how to address object initialization (§5.5 “Initialization”), by introducing special methods called initializers. Once an object is created, and before it is returned to the creator, its initializer is called.

In the presence of inheritance, the process is a bit more subtle, because if initializers override each other, some necessary initialization work can be missed. The work of an initializer can be quite specific, and we want to avoid subclasses to have to deal with all the details. One can simply assume the normal semantics of method dispatch, whereby `initialize` in a subclass can call the super initializer if needed. This is the approach adopted by Python: if one does not call `super () .__init__ (...)` the initializer of the superclass(es) will not execute. The problem with this freedom is that the initializer in the subclass may start doing things with the object when inherited fields are not yet consistently initialized, or simply skip the initialization defined in the superclasses.

To avoid this issue, in Java, the first thing a constructor must do is to call a super constructor (it may have to compute arguments for that call, but that’s all it is allowed to do). Even if the call is not in the source code, the compiler adds it. Actually, this restriction is also enforced at the VM level by the bytecode verifier: low-level bytecode surgery can therefore not be used to avoid the super constructor call.

However, this protocol does not ensure by itself that uninitialized fields are never accessed. Can you think of how such a problem can occur? Think about the possibility to call any method from a constructor.

## 7 A World of Possibilities

In this brief step-by-step construction of object systems in Scheme, we have only illustrated some fundamental concepts of object-oriented programming languages. As always in language design, there is a world of possibilities to explore, variations on the same ideas, or extensions of these ideas.

Here is just a (limited/arbitrary) list of features and mechanisms that you will find in some existing object-oriented programming languages, which were not covered in our tour, and that you may want to try to integrate in the object system. It is—of course—very interesting to think of other features by yourself, as well as to study existing languages and figure out how to integrate their distinctive features.

- Visibility for methods: public/private
- Declare methods that override a method in a superclass: override
- Declare methods that cannot be overridden: final
- Declare methods that are expected to be inherited: inherit
- Augmentable methods: inner
- Interfaces: sets of messages to understand
- Protocol to check if an object is instance of a class, if a class implements an interface, ...
- Proper initialization protocol with superclasses, named initialization attributes
- Multiple inheritance
- Mixins
- Traits
- Classes as objects, metaclasses, ...

There are also many optimizations, such as:

- compute fields offset for direct field access
- vtables & indices for direct method invocation

We now only briefly outline two mechanisms, interfaces and mixins, as well as their combination (ie. using interfaces in the specification of mixins).

## 7.1 Interfaces

Introduce a form to define interfaces (which can extend super interfaces):

```
(interface (superinterface-expr ...) id ...)
```

Introduce a new class form that expects a list of interfaces:

```
(CLASS* super-expr (interface-expr ...) decls ...)
```

Example:

```
(define positionable-interface
  (interface () get-pos set-pos move-by))

(define Figure
  (CLASS* Root (positionable-interface)
    ....))
```

Extend the protocol of a class so as to be able to check if it implements a given interface:

```
> (implements? Figure positionable-interface)
#t
```

## 7.2 Mixins

A mixin is a class declaration parameterized over its superclass. Mixins can be combined to create new classes whose implementation sharing does not fit into a single-inheritance hierarchy.

Mixins "come for free" by the mere fact of having classes be first-class values integrated with functions.

```
(define (foo-mixin cl)
  (CLASS cl (....) (....)))

(define (bar-mixin cl)
  (CLASS cl (....) (....)))

(define Point (CLASS () ....))

(define foobarPoint
  (foo-mixin (bar-mixin Point)))
(define fbp (foobarPoint 'create))
....
```

Combine with interfaces to check that the given base class implements a certain set of interfaces. Define a `MIXIN` form for that:

```
(MIXIN (interface-expr ...) decl ...)
```

Defines a function that takes a base class, checks that it implements all the specified interfaces, and return a new class that extends the base class with the given declarations.