
Part I:
Foundations of object orientation

Chapter 1 Objects and classes

Main concepts discussed in this chapter:

- objects
- classes
- methods
- parameters

This chapter is the start of our journey into the world of object-oriented programming. Here we introduce the most important concepts you will learn about: *objects* and *classes*. At the end of the chapter, you should have an understanding of what objects and classes are, what they are used for, and how to interact with them. This chapter forms the basis of all other explorations in the book.

1.1 Objects and classes

[concept box: objects] If you write a computer program in an object-oriented language, you are creating, in your computer, a model of some part of the world. The parts that the model is built up from are the objects that appear in the problem domain. These objects must be represented in the computer model being created.

Objects may be categorized and a class describes, in an abstract way, all objects of a particular kind.

We can make these abstract notions clearer by looking at an example. Assume you want to model a traffic simulation. One kind of entity you then have to deal with is cars. What is a car in our context: is it a class or an object? A few questions may help us to make a decision.

What color is a car? How fast can it go? Where is it right now?

[concept box: classes] You will notice that we cannot answer these questions until we talk about one specific car. The reason is that the word “car” in this context refers to the *class* car – we are talking about cars in general, not about one particular car.

If I say, “My old car that is parked at home in my garage,” we can answer the questions above. That car is red, it doesn’t go very fast, and it is in my garage. Now I am talking about an object – about one particular example of a car.

We usually refer to a particular object as an *instance*. We will use the term “instance” quite regularly from now on. Instance is roughly synonymous with object – we refer to objects as instances when we want to emphasize that they are of a particular class (such as, “this object is an instance of class car”).

Before we continue this rather theoretical discussion, let us look at an example.

1.2 Creating objects

Start BlueJ and open the example named *shapes*.² You should see a window similar to that shown in Figure 1.

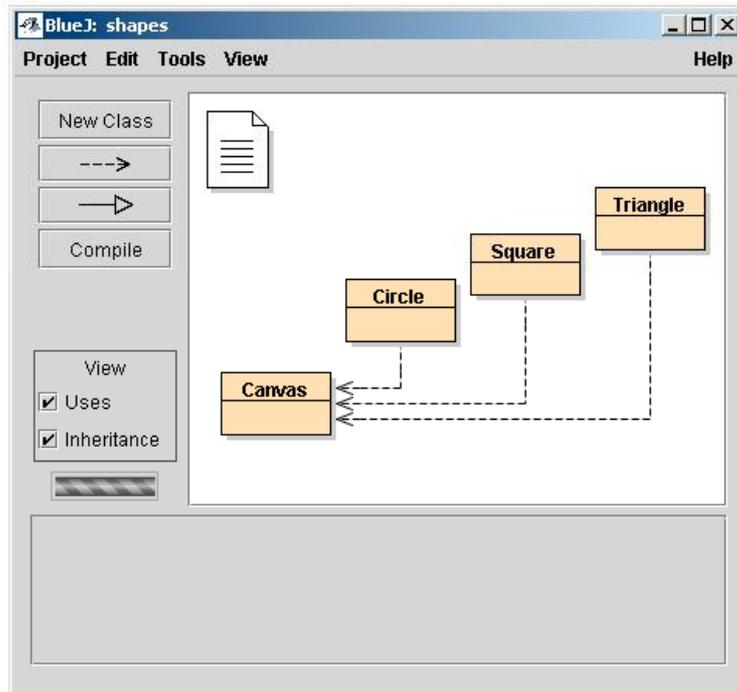


Figure 1: The shapes project in BlueJ

In this window, a diagram should become visible. Every one of the colored rectangles in the diagram represents a class in our project. In this project we have classes named Circle, Square, Triangle and Canvas.

Right-click on the Circle class and choose

```
new Circle()
```

from the popup menu. The system asks you for a “name of the instance” – click Ok, the default name supplied is good enough for now. You will see a red rectangle towards the bottom of the screen labeled “circle_1” (Figure 2).

² We regularly expect you to undertake some activities and exercises while reading this book. At this point we assume that you already know how to start BlueJ and open the example projects. If not, read [Appendix A] first.

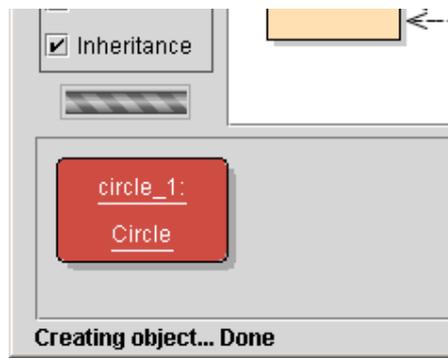


Figure 2: An object on the object bench

You have just created your first object! “Circle”, the rectangular icon in Figure 1, represents the class `Circle`, “circle_1” is an object created from this class. The area at the bottom of the screen where the object is shown is called the *object bench*.

Convention: We start names of classes with capital letters (such as `Circle`) and names of objects with lowercase letters (such as `circle_1`). This helps to distinguish what we are talking about.

Exercise:

1-1 Create another circle. Then create a square.

1.3 Calling methods

Right-click on one of the circle objects (not the class!) and you will see a popup menu with several operations. Choose `makeVisible` from the menu – this will draw a representation of this circle in a separate window (Figure 3).

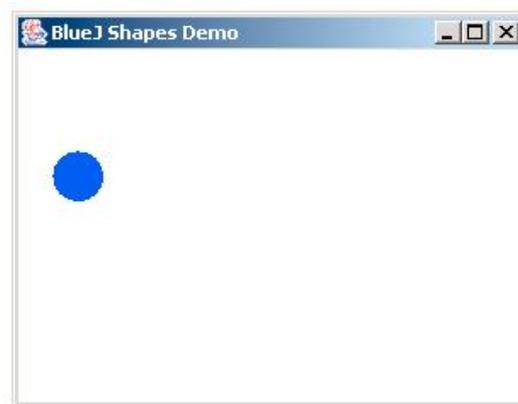


Figure 3: A drawing of a circle

You will notice several other operations in the circle’s menu. Try invoking `moveRight` and `moveDown` a few times to move the circle closer to the center of the screen. You may also like to try `makeInvisible` and `makeVisible` to hide and show the circle.

Exercise:

- 1-2 What happens if you call `moveDown` twice? Or three times? What happens if you call `makeInvisible` twice?

[concept box: methods] The entries in the circle’s menu represent operations that you can use to manipulate the circle. These are called *methods* in Java. Using common terminology, we say that these methods are *called* or *invoked*. We will use this proper terminology from now on. We might ask you to “invoke the `moveRight` method of `circle_1`”.

1.4 Parameters

[concept box: parameters] Now invoke the `moveHorizontal` method. You will see a dialog appear that prompts you for some input (Figure 4). Type in 50 and click Ok. You will see the circle move 50 pixels to the right.³

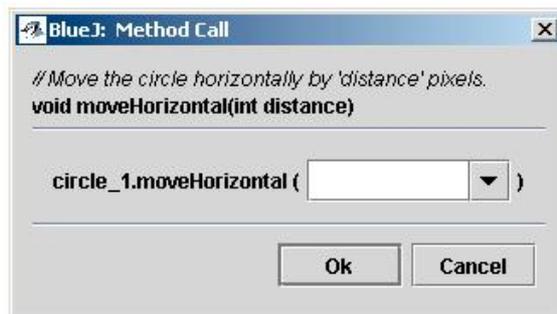


Figure 4: A method call dialog

The `moveHorizontal` method that was just called is written in such a way that it requires some more information to execute. In this case, the information required is the distance – how far the circle should be moved. Thus, the `moveHorizontal` method is more flexible than the `moveRight` or `moveLeft` methods. The latter always move the circle a fixed distance, whereas `moveHorizontal` lets you specify how far you want to move the circle.

Exercise:

- 1-3 Try invoking the `moveVertical`, `slowMoveVertical` and `changeSize` methods before you read on. Find out how you can use `moveHorizontal` to move the circle 70 pixels to the left.

The additional values that some methods require are called *parameters*. A method indicates what kinds of parameters it requires. When calling, for example, the `moveHorizontal` method as shown in Figure 4, the dialog displays the line

```
void moveHorizontal(int distance)
```

³ A pixel is a single dot on your screen. Your whole screen is made up of a grid of single pixels.

near the top. This is called the *signature* of the method. [concept box: signature] The signature provides some information about the method in question. The part between the parenthesis (`int distance`) is the information about the required parameter. For each parameter, it defines a *type* and a *name*. The signature above states that the method requires one parameter of type *int* named *distance*. The name gives a hint about the meaning of the data expected.

1.5 Data types

A type specifies what kind of data can be passed to a parameter. [concept box: types] The type "*int*" signifies whole numbers (also called "integer" numbers, therefore the abbreviation "int").

In the example above, the signature of the `moveHorizontal` method states that, before the method can execute, we need to supply a whole number specifying the distance to move. The data entry field shown in Figure 4 then lets you enter that number.

In the examples so far, the only data type we have seen is *int*. The parameters of the move methods and the `changeSize` method are all of that type.

Closer inspection of the object's popup menu shows that the method entries in the menu include the parameter types. If a method has no parameter, the method name is followed by an empty set of parentheses. If it has a parameter, the type of that parameter is displayed. In the list of methods for a circle, you will see one method with a different parameter type: the `changeColor` method has a parameter of type `String`.

The `string` type indicates that a section of text (for example a word or a sentence) is expected. Strings are always enclosed within double quotes. For example, to enter the word *red* as a string, type

```
"red"
```

The method call dialog also includes a section of text called a *comment* above the method signature. Comments are included to provide information to the (human) reader and are described in [Chapter 2]. The comment of the `changeColor` method describes what color names the system knows about.

Exercise:

- 1-4 Invoke the `changeColor` method on one of your circle objects and enter the String "red". This should change the color of the circle. Try other colors.
- 1-5 This is a very simple example, and not many colors are supported. See what happens when you specify a color that is not known.
- 1-6 Invoke the `changeColor` method, and write the color into the parameter field *without* the quotes. What happens?

Pitfall: A common error for beginners is to forget the double quotes when typing in a data value of type `String`. If you type *green* instead of “*green*”, you will get an error message saying something like “Error: undefined variable” (or something similar).

Java supports several other data types including, for example, decimal numbers and characters. We will not discuss all of them right now, but rather come back to this issue later. If you want to find out about them now, look at [Appendix B].

1.6 Multiple instances

Exercise:

1-7 Create several circle objects on the object bench. You can do so by selecting `new Circle()` from the popup menu of the Circle class. Make them visible, then move them around on the screen using the “move” methods. Make one big and yellow, make another one small and green. Try the other shapes too: create a few triangles and squares. Change their positions, sizes and colors.

[concept box: multiple instances] Once you have a class, you can create as many objects (or instances) of that class as you like. From the class `Circle`, you can create many circles. From `square`, you can create many squares.

Every one of those objects has its own position, color and size. You change an attribute of an object (such as its size) by calling a method on that object. This will affect this particular object, but not others.

You may also notice an additional detail about parameters. Have a look at the `changeSize` method of the triangle. Its signature is

```
void changeSize(int newHeight, int newWidth)
```

Here is an example of a method with more than one parameter. This method has two, and a comma separates them in the signature. Methods can, in fact, have any number of parameters.

1.7 State

[concept box: state] The set of values of all attributes defining an object (such as x-position, y-position, color, diameter and visibility status for a circle) is also referred to as the object’s *state*. This is another example of common terminology that we will use from now on.

In BlueJ, the state of an object can be inspected by selecting the *Inspect* function from the object’s popup menu. When an object is inspected, a window similar to that shown in Figure 5 is displayed. This window is called the *object inspector*.

Exercise:

1-8 Make sure you have several objects on the object bench and then inspect each of them in turn. Try changing the state of an object (for example by calling the `moveLeft` method) while the object inspector is open. You should see the values in the object inspector change.

Some methods, when called, change the state of an object. For example, `moveLeft` changes the `xPosition` attribute. Java refers to these object attributes as *fields*.

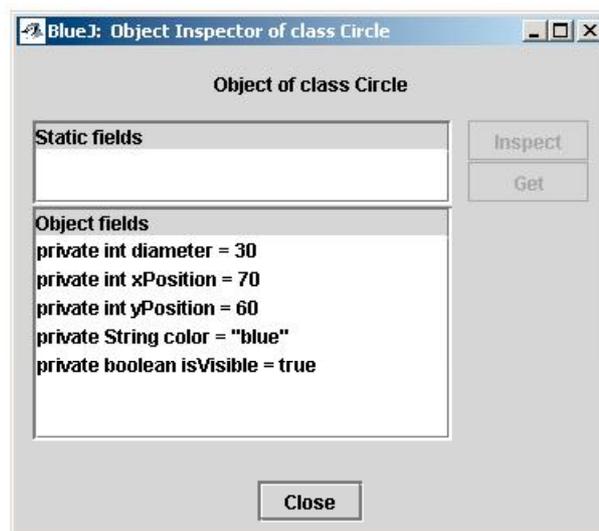


Figure 5: An object inspection dialog

1.8 What is in an object?

On inspecting different objects you will notice that objects of the *same* class all have the same fields. That is, the number, type and names of the fields are the same, while the actual value of a particular field in each object may be different. In contrast, objects of a *different* class may have different fields. A circle, for example, has a field “diameter”, while a triangle has fields for “width” and “height”.

The reason is that the number, types, and names of fields are defined in a class, not in an object. So the class `circle` defines that each circle object will have five fields, named `diameter`, `xPosition`, `yPosition`, `color` and `isVisible`. It also defines the types for these fields. That is, it specifies that the first three are of type `int`, while the `color` is of type `String` and the `isVisible` flag is of type `boolean`. (Boolean is a type that can represent two values: `true` and `false`. We will discuss it in more detail later.)

When an object of class `circle` is created, the object will automatically have these fields. The values of these fields are stored in the object. That ensures that each circle has a color, for instance, and each can have a different color (Figure 6).

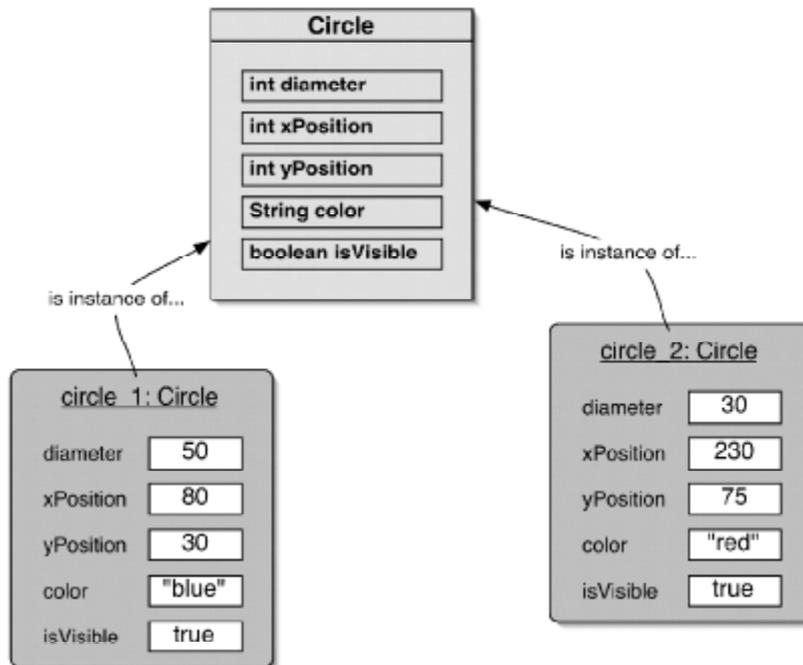


Figure 6: A class and its objects with fields and values

The story is similar for methods. Methods are defined in the class of the object. As a result, all objects of a given class have the same methods. However, the methods are invoked on objects. This makes it clear which object to change when, for example, a `moveRight` method is invoked.

Exercise:

- 1-9 Use the shapes from the *shapes* project to create an image of a house and a sun, similar to that shown in Figure 7. While you are doing this, write down what you have to do to achieve this. Could it be done in different ways?

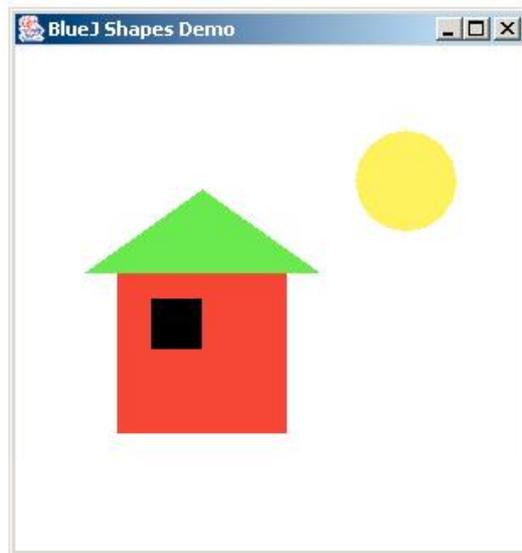


Figure 7: An image created from a set of shape objects

1.9 Object interaction

For the next section, we will work with a different example project. Close the *shapes* project if you still have it open, and open the project called *picture*.

Exercises:

- 1-10 Open the *picture* project. Create an instance of class `Picture` and invoke its `draw` method. Also, try out the `setBlackAndWhite` and `setColor` methods.
- 1-11 How do you think the `Picture` class draws the picture?

Four of the classes in the project are identical to the classes in the *shapes* project. But we now have an additional class: `Picture`. This class is programmed to do exactly what we have done by hand in exercise 1-9.

In reality, if we want a sequence of tasks done in Java, we would not normally do it by hand as in exercise 1-9. Instead, we create a class that does it for us. This is the `Picture` class.

The `Picture` class is written in a way that, when you create an instance, that instance creates two square objects (one for the wall, one for the window), a triangle and a circle, moves them around, changes their color and size, until it looks like the picture we see in Figure 7.

[concept box: method-calling] The important point here is: objects can create other objects and they can call each other's methods. In a normal Java program you may well have hundreds or thousands of objects. The user of a program just starts the program (which typically creates a first object), and all other objects are created – directly or indirectly – by that object.

The big question now is: how do you write the class for such an object?

1.10 Source code

[concept box: source code] Each class has some *source code* associated with it. The source code is text that defines the details of the class. In BlueJ, the source code of a class can be viewed by selecting the *Open Editor* function from the class's popup menu, or by double-clicking the class icon.

Exercise:

1-12 Look at the popup menu of class `Picture` again. You will see an option labeled *Open Editor*. Select it. This will open a text editor displaying the source code of the class.

The source code is text written in the Java programming language. It defines what fields and methods a class has, and precisely what happens when a method is invoked. In the next chapter we will discuss exactly what the source code of a class contains and how it is structured.

A large part of learning the art of programming is learning how to write these class definitions. To do this, we will learn to use the Java language (although there are many other programming languages that could be used to write code).

When you make a change to the source code and close the editor,⁴ the icon for that class appears striped in the diagram. The stripes indicate that the source has been changed. The class now needs to be compiled by clicking the *Compile* button. (You may like to read the "About compilation" sidebar for more information about what is happening when you compile a class.) Once a class has been compiled, objects can be created again and you can try out your change.

About compilation

When people write computer programs, they typically use a "higher level" programming language, such as Java. A problem with that is that a computer cannot execute Java source code directly. Java was designed to be reasonably easy to read for humans, not for computers. Computers, internally, work with a binary representation of a machine code, which looks quite different from Java. The problem for us is: it looks so complex that we do not want to write it directly. We prefer to write Java. What can we do about this?

The solution is a program called the compiler. The compiler translates the Java code into machine code. We can write Java, run the compiler – which generates the machine code – and the computer can then read the machine code. As a result, every time we change the source code we must first run the compiler before we can use the class again to create an object. Otherwise the machine code version that the computer needs does not exist.

Exercises:

- 1-13 In the source code of class `Picture`, find the part that actually draws the picture. Change it so that the sun will be blue rather than yellow.
- 1-14 Add a second sun to the picture. To do this, pay attention to the field definitions close to the top of the class. You will find this code:

⁴ In BlueJ, there is no need to save the text in the editor explicitly before closing. If you close the editor, the source code will automatically be saved.

```
private Square wall;  
private Square window;  
private Triangle roof;  
private Circle sun;
```

You need to add a line here for the second sun, For example:

```
private Circle sun2;
```

Then write the appropriate code for creating the second sun.

1-15 *Challenge exercise:* (This means that this exercise might not be solved quickly. We do not expect everyone to be able to solve this at the moment. If you do – great. If you don't, then don't worry. Things will become clearer as you read on. Come back to this exercise later.) Add a sunset to the single-sun version of `Picture`. That is: make the sun go down slowly. Remember: The circle has a method `slowMoveVertical` that you can use to do this.

1-16 *Challenge exercise:* If you added your sunset to the end of the `draw` method (so that the sun goes down automatically when the picture is drawn), change this now. We now want the sunset in a separate method, so that we can call `draw` and see the picture with the sun up, and then call `sunset` (a separate method!) to make the sun go down.

1.11 Another example

In this chapter, we have already discussed a large number of new concepts. To help in understanding these concepts, we will now revisit them in a different context. For this, we use a different example. Close the *picture* project if you still have it open, and open the *lab-classes* project.

This project is a simplified part of a student database designed to keep track of students in laboratory classes and to print class lists.

Exercise:

1-17 Create an object of class `Student`. You will notice that this time you are not only prompted for a name of the instance, but also for some other parameters. Fill them in before clicking Ok. (Remember that parameters of type `String` must be written in double-quotes.)

1.12 Return values

As before, you can create multiple objects. And again, as before, the objects have methods which you can call from their popup menu.

Exercise:

1-18 Create some student objects. Call the `getName` method on each object. Explain what is happening.

When calling the `getName` method of the `Student` class, we notice something new: methods may return a result value. [concept box: results] In fact, the signature of each method tells us whether or not it returns a result, and what the type of the result is. The signature of `getName` (as shown in the object's popup menu) is defined as:

```
String getName()
```

The word `String` before the method name specifies the return type. In this case it states that calling this method will return a result of type `String`. The signature of `changeName` states:

```
void changeName(String)
```

The word `void` indicates that this method does not return any result.

Methods with return values enable us to get information from an object via a method call. This means that we can use methods either to change an object's state or to find out about its state.

1.13 Objects as parameters

Exercises:

- 1-19 Create an object of class `LabClass`. As the signature indicates, you need to specify the maximum number of students in that class (an integer).
- 1-20 Call the `numberOfStudents` method of that class. What does it do?
- 1-21 Look at the signature of the `enrolStudent` method. You will notice that the type of the expected parameter is `Student`. Make sure you have two or three students and a `LabClass` object on the object bench, then call the `enrolStudent` method of the `LabClass` object. With the input cursor in the dialog entry field, click on one of the student objects – this enters the name of the student object into the parameter field of the `enrolStudent` method (Figure 8). Click Ok, and you have added the student to the `LabClass`. Add one or more other students as well.
- 1-22 Call the `printList` method of the `LabClass` object. You will see a list of all the students in that class printed to the BlueJ terminal window (Figure 9).

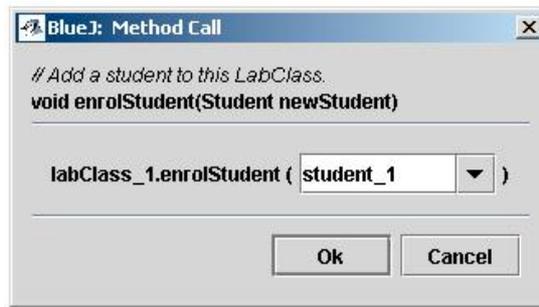


Figure 8: Adding a student to a tutorial

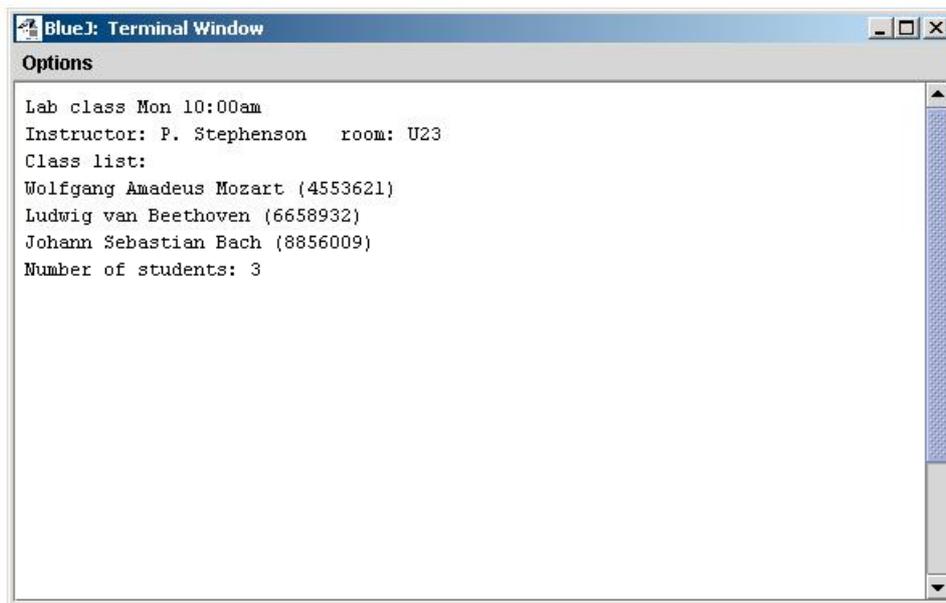


Figure 9: Output of the LabClass project's class listing

As the exercises show, objects can be passed as parameters to methods of other objects. In the case where a method expects an object as a parameter, the expected object's class name is specified as the parameter type in the method signature.

Explore this project a bit more. Try to identify the concepts discussed in the *shapes* example in this context.

Exercises:

- 1-23 Create three students with the following details:
Snow White, student ID: *100234*, credits: *24*
Lisa Simpson, student ID: *122044*, credits: *56*
Charlie Brown, student ID: *12003P*, credits: *6*
 Then enter all three into a lab and print a list to the screen.
- 1-24 Use the inspector on a `LabClass` object to discover what fields it has.
- 1-25 Set the instructor, room, and time for a lab and print the list to the terminal window to check that these new details appear.

1.14 Summary

In this chapter, we have explored the basics of classes and objects. We have discussed the fact that objects are specified by classes. Classes represent the general concept of a thing, while objects represent concrete instances of a class. We can have many objects of any class.

Objects have methods which we use to communicate with them. We can use a method to make a change to the object or to get information from the object. Methods can have parameters, and parameters have types. Methods have return types, which specify what type of data they return. If the return type is *void*, they do not return anything.

Objects store data in fields (which also have types). All the data values of an object together are referred to as the object's state.

Objects are created from class definitions that have been written in a particular programming language. Much of programming in Java is about learning to write class definitions. A large Java program will have many classes, each with many methods which call each other in many different ways.

To learn to develop Java programs, we need to learn how to write class definitions, including fields and methods, and how to put these classes together well. The rest of this book deals with these issues.

Terms introduced in this chapter

object, class, instance, method, signature, parameter, type, state, source code, return value, compiler.

Concept summary

Java objects model objects from a problem domain. [objects]

Objects are created from classes. The class describes the kind of object; the objects represent individual instantiations of the class. [classes]

We can communicate with objects by invoking methods on them. Objects usually do something if we invoke a method. [methods]

Methods can have parameters to provide additional information for a task. [parameters]

The header of a method is called its signature. It provides information needed to invoke that method. [signatures]

Parameters have types. The type defines what kinds of values a parameter can take. [types]

Many similar objects can be created from a single class. [multiple instances]

Objects have state. The state is represented by storing values in fields. [state]

Objects can communicate by calling each other's methods. [method-calling]

The source code of a class determines the structure and the behavior (the fields and methods) of each of the objects in that class. [source code]

Methods may return information about an object via a return value. [results]

Exercise:

1-26

In this chapter, we have mentioned the data types *int* and *String*. Java has more predefined data types. Find out what they are and what they are used for. To do this, you can check [Appendix B], look it up in another Java book or in an online Java language manual. One such manual is at

<http://java.sun.com/docs/books/tutorial/java/nutsandbolts/datatypes.html>