

Your Brain on Java—A Learner's Guide

2nd Edition - Covers Java 5.0

# Head First Java

Learn how threads  
can change your life



Make Java concepts  
stick to your brain



Avoid embarrassing  
OO mistakes

Fool around in  
the Java Library



Bend your mind  
around 42  
Java puzzles



Make attractive  
and useful GUIs



O'REILLY®

Kathy Sierra & Bert Bates

## A Trip to Objectville



**I was told there would be objects.** In chapter 1, we put all of our code in the `main()` method. That's not exactly object-oriented. In fact, that's not object-oriented *at all*. Well, we did *use* a few objects, like the String arrays for the Phrase-O-Matic, but we didn't actually develop any of our own object *types*. So now we've got to leave that procedural world behind, get the heck out of `main()`, and start making some objects of our own. We'll look at what makes object-oriented (OO) development in Java so much fun. We'll look at the difference between a *class* and an *object*. We'll look at how objects can give you a better life (at least the programming part of your life. Not much we can do about your fashion sense). Warning: once you get to Objectville, you might never go back. Send us a postcard.

once upon a time in Objectville

# Chair Wars

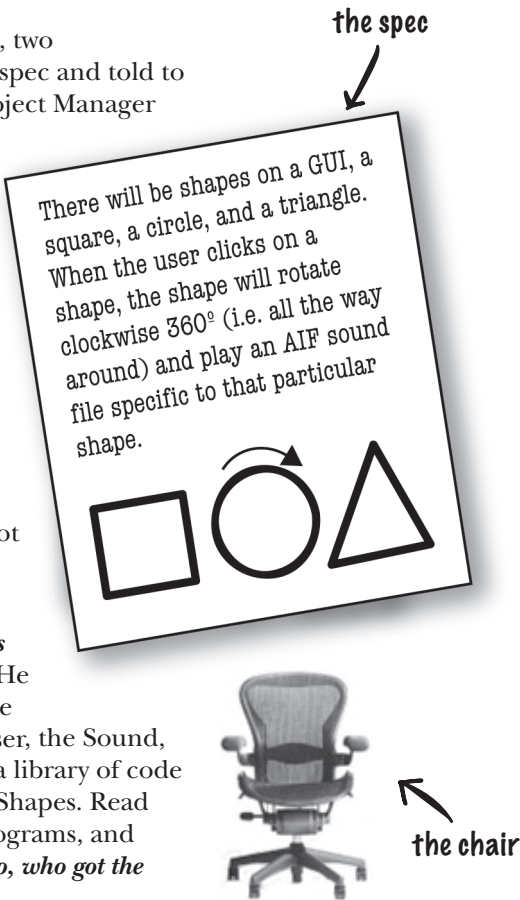
(or How Objects Can Change Your Life)

Once upon a time in a software shop, two programmers were given the same spec and told to “build it”. The Really Annoying Project Manager forced the two coders to compete,

by promising that whoever delivers first gets one of those cool Aeron™ chairs all the Silicon Valley guys have. Larry, the procedural programmer, and Brad, the OO guy, both knew this would be a piece of cake.

Larry, sitting in his cube, thought to himself, “What are the things this program has to *do*? What *procedures* do we need?”. And he answered himself, “**rotate** and **playSound**.” So off he went to build the procedures. After all, what *is* a program if not a pile of procedures?

Brad, meanwhile, kicked back at the cafe and thought to himself, “What are the *things* in this program... who are the key *players*?” He first thought of **The Shapes**. Of course, there were other objects he thought of like the User, the Sound, and the Clicking event. But he already had a library of code for those pieces, so he focused on building Shapes. Read on to see how Brad and Larry built their programs, and for the answer to your burning question, “*So, who got the Aeron?*”



## In Larry’s cube

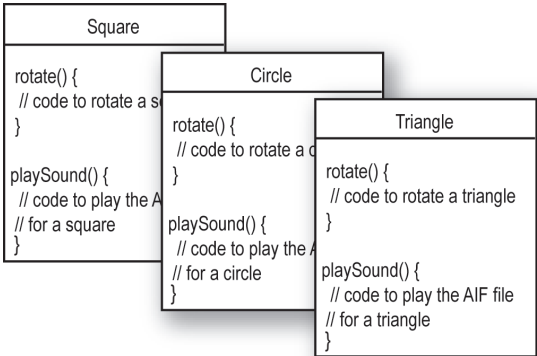
As he had done a gazillion times before, Larry set about writing his **Important Procedures**. He wrote **rotate** and **playSound** in no time.

```
rotate(shapeNum) {
    // make the shape rotate 360°
}

playSound(shapeNum) {
    // use shapeNum to lookup which
    // AIF sound to play, and play it
}
```

## At Brad’s laptop at the cafe

Brad wrote a **class** for each of the three shapes

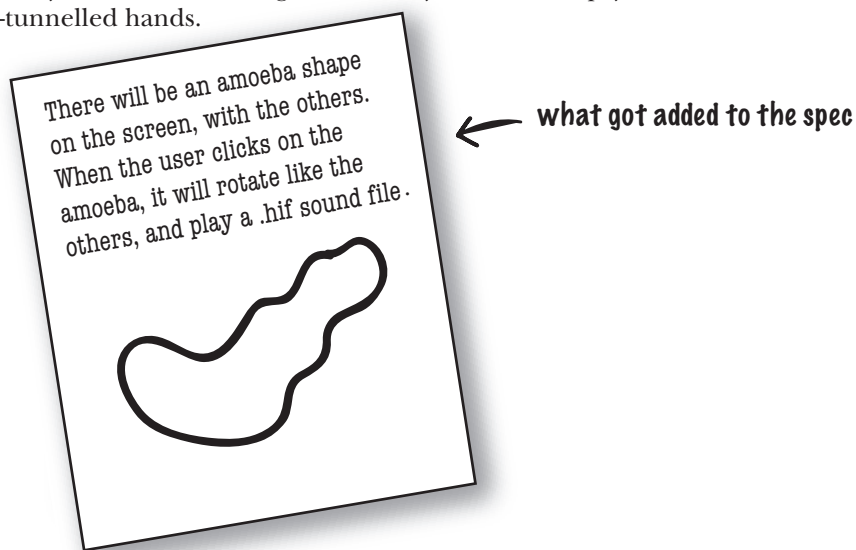


Larry thought he'd nailed it. He could almost feel the rolled steel of the Aeron beneath his...

### But wait! There's been a spec change.

"OK, *technically* you were first, Larry," said the Manager, "but we have to add just one tiny thing to the program. It'll be no problem for crack programmers like you two."

"If I had a dime for every time I've heard that one", thought Larry, knowing that spec-change-no-problem was a fantasy. "And yet Brad looks strangely serene. What's up with that?" Still, Larry held tight to his core belief that the OO way, while cute, was just slow. And that if you wanted to change his mind, you'd have to pry it from his cold, dead, carpal-tunnelled hands.



### Back in Larry's cube

The rotate procedure would still work; the code used a lookup table to match a shapeNum to an actual shape graphic. But *playSound would have to change*. And what the heck is a .hif file?

```
playSound(shapeNum) {
    // if the shape is not an amoeba,
    // use shapeNum to lookup which
    // AIF sound to play, and play it
    // else
    // play amoeba .hif sound
}
```

It turned out not to be such a big deal, but *it still made him queasy to touch previously-tested code*. Of all people, *he* should know that no matter what the project manager says, *the spec always changes*.

### At Brad's laptop at the beach

Brad smiled, sipped his margarita, and *wrote one new class*. Sometimes the thing he loved most about OO was that he didn't have to touch code he'd already tested and delivered. "Flexibility, extensibility,..." he mused, reflecting on the benefits of OO.

Amoeba
<pre>rotate() {     // code to rotate an amoeba }  playSound() {     // code to play the new     // .hif file for an amoeba }</pre>

once upon a time in Objectville

Larry snuck in just moments ahead of Brad.



(Hah! So much for that foofy OO nonsense). But the smirk on Larry’s face melted when the Really Annoying Project Manager said (with that tone of disappointment), “Oh, no, *that’s* not how the amoeba is supposed to rotate...”

Turns out, both programmers had written their rotate code like this:

- 1) determine the rectangle that surrounds the shape
- 2) calculate the center of that rectangle, and rotate the shape around that point.

But the amoeba shape was supposed to rotate around a point on one *end*, like a clock hand.

“I’m toast.” thought Larry, visualizing charred Wonderbread™. “Although, hmmm. I could just add another if/else to the rotate procedure, and then just hard-code the rotation point code for the amoeba. That probably won’t break anything.” But the little voice at the back of his head said, “*Big Mistake. Do you honestly think the spec won’t change again?*”



What the spec conveniently forgot to mention

Back in Larry’s cube

He figured he better add rotation point arguments to the rotate procedure. *A lot of code was affected.* Testing, recompiling, the whole nine yards all over again. Things that used to work, didn’t.

```
rotate(shapeNum, xPt, yPt) {
  // if the shape is not an amoeba,
  // calculate the center point
  // based on a rectangle,
  // then rotate
// else
  // use the xPt and yPt as
  // the rotation point offset
  // and then rotate
}
```

At Brad’s laptop on his lawn chair at the Telluride Bluegrass Festival

Without missing a beat, Brad modified the rotate **method**, but only in the Amoeba class. *He never touched the tested, working, compiled code* for the other parts of the program. To give the Amoeba a rotation point, he added an **attribute** that all Amoebas would have. He modified, tested, and delivered (wirelessly) the revised program during a single Bela Fleck set.

Amoeba
int xPoint
int yPoint
rotate() {
// code to rotate an amoeba
// using amoeba's x and y
}
playSound() {
// code to play the new
// .hif file for an amoeba
}



## So, Brad the OO guy got the chair, right?

*Not so fast.* Larry found a flaw in Brad's approach. And, since he was sure that if he got the chair he'd also get Lucy in accounting, he had to turn this thing around.

**LARRY:** You've got duplicated code! The rotate procedure is in all four Shape things.


**BRAD:** It's a *method*, not a *procedure*. And they're *classes*, not *things*.

**LARRY:** Whatever. It's a stupid design. You have to maintain *four* different rotate "methods". How can that ever be good?

**BRAD:** Oh, I guess you didn't see the final design. Let me show you how OO **inheritance** works, Larry.



What Larry wanted ↗  
(figured the chair would impress her)




1

I looked at what all four classes have in common.


2

They're Shapes, and they all rotate and playSound. So I abstracted out the common features and put them into a new class called Shape.



3

Then I linked the other four shape classes to the new Shape class, in a relationship called inheritance.



You can read this as, "Square inherits from Shape", "Circle inherits from Shape", and so on. I removed rotate() and playSound() from the other shapes, so now there's only one copy to maintain.

The Shape class is called the **superclass** of the other four classes. The other four are the **subclasses** of Shape. The subclasses inherit the methods of the superclass. In other words, if the Shape class has the functionality, then the subclasses automatically get that same functionality.

once upon a time in Objectville

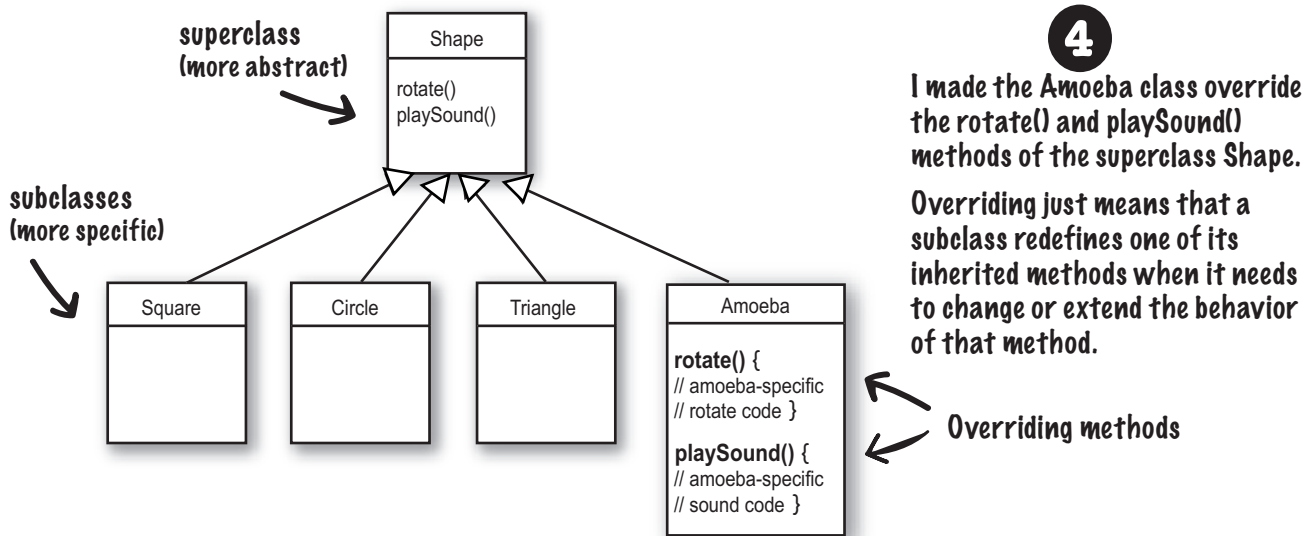
## What about the Amoeba rotate()?

**LARRY:** Wasn't that the whole problem here — that the amoeba shape had a completely different rotate and playSound procedure?

**BRAD:** Method.

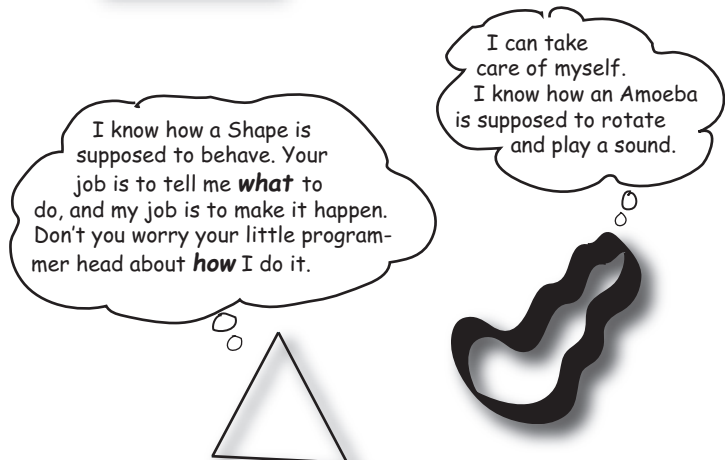
**LARRY:** Whatever. How can amoeba do something different if it “inherits” its functionality from the Shape class?

**BRAD:** That's the last step. The Amoeba class **overrides** the methods of the Shape class. Then at runtime, the JVM knows exactly which rotate() method to run when someone tells the Amoeba to rotate.



**LARRY:** How do you “tell” an Amoeba to do something? Don’t you have to call the procedure, sorry—*method*, and then tell it *which* thing to rotate?

**BRAD:** That's the really cool thing about OO. When it's time for, say, the triangle to rotate, the program code invokes (calls) the rotate() method *on the triangle object*. The rest of the program really doesn't know or care *how* the triangle does it. And when you need to add something new to the program, you just write a new class for the new object type, so the **new objects will have their own behavior**.



## The suspense is killing me. Who got the chair?



**Amy from the second floor.**

(unbeknownst to all, the Project Manager had given the spec to *three* programmers.)

### What do you like about OO?

"It helps me design in a more natural way. Things have a way of evolving."

-Joy, 27, software architect

"Not messing around with code I've already tested, just to add a new feature."

-Brad, 32, programmer

"I like that the data and the methods that operate on that data are together in one class."

-Josh, 22, beer drinker

"Reusing code in other applications. When I write a new class, I can make it flexible enough to be used in something new, later."

-Chris, 39, project manager

"I can't believe Chris just said that. He hasn't written a line of code in 5 years."

-Daryl, 44, works for Chris

"Besides the chair?"

-Amy, 34, programmer



### Time to pump some neurons.

You just read a story about a procedural programmer going head-to-head with an OO programmer. You got a quick overview of some key OO concepts including classes, methods, and attributes. We'll spend the rest of the chapter looking at classes and objects (we'll return to inheritance and overriding in later chapters).

Based on what you've seen so far (and what you may know from a previous OO language you've worked with), take a moment to think about these questions:

What are the fundamental things you need to think about when you design a Java class? What are the questions you need to ask yourself? If you could design a checklist to use when you're designing a class, what would be on the checklist?



### metacognitive tip

If you're stuck on an exercise, try talking about it out loud. Speaking (and hearing) activates a different part of your brain. Although it works best if you have another person to discuss it with, pets work too. That's how our dog learned polymorphism.



When you design a class, think about the objects that will be created from that class type. Think about:

- things the object **knows**
- things the object **does**

ShoppingCart
cartContents
addToCart() removeFromCart() checkOut()

**knows**

**does**

Button
label color
setColor() setLabel() dePress() unDepress()

**knows**

**does**

Alarm
alarmTime alarmMode
setAlarmTime() getAlarmTime() isAlarmSet() snooze()

**knows**

**does**

Things an object *knows* about itself are called

- instance variables

Things an object can *do* are called

- methods

**instance  
variables**  
(state)

**methods**  
(behavior)

Song
title artist
setTitle() setArtist() play()

**knows**

**does**

Things an object *knows* about itself are called **instance variables**. They represent an object's state (the data), and can have unique values for each object of that type.

Think of **instance** as another way of saying **object**.


Things an object can *do* are called **methods**. When you design a class, you think about the data an object will need to know about itself, and you also design the methods that operate on that data. It's common for an object to have methods that read or write the values of the instance variables. For example, Alarm objects have an instance variable to hold the alarmTime, and two methods for getting and setting the alarmTime.

So objects have instance variables and methods, but those instance variables and methods are designed as part of the class.



Fill in what a television object might need to know and do.

Television



**instance  
variables**

**methods**

## Better Living in Objectville

We were underpaid, overworked coders 'till we tried the Polymorphism Plan. But thanks to the Plan, our future is bright. Yours can be too!



**Plan your programs with the future in mind.** If there were a way to write Java code such that you could take more vacations, how much would it be worth to you? What if you could write code that someone *else* could extend, **easily**? And if you could write code that was flexible, for those pesky last-minute spec changes, would that be something you're interested in? Then this is your lucky day. For just three easy payments of 60 minutes time, you can have all this. When you get on the Polymorphism Plan, you'll learn the 5 steps to better class design, the 3 tricks to polymorphism, the 8 ways to make flexible code, and if you act now—a bonus lesson on the 4 tips for exploiting inheritance. Don't delay, an offer this good will give you the design freedom and programming flexibility you deserve. It's quick, it's easy, and it's available now. Start today, and we'll throw in an extra level of abstraction!

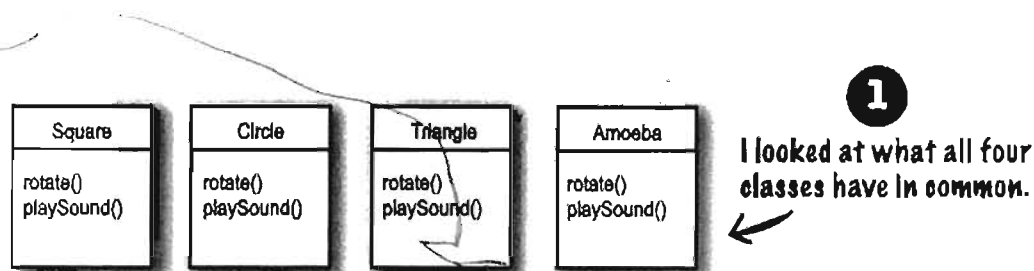
the power of inheritance

## Chair Wars Revisited...

*Remember way back in chapter 2, when Larry (procedural guy) and Brad (OO guy) were vying for the Aeron chair? Let's look at a few pieces of that story to review the basics of inheritance.*


**LARRY:** You've got duplicated code! The rotate procedure is in all four Shape things. It's a stupid design. You have to maintain four different rotate "methods". How can that ever be good?

**BRAD:** Oh, I guess you didn't see the final design. Let me show you how OO inheritance works, Larry.



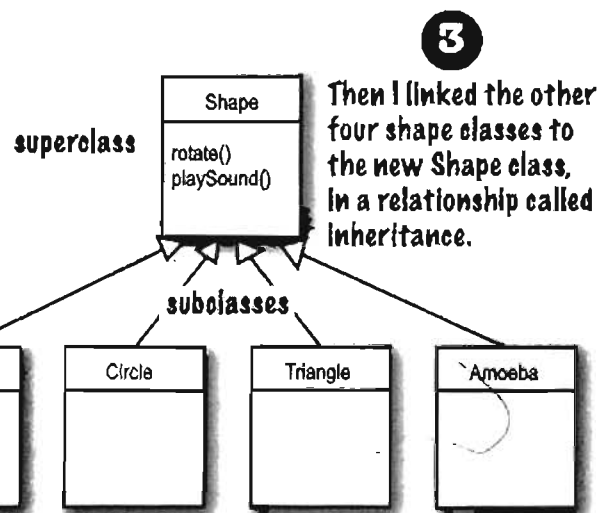
2

They're Shapes, and they all rotate and playSound. So I abstracted out the common features and put them into a new class called Shape.



You can read this as, "Square inherits from Shape", "Circle inherits from Shape", and so on. I removed rotate() and playSound() from the other shapes, so now there's only one copy to maintain.

The Shape class is called the **superclass** of the other four classes. The other four are the **subclasses** of Shape. The subclasses inherit the methods of the superclass. In other words, if the Shape class has the functionality, then the subclasses automatically get that same functionality.


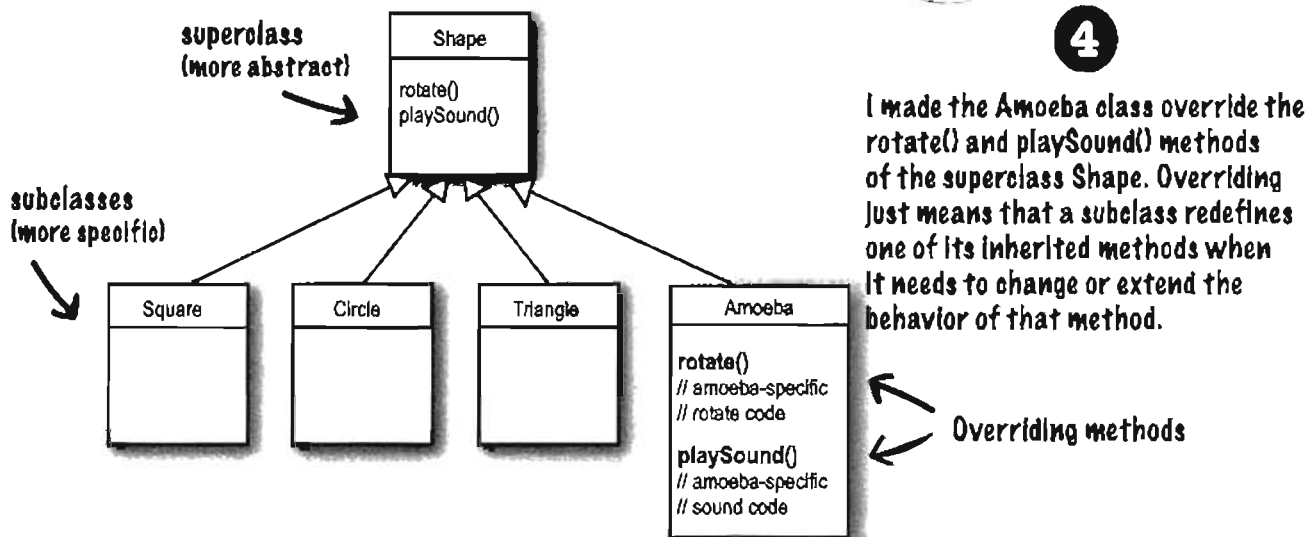


## What about the Amoeba rotate()?

**LARRY:** Wasn't that the whole problem here — that the amoeba shape had a completely different rotate and playSound procedure?

How can amoeba do something different if it *inherits* its functionality from the Shape class?

**BRAD:** That's the last step. The Amoeba class *overrides* the methods of the Shape class. Then at runtime, the JVM knows exactly which *rotate()* method to run when someone tells the Amoeba to rotate.



How would you represent a house cat and a tiger, in an inheritance structure. Is a domestic cat a specialized version of a tiger? Which would be the subclass and which would be the superclass? Or are they both subclasses to some *other* class?

How would you design an Inheritance structure? What methods would be overridden?

Think about it. *Before* you turn the page.

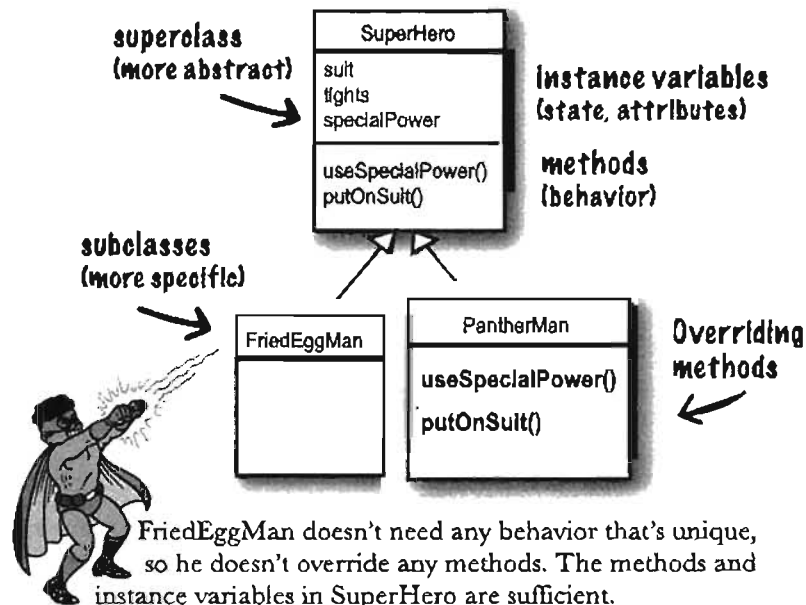
## Understanding Inheritance

When you design with inheritance, you put common code in a class and then tell other more specific classes that the common (more abstract) class is their superclass. When one class inherits from another, **the subclass inherits from the superclass**.

In Java, we say that the **subclass extends the superclass**.

An inheritance relationship means that the subclass inherits the **members** of the superclass. When we say “members of a class” we mean the instance variables and methods.

For example, if `PantherMan` is a subclass of `SuperHero`, the `PantherMan` class automatically inherits the instance variables and methods common to all superheroes including `suit`, `tights`, `specialPower`, `useSpecialPower()` and so on. But the `PantherMan` subclass can add new methods and instance variables of its own, and it can override the methods it inherits from the superclass `SuperHero`.



**Instance variables are not overridden** because they don't need to be. They don't define any special behavior, so a subclass can give an inherited instance variable any value it chooses. `PantherMan` can set his inherited `tights` to purple, while `FriedEggMan` sets his to white.



**An inheritance example:**

```

public class Doctor {

    boolean worksAtHospital;

    void treatPatient() {
        // perform a checkup
    }
}

public class FamilyDoctor extends Doctor {

    boolean makesHouseCalls;
    void giveAdvice() {
        // give homespun advice
    }
}

public class Surgeon extends Doctor{

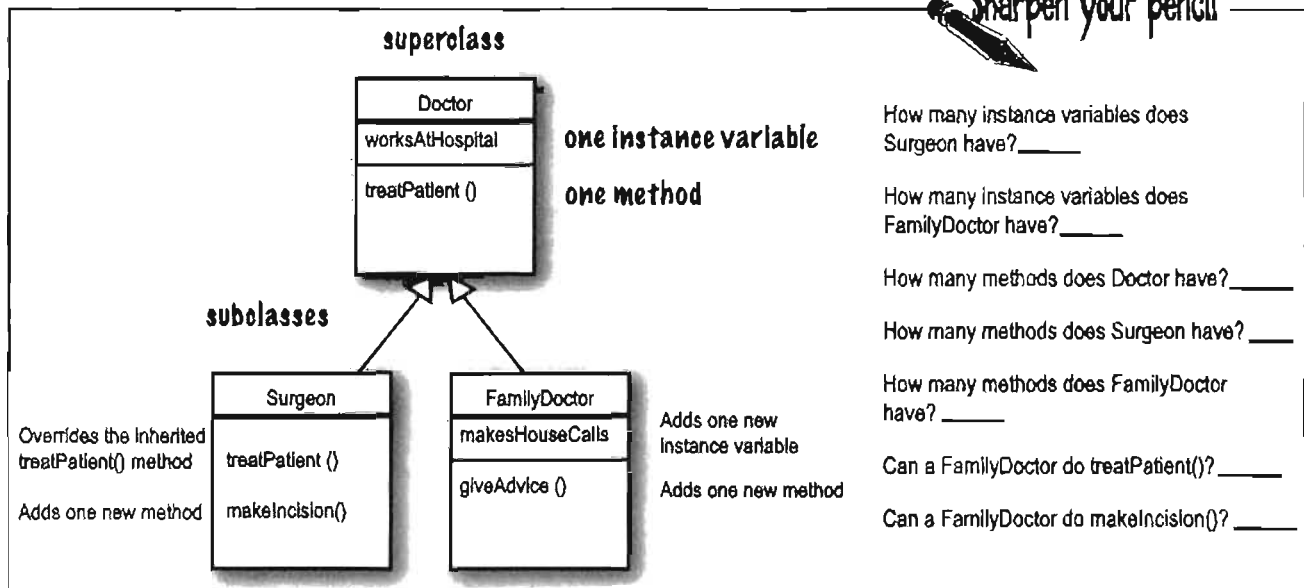
    void treatPatient() {
        // perform surgery
    }

    void makeIncision() {
        // make incision (yikes!)
    }
}

```



*Sharpen your pencil*



## Let's design the inheritance tree for an Animal simulation program

Imagine you're asked to design a simulation program that lets the user throw a bunch of different animals into an environment to see what happens. We don't have to code the thing now, we're mostly interested in the design.

We've been given a list of *some* of the animals that will be in the program, but not all. We know that each animal will be represented by an object, and that the objects will move around in the environment, doing whatever it is that each particular type is programmed to do.

*And we want other programmers to be able to add new kinds of animals to the program at any time.*

First we have to figure out the common, abstract characteristics that all animals have, and build those characteristics into a class that all animal classes can extend.

- 1 Look for objects that have common attributes and behaviors.

What do these six types have in common? This helps you to abstract out behaviors. (step 2)

How are these types related? This helps you to define the inheritance tree relationships (step 4-5)



## Using inheritance to avoid duplicating code in subclasses

We have five *instance variables*:

*picture* – the file name representing the JPEG of this animal

*food* – the type of food this animal eats. Right now, there can be only two values: *meat* or *grass*.

*hunger* – an int representing the hunger level of the animal. It changes depending on when (and how much) the animal eats.

*boundaries* – values representing the height and width of the 'space' (for example, 640 x 480) that the animals will roam around in.

*location* – the X and Y coordinates for where the animal is in the space.

We have four *methods*:

*makeNoise()* – behavior for when the animal is supposed to make noise.

*eat()* – behavior for when the animal encounters its preferred food source, *meat* or *grass*.

*sleep()* – behavior for when the animal is considered asleep.

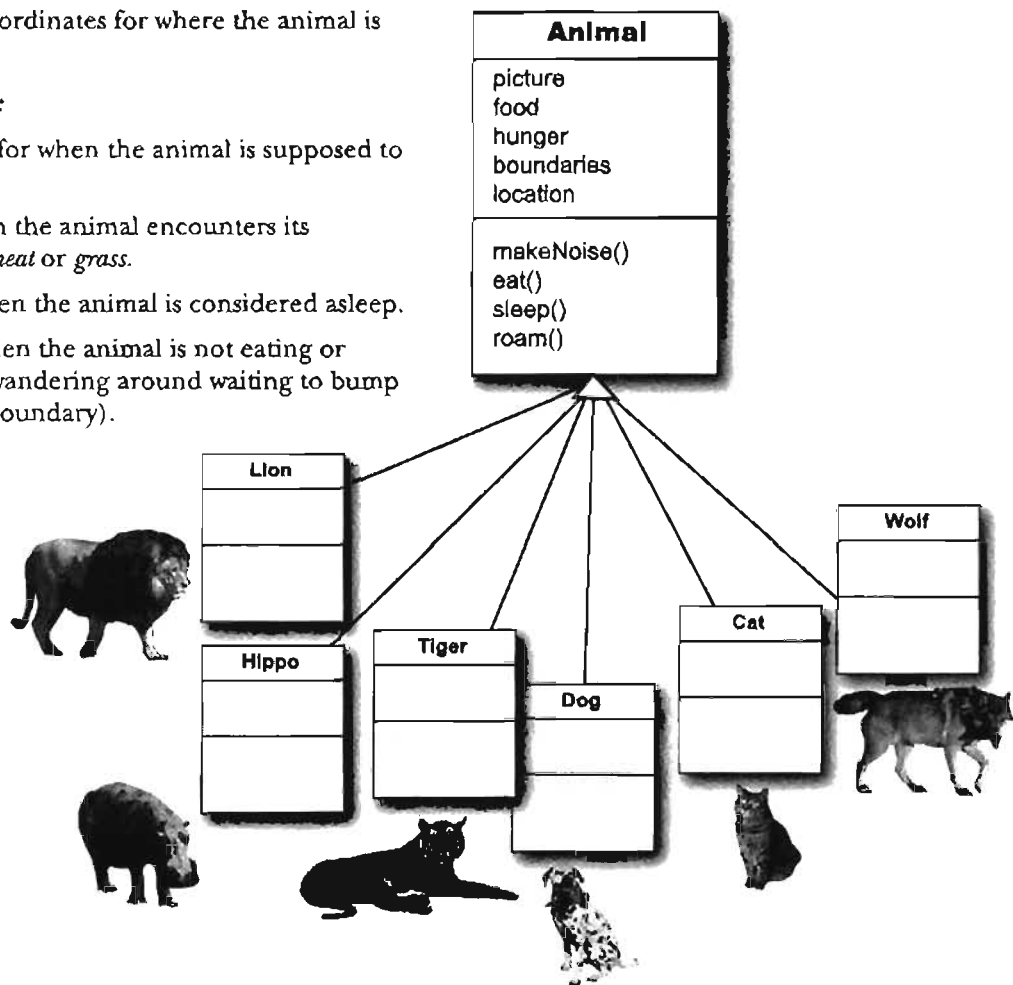
*roam()* – behavior for when the animal is not eating or sleeping (probably just wandering around waiting to bump into a food source or a boundary).

2

Design a class that represents the common state and behavior.

These objects are all animals, so we'll make a common superclass called **Animal**.

We'll put in methods and instance variables that all animals might need.



## Do all animals eat the same way?

Assume that we all agree on one thing: the instance variables will work for *all* Animal types. A lion will have his own value for picture, food (we're thinking *meat*), hunger, boundaries, and location. A hippo will have different *values* for his instance variables, but he'll still have the same variables that the other Animal types have. Same with dog, tiger, and so on. But what about *behavior*?

- 3 Decide if a subclass needs behaviors (method implementations) that are specific to that particular subclass type.

## Which methods should we override?


Does a lion make the same noise as a dog? Does a cat *eat* like a hippo? Maybe in *your* version, but in ours, eating and making noise are Animal-type-specific. We can't figure out how to code those methods in such a way that they'd work for any animal. OK, that's not true. We could write the `makeNoise()` method, for example, so that all it does is play a sound file defined in an instance variable for that type, but that's not very specialized. Some animals might make different noises for different situations (like one for eating, and another when bumping into an enemy, etc.)

So just as with the Amoeba overriding the Shape class `rotate()` method, to get more amoeba-specific (in other words, *unique*) behavior, we'll have to do the same for our Animal subclasses.

I'm one bad\*ss plant-eater.



In the dog community, barking is an important part of our cultural identity. We have a unique sound, and we want that diversity to be recognized and respected.



We better override these two methods, `eat()` and `makeNoise()`, so that each animal type can define its own specific behavior for eating and making noise. For now, it looks like `sleep()` and `roam()` can stay generic.

## Looking for more inheritance opportunities


The class hierarchy is starting to shape up. We have each subclass override the *makeNoise()* and *eat()* methods, so that there's no mistaking a Dog bark from a Cat meow (quite insulting to both parties). And a Hippo won't eat like a Lion.

But perhaps there's more we can do. We have to look at the subclasses of Animal, and see if two or more can be grouped together in some way, and given code that's common to only *that* new group. Wolf and Dog have similarities. So do Lion, Tiger, and Cat.

### 4

Look for more opportunities to use abstraction, by finding two or more *subclasses* that might need common behavior.

We look at our classes and see that Wolf and Dog might have some behavior in common, and the same goes for Lion, Tiger, and Cat.






## 5 Finish the class hierarchy

Since animals already have an organizational hierarchy (the whole kingdom, genus, phylum thing), we can use the level that makes the most sense for class design. We'll use the biological "families" to organize the animals by making a *Feline* class and a *Canine* class.

We decide that *Canines* could use a common `roam()` method, because they tend to move in packs. We also see that *Felines* could use a common `roam()` method, because they tend to avoid others of their own kind. We'll let *Hippo* continue to use its inherited `roam()` method—the generic one it gets from *Animal*.

So we're done with the design for now; we'll come back to it later in the chapter.



## Which method is called?


The Wolf class has four methods. One inherited from Animal, one inherited from Canine (which is actually an overridden version of a method in class Animal), and two overridden in the Wolf class. When you create a Wolf object and assign it to a variable, you can use the dot operator on that reference variable to invoke all four methods. But which *version* of those methods gets called?

make a new Wolf object	<code>Wolf w = new Wolf();</code>
calls the version in Wolf	<code>w.makeNoise();</code>
calls the version in Canine	<code>w.roam();</code>
calls the version in Wolf	<code>w.eat();</code>
calls the version in Animal	<code>w.sleep();</code>

When you call a method on an object reference, you're calling the most specific version of the method for that object type.

In other words, *the lowest one wins!*

"Lowest" meaning lowest on the inheritance tree. Canine is lower than Animal, and Wolf is lower than Canine, so invoking a method on a reference to a Wolf object means the JVM starts looking first in the Wolf class. If the JVM doesn't find a version of the method in the Wolf class, it starts walking back up the inheritance hierarchy until it finds a match.




practice designing an inheritance tree

## Designing an Inheritance Tree

Class	Superclasses	Subclasses
Clothing	--	Boxers, Shirt
Boxers	Clothing	
Shirt	Clothing	

Inheritance Table



Inheritance Class Diagram



Sharpen your pencil

Find the relationships that make sense. Fill in the last two columns

Class	Superclasses	Subclasses
Musician		
Rock Star		
Fan		
Bass Player		
Concert Pianist		

Hint: not everything can be connected to something else.  
Hint: you're allowed to add to or change the classes listed.

Draw an inheritance diagram here.

## there are no Dumb Questions

**Q:** You said that the JVM starts walking up the inheritance tree, starting at the class type you invoked the method on (like the Wolf example on the previous page). But what happens if the JVM doesn't ever find a match?

**A:** Good question! But you don't have to worry about that. The compiler guarantees that a particular method is callable for a specific reference type, but it doesn't say (or care) from which class that method actually comes from at runtime. With the Wolf example, the compiler checks for a sleep() method, but doesn't care that sleep() is actually defined in (and inherited from) class Animal. Remember that if a class inherits a method, it has the method.

Where the inherited method is defined (in other words, in which superclass it is defined) makes no difference to the compiler. But at runtime, the JVM will always pick the right one. And the right one means, the most specific version for that particular object.

## Using IS-A and HAS-A

Remember that when one class inherits from another, we say that the subclass *extends* the superclass. When you want to know if one thing should extend another, apply the IS-A test.

Triangle IS-A Shape, yeah, that works.

Cat IS-A Feline, that works too.

Surgeon IS-A Doctor, still good.

Tub extends Bathroom, sounds reasonable.



*Until you apply the IS-A test.*

To know if you've designed your types correctly, ask, "Does it make sense to say type X IS-A type Y?" If it doesn't, you know there's something wrong with the design, so if we apply the IS-A test, Tub IS-A Bathroom is definitely false.

What if we reverse it to Bathroom extends Tub? That still doesn't work, Bathroom IS-A Tub doesn't work.

Tub and Bathroom *are* related, but not through inheritance. Tub and Bathroom are joined by a HAS-A relationship. Does it make sense to say "Bathroom HAS-A Tub"? If yes, then it means that Bathroom has a Tub instance variable. In other words, Bathroom has a *reference* to a Tub, but Bathroom does not *extend* Tub and vice-versa.

Does it make sense to say a Tub IS-A Bathroom? Or a Bathroom IS-A Tub? Well it doesn't to me. The relationship between my Tub and my Bathroom is HAS-A. Bathroom HAS-A Tub. That means Bathroom has a Tub instance variable.



Bathroom HAS-A Tub and Tub HAS-A Bubbles.  
But nobody inherits from (extends) anybody else.


## But wait! There's more!

The IS-A test works *anywhere* in the inheritance tree. If your inheritance tree is well-designed, the IS-A test should make sense when you ask *any* subclass if it IS-A *any* of its supertypes.

**If class B extends class A, class B IS-A class A.  
This is true anywhere in the inheritance tree. If  
class C extends class B, class C passes the IS-A  
test for both B and A.**

Canine extends Animal  
Wolf extends Canine  
Wolf extends Animal

Canine IS-A Animal  
Wolf IS-A Canine  
Wolf IS-A Animal



With an inheritance tree like the one shown here, you're *always* allowed to say “Wolf extends Animal” or “Wolf IS-A Animal”. It makes no difference if Animal is the superclass of the superclass of Wolf. In fact, as long as Animal is *somewhere* in the inheritance hierarchy above Wolf, Wolf IS-A Animal will always be true.

The structure of the Animal inheritance tree says to the world: “Wolf IS-A Canine, so Wolf can do anything a Canine can do. And Wolf IS-A Animal, so Wolf can do anything an Animal can do.”

It makes no difference if Wolf overrides some of the methods in Animal or Canine. As far as the world (of other code) is concerned, a Wolf can do those four methods. *How* he does them, or *in which class they're overridden* makes no difference. A Wolf can makeNoise(), eat(), sleep(), and roam() because a Wolf extends from class Animal.



## How do you know if you've got your inheritance right?


There's obviously more to it than what we've covered so far, but we'll look at a lot more OO issues in the next chapter (where we eventually refine and improve on some of the design work we did in *this* chapter).

For now, though, a good guideline is to use the IS-A test. If "X IS-A Y" makes sense, both classes (X and Y) should probably live in the same inheritance hierarchy. Chances are, they have the same or overlapping behaviors.

### Keep in mind that the inheritance IS-A relationship works in only one direction!

Triangle IS-A Shape makes sense, so you can have Triangle extend Shape.

But the reverse—Shape IS-A Triangle—does *not* make sense, so Shape should not extend Triangle. Remember that the IS-A relationship implies that if X IS-A Y, then X can do anything a Y can do (and possibly more).



### Sharpen your pencil

Put a check next to the relationships that make sense.

- ☐ Oven extends Kitchen
- ☐ Guitar extends Instrument
- ☐ Person extends Employee
- ☐ Ferrari extends Engine
- ☐ FriedEgg extends Food
- ☐ Beagle extends Pet
- ☐ Container extends Jar
- ☐ Metal extends Titanium
- ☐ GratefulDead extends Band
- ☐ Blonde extends Smart
- ☐ Beverage extends Martini

Hint: apply the IS-A test

who inherits what

## there are no Dumb Questions

**Q:** So we see how a subclass gets to inherit a superclass method, but what if the superclass wants to use the subclass version of the method?

**A:** A superclass won't necessarily know about any of its subclasses. You might write a class and much later someone else comes along and extends it. But even if the superclass creator does know about (and wants to use) a subclass version of a method, there's no sort of *reverse* or *backwards* inheritance. Think about it, children inherit from parents, not the other way around.

**Q:** In a subclass, what if I want to use **BOTH** the superclass version and my overriding subclass version of a method? In other words, I don't want to completely *replace* the superclass version, I just want to add more stuff to it.

**A:** You can do this! And it's an important design feature. Think of the word "extends" as meaning, "I want to extend the functionality of the superclass".

```
public void roam() {  
    super.roam();  
    // my own roam stuff  
}
```

You can design your superclass methods in such a way that they contain method implementations that will work for any subclass, even though the subclasses may still need to 'append' more code. In your subclass overriding method, you can call the superclass version using the keyword **super**. It's like saying, "first go run the superclass version, then come back and finish with my own code..."

this calls the inherited version of roam(), then comes back to do your own subclass-specific code

## Who gets the Porsche, who gets the porcelain? (how to know what a subclass can inherit from its superclass)



A subclass inherits members of the superclass. Members include instance variables and methods, although later in this book we'll look at other inherited members. A superclass can choose whether or not it wants a subclass to inherit a particular member by the level of access the particular member is given.

There are four access levels that we'll cover in this book. Moving from most restrictive to least, the four access levels are:

**private   default   protected   public**

Access levels control *who sees what*, and are crucial to having well-designed, robust Java code. For now we'll focus just on public and private. The rules are simple for those two:

**public members are inherited**  
**private members are not inherited**

When a subclass inherits a member, it is *as if the subclass defined the member itself*. In the Shape example, Square inherited the `rotate()` and `playSound()` methods and to the outside world (other code) the Square class simply *has* a `rotate()` and `playSound()` method.

The members of a class include the variables and methods defined in the class plus anything inherited from a superclass.

*Note: get more details about default and protected in chapter 16 (deployment) and appendix B.*

## When designing with inheritance, are you **using** or **abusing**?

Although some of the reasons behind these rules won't be revealed until later in this book, for now, simply *knowing* a few rules will help you build a better inheritance design.

**DO** use inheritance when one class is a more specific type of a superclass. Example: Willow *is a* more specific type of Tree, so Willow *extends* Tree makes sense.

**DO** consider inheritance when you have behavior (implemented code) that should be shared among multiple classes of the same general type. Example: Square, Circle, and Triangle all need to rotate and play sound, so putting that functionality in a superclass Shape might make sense, and makes for easier maintenance and extensibility. Be aware, however, that while inheritance is one of the key features of object-oriented programming, it's not necessarily the best way to achieve behavior reuse. It'll get you started, and often it's the right design choice, but design patterns will help you see other more subtle and flexible options. If you don't know about design patterns, a good follow-on to this book would be *Head First Design Patterns*.

**DO NOT** use inheritance just so that you can reuse code from another class, if the relationship between the superclass and subclass violate either of the above two rules. For example, imagine you wrote special printing code in the Alarm class and now you need printing code in the Piano class, so you have Piano extend Alarm so that Piano inherits the printing code. That makes no sense! A Piano is *not* a more specific type of Alarm. (So the printing code should be in a Printer class, that all printable objects can take advantage of via a HAS-A relationship.)

**DO NOT** use inheritance if the subclass and superclass do not pass the IS-A test. Always ask yourself if the subclass IS-A more specific type of the superclass. Example: Tea IS-A Beverage makes sense. Beverage IS-A Tea does not.



### BULLET POINTS

- A subclass *extends* a superclass.
- A subclass inherits all *public* instance variables and methods of the superclass, but does not inherit the *private* instance variables and methods of the superclass.
- Inherited methods *can* be overridden; instance variables *cannot* be overridden (although they can be *redefined* in the subclass, but that's not the same thing, and there's almost never a need to do it.)
- Use the IS-A test to verify that your inheritance hierarchy is valid. If X *extends* Y, then X IS-A Y must make sense.
- The IS-A relationship works in only one direction. A Hippo is an Animal, but not all Animals are Hippos.
- When a method is overridden in a subclass, and that method is invoked on an instance of the subclass, the overridden version of the method is called. (*The lowest one wins.*)
- If class B extends A, and C extends B, class B IS-A class A, and class C IS-A class B, and class C also IS-A class A.

## So what does all this inheritance really buy you?

You get a lot of OO mileage by designing with inheritance. You can get rid of duplicate code by abstracting out the behavior common to a group of classes, and sticking that code in a superclass. That way, when you need to modify it, you have only one place to update, and *the change is magically reflected in all the classes that inherit that behavior*. Well, there's no magic involved, but it is pretty simple: make the change and compile the class again. That's it. **You don't have to touch the subclasses!**

**Just deliver the newly-changed superclass, and all classes that extend it will automatically use the new version.**

A Java program is nothing but a pile of classes, so the subclasses don't have to be recompiled in order to use the new version of the superclass. As long as the superclass doesn't *break* anything for the subclass, everything's fine. (We'll discuss what the word 'break' means in this context, later in the book. For now, think of it as modifying something in the superclass that the subclass is depending on, like a particular method's arguments or return type, or method name, etc.)

### ❶ You avoid duplicate code.

Put common code in one place, and let the subclasses inherit that code from a superclass. When you want to change that behavior, you have to modify it in only one place, and everybody else (i.e. all the subclasses) see the change.

### ❷ You define a common protocol for a group of classes.



**Inheritance lets you guarantee that all classes grouped under a certain supertype have all the methods that the supertype has.\***

In other words, you define a common protocol for a set of classes related through inheritance.

When you define methods in a superclass, that can be inherited by subclasses, you're announcing a kind of protocol to other code that says, "All my subtypes (i.e. subclasses) can do these things, with these methods that look like this..."

In other words, you establish a *contract*.

Class `Animal` establishes a common protocol for all `Animal` subtypes:

Animal
<pre>makeNoise() eat() sleep() roam()</pre>

You're telling the world that *any* `Animal` can do these four things. That includes the method arguments and return types.

And remember, when we say *any* `Animal`, we mean `Animal` and *any class that extends from* `Animal`. Which again means, *any class that has* `Animal` *somewhere above it in the inheritance hierarchy*.

But we're not even at the really cool part yet, because we saved the best—*polymorphism*—for last.

When you define a supertype for a group of classes, *any subclass of that supertype can be substituted where the supertype is expected*.

Say, what?

Don't worry, we're nowhere near done explaining it. Two pages from now, you'll be an expert.

\*When we say "all the methods" we mean "all the *inheritable* methods", which for now actually means, "all the *public* methods", although later we'll refine that definition a bit more.

## And I care because...

Because you get to take advantage of polymorphism.

## Which matters to me because...

Because you get to refer to a subclass object using a reference declared as the supertype.

## And that means to me...

You get to write really flexible code. Code that's cleaner (more efficient, simpler). Code that's not just easier to *develop*, but also much, much easier to *extend*, in ways you never imagined at the time you originally wrote your code.

That means you can take that tropical vacation while your co-workers update the program, and your co-workers might not even need your source code.

You'll see how it works on the next page.

We don't know about you, but personally, we find the whole tropical vacation thing particularly motivating.






## Keeping the contract: rules for overriding

When you override a method from a superclass, you're agreeing to fulfill the contract. The contract that says, for example, "I take no arguments and I return a boolean." In other words, the arguments and return types of your overriding method must look to the outside world *exactly* like the overridden method in the superclass.

### The methods *are* the contract.

If polymorphism is going to work, the Toaster's version of the overridden method from Appliance has to work at runtime. Remember, the compiler looks at the reference type to decide whether you can call a particular method on that reference. With an Appliance reference to a Toaster, the compiler cares only if class *Appliance* has the method you're invoking on an Appliance reference. But at runtime, the JVM looks not at the *reference* type (Appliance) but at the actual *Toaster* object on the heap. So if the compiler has already *approved* the method call, the only way it can work is if the overriding method has the same arguments and return types. Otherwise, someone with an Appliance reference will call `turnOn()` as a no-arg method, even though there's a version in Toaster that takes an `int`. Which one is called at runtime? The one in Appliance. In other words, *the `turnOn(int level)` method in Toaster is not an override!*




### Arguments must be the same, and return types must be compatible.

The contract of superclass defines how other code can use a method. Whatever the superclass takes as an argument, the subclass overriding the method must use that same argument. And whatever the superclass declares as a return type, the overriding method must declare either the same type, or a subclass type. Remember, a subclass object is guaranteed to be able to do anything its superclass declares, so it's safe to return a subclass where the superclass is expected.

### The method can't be less accessible.

That means the access level must be the same, or friendlier. That means you can't, for example, override a public method and make it private. What a shock that would be to the code invoking what it *thinks* (at compile time) is a public method, if suddenly at runtime the JVM slammed the door shut because the overriding version called at runtime is private!

So far we've learned about two access levels: private and public. The other two are in the deployment chapter (Release your Code) and appendix B. There's also another rule about overriding related to exception handling, but we'll wait until the chapter on exceptions (Risky Behavior) to cover that.







He treats me like an Object. But I can do so much more...if only he'd see me for what I really am.

## Get in touch with your inner Object.

An object contains *everything* it inherits from each of its superclasses. That means *every* object—regardless of its actual class type—is *also* an instance of class `Object`. That means any object in Java can be treated not just as a `Dog`, `Button`, or `Snowboard`, but also as an `Object`. When you say `new Snowboard()`, you get a single object on the heap—a `Snowboard` object—but that `Snowboard` wraps itself around an inner core representing the `Object` (capital “O”) portion of itself.



Snowboard inherits methods from superclass `Object`, and adds four more.



There is only *ONE* object on the heap here. A `Snowboard` object. But it contains both the `Snowboard` class parts of itself and the `Object` class parts of itself.

## ‘Polymorphism’ means ‘many forms’.

### You can treat a Snowboard as a Snowboard or as an Object.

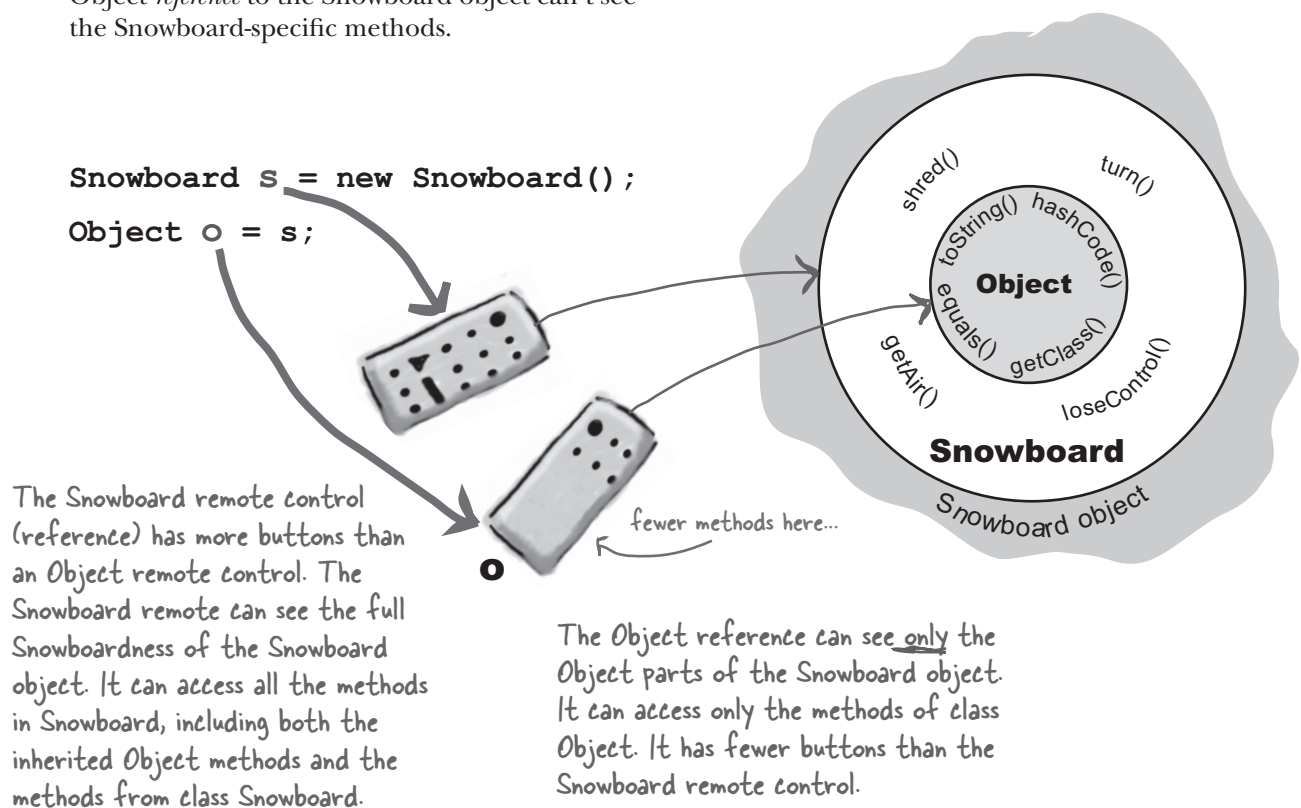
If a reference is like a remote control, the remote control takes on more and more buttons as you move down the inheritance tree. A remote control (reference) of type `Object` has only a few buttons—the buttons for the exposed methods of class `Object`. But a remote control of type `Snowboard` includes all the buttons from class `Object`, plus any new buttons (for new methods) of class `Snowboard`. The more specific the class, the more buttons it may have.

Of course that’s not always true; a subclass might not add any new methods, but simply override the methods of its superclass. The key point is that even if the *object* is of type `Snowboard`, an *Object* *reference* to the `Snowboard` object can’t see the `Snowboard`-specific methods.

**When you put an object in an `ArrayList<Object>`, you can treat it only as an `Object`, regardless of the type it was when you put it in.**

**When you get a reference from an `ArrayList<Object>`, the reference is always of type `Object`.**

**That means you get an `Object` remote control.**



## What if you need to change the contract?

OK, pretend you're a Dog. Your Dog class isn't the *only* contract that defines who you are. Remember, you inherit accessible (which usually means *public*) methods from all of your superclasses.

True, your Dog class defines a contract.

But not *all* of your contract.

**Everything in class *Canine* is part of your contract.**

**Everything in class *Animal* is part of your contract.**

**Everything in class *Object* is part of your contract.**

According to the IS-A test, you *are* each of those things—Canine, Animal, and Object.

But what if the person who designed your class had in mind the Animal simulation program, and now he wants to use you (class Dog) for a Science Fair Tutorial on Animal objects.

That's OK, you're probably reusable for that.

But what if later he wants to use you for a PetShop program? *You don't have any **Pet** behaviors.* A Pet needs methods like *beFriendly()* and *play()*.

OK, now pretend you're the Dog class programmer. No problem, right? Just add some more methods to the Dog class. You won't be breaking anyone else's code by *adding* methods, since you aren't touching the *existing* methods that someone else's code might be calling on Dog objects.

Can you see any drawbacks to that approach (adding Pet methods to the Dog class)?



Think about what **YOU** would do if **YOU** were the Dog class programmer and needed to modify the Dog so that it could do Pet things, too. We know that simply adding new Pet behaviors (methods) to the Dog class will work, and won't break anyone else's code.

But... this is a PetShop program. It has more than just Dogs! And what if someone wants to use your Dog class for a program that has *wild* Dogs? What do you think your options might be, and without worrying about how Java handles things, just try to imagine how you'd *like* to solve the problem of modifying some of your Animal classes to include Pet behaviors.

Stop right now and think about it, **before you look at the next page** where we begin to reveal *everything*.

(thus rendering the whole exercise completely useless, robbing you of your One Big Chance to burn some brain calories)

## Let's explore some design options for reusing some of our existing classes in a PetShop program.

On the next few pages, we're going to walk through some possibilities. We're not yet worried about whether Java can actually *do* what we come up with. We'll cross that bridge once we have a good idea of some of the tradeoffs.

### ① Option one

We take the easy path, and put pet methods in class `Animal`.

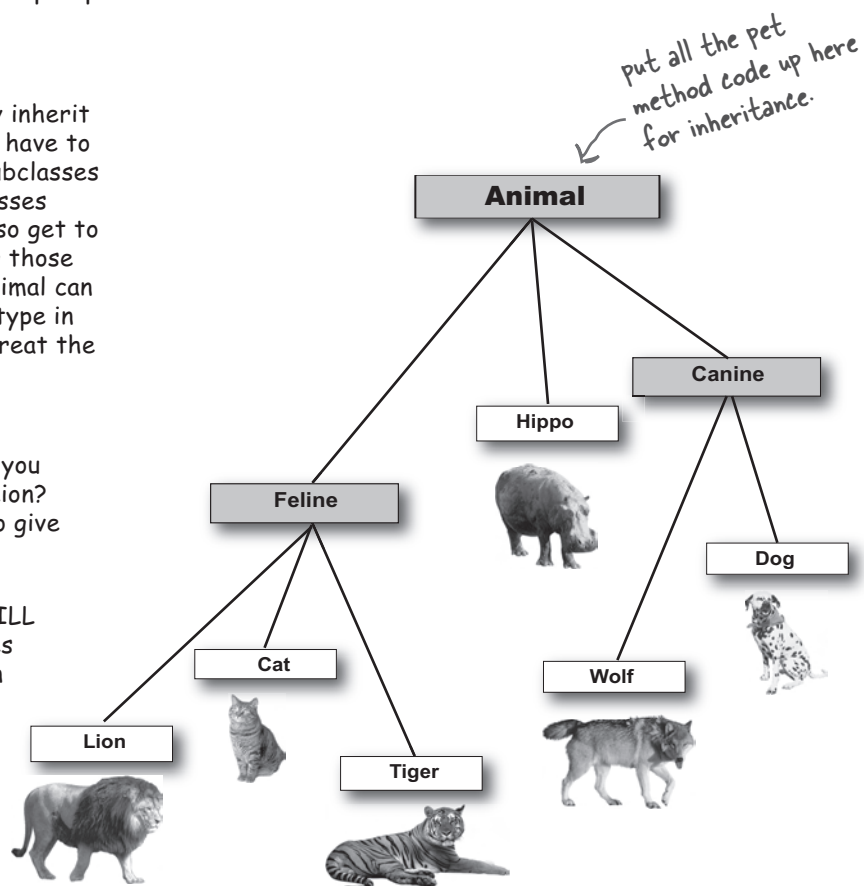
#### Pros:

All the `Animals` will instantly inherit the pet behaviors. We won't have to touch the existing `Animal` subclasses at all, and any `Animal` subclasses created in the future will also get to take advantage of inheriting those methods. That way, class `Animal` can be used as the polymorphic type in any program that wants to treat the `Animals` as pets.

#### Cons:

So... when was the last time you saw a Hippo at a pet shop? Lion? Wolf? Could be dangerous to give non-pets pet methods.

Also, we almost certainly **WILL** have to touch the pet classes like `Dog` and `Cat`, because (in our house, anyway) `Dogs` and `Cats` tend to implement pet behaviors **VERY** differently.



## ② Option two

We start with Option One, putting the pet methods in class `Animal`, but we make the methods abstract, forcing the `Animal` subclasses to override them.


### Pros:

That would give us all the benefits of Option One, but without the drawback of having non-pet `Animals` running around with pet methods (like `beFriendly()`). All `Animal` classes would have the method (because it's in class `Animal`), but because it's abstract the non-pet `Animal` classes won't inherit any functionality. All classes **MUST** override the methods, but they can make the methods "do-nothings".

### Cons:

Because the pet methods in the `Animal` class are all abstract, the concrete `Animal` subclasses are forced to implement all of them. (Remember, all abstract methods **MUST** be implemented by the first concrete subclass down the inheritance tree.) What a waste of time! You have to sit there and type in each and every pet method into each and every concrete non-pet class, and all future subclasses as well. And while this does solve the problem of non-pets actually **DOING** pet things (as they would if they inherited pet functionality from class `Animal`), the contract is bad. Every *non-pet* class would be announcing to the world that it, too, has those pet methods, even though the methods wouldn't actually **DO** anything when called.

This approach doesn't look good at all. It just seems wrong to stuff everything into class `Animal` that more than one `Animal` type might need, **UNLESS** it applies to **ALL** `Animal` subclasses.



### ③ Option three

Put the pet methods **ONLY** in the classes where they belong.

#### Pros:


No more worries about Hippos greeting you at the door or licking your face. The methods are where they belong, and **ONLY** where they belong. Dogs can implement the methods and Cats can implement the methods, but nobody else has to know about them.

#### Cons:

Two Big Problems with this approach. First off, you'd have to agree to a protocol, and all programmers of pet Animal classes now and in the future would have to **KNOW** about the protocol. By protocol, we mean the exact methods that we've decided all pets should have. The pet contract without anything to back it up. But what if one of the programmers gets it just a tiny bit wrong? Like, a method takes a String when it was supposed to take an int? Or they named it `doFriendly()` instead of `beFriendly()`? Since it isn't in a contract, the compiler has no way to check you to see if you've implemented the methods correctly. Someone could easily come along to use the pet Animal classes and find that not all of them work quite right.

And second, you don't get to use polymorphism for the pet methods. Every class that needs to use pet behaviors would have to know about each and every class! In other words, you can't use Animal as the polymorphic type now, because the compiler won't let you call a Pet method on an Animal reference (even if it's really a Dog object) because class Animal doesn't have the method.

★ Put the pet methods **ONLY** in the Animal classes that can be pets, instead of in Animal.





multiple inheritance?


### So what we REALLY need is:

- ✱ A way to have pet behavior in **just** the pet classes
- ✱ A way to guarantee that all pet classes have all of the same methods defined (same name, same arguments, same return types, no missing methods, etc.), without having to cross your fingers and hope all the programmers get it right.
- ✱ A way to take advantage of polymorphism so that all pets can have their pet methods called, without having to use arguments, return types, and arrays for each and every pet class.

**It looks like we need TWO superclasses at the top**

*We make a new abstract superclass called Pet, and give it all the pet methods.*

*Cat now extends from both Animal AND Pet, so it gets the methods of both.*




*Dog extends both Pet and Animal*

*The non-pet Animals don't have any inherited Pet stuff.*



## Fireside Chats



Tonight's Talk: **An instance variable and a local variable discuss life and death (with remarkable civility)**

### Instance Variable

I'd like to go first, because I tend to be more important to a program than a local variable. I'm there to support an object, usually throughout the object's entire life. After all, what's an object without *state*? And what is state? Values kept in *instance variables*.

No, don't get me wrong, I do understand your role in a method, it's just that your life is so short. So temporary. That's why they call you guys "temporary variables".

My apologies. I understand completely.

I never really thought about it like that. What are you doing while the other methods are running and you're waiting for your frame to be the top of the Stack again?

### Local Variable

I appreciate your point of view, and I certainly appreciate the value of object state and all, but I don't want folks to be misled. Local variables are *really* important. To use your phrase, "After all, what's an object without *behavior*?" And what is behavior? Algorithms in methods. And you can bet your bits there'll be some *local variables* in there to make those algorithms work.

Within the local-variable community, the phrase "temporary variable" is considered derogatory. We prefer "local", "stack", "automatic", or "Scope-challenged".

Anyway, it's true that we don't have a long life, and it's not a particularly *good* life either. First, we're shoved into a Stack frame with all the other local variables. And then, if the method we're part of calls another method, another frame is pushed on top of us. And if *that* method calls *another* method... and so on. Sometimes we have to wait forever for all the other methods on top of the Stack to complete so that our method can run again.

Nothing. Nothing at all. It's like being in stasis—that thing they do to people in science fiction movies when they have to travel long distances. Suspended animation, really. We just sit there on hold. As long as our frame is still there, we're safe and the value we hold is secure, but it's a mixed blessing when our

**Instance Variable**

We saw an educational video about it once. Looks like a pretty brutal ending. I mean, when that method hits its ending curly brace, the frame is literally *blown* off the Stack! Now *that's* gotta hurt.

I live on the Heap, with the objects. Well, not *with* the objects, actually *in* an object. The object whose state I store. I have to admit life can be pretty luxurious on the Heap. A lot of us feel guilty, especially around the holidays.

OK, hypothetically, yes, if I'm an instance variable of the Collar and the Collar gets GC'd, then the Collar's instance variables would indeed be tossed out like so many pizza boxes. But I was told that this almost never happens.

They let us *drink*?

**Local Variable**

frame gets to run again. On the one hand, we get to be active again. On the other hand, the clock starts ticking again on our short lives. The more time our method spends running, the closer we get to the end of the method. We *all* know what happens then.

Tell me about it. In computer science they use the term *popped* as in "the frame was popped off the Stack". That makes it sound fun, or maybe like an extreme sport. But, well, you saw the footage. So why don't we talk about you? I know what my little Stack frame looks like, but where do *you* live?

But you don't *always* live as long as the object who declared you, right? Say there's a Dog object with a Collar instance variable. Imagine *you're* an instance variable of the *Collar* object, maybe a reference to a Buckle or something, sitting there all happy inside the *Collar* object who's all happy inside the *Dog* object. But... what happens if the Dog wants a new Collar, or *nulls* out its Collar instance variable? That makes the Collar object eligible for GC. So... if *you're* an instance variable inside the Collar, and the whole *Collar* is abandoned, what happens to *you*?

And you believed it? That's what they say to keep us motivated and productive. But aren't you forgetting something else? What if you're an instance variable inside an object, and that object is referenced *only* by a *local* variable? If I'm the only reference to the object you're in, when I go, you're coming with me. Like it or not, our fates may be connected. So I say we forget about all this and go get drunk while we still can. Carpe RAM and all that.