

ADT Implementation in Java (1st pass)

As I said in my introductory lecture, we will be programming in Java in this course. In part, this is because Java is a reasonable language, and also because you already have seen some Java in CSU 213. In fact, I will assume you know the basics of the language: how to define a basic class, derive some objects from the class, what are methods, what are fields (which we'll also call variables to confuse things sometimes), and what are the basic types such as integers, booleans, and some fundamental operations on values of those types.

Now, working in a specific language such as Java will lead to terminological problems. In particular, many of the terms we use when discussing design in a language-independent way (as we have been doing until now) also have a meaning in Java, a meaning which is often quite similar to the generic language-independent meaning, but often much more specific. For instance, the term *constructor* has a specific meaning in Java which is along the lines of the meaning we gave it last lecture, but actually means something much more specific, viz. a method automatically invoked when using `new` to create a new object. Another example is *interface*, which means something specific in Java that we will cover in a couple of lectures. There are many others. Generally, the context will make it clear whether I am talking about the generic term, or a specific term that has a technical meaning in Java. It is your job to try to keep the two apart. Please ask when you are unsure.

Some remarks on the language: Java is a class-based object-oriented language. “Object-oriented” means that (almost) every value in the language is an object. “Class-based” means that objects are created from classes. A class is essentially a template, a description of how objects of the class are created.

it helps to think of a program has having two parts. One part, the part that correspond to the source code, is static. (Static means non-moving, which we take to mean non-executing). It represents what information about the program we have before anything executes. In Java, the only thing we know before a program executes are what classes are defined. Thus, the classes are static. Classes exist, in some sense, even before programs start executing. The other part is the dynamic part, which corresponds to program execution. During execution, instances of the classes, that is, objects, get created, updated, destroyed. Thus, objects are dynamic.

You should probably be aware that not every object-oriented language is class-based. Self, for example, has no concept of class, but still has objects. Java is also statically typed - types are associated with variables, and before a program is run those types are checked.

not every object-oriented language is statically typed. Smalltalk, for example, checks types dynamically, like Scheme does. We'll talk more about types in a few lectures.

Object-Oriented Signatures and Specifications

To help us devise implementation for ADTs in Java (or in fact, any object-oriented language), let us consider a slight variant way of writing signatures and specifications.

The plan today is to give a quick and dirty implementation of the Drawing ADT, to illustrate the basic uses of Java, and also to introduce the conventions we will use. Recall the signature we had:

```
CREATORS    empty    :    -> Drawing
            oneLine  : Line -> Drawing
            merge    : Drawing Drawing -> Drawing

ACCESSORS   isEmpty  :    Drawing -> boolean
            firstLine :    Drawing -> Line
            restLines :    Drawing -> Drawing
```

(I am making a small change in the specification, `oneLine` and `firstLine` taking and returning a `Line` instead of a pair of `Point`.)

An object-oriented signature, to a first approximation, consider that accessors (not creators), which must take at least one argument of the type of the ADT¹ take that value on which they act as an *implicit* argument, as opposed to an *explicit* argument that appears in the argument list. Thus, to a first approximation, the Drawing object-oriented signature looks as follows:

```
CREATORS    empty    :    -> Drawing
            oneLine  : Line -> Drawing
            merge    : Drawing Drawing -> Drawing

ACCESSORS   isEmpty  :    -> boolean
            firstLine :    -> Line
            restLines :    -> Drawing
```

This looks a bit weird, so let's actually use a Java-like syntax:

```
CREATORS    Drawing empty ();
            Drawing oneLine (Line);
```

¹otherwise, I would claim that such an accessor has no business being part of the ADT.

```

        Drawing merge (Drawing, Drawing);

ACCESSORS  boolean isEmpty ();
           Line firstLine ();
           Drawing restLines ();

```

Let's change how those operations are invoked, because we need to understand where the implicit argument to the accessors is coming from. Creators are invoked as before, e.g., `oneLine(1)`. Accessors, on the other hand, are invoked on an expression yielding a `Drawing` value, e.g., `v.firstLine()`, where `v` is a `Drawing` value, or `merge(d1,d2).firstLine()`.

We can now similarly convert the specification using the above invocation forms:

```

empty().isEmpty() = true
oneLine(1).isEmpty() = false
merge(d1,d2).isEmpty() = isEmpty (d1) & isEmpty (d2)

oneLine(1).firstLine() = 1
merge(d1,d2).firstLine() =
    d2.firstLine()      if d1.isEmpty()
    d1.firstLine()      otherwise

oneLine(1).restLines() = empty ()
merge(d1,d2).restLines() =
    d2.restLines()      if d1.isEmpty()
    merge(d1.restLines(),d2) otherwise

```

This presentation of signatures and specifications is easier to implement, as we now confirm.

Implementation of Drawing ADT

Let's implement the above signature, then. The one decision we have to make, at this point, is how to represent drawings. We did not have to worry about this when designing the specification, because there, the only thing we cared about is how drawings behave, that is, how their operations interact. But now that we have to write code to actual create and manipulate drawings, we need to choose a representation. Later, we will see that there is a very natural representation that comes out of various principles of object-orientation. But for now, we will focus on a naive and somewhat ugly representation, partly for contrasting purposes later on. The representation we choose is to take a drawing to be a linked list of lines.

With this in mind, we define a class `Drawing` in a file `Drawing.java`. Recall that in Java, we can only put one public class (that is, a world-accessible class) per file, and the name of

the file should be the same name as the class, with `.java` appended.

```
// Class implementing the Drawing ADT given in lecture
public class Drawing {

    ...

}
```

Let's fill in the body of this class.

There are several rules that we will follow when writing classes in this course. Six, to be precise, to be introduced throughout this example. Here is the first one:

- (1) The only methods in the class that we should be able to invoke are those given in the signature.

Now, the class can define other methods, we just have to make sure they are not accessible from outside the class.

Java's lets us restrict accessibility to methods (and to fields) using the `private` keyword. In contrast, accessible methods are specified using the `public` keyword.

We decided to represent drawings as lists of lines. We have a few ways of doing this. Here is one. We define fields in a drawing holding the first line of the drawing, and the rest of the drawing.

```
private Line first;           // the first line of the drawing
private Drawing rest;        // the rest of the lines of the drawing
private boolean emptyFlag;   // whether the drawing is empty
```

The second rule explains some of the choices:

- (2) All fields are private.

Fields are not part of the signature, so the spirit of rule (1) says that they should be private. This is not a big restrictions, as we can always use methods (if the signature tells us to) to read and update fields. Mostly, this is to make sure that the rest of the code does not depend on there being a particular field in the object, so that we can, for instance, change the representation of an ADT without worrying about breaking code elsewhere in our program.

From 213 at least, you know that every Java class requires a *constructor*, the method invoked when the `new` operator is used to create an instance of the class. The Java constructor for the `Drawing` class is simple:

```

private Drawing (Line f, Drawing r, boolean e) {
    first = f;
    rest = r;
    emptyFlag = f;
}

```

The only thing the constructor does is initialize the fields. That's always what a constructor will do:

(3) The class constructor takes arguments initial values for all the fields and sets them.

More interestingly, note that the constructor is private. This is in keeping with rule (1), because the constructor is not listed in the signature of the ADT, and never will be:

(4) The class constructor is private.

But that's not a big problem, because we have creators, which are used to create new instances. The creators, being defined inside the class, will be able to invoke the constructor.

So let's look at the creators. Creators, the way we will use them, are sometimes called *factory methods*, and the technique of only using creators to create new objects instead of Java constructors is called the *factory design pattern*.

```

// Creator for an empty drawing
public static Drawing empty () {
    return new Drawing (null,null,true);
}

// Creator for a drawing made up of one line only
public static Drawing oneLine (Line l) {
    return new Drawing (l,empty(),false);
}

// Creator for a merged drawing
public static Drawing merge (Drawing d1, Drawing d2) {
    if (d1.isEmpty()) {
        return d2;
    } else {
        return new Drawing (d1.firstLine(), merge(d1.restLines(),d2));
    }
}

```

This should all be fairly clear, except perhaps the `static` annotation on the creators. This is because we have painted ourselves in a corner, in a sense. Consider this. We would like to

write a creator such as `empty` as a public method that calls the constructor. But how will we be able to call the `empty` method? Recall that in general, to invoke a method, we need an instance of the class. And we just made sure that the constructor of the class was private, that is, not invocable from outside that class. So there is no way to have an object that would let us invoke any of the creators in the first place.

The key is to somehow make sure that the creator methods are available even when no instance of the class exist. That's what the `static` annotation does. It says that the method does not require an instance of the class to be invoked. There is no implicit object passed to the method. It acts, essentially, just like a normal function in other language. (Why static? Recall the discussion at the beginning of this lecture, and try to figure it out.)

How do you use static methods? Let's examine more carefully how to call methods:

- (a) A non-static method `m` in class `T` is called with arguments `a` and `b` (say) by writing:

`obj.m(a,b)`

where `obj` is some expression denoting an object of class `T`. In the body of method `m`, the keyword `this` denotes the same object as `obj`.

- (b) A static method `m` in class `T` is called with arguments `a` and `b` (say) by writing:

`T.m(a,b)`

and the keyword `this` cannot be used in the body of method `m`.²

Following this discussion, we add the following rule to our list:

- (5) Creators are static methods.

Finally, we can implement the accessors, which in this case simply query the appropriate fields of a drawing:

```
// Predicate for an empty drawing
public boolean isEmpty () {
    return (this.emptyFlag);
}

// Accessor for the first line of a drawing
public Line firstLine () {
    if (this.isEmpty())
```

²There are rules about static methods; in particular, a static method cannot refer to methods or fields that are defined in objects, which makes sense because a static method exists even when there are no objects around.

```
        throw new Error ("firstLine() of empty drawing");
    return this.first;
}

// Accessor for the rest of the drawing
public Drawing restLines () {
    if (this.isEmpty())
        throw new Error ("restLines() of empty drawing");
    return this.rest;
}
```

A couple of things to notice. First off, I explicitly use `this.` as a qualifier when invoking a method of the current object, or accessing a field of the current object. Also, to report an error, I throw an exception (in this case, the `Error` exception). We will see exceptions in more detail later on in the course.

Fields are never updated. This is important enough that I will make it a rule that we will only break towards the end of the course:

- (6) Fields, once initialized by the class constructor, are never updated.

In combination with rule (2) that makes every field private, this makes every instance of the class *immutable*—once created, it cannot be changed. Immutable instances have a host of advantages: the code is easier to reason about, it is easier to replace the code or debug it, etc. As we will see when we look at mutation, understanding what actually happens when a field is updated can get very tricky when a program uses all the features of Java. Because of this, and other reasons that we will return to in the course of the semester, we will restrict our attention to immutable instances.

I have also followed a couple of conventions for naming, which you should follow as well. The Java compiler will not enforce them, but your brain will learn to recognize them and use them to spot some errors some times. Class names are capitalized, like `Drawing`, or the hypothetical `LargeDrawing`. Method names and variable names are capitalized but for the first letter, like `oneLine`, or `first`.

All code should be commented. Not putting any comments is a sin that I will not permit you to indulge in. Every class should have a comment at the top indicating the purpose of the class, and every method and variable should have a comment indicating the role of the method or variable.