


# COMP6771



# Dynamic Polymorphism

## Lecture 7.1

Author(s): Hayden Smith



[\(Download as PDF\)](#)



# Key Concepts

- **Inheritance:** ability to create new classes based on existing ones
  - Supported by class derivation
- **Polymorphism:** allows objects of a subclass to be used as if they were objects of a base class
  - Supported via virtual functions
- **Dynamic binding:** run-time resolution of the appropriate function to invoke based on the type of the object
  - Closely related to polymorphism
  - Supported via virtual functions



# OO Tenets Of C++

- Don't pay for what you don't use
  - C++ Supports OOP
    - No runtime performance penalty
  - C++ supports generic programming with the STL and templates
    - No runtime performance penalty
  - Polymorphism is extremely powerful, and we need it in C++
    - Do we need polymorphism at all when using inheritance?
      - Answer: sometimes
      - But how do we do so, considering that we don't want to make anyone who doesn't use it pay a performance penalty
- One of the guiding principles of C++ is "You don't pay for what you don't use"

# Inheritance In C++

```
1 class BaseClass {
2     public:
3         int get_int_member() { return int_member_; }
4         std::string get_class_name() {
5             return "BaseClass"
6         };
7
8     private:
9         int int_member_;
10        std::string string_member_;
11 }
```

```
1 class SubClass: public BaseClass {
2     public:
3         std::string get_class_name() {
4             return "SubClass";
5         }
6
7     private:
8         std::vector<int> vector_member_;
9         std::unique_ptr<int> ptr_member_;
10 }
```

# Inheritance In C++

- To inherit off classes in C++, we use "class DerivedClass: public BaseClass"
- Visibility can be one of:
  - **public** (generally use this unless you have good reason not to)
    - If you don't want public, you should (usually) use composition
  - **protected**
  - **private**
- Visibility is the maximum visibility allowed
  - If you specify ": private BaseClass", then the maximum visibility is private
    - Any BaseClass members that were public or protected are now private

# Inheritance In C++

## Memory Layout

BaseClass object		SubClass object
int_member_	BaseClass subobject	int_member_
string_member_		string_member_
	SubClass subobject	vector_member_
		ptr_member_



# Code Exploration

```
1 #include <string>
2 #include <iostream>
3
4 class BaseClass {
5     public:
6     int get_member() { return member_; }
7     std::string get_class_name() {
8         return "BaseClass";
9     };
10
11     private:
12     int member_;
13 };
14
15 class SubClass: public BaseClass {
16     public:
17     std::string get_class_name() {
18         (void) subclass_data_;
19         return "SubClass";
20     }
21
22     private:
23     int subclass_data_;
24 };
25
26 void print_class_name(BaseClass base) {
27     std::cout << base.get_class_name()
28         << ' ' << base.get_member()
29         << '\n';
30 }
31
32 int main() {
33     BaseClass base_class;
34     SubClass subclass;
35     print_class_name(base_class);
36     print_class_name(subclass);
37 }
```



# Code Exploration

How many bytes is a BaseClass instance?

How many bytes is a SubClass instance?

How can the compiler allocate space for it on the stack, when it doesn't know how big it could be?





# Code Exploration

The answer: since we care about performance, a BaseClass can only store a BaseClass,  
not a SubClass

If we try to fill that value with a SubClass, then it just fills it with the BaseClass subobject,  
and drops the SubClass subobject

This is called the **object slicing problem**

## Code Exploration

The **solution** to this is to use references/pointers (preferably references) to the baseclass



# Code Exploration

```
1 #include <string>
2 #include <iostream>
3
4 class BaseClass {
5     public:
6     int get_member() { return member_; }
7     std::string get_class_name() {
8         return "BaseClass";
9     };
10
11     private:
12     int member_;
13 };
14
15 class SubClass: public BaseClass {
16     public:
17     std::string get_class_name() {
18         (void) subclass_data_;
19         return "SubClass";
20     }
21
22     private:
23     int subclass_data_;
24 };
25
26 void print_class_name(BaseClass& base) {
27     std::cout << base.get_class_name()
28         << ' ' << base.get_member()
29         << '\n';
30 }
31
32 int main() {
33     BaseClass base_class;
34     SubClass subclass;
35     print_class_name(base_class);
36     print_class_name(subclass);
37 }
```



# More Problems

```
1 class BaseClass {
2   public:
3     int get_member() { return member_; }
4     std::string get_class_name() {
5       return "BaseClass";
6     };
7
8   private:
9     int member_;
10 };
11 class SubClass: public BaseClass {
12   public:
13     std::string get_class_name() {
14       return "SubClass";
15     }
16
17   private:
18     int subclass_data_;
19 }
```

```
1 void print_class_name(BaseClass& base) {
2   std::cout << base.get_class_name()
3             << ' ' << base.get_member()
4             << '
5 ';
6 }
7
8 int main() {
9   BaseClass base_class;
10  SubClass subclass;
11  print_class_name(base_class);
12  print_class_name(subclass);
13 }
```

- How does the compiler decide which version of `get_class_name` to call?
  - When does the compiler decide this? Compile or runtime?
- How can it ensure that calling `get_member` doesn't have similar overhead?



# Virtual & Override

```
1 void print_stuff(const BaseClass& base) {
2     std::cout << base.get_class_name() << '
3     ' ;
4 }
5 int main() {
6     SubClass subclass;
7     print_stuff(subclass);
8 }
```

By default, C++ will call "get\_class\_name()" of the BaseClass.

However, if the base class has the function marked as **virtual**, it will happily look toward the derived class.



# Virtual & Override

For example:

```
1 class BaseClass {
2   public:
3     int get_member() { return member_; }
4     virtual std::string get_class_name() {
5       return "BaseClass"
6     };
7
8   private:
9     int member_;
10 }
```

```
1 class SubClass: public BaseClass {
2   public:
3     std::string GetClassName() override {
4       return "SubClass";
5     }
6
7   private:
8     int subclass_data_;
9 }
```

Note the use of the **override** keyword to.



# Virtual & Override

- While override isn't required by the compiler, you should always use it
- Override fails to compile if the function doesn't exist in the base class. This helps with:
  - Typos
  - Refactoring
  - Const / non-const methods
  - Slightly different signatures



# Virtual & Override

Let's explore some outputs with virtual members

```
1 #include <iostream>
2 #include <string>
3
4 class BaseClass {
5 public:
6     virtual std::string get_class_name() const {
7         return "BaseClass";
8     };
9
10    virtual ~BaseClass() {
11        std::cout << "Destructing base class\n";
12    }
13 };
14
15 class SubClass: public BaseClass {
16 public:
17     std::string get_class_name() const override {
18         return "SubClass";
19     }
20
21    ~SubClass() {
22        std::cout << "Destructing subclass\n";
23    }
24 };
25
26 void print_stuff(const BaseClass& base_class) {
27     std::cout << base_class.get_class_name() << '\n';
28 }
29
30 int main() {
31     auto subclass = static_cast<std::unique_ptr<BaseClass>>(
32         std::make_unique<SubClass>());
33     std::cout << subclass->get_class_name();
34 }
```

virtual.cpp



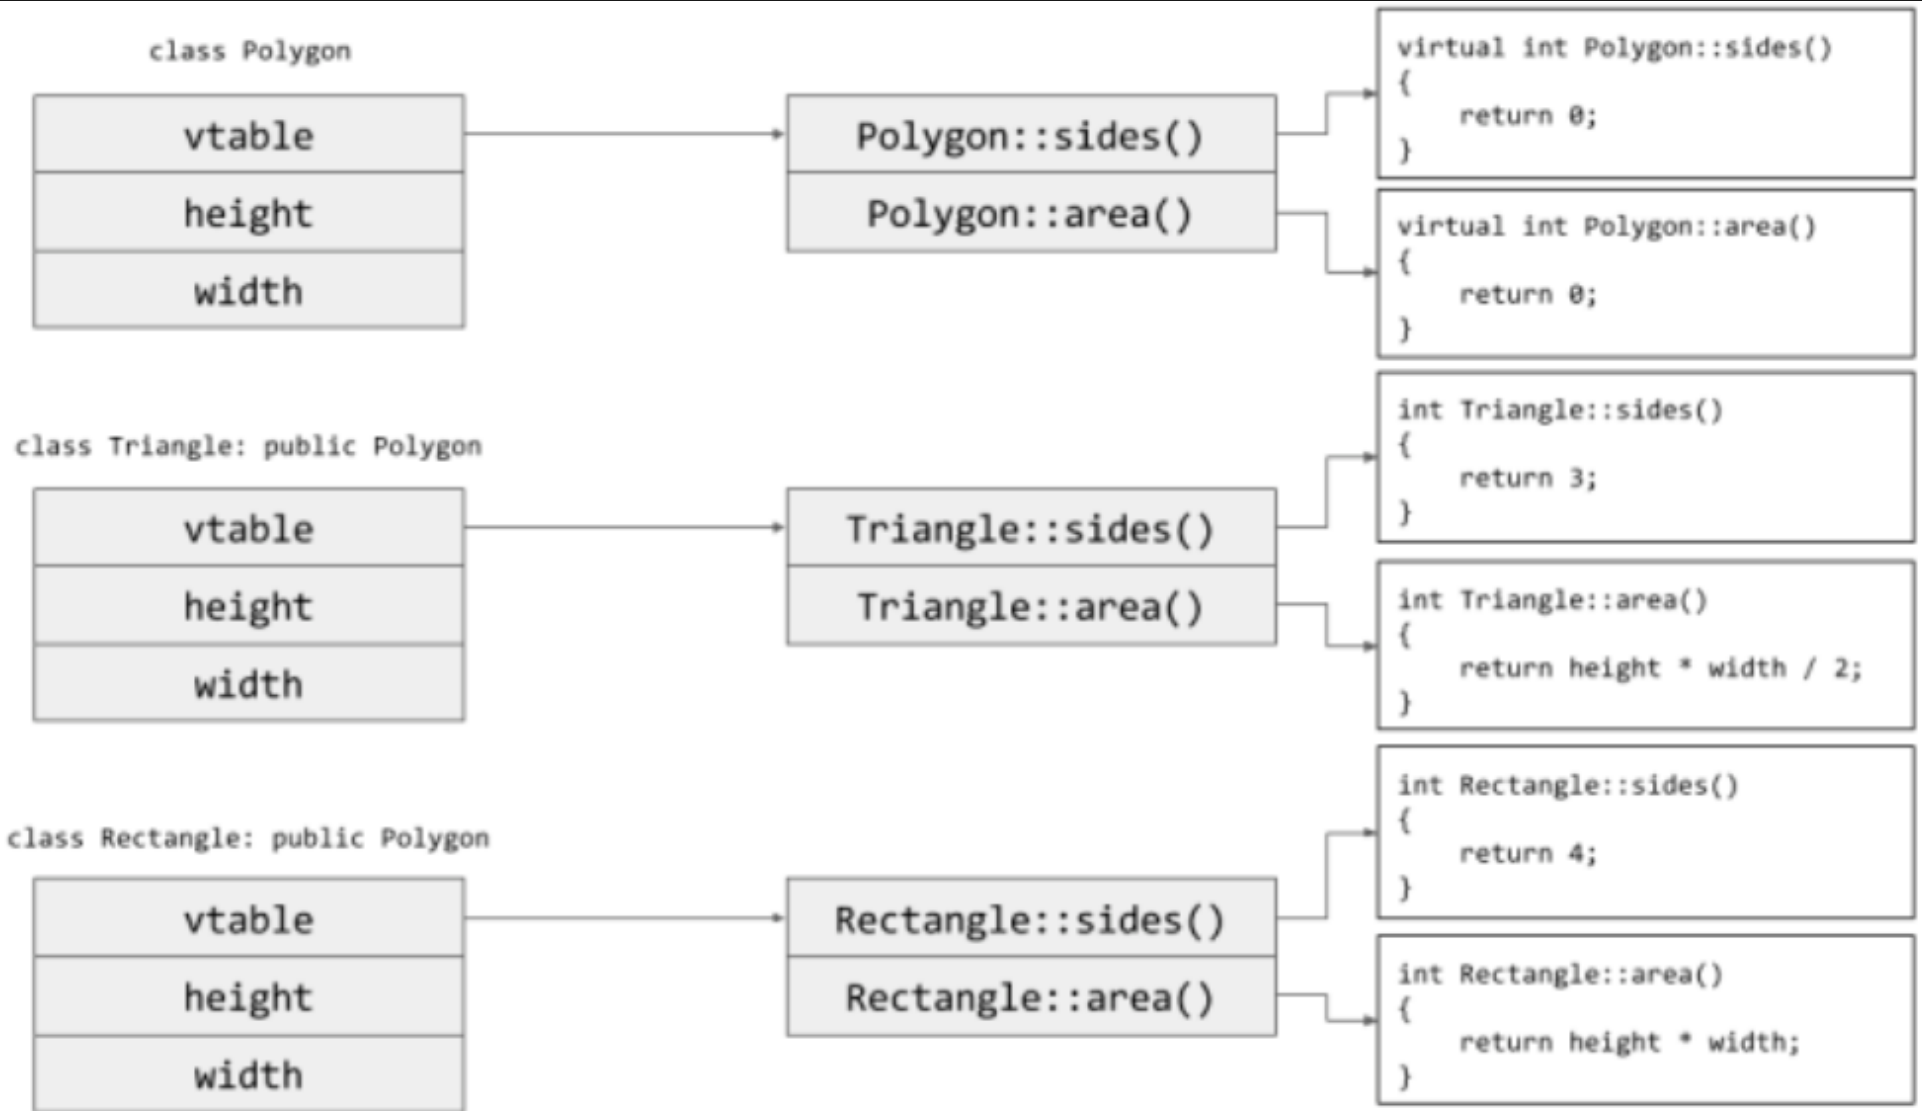


# Vtables

- Each class has a VTable stored in the data segment
  - A vtable is an array of function pointers that says which definition each virtual function points to for that class
- If the VTable for a class is non-empty, then every member of that class has an additional data member that is a pointer to the vtable
- When a virtual function is called on a reference or pointer type, then the program actually does the following
  1. Follow the vtable pointer to get to the vtable
  2. Increment by an offset, which is a constant for each function
  3. Follow the function pointer at `vtable[offset]` and call the function



# Vtables



# ✓ Final

- Specifies to the compiler "this is not virtual for any subclasses"
- If the compiler has a variable of type SubClass&, it now no longer needs to look it up in the vtable
- This means static binding if you have a SubClass&, but dynamic binding for BaseClass&

```
1 class BaseClass {
2   public:
3     int get_member() { return member_; }
4     virtual std::string get_class_name() {
5       return "BaseClass"
6     };
7
8   private:
9     int member_;
10 }
```

```
1 class SubClass: public BaseClass {
2   public:
3     std::string get_class_name() override final {
4       return "SubClass";
5     }
6
7   private:
8     int subclass_data_;
9 }
```

# ✓ Final

Syntax	Name	Meaning
<code>virtual void fn() = 0;</code>	pure virtual	Inherit interface only
<code>virtual void fn()</code>	virtual	Inherit interface with optional implementation
<code>void fn()</code>	nonvirtual	Inherit interface and mandatory implementation

- **Note: nonvirtualls can be hidden by writing a function with the same name in a subclass**
  - **DO NOT DO THIS**

# Abstract Base Classes (ABCs)

- Might want to deal with a base class, but the base class by itself is nonsense
  - What is the default way to draw a shape? How many sides by default?
  - A function takes in a "Clickable"
- Might want some default behaviour and data, but need others
  - All files have a name, but are reads done over the network or from a disk
- If a class has at least one "abstract" (pure virtual in C++) method, the class is abstract and cannot be constructed
  - It can, however, have constructors and destructors
  - These provide semantics for constructing and destructing the ABC subobject of any derived classes

# Pure Virtual Functions

- Virtual functions are good for when you have a default implementation that subclasses may want to overwrite
- Sometimes there is no default available
- A pure virtual function specifies a function that a class must override in order to not be abstract

```
1 class Shape {
2     // Your derived class "Circle" may forget to write this.
3     virtual void draw(Canvas&) {}
4
5     // Fails at link time because there's no definition.
6     virtual void draw(Canvas&);
7
8     // Pure virtual function.
9     virtual void draw(Canvas&) = 0;
10 };
```



# Creating Polymorphic Objects

- In a language like Java, everything is a pointer
  - This allows for code like on the left
  - Not possible in C++ due to objects being stored inline
    - This then leads to slicing problem
- If you want to store a polymorphic object, use a pointer

```
1 // Java-style C++ here
2 // Don't do this.
3
4 auto base = std::vector<BaseClass>();
5 base.push_back(BaseClass{});
6 base.push_back(SubClass1{});
7 base.push_back(SubClass2{});
```

```
1 // Good C++ code
2 // But there's a potential problem here.
3 // (*very* hard to spot)
4
5 auto base = std::vector<std::unique_ptr<BaseClass>>();
6 base.push_back(std::make_unique<BaseClass>());
7 base.push_back(std::make_unique<Subclass1>());
8 base.push_back(std::make_unique<Subclass2>());
```



# Inheritance And Constructors

- Every subclass constructor must call a base class constructor
  - If none is manually called, the default constructor is used
  - A subclass cannot initialise fields defined in the base class
  - Abstract classes must have constructors

```
1 class BaseClass {
2     public:
3         BaseClass(int member): int_member_{member} {}
4
5     private:
6         int int_member_;
7         std::string string_member_;
8 }
9
10 class SubClass: public BaseClass {
11     public:
12         SubClass(int member, std::unique_ptr<int>&& ptr): BaseClass(member), ptr_member_(std::move(ptr)) {}
13         // Won't compile.
14         SubClass(int member, std::unique_ptr<int>&& ptr): int_member_(member), ptr_member_(std::move(ptr)) {}
15
16     private:
17         std::vector<int> vector_member_;
18         std::unique_ptr<int> ptr_member_;
19 }
```





# Destructing Polymorphic Objects

- Which constructor is called?
- Which destructor is called?
- What could the problem be?
  - What would the consequences be?
- How might we fix it, using the techniques we've already learnt?

```
1 // Simplification of previous slides code.  
2  
3 auto base = std::make_unique<BaseClass>();  
4 auto subclass = std::make_unique<Subclass>();
```



# Destructing Polymorphic Objects

- Whenever you write a class intended to be inherited from, **always make your destructor virtual**
- Remember: When you declare a destructor, the move constructor and assignment are not synthesized

```
1 class BaseClass {  
2     BaseClass(BaseClass&&) = default;  
3     BaseClass& operator=(BaseClass&&) = default;  
4     virtual ~BaseClass() = default;  
5 }
```

Forgetting this can be a hard bug to spot



# Static & Dynamic Types

- Static type is the type it is declared as
- Dynamic type is the type of the object itself
- Static means compile-time, and dynamic means runtime
  - Due to object slicing, an object that is neither reference or pointer always has the same static and dynamic type

```
1 int main() {
2     auto base_class = BaseClass();
3     auto subclass = SubClass();
4     auto sub_copy = subclass;
5     // The following could all be replaced with pointers
6     // and have the same effect.
7     const BaseClass& base_to_base{base_class};
8     // Another reason to use auto - you can't accidentally do this.
9     const BaseClass& base_to_sub{subclass};
10    // Fails to compile
11    const SubClass& sub_to_base{base_class};
12    const SubClass& sub_to_sub{subclass};
13    // Fails to compile (even though it refers to at a sub);
14    const SubClass& sub_to_base_to_sub{base_to_sub};
15 }
```



# Static & Dynamic Types

- **Static binding:** Decide which function to call at compile time (based on static type)
- **Dynamic binding:** Decide which function to call at runtime (based on dynamic type)
- C++
  - Statically typed (types are calculated at compile time)
  - Static binding for non-virtual functions
  - Dynamic binding for virtual functions
- Java
  - Statically typed
  - Dynamic binding



# Static & Dynamic Types

## Up-casting

- Casting from a derived class to a base class is called up-casting
- This cast is always safe
  - All dogs are animals
- Because the cast is always safe, C++ allows this as an implicit cast
- One of the reasons to use auto is that it avoids implicit casts

```
1 auto dog = Dog();  
2 Animal& animal = dog;  
3 Animal* animal = &dog;
```



# Static & Dynamic Types

## Down-casting

- Casting from a base class to a derived class is called down-casting
- This cast is not safe
  - Not all animals are dogs

```
1 auto dog = Dog();
2 auto cat = Cat();
3 Animal& animal_dog{dog};
4 Animal& animal_cat{cat};
5
6 // Attempt to down-cast with references.
7 // Neither of these compile.
8 // Why not?
9 Dog& dog_ref{animal_dog};
10 Dog& dog_ref{animal_cat};
```



# Static & Dynamic Types

## How to down cast

- The compiler doesn't know if an Animal happens to be a Dog
  - If you know it is, you can use `static_cast`
  - Otherwise, you can use `dynamic_cast`
    - Returns null pointer for pointer types if it doesn't match
    - Throws exceptions for reference types if it doesn't match

```
1 auto dog = Dog();
2 auto cat = Cat();
3 Animal& animal_dog{dog};
4 Animal& animal_cat{cat};
5
6 // Attempt to down-cast with references.
7 Dog& dog_ref{static_cast<Dog&>(animal_dog)};
8 Dog& dog_ref{dynamic_cast<Dog&>(animal_dog)};
9 // Undefined behaviour (incorrect static cast).
10 Dog& dog_ref{static_cast<Dog&>(animal_cat)};
11 // Throws exception
12 Dog& dog_ref{dynamic_cast<Dog&>(animal_cat)};
```

```
1 auto dog = Dog();
2 auto cat = Cat();
3 Animal& animal_dog{dog};
4 Animal& animal_cat{cat};
5
6 // Attempt to down-cast with pointers.
7 Dog* dog_ref{static_cast<Dog*>(&animal_dog)};
8 Dog* dog_ref{dynamic_cast<Dog*>(&animal_dog)};
9 // Undefined behaviour (incorrect static cast).
10 Dog* dog_ref{static_cast<Dog*>(&animal_cat)};
11 // returns null pointer
12 Dog* dog_ref{dynamic_cast<Dog*>(&animal_cat)};
```



# Covariants

- [Read more about covariance and contravariance](#)
- **If a function overrides a base, which type can it return?**
  - If a base specifies that it returns a LandAnimal, a derived also needs to return a LandAnimal
- Every possible return type for the derived must be a valid return type for the base

```
1 class Base {
2     virtual LandAnimal& get_favorite_animal();
3 };
4
5 class Derived: public Base {
6     // Fails to compile: Not all animals are land animals.
7     Animal& get_favorite_animal() override;
8     // Compiles: All land animals are land animals.
9     LandAnimal& get_favorite_animal() override;
10    // Compiles: All dogs are land animals.
11    Dog& get_favorite_animal() override;
12 };
```





# Contravariants

- If a function overrides a base, which types can it take in?
  - If a base specifies that it takes in a LandAnimal, a LandAnimal must always be valid input in the derived
- Every possible parameter to the base must be a possible parameter for the derived

```
1 class Base {
2     virtual void use_animal(LandAnimal&);
3 };
4
5 class Derived: public Base {
6     // Compiles: All land animals are valid input (animals).
7     void use_animal(Animal&) override;
8     // Compiles: All land animals are valid input (land animals).
9     void use_animal(LandAnimal&) override;
10    // Fails to compile: Not All land animals are valid input (dogs).
11    void use_animal(Dog&) override;
12 };
```



# Default Arguments And Virtuals

- Default arguments are determined at compile time for efficiency's sake
- Hence, default arguments need to use the static type of the function
- **Avoid default arguments when overriding virtual functions**

```
1 #include <iostream>
2
3 class Base {
4 public:
5     virtual ~Base() = default;
6     virtual void print_num(int i = 1) {
7         std::cout << "Base " << i << '\n';
8     }
9 };
10
11 class Derived: public Base {
12 public:
13     void print_num(int i = 2) override {
14         std::cout << "Derived " << i << '\n';
15     }
16 };
17
18 int main() {
19     Derived derived;
20     Base* base = &derived;
21     derived.print_num(); // Prints "Derived 2"
22     base->print_num(); // Prints "Derived 1"
23 }
```

default.cpp



# Construction Of Derived Classes

- Base classes are always constructed before the derived class is constructed
  - The base class ctor never depends on the members of the derived class
  - The derived class ctor may be dependent on the members of the base class

```
1 class Animal {...}
2 class LandAnimal: public Animal {...}
3 class Dog: public LandAnimals {...}
4
5 Dog d;
6
7 // Dog() calls LandAnimal()
8 // LandAnimal() calls Animal()
9 // Animal members constructed using initialiser list
10 // Animal constructor body runs
11 // LandAnimal members constructed using initialiser list
12 // LandAnimal constructor body runs
13 // Dog members constructed using initialiser list
14 // Dog constructor body runs
```



# Virtuals In Constructors

If a class is not fully constructed, cannot perform dynamic binding

```
1 class Animal {...};
2 class LandAnimal: public Animal {
3     LandAnimal() {
4         Run();
5     }
6
7     virtual void Run() {
8         std::cout << "Land animal running
9 ";
10    }
11 };
12 class Dog: public LandAnimals {
13     void Run() override {
14         std::cout << "Dog running
15 ";
16     }
17 };
18
19 // When the LandAnimal constructor is being called,
20 // the Dog part of the object has not been constructed yet.
21 // C++ chooses to not allow dynamic binding in constructors
22 // because Dog::Run() might depend upon Dog's members.
23 Dog d;
```



# Destruction Of Derived Classes

Easy to remember order: Always opposite to construction order

```
1 class Animal {...}
2 class LandAnimal: public Animal {...}
3 class Dog: public LandAnimals {...}
4
5 auto d = Dog();
6
7 // ~Dog() destructor body runs
8 // Dog members destroyed in reverse order of declaration
9 // ~LandAnimal() destructor body runs
10 // LandAnimal members destroyed in reverse order of declaration
11 // ~Animal() destructor body runs
12 // Animal members destroyed in reverse order of declaration.
```



# Virtuals In Destructors

- If a class is partially destructed, cannot perform dynamic binding
- Unrelated to the destructor itself being virtual

```
1 class Animal {...};
2 class LandAnimal: public Animal {
3     virtual ~LandAnimal() {
4         Run();
5     }
6
7     virtual void Run() {
8         std::cout << "Land animal running
9 ";
10    }
11 };
12 class Dog: public LandAnimals {
13     void Run() override {
14         std::cout << "Dog running
15 ";
16     }
17 };
18
19 // When the LandAnimal constructor is being called,
20 // the Dog part of the object has already been destroyed.
21 // C++ chooses to not allow dynamic binding in destructors
22 // because Dog::Run() might depend upon Dog's members.
23 auto d = Dog();
```

# Feedback



Or go to the [form here](#).

