

Java for Python Programmers

Bradley N. Miller

January 27, 2008

1 Preface

Welcome to *Java for Python Programmers*. This short ebook is an ongoing project to help Computer Science students who have had one or two semesters of Python learn the Java programming language. If you are not a part of that audience you may still find this a useful way to learn about Java. This book is written using the build on what you know philosophy. In order to help you learn Java I will start with a Python example and then implement the example in Java. Along the way we will examine the strengths, weaknesses and differences between those two languages.

This book does not attempt to replace the many good Java reference books that are available, in fact I use this in my course along with Horstman's *Core Java* volumes. Please feel free to use this book for yourself, or if it fits a class you are teaching you are welcome to use this as a resource for your own class.

I have published this article using a Creative Commons license to encourage you to use it, change it, and modify it for your own purposes. I would appreciate knowing what you think if you do use this book, and I would love to see any modifications or additions you make.

Brad Miller bmiller@luther.edu January, 2008



This work is licensed under a Creative Commons Attribution 3.0 United States License. See <http://creativecommons.org>

Contents

1	Preface	2
2	Introduction	5
3	Why Learn another programming Language?	5
4	Why Learn Java? Why not C or C++?	6
4.1	Lets look at a Java Program	6
5	Java Data Types	9
5.1	Numeric	9
5.1.1	Import	10
5.1.2	Declaring Variables	11
5.1.3	Input / Output / Scanner	12
5.2	String	13
5.3	List	13
5.4	Arrays	16
5.5	Dictionary	17
6	Conditionals	18
6.1	Simple if	18
6.2	if else	19
6.3	elif	19
7	Loops and Iteration	21
7.1	Definite Loop	21
7.2	Indefinite Loops	22
8	Defining Classes in Java	23
8.1	Writing a constructor	25
8.2	Methods or Member Functions	25
8.2.1	Method Signatures and Overloading	26
8.3	Inheritance	28
8.3.1	The Object Class	28
8.3.2	Abstract Classes and Methods	29
8.4	Interfaces	30
8.5	Static member variables	32
8.6	Static Methods	32
8.7	Full Implementation of the Fraction Class	33
9	Common Mistakes	35

10 Naming Conventions	35
11 Java Documentation Online	36
12 Colophon	36
13 Shameless Plug	36

2 Introduction

This book assumes that you are already familiar with the **Python** programming language. We will use Python as a starting point for our journey into **Java**. We will begin by looking at a very simple Java program, just to see what the language looks like and how we get a program to run. Next, we will look at the main constructs that are common to most programming languages:

- Data Types
- Loops
- Reading user input
- Conditionals

Once we have the basics of Java behind us we will move on to look at the features of Java that are both unique and powerful.

- Classes
- Interfaces
- Collections
- Graphical User Interface Programming
- Generic Programming

Please note that this book is a work in progress. I will continue to update and post new versions.

3 Why Learn another programming Language?

Python is a nice language for beginning programming for several reasons. First the syntax is sparse, and clear. Second, the underlying model of how objects and variables work is very consistent. Third, you can write powerful and interesting programs without a lot of work. However, Python is representative of one kind of language, called a dynamic language. You might think of Python as being fairly informal. There are other languages, Like Java and C++ that are more formal.

These languages have some advantages of their own. First, is speed: For very large programs Java and C++ are going to give you the best performance. Second is their maintainability. A lot of what makes Python easy to use is that you must remember certain things. For example if you set variable `x` to reference a turtle, and forget later that `x` is a turtle but try to invoke a string method on it, you will get an error. Java and C++ protect you by forcing you to be upfront and formal about the kind of object each variable is going to refer to.

In one sense Python is representative of a whole class of languages, sometimes referred to as “scripting languages.” Other languages in the same category as Python are Ruby and Perl. Java is representative of what I will call industrial strength languages. Industrial strength languages are good for projects with several people working on the project where being formal and careful about what you do may impact lots of other people. Languages in this category include C++, C, C# and Ada.

Programming languages will always change. As the field of computer science advances there will be new programming languages and you will need to learn them. It is important to learn several programming languages so that you know what to expect. There are certain features that most programming languages have in common; variables, loops, conditionals, functions. And there are some features that are unique. If you know what is common in languages that is a good place to start.

4 Why Learn Java? Why not C or C++?

- Java is the most widely taught programming language.
- Java is more popular
- Java is industrial strength used for large systems by large groups of people
- If you know Java learning C++ is easy.

Java is an enormous language. There are over 3700 different classes included in the Java 6 Standard Edition. We could not begin to scratch the surface of these classes even if we devoted all 2700 minutes of class time that we have in a semester. However Java is very powerful and we will write some very powerful programs this semester.

4.1 Lets look at a Java Program

A time honored tradition in Computer Science is to write a program called “hello world.” The “hello world” program is simple and easy. There are no logic errors to make, so getting it to run relies only on understanding the syntax. To be clear lets look a a “complicated” version of hello world for Python:

```
1 def main():
2     print "Hello World!"
```

Remember that we can define this program right at the Python command line and then run it:

```
>>> main()
"Hello World!"
>>>
```

Now lets look at the same program written in Java:

```
1 public class Hello {
2
3     public static void main(String[] args) {
4         System.out.println("Hello World!");
5     }
6
7 }
```

What we see is that at the core there are a few similarities, such as a main and the string “Hello World” However there is a lot more stuff around the edges that make it harder to see the core of the program. Do not worry! An important skill for a computer scientist is to learn what to ignore and what to look at carefully. You will soon find that there are some elements of Java that will fade into the background as you become used to seeing them. One thing that will help you is to learn a little bit about Java [Naming Conventions](#).

The first question you probably have about this little program is “How do I run it?” Running a Java program is not as simple as running a Python program. The first thing you need to do with a Java program is compile it. The first big difference between Java and Python is that Python is an interpreted language. We could run our Python programs in the Python **interpreter** and we were quite happy to do that. Java makes running programs a two step process. First we must type the hello world program into a file and save that file using the name `Hello.java` The file name must be the same as the public class you define in the file. Once we have saved the file we **compile** it from the command line as follows:

```
$ javac Hello.java
$ ls -l Hello.*
-rw-r--r--  1 bmiller  bmiller  391 Jul 19 17:47 Hello.class
-rw-r--r--  1 bmiller  bmiller  117 Jul 19 17:46 Hello.java
```

The command `javac` compiles our java source code into compiled byte code and saves it in a file called `Hello.class`. `Hello.class` is a binary file so you won't learn much if you try to examine the class file with an editor. Hopefully you didn't make any mistakes, but if you did you may want to consult the [Common Mistakes](#) section for helpful hints on compiler errors.

Now that we have compiled our java source code we can run the compiled code using the `java` command.

```
$ java Hello
Hello World!
$
```

Now you may be wondering what good is that extra step? What does compiling do for us? There are a couple of important benefits we get from compiling:

- Early detection of errors
- Faster Program Execution

The job of the compiler is to turn your java code into language that the Java Virtual Machine (JVM) can understand. We call the code that the JVM understands **byte code**. The JVM interprets the byte code much like the Python interpreter interprets your Python. However since byte code is much closer to the native language of the computer it can run faster.

When the compiler does the translation it can find many different kinds of errors. For example if you make a typo the compiler will find the typo and point it out to you before you ever run the program. We will look at some examples of compiler errors shortly. Chances are you will create some on your own very soon too.

Now that we have run our hello world program, lets go back and look at it carefully to see what we can learn about the Java language. This simple example illustrates a few very important rules:

1. Every Java program must define a class, all code is inside a class.
2. Everything in Java must have a type
3. Every Java program must have a function called `public static void main(String[] args)`

Lets take the hello world example a line at a time to see how these rules are applied. On line 1 we see that we are declaring a class called `Hello`. As rule 1 says all Java code resides inside a class. Unlike Python where a program can simply be a bunch of statements in a file, Java programs must be inside a class. So, we define a class, `Hello` is not a very useful class it has no instance variables, and only one method. You will also notice the curly brace `{` In Java blocks of code are identified by pairs of curly braces. The block starts with a `{` and ends with a `}`. You will notice that I indented my code that followed the left brace, but in Java this is only done by convention it is not enforced.

On the next line we start our method definition. The name of this method is:

```
public static void main(String[] args)!
```

Everything on this line is significant, and helps in the identification of this method. For example the following lines look similar but are in fact treated by Java as completely different methods:

- `public void main(String[] args)`
- `public static void main(String args)`
- `public static void main()`
- `void main(String args)`

Just digging in to this one line will take us deep into the world of Java, so we are going to start digging but we are not going to dig too deeply right away. Much of what could be revealed by this one line is better understood through other examples, so be patient.

The first word, **public** indicates to the Java compiler that this is a method that anyone can call. We will see that Java enforces several levels of security on the methods we write, including **public**, **protected**, and **private** methods.

The next word, **static** tells Java that this is a method that is part of the class, but is not a method for any one instance of the class. The kind of methods we typically wrote in Python required an instance in order for the method to be called. With a static method, the object to the left of the `.` is a class, not an instance of the class. For example the way that we would call the `main` method directly is: `Hello.main(parameter1)`. For now you can think of static methods the same way you think of methods in Python modules that don't require an instance, for example the `math` module contains many methods: `sin`, `cos`, etc. You probably evaluated these methods using the names `math.cos(90)` or `math.sin(60)`.

The next word, `void` tells the Java compiler that the method `main` will not return a value. This is roughly analogous to omitting the return statement in a Python method. In other words the method will run to completion and exit but will not return a value that you can use in an assignment statement. As we look at other examples we will see that every Python function must tell the compiler what kind of an object it will return. This is in keeping with the rule that says everything in Java must have a type. In this case we use the special type called `void` which means no type.

Next we have the proper name for the method: `main`. The rules for names in Java are similar to the rules in Python. Names can include letters, numbers, and the `_`. Names in Java must start with a letter.

Finally we have the parameter list for the method. In this example we have one parameter. The name of the parameter is `args` however, because everything in Java must have a type we also have to tell the compiler that the value of `args` is an array of strings. For the moment You can just think of an array as being the same thing as a list in Python. The practical benefit of declaring that the method `main` must accept one parameter and the parameter must be a an array of strings is that if you call `main` somewhere else in your code and and pass it an array of integers or even a single string, the compiler will flag it as an error.

That is a lot of new material to digest in only a single line of Java. Lets press on and look at the next line: `System.out.println("Hello World!");`. This line should look a bit more familiar to you. Python and Java both use the dot notation for finding names. In this example we start with `System`. `System` is a class. Within the `System` class we find the object named `out`. The `out` object is the standard output stream for this program. Having located the `out` object Java will now call the method named `println(String s)` on that object. The `println` method prints a string and adds a newline character at the end. Anywhere in Python that you used the `print` function you will use the `System.out.println` method in Java.

Now there is one more character on this line that is significant and that is the `;` at the end. In Java the `;` signifies the end of a statement. Unlike Python where statements are almost always only one line long java statements can spread across many lines. The compiler knows it has reached the end of a statement when it encounters a `;`. This is a very important difference to remember. In Java the following statements are all legal and equivalent. I would not encourage you to write your code like this, but you should know that it is legal.

```
System.out.println("Hello World");
System.out.println("Hello World")
```



```

;
System.out.println
(
    "Hello World"
)
;
System.
out.
println("Hello World")
;

```

The last two lines of the hello world program simply close the two blocks. The first or outer block is the class definition. The second or inner block is the function definition.

If we wanted to translate the Java back to Python we would have something like the following class definition.

```

1 class Hello(object):
2     @staticmethod
3     def main(args):
4         print "Hello World!"

```

Notice that we used the decorator `@staticmethod` to tell the Python interpreter that `main` is going to be a static method. The impact of this is that we don't have to, indeed we should not, use `self` as the first parameter of the main method! Using this definition we can call the main method in a Python session like this:

```

>>> Hello.main("")
Hello World!
>>>

```

5 Java Data Types

5.1 Numeric

One of the great things about Python is that all of the basic data types are objects. Integers are objects, floating point numbers are objects, lists are objects, everything. In Java that is not the case. In Java some of the most basic data types like integers and floating point numbers are not objects. The benefit of having these primitive data types be non-objects is that operations on the primitives are fast. The problem is that it became difficult for programmers to combine objects and non-objects in the way that we do in Python. So, eventually all the non-object primitives ended up with Objectified versions.

Primitive	Object
int	Integer
float	Float
double	Double
char	Char
boolean	Boolean

In older versions of Java it was the programmers responsibility to convert back and forth from a primitive to an object whenever necessary. This processing of converting a primitive to an object was called "boxing." The reverse process is called "unboxing." In Java 5, the compiler became smart enough to know when to convert back and forth and is called "autoboxing." In this book we will typically use the Object version of all the numeric data types and let the compiler do its thing.

Lets go back in time and look at another of our very early Python programs. Here is a simple Python function to convert a Fahrenheit temperature to Celsius.

```

1 def main():
2     fahr = input("Enter the temperature in F: ")
3     cel = (fahr - 32) * 5.0/9.0
4     print "the temperature in C is: ", cel

```

Next, lets look at the Java Equivalent.

```

1 import java.util.Scanner;
2
3 public class TempConv {
4     public static void main(String[] args) {
5         Double fahr;
6         Double cel;
7         Scanner in;
8
9         in = new Scanner(System.in);
10        System.out.println("Enter the temperature in F: ");
11        fahr = in.nextDouble();
12
13        cel = (fahr - 32) * 5.0/9.0;
14        System.out.println("The temperature in C is: " + cel);
15
16        System.exit(0);
17    }
18
19 }

```

There are several new concepts introduced in this example. We will look at them in the following order:

- Import
- Variable Declaration
- Input/Output and the Scanner Class

5.1.1 Import

In Java you can use any class that is available without having to import the class subject to two very important conditions:

1. The javac and java must know that the class exists.
2. You must use the full name of the class

Your first question might be how do the `java` and `javac` commands know that certain classes exist. The answer is the following:

1. Java knows about all the classes that are defined in `.java` and `.class` files in your current working directory.
2. Java knows about all the classes that are shipped with java.
3. Java knows about all the classes that are included in your `CLASSPATH` environment variable. Your `CLASSPATH` environment variable can name two kinds of structures.
 1. A jar file that contains java classes
 2. Another unix directory that contains java class files

You can think of the import statement in Java as working a little bit like the `from module import xxx` statement in Python. However, behind the scenes the two statements actually do very different things. The first important difference to understand is that the class naming system in Java is very hierarchical. The *full* name of the Scanner class is really `java.util.Scanner`. You can think of this name as having two parts: The first part `java.util` is called the **package** and the last part is the class. We'll talk more about the class naming system a bit later. The second important difference is that it is the Java class loader's responsibility to load classes into memory, not the import statement's.

So, what exactly does the import statement do? What it does is tell the compiler that we are going to use a shortened version of the class's name. In this example we are going to use the class `java.util.Scanner` but we can refer to it as just `Scanner`. We could use the `java.util.Scanner` class without any problem and without any import statement provided that we always referred to it by its full name. As an Experiment you may want to try this yourself. Remove the import statement and change the string `Scanner` to `java.util.Scanner` in the rest of the code. The program should still compile and run.

5.1.2 Declaring Variables

Here is where we run into one of the most important differences between Java and Python. Python is a **dynamically typed** language. In a dynamically typed language a variable can refer to any kind of object at any time. When the variable is used, the interpreter figures out what kind of object it is. Java is a **statically typed** language. In a statically typed language the association between a variable and the type of object the variable can refer to is determined when the variable is **declared**. Once the declaration is made it is an error for a variable to refer to an object of any other type.

In the example above, lines 5—7 contain variable declarations. Specifically we are saying that `fahr` and `cel` are going to reference objects that are of type `Double`. The variable `in` will reference a `Scanner` object. This means that if we were to try an assignment like `fahr = "xyz"` the compiler would generate an error because `"xyz"` is a string and `fahr` is supposed to be a double.

For Python programmers the following error is likely to be even more common. Suppose we forgot the declaration for `cel` and instead left line 6 blank. What would happen when we type `javac TempConv.java` on the command line?

```
TempConv.java:13: cannot find symbol
symbol   : variable cel
location: class TempConv
    cel = (fahr - 32) * 5.0/9.0;
    ~
TempConv.java:14: cannot find symbol
symbol   : variable cel
location: class TempConv
    System.out.println("The temperature in C is: " + cel);
    ~
2 errors
```

When you see the first kind of error, where the symbol is on the left side of the equals sign it usually means that you have not declared the variable. If you have ever tried to use a Python variable that you have not initialized the second error message will be familiar to you. The difference here is that we see the message before we ever try to test our program. More common error messages are discussed in the section [Common Mistakes](#).

The general rule in Java is that you must decide what kind of an object your variable is going to reference and then you must declare that variable before you use it. There is much more to say about the static typing of Java but for now this is enough.

5.1.3 Input / Output / Scanner

In the previous section you saw that we created a `Scanner` object. In Java `Scanner` objects make getting input from the user, a file, or even over the network relatively easy. In our case we simply want to ask the user to type in a number at the command line, so in line 9 we construct a `Scanner` by calling the constructor and passing it the `System.in` object. Notice that this `Scanner` object is assigned to the name `in`, which we declared to be a `Scanner` on line 7. `System.in` is similar to `System.out` except of course it is used for input. If you are wondering why we must create a `Scanner` to read data from `System.in` when we can write data directly to `System.out` using `println`, you are not alone. We will talk about the reasons why this is so later when we talk in depth about Java streams. You will also see in other examples that we can create a `Scanner` by passing the `Scanner` a `File` object. You can think of a scanner as a kind of “adapter” that makes low level objects easier to use.

On line 11 we use the `Scanner` object to read in a number. Here again we see the implications of Java being a strongly typed language. Notice that we must call the method `nextDouble`. Because the variable `fahr` was declared as a `double`. So, we must have a function that is guaranteed to return each kind of object we might want to read. In this case we need to read a `Double` so we call the function `nextDouble`. The compiler matches up these assignment statements and if you try to assign the results of a method call to the wrong kind of variable it will be flagged as an error.

Table 2 shows you some commonly used methods of the scanner class. There are many more methods supported by this class and we will talk about how to find them in the next chapter.

Return type	Method name	Description
boolean	<code>hasNext()</code>	returns true if more data is present
boolean	<code>hasNextInt()</code>	returns true if the next thing to read is an integer
boolean	<code>hasNextFloat()</code>	returns true if the next thing to read is a float
boolean	<code>hasNextDouble()</code>	returns true if the next thing to read is a double
Integer	<code>nextInt()</code>	returns the next thing to read as an integer
Float	<code>nextFloat()</code>	returns the next thing to read as a float
Double	<code>nextDouble()</code>	returns the next thing to read as a Double
String	<code>next()</code>	returns the next thing to read as a String

Of course Java is more well known for producing applications that have more of a user interface to them than reading and writing from the command line. Lets look at a version of our temperature control application that uses dialog boxes for input and output.

```
1 import javax.swing.*;
2
3 public class TempConvGUI {
4
5     public static void main(String[] args) {
6         String fahrString;
7         Double fahr, cel;
8
9         fahrString = JOptionPane.showInputDialog("Enter the temperature in F");
10        fahr = new Double(fahrString);
11        cel = (fahr - 32) * 5.0/9.0;
12
13        JOptionPane.showMessageDialog(null,"The temperature in C is, " + cel);
14    }
15
16 }
```

This example illustrates a couple of interesting points:

First, the function call `JOptionPane.showInputDialog` pops up a dialog box to allow you to enter a temperature. But, since you could enter anything into the text input box it returns a `String`.

On the next line the string is converted into a Double by the Double constructor. This is similar to what happens in Python when you call `float()` with either a string or an integer as the argument.

The next dialog box is `JOptionPane.showMessageDialog`. Notice that the first parameter is `null`. In Java `null` serves the same purpose as `None` in Python. The first parameter is null because we do not have a ‘main window’ for this little application. When we look at creating full blown java programs with user interfaces, we will learn more about this parameter.

The second parameter is `"The temperature in C is, " + cel`. Now you may be thinking to yourself that this must surely be a violation of the strong typing I have been describing to you. After all you should not be able to add together a string and a Double right? You are correct, however, all java objects have a method called `toString`. The `toString` method acts much like the Python method `__str__()` and is called automatically by the compiler whenever it makes sense to convert a Java object to a string.

5.2 String

Strings in Java and Python are quite similar. Like Python, Java strings are immutable. However, manipulating strings in Java is not quite as obvious since Strings do not support an indexing or slicing operator. That is not to say that you can’t index into a Java string, you can. You can also pull out a substring just as you can with slicing. The difference is that Java uses method calls where Python uses Operators.

In fact this is the first example of another big difference between Java and Python. Java does not support any operator overloading. Table 3 maps common Python string operations to their Java counterparts. For the examples shown in the table we will use a string variable called “str”

Python	Java	Description
<code>str[3]</code>	<code>str.charAt(3)</code>	Return character in 3rd position
<code>str[2:5]</code>	<code>str.substring(2,4)</code>	Return substring from 2nd to 4th
<code>len(str)</code>	<code>str.length()</code>	Return the length of the string
<code>str.find('x')</code>	<code>str.indexOf('x')</code>	Find the first occurrence of x
<code>str.split()</code>	<code>str.split('\s')</code>	Split the string on whitespace into a list/array of strings
<code>str.split(',')</code>	<code>str.split(',')</code>	Split the string at ',' into a list/array of strings
<code>str + str</code>	<code>str.concat(str)</code>	Concatenate two strings together
<code>str.strip()</code>	<code>str.trim()</code>	Remove any whitespace at the beginning or end

5.3 List

Lets look at another early Python program. We are going to read numbers from a file and produce a histogram that shows the frequency of the various numbers. The data file we will use has one number between 0 and 9 on each line of the file. Here is a simple Python program that creates and prints a histogram.

```
def main():
    count = [0]*10
    data = open('test.dat')

    for line in data:
        count[int(line)] = count[int(line)] + 1

    idx = 0
    for num in count:
        print idx, " occured ", num, " times."
```

```
    idx += 1
```

Now if we run this program on a data file that looks like this:

```
9 8 4 5 3 5 2 1 5
```

We will get output that looks like this:

```
0 occurred 0 times
1 occurred 1 times
2 occurred 1 times
3 occurred 1 times
4 occurred 1 times
5 occurred 3 times
6 occurred 0 times
7 occurred 0 times
8 occurred 1 times
9 occurred 1 times
```

Lets review what is happening in this little program. In the first line we create a list and initialize the first 10 positions in the list to be 0. Next we open the data file called 'test.dat' Third, we have a loop that reads each line of the file. As we read each line we convert it to an integer and increment the counter at the position in the list indicated by the number on the line we just read. Finally we iterate over each element in the list printing out both the position in the list and the total value stored in that position.

To write the Java version of this program we will have to introduce several new Java concepts. First, you will see the Java equivalent of a list, called an `ArrayList`. Next you will see three different kinds of loops used in Java. Two of the loops we will use are going to be very familiar, the third one is different from what you are used to in Python but is easy when you understand the syntax:

while Used with boolean expression for loop exit condition.

for Used to iterate over a sequence. This is very similar to `for i in xxx` where xxx is a list or string or file.

for Used to iterate through a sequence of numbers. This is most similar to `for i in range()`, except the syntax is different.

Here is the Java code needed to write the exact same program:

```
1 import java.util.Scanner;
2 import java.util.ArrayList;
3 import java.io.File;
4 import java.io.IOException;
5
6 public class Histo {
7
8     public static void main(String[] args) {
9         Scanner data = null;
10        ArrayList<Integer> count;
11        Integer idx;
12
13        try {
14            data = new Scanner(new File("test.dat"));
15        }
16        catch ( IOException e) {
17            System.out.println("Sorry but I was unable to open your data file");
```

```

18         e.printStackTrace();
19         System.exit(0);
20     }
21
22     count = new ArrayList<Integer>(10);
23     for (Integer i =0; i<10;i++) {
24         count.add(i,0);
25     }
26
27     while(data.hasNextInt()) {
28         idx = data.nextInt();
29         count.set(idx,count.get(idx)+1);
30     }
31
32     idx = 0;
33     for(Integer i : count) {
34         System.out.println(idx + " occurred " + i + " times.");
35         idx++;
36     }
37 }
38 }

```

Before going any further, I suggest you try to compile the above program and run it on some test data that you create.

Now, lets look at what is happening in the Java source. As usual we declare the variables we are going to use at the beginning of the method. In this example we are declaring a Scanner variable called data, an integer called idx and an ArrayList called count. However, there is a new twist to the ArrayList declaration. Unlike Python where lists can contain just about anything, in Java we let the compiler know what kind of objects our array list is going to contain. In this case the ArrayList will contain Integers. The syntax we use to declare what kind of object the list will contain is the `<Type>` syntax.

Technically, you don't *have* to declare what is going to be on an array list. The compiler will allow you to leave the `<Type>` off the declaration. If you don't tell Java what kind of object is going to be on the list Java will give you a warning message like this:

Note: Histo.java uses unchecked or unsafe operations.

Note: Recompile with -Xlint:unchecked for details.

Without the `<Integer>` part of the declaration Python simply assumes that *any* object can be on the list. However, without resorting to an ugly notation called casting, you cannot do anything with the objects on a list like this! So, if you forget you will surely see more errors later in your code. (Try it and see what you get)

Lines 13—20 are required to open the file. Why so many lines to open a file in Java? The additional code mainly comes from the fact that Java forces you to reckon with the possibility that the file you want to open is not going to be there. If you attempt to open a file that is not there you will get an error. A try/catch construct allows us to try things that are risky, and gracefully recover from an error if one occurs. The following example shows the general structure of a try catch block.

```

try {
    Put some risky code in here.... like opening a file
}
catch (Exception e) {
    If an error happens in the try block an exception is thrown.
    We will catch that exception here!
}

```

Notice that in line 16 we are catching an `IOException`. In fact we will see later that we can have multiple catch blocks to catch different types of exceptions. If we want to be lazy and catch any old exception we can catch an `Exception` which is the parent of all exceptions.

On line 22 we create our array list and give it an initial size of 10. Strictly speaking it is not necessary to give the `ArrayList` any size. It will grow or shrink dynamically as needed just like a list in Python. On line 23 we start the first of three loops. The for loop on lines 23–25 serves the same purpose as the Python statement `count = [0]*10`, that is it initializes the first 10 positions in the `ArrayList` to hold the value 0.

The syntax of this for loop probably looks very strange to you, but in fact it is not too different from what happens in Python using `range`. In fact `for(Integer i = 0; i < 10; i++)` is exactly equivalent to the Python `for i in range(10)`. The first statement inside the parenthesis declares and initializes a loop variable `i`. The second statement is a Boolean expression that is our exit condition. In other words we will keep looping as long as this expression evaluates to true. The third clause is used to increment the value of the loop variable at the end of iteration through the loop. In fact `i++` is Java shorthand for `i = i + 1`. Java also supports the shorthand `i--` to decrement the value of `i`. Like Python you can also write `i += 2` as shorthand for `i = i + 2`. Try to rewrite the following Python for loops as Java for loops:

- `for i in range(2,101,2)`
- `for i in range(1,100)`
- `for i in range(100,0,-1)`
- `for x,y in zip(range(10),range(0,20,2))` [hint, you can separate statements in the same clause with a ,]

The next loop (lines 27–30) shows a typical Java pattern for reading data from a file. Java while loops and Python while loops are identical in their logic. In this case we will continue to process the body of the loop as long as `data.hasNextInt()` returns true.

Line 29 illustrates another important difference between Python and Java. Notice that in Java we can not write `count[idx] = count[idx] + 1`. This is because in Java there is no overloading of operators. Everything except the most basic math and logical operations is done using methods. So, to set the value of an `ArrayList` element we use the `set` method. The first parameter of `set` indicates the index or position in the `ArrayList` we are going to change. The next parameter is the value we want to set. Notice that once again we cannot use the indexing square bracket operator to retrieve a value from the list, but we must use the `get` method.

The last loop in this example is similar to the Python for loop where the object of the loop is a Sequence. In Java we can use this kind of for loop over all kinds of sequences, which are called Collection classes in Java. The for loop on line 33 `for(Integer i : count)` is equivalent to the Python loop `for i in count`: This loop iterates over all of the elements in the `ArrayList` called `count`. Each time through the loop the Integer variable `i` is bound to the next element of the `ArrayList`. If you tried the experiment of removing the `<Integer>` part of the `ArrayList` declaration you probably noticed that you had an error on this line. Why?

5.4 Arrays

As I said at the outset of this Section we are going to use Java `ArrayLists` because they are easier to use and more closely match the way that Python lists behave. However, if you look at Java code on the internet or even in your Core Java books you are going to see examples of something called arrays. In fact you have already seen one example of an array declared in the ‘Hello World’ program. Lets rewrite this program to use primitive arrays rather than array lists.

```
1 import java.util.Scanner;  
2 import java.io.File;  
3 import java.io.IOException;
```



```

4
5 public class HistoArray {
6     public static void main(String[] args) {
7         Scanner data = null;
8         Integer[] count = {0,0,0,0,0,0,0,0,0,0};
9         Integer idx;
10
11
12
13         try {
14             data = new Scanner(new File("test.dat"));
15         }
16         catch ( IOException e) {
17             System.out.println("Sorry but I was unable to open your data file");
18             e.printStackTrace();
19             System.exit(0);
20         }
21
22         while(data.hasNextInt()) {
23             idx = data.nextInt();
24             count[idx] = count[idx] + 1;
25         }
26
27         idx = 0;
28         for(Integer i : count) {
29             System.out.println(idx + " occurred " + i + " times.");
30             idx++;
31         }
32     }
33 }

```

The main difference between this example and the previous example is that we declare count to be an Array of integers. We also can initialize short arrays directly using the syntax shown on line 8. Then notice that on line 24 we can use the square bracket notation to index into an array.

5.5 Dictionary

Just as Python provides the dictionary when we want to have easy access to key, value pairs, Java also provides us a similar mechanism. Rather than the dictionary terminology, Java calls these objects Maps. Java provides two different implementations of a map, one is called the `TreeMap` and the other is called a `HashMap`. As you might guess the `TreeMap` uses a balanced binary tree behind the scenes, and the `HashMap` uses a hash table.

Lets stay with a simple frequency counting example, only this time we will count the frequency of words in a document. A simple Python program for this job could look like this:

```

1 def main():
2     data = open('alice30.txt')
3     wordList = data.read().split()
4     count = {}
5     for w in wordList:
6         w = w.lower()
7         count[w] = count.setdefault(w,0) + 1
8
9     keyList = count.keys()
10    keyList.sort()
11    for k in keyList:
12        print "%-20s occurred %4d times"%(k, count[k])
13
14    main()

```

Notice that the structure of the program is very similar to the numeric histogram program.

```
1 import java.util.Scanner;
2 import java.util.ArrayList;
3 import java.io.File;
4 import java.io.IOException;
5 import java.util.TreeMap;
6
7 public class HistoMap {
8
9     public static void main(String[] args) {
10         Scanner data = null;
11         TreeMap<String,Integer> count;
12         Integer idx;
13         String word;
14         Integer wordCount;
15
16         try {
17             data = new Scanner(new File("alice30.txt"));
18         }
19         catch ( IOException e) {
20             System.out.println("Sorry but I was unable to open your data file");
21             e.printStackTrace();
22             System.exit(0);
23         }
24
25         count = new TreeMap<String,Integer>();
26
27         while(data.hasNext()) {
28             word = data.next().toLowerCase();
29             wordCount = count.get(word);
30             if (wordCount == null) {
31                 wordCount = 0;
32             }
33             count.put(word,++wordCount);
34         }
35
36         for(String i : count.keySet()) {
37             System.out.printf("%-20s occured %5d times\n", i, count.get(i) );
38         }
39     }
40 }
```

6 Conditionals

Conditional statements in Python and Java are very similar. In Python we have three patterns:

6.1 Simple if

```
if condition:
    statement1
    statement2
    ...
```

In Java this same pattern is simply written as:

```
if (condition) {
```

```
    statement1
    statement2
    ...
}
```

Once again you can see that in Java the curly braces define a block rather than indentation. In Java the parenthesis around the condition are required because if is technically a function that evaluates to True or False.

6.2 if else

```
if condition:
    statement1
    statement2
    ...
else:
    statement1
    statement2
    ...
```

In Java this is written as:

```
if (condition) {
    statement1
    statement2
    ...
} else {
    statement1
    statement2
    ...
}
```

6.3 elif

Java does not have an elif pattern like Python. In Java you can get the functionality of an elif statement by nesting if and else. Here is a simple example in both Python and Java.

```
if grade < 60:
    print 'F'
elif grade < 70:
    print 'D'
elif grade < 80:
    print 'C'
elif grade < 90:
    print 'B'
else:
    print 'A'
```

In Java we have a couple of ways to write this

```
if (grade < 60) {
    System.out.println('F');
} else {
    if (grade < 70) {
```

```

        System.out.println('F');
    } else {
        if (grade < 80) {
            System.out.println('F');
        } else {
            if (grade < 90) {
                System.out.println('F');
            } else {
                System.out.println('F');
            }
        }
    }
}

```

We can get even closer to the `elif` statement by taking advantage of the Java rule that a single statement does not need to be enclosed in curly braces. Since the `if` is the only statement used in each `else` we can get away with the following.

```

if (grade < 60) {
    System.out.println('F');
} else if (grade < 70) {
    System.out.println('D');
} else if (grade < 80) {
    System.out.println('C');
} else if (grade < 90) {
    System.out.println('B');
} else System.out.println('A');

```

Java also supports a `switch` statement that acts something like the `elif` statement of Python under certain conditions. To write the grade program using a `switch` statement we would use the following:

```

int tempgrade = grade / 10;
switch(tempgrade) {
case 10:
case 9:
    System.out.println('A');
    break;
case 8:
    System.out.println('B');
    break;
case 7:
    System.out.println('C');
    break;
case 6:
    System.out.println('A');
    break;
default:
    System.out.println('F');
}

```

The `switch` statement is not used very often, and I recommend you do not use it! First, it is not as powerful as the `else if` model because the `switch` variable can only be compared for equality with an integer or enumerated constant. Second it is very easy to forget to put in the `break` statement. If the `break` statement is left out then then the next alternative will be automatically executed. For example if the grade was 95 and the `break` was omitted from the `case 9`: alternative then the program would print out both A and B.

7 Loops and Iteration

You have already seen a couple of examples of iteration and looping in Java. So this section will just serve as a reference for the differences in Snytax.

7.1 Definite Loop

In Python the easiest way to write a definite loop is using the for loop in conjunction with the range function. For example:

```
for i in range(10):
    print i
```

In Java we would write this as:

```
for (Integer i = 0; i < 10; i++ ) {
    System.out.println(i);
}
```

Recall that the `range` function provides you with a wide variety of options for controlling the value of the loop variable.

```
range(stop)
range(start,stop)
range(start,stop,step)
```

The Java for loop is really analogous to the last option giving you explicit control over the starting, stopping, and stepping in the three clauses inside the parenthesis. You can think of it this way:

```
for (start clause; stop clause; step clause) {
    statement1
    statement2
    ...
}
```

If you want to start at 100, stop at 0 and count backward by 5 the Python loop would be written as:

```
for i in range(100,-1,-5):
    print i
```

In Java we would write this as:

```
for (Integer i = 100; i >= 0; i -= 5)
    System.out.println(i);
```

In Python the for loop can also iterate over any sequence such as a list, a string, or a tuple. Java also provides a variation of its for loop that provides the same functionality in its so called `for each` loop.

In Python we can iterate over a list as follows:

```
l = [1, 1, 2, 3, 5, 8, 13, 21]
for fib in l:
    print fib
```

In Java we can iterate over an ArrayList of integers too:

```
ArrayList<Integer> l = new ArrayList<Integer>();
l.add(1); l.add(1); l.add(2); l.add(3);
for (Integer i : l) {
    System.out.println(i)
}
```

This example stretches the imagination a bit, and in fact points out one area where Java's primitive arrays are easier to use than an array list. In fact all primitive arrays can be used in a `for each` loop.

```
int l[] = {1,1,2,3,5,8,13,21};
for(int i : l) {
    System.out.println(i);
}
```

To iterate over the characters in a string in Java do the following:

```
String t = "Hello World";
for (char c : t.toCharArray()) {
    System.out.println(c);
}
```

7.2 Indefinite Loops

Both Python and Java support the while loop. Recall that in Python the while loop is written as:

```
while condition:
    statement1
    statement2
    ...
```

In Java we add parenthesis and curly braces to get:

```
while (condition) {
    statement1
    statement2
    ...
}
```

Java adds an additional, if seldom used variation of the while loop called the do loop. The do loop is very similar to while except that the condition is evaluated at the end of the loop rather than the beginning. This ensures that a loop will be executed at least one time. Some programmers prefer this loop in some situations because it avoids an additional assignment prior to the loop. For example:

```
do {
    statement1
    statement2
    ...
} while (condition);
```

8 Defining Classes in Java

You have already seen how to define classes in Java. Its unavoidable for even the simplest of programs. In this section we will look at how we define classes to create our own data types. Lets start by creating a fraction class to extend the set of numeric data types provided by our language. The requirements for this new data type are as follows:

- Given a numerator and a denominator create a new Fraction.
- When a fraction is printed it should be simplified.
- Two fractions can be added or subtracted
- Two fractions can be multiplied or divided
- Two fractions can be compared
- A fraction and an integer can be added together.
- Given a list of Fractions that list should be sortable by the default sorting function.

Here is a mostly complete implementation of a Fraction class in Python that we will refer to throughout this section:

```
1
2     def __init__(self,top,bottom):
3
4         self.num = top           #the numerator is on top
5         self.den = bottom       #the denominator is on the bottom
6
7
8     def __repr__(self):
9         if self.num > self.den:
10            retWhole = self.num / self.den
11            retNum = self.num - (retWhole * self.den)
12            return str(retWhole) + " " + str(retNum)+"/"+str(self.den)
13        else:
14            return str(self.num)+"/"+str(self.den)
15
16    def show(self):
17        print self.num,"/",self.den
18
19    def __add__(self,otherfraction):
20        # convert to a fraction
21        otherfraction = self.toFract(otherfraction)
22
23        newnum = self.num*otherfraction.den + self.den*otherfraction.num
24        newden = self.den * otherfraction.den
25
26        common = gcd(newnum,newden)
27
28        return Fraction(newnum/common,newden/common)
29
30    def __radd__(self,leftNum):
31        otherfraction = self.toFract(leftNum)
32        newnum = self.num*otherfraction.den + self.den*otherfraction.num
33        newden = self.den * otherfraction.den
34
35        common = gcd(newnum,newden)
36
37        return Fraction(newnum/common,newden/common)
38
```

```

39     def __cmp__(self, otherfraction):
40
41         num1 = self.num*otherfraction.den
42         num2 = self.den*otherfraction.num
43
44         if num1 < num2:
45             return -1
46         else:
47             if num1 == num2:
48                 return 0
49             else:
50                 return 1
51
52     def toFract(self, n):
53         if isinstance(n, int):
54             otherfraction = Fraction(n, 1)
55         elif isinstance(n, float):
56             wholePart = int(n)
57             fracPart = n - wholePart
58             # convert to 100ths???
59             fracNum = int(fracPart * 100)
60             newNum = wholePart * 100 + fracNum
61             otherfraction = Fraction(newNum, 100)
62         elif isinstance(n, Fraction):
63             otherfraction = n
64         else:
65             print "Error: cannot add a fraction to a ", type(n)
66             return None
67         return otherfraction
68
69     #gcd is a helper function for Fraction
70
71     def gcd(m, n):
72         while m%n != 0:
73             oldm = m
74             oldn = n
75
76             m = oldn
77             n = oldm%oldn
78
79     return n

```

The instance variables (data members) we will need for our fraction class are the numerator and denominator. Of course in Python we can add instance variables to a class at any time by simply assigning a value to `objectReferenc.variableName`. In Java all data members must be declared up front.

The declarations of instance variables can come at the beginning of the class definition or the end. Cay Horstman, Author of the *Core Java* books puts the declarations at the end of the class. I like them at the very beginning so you see the variables that are declared before you begin looking at the code that uses them. With that in mind the first part of the Fraction class definition is as follows:

```

public class Fraction {
    private Integer numerator;
    private Integer denominator;

}

```

Notice that we have declared the numerator and denominator to be private. This means that the compiler will generate an error if another method tries to write code like the following:


```
Fraction f = new Fraction(1,2);
Integer y = f.numerator * 10;
```

Direct access to instance variables is not allowed. Therefore if we legitimately want to be able to access information such as the numerator or denominator for a particular fraction we must have getter functions. It is very common programming practice to provide getter and setter functions for instance variables in Java.

```
public Integer getNumerator() {
    return numerator;
}

public void setNumerator(Integer numerator) {
    this.numerator = numerator;
}

public Integer getDenominator() {
    return denominator;
}

public void setDenominator(Integer denominator) {
    this.denominator = denominator;
}
```

8.1 Writing a constructor

Once you have identified the instance variables for your class the next thing to consider is the constructor. In Java, constructors have the same name as the class and are declared public. They are declared without a return type. So any function that is named the same as the class and has no return type is a constructor. Our constructor will take two parameters the numerator and the denominator.

```
public Fraction(Integer top, Integer bottom) {
    num = top;
    den = bottom;
}
```

There are a couple of important things to notice here. First, you will notice that the constructor does not have a self parameter. You will also notice that we can simply refer to the instance variables by name without the self prefix, because they have already been declared. This allows the Java compiler to do the work of dereferencing the current Java object. Java does provide a special variable called `this` that works like the self variable. In Java, `this` is typically only used when it is needed to differentiate between a parameter or local variable and an instance variable. For example this alternate definition of the the Fraction constructor uses `this` to differentiate between parameters and instance variables.

```
public Fraction(Integer num, Integer den) {
    this.num = num;
    this.den = den;
}
```

8.2 Methods or Member Functions

Now we come to one of the major differences between Java and Python. The Python class definition used the special methods for addition, and comparison that have the effect of redefining how the

standard operators behave. In Java there is **no operator overloading**. So we will have to write member functions to do addition, subtraction, multiplication, and division. Lets begin with addition.

```
1 public Fraction add(Fraction otherFrac) {
2     Integer newNum, newDen, common;
3
4     newNum = otherFrac.getDenominator()*this.numerator +
5             this.denominator*otherFrac.getNumerator();
6     newDen = this.denominator * otherFrac.getDenominator();
7     common = gcd(newNum,newDen);
8     return new Fraction(newNum/common, newDen/common );
9 }
```

First you will notice that the add member function is declared as `public Fraction`. The `public` part means that any other method may call the add method. The `Fraction` part means that add will return a fraction as its result.

Second, you will notice that on line two all of the local variables used in the function are declared. In this case there are three local variables: `newNum`, `newDen`, and `common`. It is a good idea for you to get in the habit of declaring your local variables at the beginning of your function. This declaration section provides a simple road map for the function in terms of the data that will be used. The listing above also makes use of the `this` variable, you may find it useful to use `this` until you are comfortable with abandoning your Pythonic ideas about `self`.

Declaring your variables at the top is not a requirement, it is just a recommended practice for you. Java only requires that you declare your variables before they are used. The following version of `Fraction` is also legal Java, but may be somewhat less readable.

```
1 public Fraction add(Fraction otherFrac) {
2     Integer newNum = otherFrac.getDenominator()*numerator +
3             denominator*otherFrac.getNumerator();
4     Integer newDen = denominator * otherFrac.getDenominator();
5     Integer common = gcd(newNum,newDen);
6     return new Fraction(newNum/common, newDen/common );
7 }
```

The addition takes place by multiplying each numerator by the opposite denominator before adding. This procedure ensures that we are adding two fractions with common denominators. Using this approach the denominator is computed by multiplying the two denominators. The greatest common divisor function is used to find a common divisor to simplify the numerator and denominator in the result.

Finally on line 8 a new fraction is returned as the result of the computation. The value that is returned by the return statement must match the value that is specified as part of the declaration. So, in this case the return value on line 8 must match the declared value on line 1.

8.2.1 Method Signatures and Overloading

Our specification for this project said that we need to be able to add a `Fraction` to an `Integer`. In Python we can do this by checking the type of the parameter using the `isinstance` function at runtime. Recall that `isinstance(1,int)` returns `True` to indicate that 1 is indeed an instance of the `int` class. See lines 22 and 53—68 of the Python version of the `Fraction` class to see how our Python implementation fulfills this requirement.

In Java we can do runtime type checking, but the compiler will not allow us to pass an `Integer` to the add function since the parameter has been declared to be a `Fraction`. The way that we solve this problem is by writing another `add` method with a different set of parameters. In Java this practice is legal and common we call this practice **overloading**.

This idea of overloading raises a very important difference between Python and Java. In Python a method is known by its name only. In Java a method is known by its signature. The signature of

a method includes its name, and the types of all of its parameters. The name and the types of the parameters are enough information for the Java compiler to decide which method to call at runtime.

To solve the problem of adding an `Integer` and a `Fraction` in Java we will overload both the constructor and the add function. We will overload the constructor so that if it only receives a single `Integer` it will convert the `Integer` into a `Fraction`. We will also overload the add method so that if it receives an `Integer` as a parameter it first construct a `Fraction` from that integer and then add the two `Fractions` together. The new methods that accomplish this task are as follows:

```
public Fraction(Integer num) {
    this.numerator = num;
    this.denominator = 1;
}

public Fraction add(Integer other) {
    return add(new Fraction(other));
}
```

Notice that the overloading approach can provide us with a certain elegance to our code. Rather than utilizing if statements to check the types of parameters we just overload functions ahead of time which allows us to call the method we want and allow the compiler to make the decisions for us. This way of thinking about programming takes some practice.

Our full `Fraction` class to this point would look like the following. You may want to try to compile and run the short test program provided just to see what happens.

```
1 public class Fraction {
2
3     private Integer numerator;
4     private Integer denominator;
5
6     public Fraction(Integer num, Integer den) {
7         this.numerator = num;
8         this.denominator = den;
9     }
10
11    public Fraction(Integer num) {
12        this.numerator = num;
13        this.denominator = 1;
14    }
15
16    public Fraction add(Fraction other) {
17        Integer newNum, newDen, common;
18
19        newNum = other.getDenominator()*this.numerator + this.denominator*other.getNumerator();
20        newDen = this.denominator * other.getDenominator();
21        common = gcd(newNum, newDen);
22        return new Fraction(newNum/common, newDen/common );
23    }
24
25    public Fraction add(Integer other) {
26        return add(new Fraction(other));
27    }
28
29    private static Integer gcd(Integer m, Integer n) {
30        while (m % n != 0) {
31            Integer oldm = m;
32            Integer oldn = n;
33            m = oldn;
34            n = oldm%oldn;
35        }
36    }
37 }
```

```

36         return n;
37     }
38
39     public static void main(String[] args) {
40         Fraction f1 = new Fraction(1,2);
41         Fraction f2 = new Fraction(2,3);
42
43         System.out.println(f1.mul(f2));
44         System.out.println(f1.add(1));
45     }
46
47 }

```

8.3 Inheritance

If you ran the program above you probably noticed that the output is not very satisfying. Chances are your output looked something like this:

```

Fraction@7b11a3ac
Fraction@6c22c95b

```

The reason is that we have not yet provided a friendly string representation for our Fraction objects. The truth is that, just like in Python, whenever an object is printed by the `println` method it must be converted to string format. In Python you can control how that looks by writing an `__str__` method for your class. If you do not then you will get the default, which looked something like the above.

8.3.1 The Object Class

In Java, the equivalent of `__str__` is the `toString` method. Every object in Java already has a `toString` method defined for it because every class in Java automatically inherits from the Object class. The object class provides default implementations for the following functions.

- clone
- equals
- finalize
- getClass
- hashCode
- notify
- notifyAll
- toString
- wait

We are not interested in most of the functions on that list, and many Java programmers live happy and productive lives without knowing much about most of the functions on that list. However, to make our output nicer we will implement the `toString` method for the `Fraction` class. A simple version of the method is provided below.

```

public String toString() {
    return numerator.toString() + "/" + denominator.toString();
}

```

The other important class for us to implement from the list of methods inherited from Object is the `equals` method. When two objects are compared in Java using the `==` operator they are tested to see if they are exactly the same object, that is do the two objects occupy the same exact space in the computers memory. This is the default behavior of the `equals` method provided by Object. The `equals` method allows us to decide if two objects are equal by looking at their instance variables. However it is important to remember that since Java does not have operator overloading if you want to use your `equals` method you must call it directly. Therefore once you write your own `equals` method:

```
object1 == object2
```

is NOT the same as

```
object1.equals(object2)
```

Here is an `equals` method for the Fraction class:

```

1 public boolean equals(Fraction other) {
2     Integer num1 = this.numerator * other.getDenominator();
3     Integer num2 = this.denominator * other.getNumerator();
4     if (num1 == num2)
5         return true;
6     else
7         return false;
8 }

```

One important thing to remember about `equals` is that it only checks to see if two objects are equal it does not have any notion of less than or greater than. We'll see more about that shortly.

8.3.2 Abstract Classes and Methods

If we want to make our Fraction class behave like Integer, Double, and the other numeric classes in Java We need to make a couple of additional modifications to the class. The first thing we will do is plug `Fraction` into the Java class hierarchy at the same place as Integer and its siblings. If you look at the documentation for Integer you will see that Integer's parent class is `Number`. `Number` is an **abstract class** that specifies several methods that all of its children must implement. In Java an abstract class is more than just a placeholder for common functions. In Java an abstract class has the power to specify certain functions that all of its children **must** implement. You can trace this power back to the strong typing nature of Java.

The that makes the Fraction class a child of Number is as follows:

```

public class Fraction extends Number {
    ...
}

```

The keyword `extends` tells the compiler that the class `Fraction` extends, or adds new functionality to the `Number` class. A child class always extends its parent.

The methods we must implement if `Fraction` is going to be a child of `Number` are:

- `longValue`

- intValue
- floatValue
- doubleValue

This really isn't much work for us to implement these functions as all we have to do is some conversion of our own and some division. The implementation of these methods is as follows:

```
public double doubleValue() {
    return numerator.doubleValue() / denominator.doubleValue();
}
```

```
public float floatValue() {
    return numerator.floatValue() / denominator.floatValue();
}
```

```
public int intValue() {
    return numerator.intValue() / denominator.intValue();
}
```

```
public long longValue() {
    return numerator.longValue() / denominator.longValue();
}
```

By having the `Fraction` class extend the `Number` class we can now pass a `Fraction` to any Java function that specifies it can receive a `Number` as one of its parameters. For example many Java user interface methods accept any object that is a subclass of `Number` as a parameter. In Java the class hierarchy and the IS-A relationships are very important. Whereas in Python you can pass any kind of object as a parameter to any function the strong typing of Java makes sure that you only pass an object as a parameter that is of the type specified in the function call or one of its children. So, in this case when you see a parameter of type `Number` its important to remember that an `Integer` *is-a* `Number` and a `Double` *is-a* `Number` and a `Fraction` *is-a* `Number`.

However, and this is a big however, it is also important to remember that if you specify `Number` as the type on a particular parameter then the Java compiler will only let you use the methods of a `Number`. In this case `longValue`, `intValue`, `floatValue`, and `doubleValue`.

Lets suppose you define a method in some class as follows:

```
public void test(Number a, Number b) {
    a.add(b);
}
```

The Java compiler would give an error because `add` is not a defined method of the `Number` class. Even if you called the `add` method and passed two `Fractions` as parameters.

8.4 Interfaces

Lets turn our attention to making a list of fractions sortable by the standard Java sorting method `Collections.sort`. In Python all we would need to do is implement the `__cmp__` method. But in Java we cannot be that informal. In Java Things that are sortable must be `Comparable`. Your first thought might be that `Comparable` is Superclass of `Number`. That would be a good thought but it would not be correct. Java only supports single inheritance, that is, a class can have only one

parent. Although it would be possible to add an additional Layer to the class hierarchy it would also complicate things dramatically. Because Not only are Numbers comparable, but Strings are also Comparable as would many other types. For example we might have a `Student` class and we want to be able to sort Students by their gpa. But `Student` already extends the class `Person` for which we have no natural comparison function.

Java's answer to this problem is the `Interface` mechanism. Interfaces are like a combination of Inheritance and contracts all rolled into one. An interface is a *specification* that says any object that claims it implements this interface must provide the following methods. It sounds a little bit like an abstract class, however it is outside the inheritance mechanism. You can never create an instance of `Comparable`. Many objects, however, do implement the `Comparable` interface. What does the `Comparable` interface specify?

The `Comparable` interface says that any object that claims to be `Comparable` must implement the `compareTo` method. The following is the documentation for the `compareTo` method as specified by the `Comparable` interface.

```
int compareTo(T o)
```

Compares this object with the specified object for order. Returns a negative integer, zero, or a positive integer as this object is less than, equal to, or greater than the specified object. The implementor must ensure `sgn(x.compareTo(y)) == -sgn(y.compareTo(x))` for all `x` and `y`. (This implies that `x.compareTo(y)` must throw an exception iff `y.compareTo(x)` throws an exception.)

The implementor must also ensure that the relation is transitive: `(x.compareTo(y)>0 && y.compareTo(z)>0)` implies `x.compareTo(z)>0`.

Finally, the implementor must ensure that `x.compareTo(y)==0` implies that `sgn(x.compareTo(z)) == sgn(y.compareTo(z))`, for all `z`.

It is strongly recommended, but not strictly required that `(x.compareTo(y)==0) == (x.equals(y))`. Generally speaking, any class that implements the `Comparable` interface and violates this condition should clearly indicate this fact. The recommended language is "Note: this class has a natural ordering that is inconsistent with equals."

In the foregoing description, the notation `sgn(expression)` designates the mathematical signum function, which is defined to return one of `-1`, `0`, or `1` according to whether the value of `expression` is negative, zero or positive.

To make our `Fraction` class `Comparable` we must modify the class declaration line as follows:

```
public class Fraction extends Number implements Comparable<Fraction> {  
    ...  
}
```

The specification `Comparable<Fraction>` makes it clear that `Fraction` is only comparable with another `Fraction`. The `compareTo` method could be implemented as follows:

```
public int compareTo(Fraction other) {  
    Integer num1 = this.numerator * other.getDenominator();  
    Integer num2 = this.denominator * other.getNumerator();  
    return num1 - num2;  
}
```

8.5 Static member variables

Suppose that you wanted to write a `Student` class so that the class could keep track of the number of students it had created. Although you could do this with a global counter variable that is an ugly solution. The right way to do it is to use a static variable. In Python we could do this as follows:

```
1     numStudents = 0
2
3     def __init__(self, id, name):
4         self.id = id
5         self.name = name
6
7         Student.numStudents = Student.numStudents + 1
8
9 def main():
10     for i in range(10):
11         s = Student(i,"Student-"+str(i))
12     print 'The number of students is: ', Student.numStudents
```

```
>>> main()
The number of students is:  10
>>>
```

In Java we would write this same example using a static declaration.

```
1 public class Student {
2     public static Integer numStudents = 0;
3
4     private int id;
5     private String name;
6
7     public Student(Integer id, String name) {
8         this.id = id;
9         this.name = name;
10
11     numStudents = numStudents + 1;
12 }
13
14     public static void main(String[] args) {
15         for(Integer i = 0; i < 10; i++) {
16             Student s = new Student(i,"Student"+i.toString());
17         }
18         System.out.println("The number of students: "+Student.numStudents.toString());
19     }
20 }
```

In this example notice that we create a static member variable by using the static modifier on the variable declaration. Once a variable has been declared static in Java it can be access from inside the class without prefixing the name of the class as we had to do in Python.

8.6 Static Methods

We have already discussed the most common static method of all, `main`. However in our `Fraction` class we also implemented a method to calculate the greatest common divisor for two fractions (`gcd`). There is no reason for this method to be a member method since it takes two `Integer` values as its parameters. Therefore we declare the method to be a static method of the class. Furthermore since we are only going to use this `gcd` method for our own purposes we can make it private.


```

private static Integer gcd(Integer m, Integer n) {
    while (m % n != 0) {
        Integer oldm = m;
        Integer oldn = n;
        m = oldn;
        n = oldm%oldn;
    }
    return n;
}

```

8.7 Full Implementation of the Fraction Class

A final version of the Fraction class that exercises all of the features we have discussed is as follows.

```

1  import java.util.ArrayList;
2  import java.util.Collections;
3
4
5  public class Fraction extends Number implements Comparable<Fraction> {
6
7      private Integer numerator;
8      private Integer denominator;
9
10     /** Creates a new instance of Fraction */
11     public Fraction(Integer num, Integer den) {
12         this.numerator = num;
13         this.denominator = den;
14     }
15
16     public Fraction(Integer num) {
17         this.numerator = num;
18         this.denominator = 1;
19     }
20
21     public Fraction add(Fraction other) {
22         Integer newNum = other.getDenominator()*this.numerator + this.denominator*other.getN
23         Integer newDen = this.denominator * other.getDenominator();
24         Integer common = gcd(newNum,newDen);
25         return new Fraction(newNum/common, newDen/common );
26     }
27
28     public Fraction add(Integer other) {
29         return add(new Fraction(other));
30     }
31
32     public Integer getNumerator() {
33         return numerator;
34     }
35
36     public void setNumerator(Integer numerator) {
37         this.numerator = numerator;
38     }
39
40     public Integer getDenominator() {
41         return denominator;
42     }
43
44     public void setDenominator(Integer denominator) {
45         this.denominator = denominator;
46     }
47

```

```

48     public String toString() {
49         return numerator.toString() + "/" + denominator.toString();
50     }
51
52     public boolean equals(Fraction other) {
53         Integer num1 = this.numerator * other.getDenominator();
54         Integer num2 = this.denominator * other.getNumerator();
55         if (num1 == num2)
56             return true;
57         else
58             return false;
59     }
60
61     public double doubleValue() {
62         return numerator.doubleValue() / denominator.doubleValue();
63     }
64
65     public float floatValue() {
66         return numerator.floatValue() / denominator.floatValue();
67     }
68
69     public int intValue() {
70         return numerator.intValue() / denominator.intValue();
71     }
72
73     public long longValue() {
74         return numerator.longValue() / denominator.longValue();
75     }
76
77     public int compareTo(Fraction other) {
78         Integer num1 = this.numerator * other.getDenominator();
79         Integer num2 = this.denominator * other.getNumerator();
80         return num1 - num2;
81     }
82
83     private static Integer gcd(Integer m, Integer n) {
84         while (m % n != 0) {
85             Integer oldm = m;
86             Integer oldn = n;
87             m = oldn;
88             n = oldm%oldn;
89         }
90         return n;
91     }
92
93
94     public static void main(String[] args) {
95         Fraction f1 = new Fraction(1,2);
96         Fraction f2 = new Fraction(2,3);
97         Fraction f3 = new Fraction(1,4);
98
99         System.out.println(f1.add(1));
100        System.out.println(f1.intValue());
101        System.out.println(f1.doubleValue());
102
103        ArrayList<Fraction> myFrac = new ArrayList<Fraction>();
104        myFrac.add(f1);
105        myFrac.add(f2);
106        myFrac.add(f3);
107        Collections.sort(myFrac);
108
109        for(Fraction f : myFrac) {

```

```

110         System.out.println(f);
111     }
112 }
113
114
115 }
```

9 Common Mistakes

- **Forgetting to declare your variables.**

```

Histo.java:21: cannot find symbol
symbol   : variable count
location: class Histo
    count = new ArrayList<Integer>(10);
    ^
```

- **Not importing a class:**

```

Histo.java:9: cannot find symbol
symbol   : class Scanner
location: class Histo
    Scanner data = null;
    ^
```

- **Forgetting to use the new keyword to create an object.** Here's an example of the error message that occurs when you forget to use the new keyword. Notice that the message is pretty unhelpful. Java *thinks* you are trying to call the Method Scanner, but there are two problems. First Scanner is not really a method it is a constructor.:

```

Histo.java:14: cannot find symbol
symbol   : method Scanner(java.io.File)
location: class Histo
    data = Scanner(new File("test.dat"));
    ^
```

- **Forgetting a Semicolon:**

```

Histo.java:19:
';' expected
    System.exit(0);
    ^
```

- Forgetting to declare the kind of object in a container.:

Note: Histo.java uses unchecked or unsafe operations. Note: Recompile with `-Xlint:unchecked` for details.

10 Naming Conventions

Java has some very handy naming conventions.

- Class names always start with an upper case letter. For example, `Scanner`, `System`, `Hello`

- Method names always start with a lower case letter, and use camelCase to represent multiword method names. for example `nextInt()`
- Instance variables of a class start with a lower case letter and use camelCase
- Constants are in all upper case letters. for example `Math.MAXINT`

11 Java Documentation Online

All Java class libraries are documented and available online. Here are two good resources for you to use:

- **JavaDoc** The Javadoconline website provides a nice searchable interface. Search for a classname and you will get the documentation you are looking for.
- **JavaAPI** contains the same information but in a browsable format. If you don't know the class name exactly this is a good way to see what is close.

In general the Javadoc page for any class contains information about:

- Where this class falls in the class hierarchy. What classes are its parents and what classes are its decedents.
- A summary and some examples of using the class.
- A summary listing of instance variables
- A summary listing of Constructors
- A summary listing of Methods
- Detailed documentation on constructors and methods.

Typically the Javadoc pages are constructed from the source code where the class is implemented. This encourages Java programmers to do a good job of documenting their code, while providing a user friendly way to read the documentation without looking at the code directly.

12 Colophon

I started writing this article in 2006 using Python's restructured text. In 2007 I switched to markdown since everything else I was writing used markdown. In particular I switched to the variant of markdown used by the excellent **pandoc** program. Pandoc does an excellent job of converting markdown to html, latex, pdf, and other formats. The markdown source for this article is available [here](#). Finally the document became so big that I used **pandoc** to convert it to latex for preparing the final version. I used The **htlatex** program to generate html for the online version of this document.

13 Shameless Plug

If you got this far, I would also like to use this space as a shameless plug for two books. At Luther college we use Python for CS1 and CS2. When we decided to make the switch to introducing CS with Python we wanted to provide our students with two semesters of Python. The reason is that after one semester students are just getting comfortable with the language and it does not make sense to push them into a brand new language just as they are getting some comfort. The second reason is that Python really is a great language for teaching data structures. The language makes

the algorithms evident by clearing away all of the syntactic clutter. So we (David Ranum and I) wrote a CS2 book called [Problem Solving with Algorithms and Data Structures using Python](#). After we finished the CS2 book we decided to write our own CS1 book as well. This book will be available for Fall 2008 classes from Jones and Bartlett publishers. The tentative title is *Python in Context*. Please feel free to contact me for more information.