

In this module

Advanced Rust syntax

Ownership

We previously talked about ownership

- In rust there is always a single owner for each stack value
- Once the owner goes out of scope any associated values should be cleaned up
- Copy types creates copies, all other types are *moved*

Moving out of a function

We have previously seen this example

```
1 fn main() {
2     let s1 = String::from("hello");
3     let len = calculate_length(s1);
4     println!("The length of '{}' is {}.", s1, len);
5 }
6 fn calculate_length(s: String) -> usize {
7     s.len()
8 }
```

- This does not compile because ownership of `s1` is moved into `calculate_length`, meaning it is no longer available in `main` afterwards
- We can use `Clone` to create an explicit copy
- We can give ownership back by returning the value
- What about other options?

Borrowing

- We can make an analogy with real life: if somebody owns something you can borrow it from them, but eventually you have to give it back
- If a value is borrowed, it is not moved and the ownership stays with the original owner
- To borrow in rust, we create a *reference*

```
1  fn main() {
2      let x = String::from("hello");
3      let len = get_length(&x);
4      println!("{}", x, len);
5  }
6
7  fn get_length(arg: &String) -> usize {
8      arg.len()
9  }
```


References (mutable)

```
1 fn main() {
2     let mut s = String::from("hello");
3     change(&mut s);
4     println!("{}", s);
5 }
6
7 fn change(some_string: &mut String) {
8     some_string.push_str(", world");
9 }
```

```
Compiling playground v0.0.1 (/playground)
Finished dev [unoptimized + debuginfo] target(s) in 2.55s
Running `target/debug/playground`
hello, world
```

- A mutable reference can even fully replace the original value
- To do this, you can use the dereference operator (`*``) to modify the value:

```
1 *some_string = String::from("Goodbye");
```

Rules for borrowing and references

- You may only ever have one mutable reference at the same time
- You may have any number of immutable references at the same time as long as there is no mutable reference
- References cannot *live* longer than their owners
- A reference will always at all times point to a valid value

These rules can be checked by the Rust compiler.

Borrowing and memory safety

Combined with the ownership model we can be sure that whole classes of errors cannot occur.

- Rust is memory safe without having to use any runtime background proces such as a garbage collector
- But we still get the performance of a language that would normally let you manage memory manually

Reference example

```
1 fn main() {
2     let mut s = String::from("hello");
3     let s1 = &s;
4     let s2 = &s;
5     let s3 = &mut s;
6     println!("{}", s1, s2, s3);
7 }
```

Compiling playground v0.0.1 (/playground)

error[E0502]: cannot borrow `s` as mutable because it is also borrowed as immutable

```
--> src/main.rs:5:14
|
3 |     let s1 = &s;
|               -- immutable borrow occurs here
4 |     let s2 = &s;
5 |     let s3 = &mut s;
|               ^^^^^^ mutable borrow occurs here
6 |     println!("{}", s1, s2, s3);
|                                   -- immutable borrow later used here
```

For more information about this error, try `rustc --explain E0502`.

error: could not compile `playground` due to previous error

Returning references

You can return references, but the value borrowed from must exist at least as long

```
1 fn give_me_a_ref() -> &String {
2     let s = String::from("Hello, world!");
3     &s
4 }
```

```
Compiling playground v0.0.1 (/playground)
error[E0106]: missing lifetime specifier
--> src/lib.rs:1:23
|
|
1 | fn give_me_a_ref() -> &String {
|                        ^ expected named lifetime parameter
|
= help: this function's return type contains a borrowed value, but there is no value for it to be borrowed from
help: consider using the `static` lifetime
|
1 | fn give_me_a_ref() -> &'static String {
|                        ~~~~~~
```

For more information about this error, try `rustc --explain E0106`.
error: could not compile `playground` due to previous error

Returning references

You can return references, but the value borrowed from must exist at least as long

```
1 fn give_me_a_ref(input: &(String, i32)) -> &String {  
2     &input.0  
3 }
```

```
1 fn give_me_a_value() -> String {  
2     let s = String::from("Hello, world!");  
3     s  
4 }
```

Types redux

We have previously looked at some of the basic types in the Rust typesystem

- Primitives (integers, floats, booleans, characters)
- Compounds (tuples, arrays)
- Most of the types we looked at were `Copy`
- Borrowing will make more sense when we look at some more ways we can type our data

Structuring data

Rust has two important ways to structure data

- structs
- enums
- unions

Structs

A struct is similar to a tuple, but this time the combined type gets its own name

```
1 struct ControlPoint(f64, f64, bool);
```

This is an example of a *tuple struct*. You can access the fields in the struct the same way as with tuples:

```
1 fn main() {  
2     let cp = ControlPoint(10.5, 12.3, true);  
3     println!("{}", cp.0); // prints 10.5  
4 }
```

Structs

Much more common though are structs with named fields

```
1 struct ControlPoint {
2     x: f64,
3     y: f64,
4     enabled: bool,
5 }
```

- We can add a little more purpose to each field
- No need to keep our indexing up to date when we add or remove a field

```
1 fn main() {
2     let cp = ControlPoint {
3         x: 10.5,
4         y: 12.3,
5         enabled: true,
6     };
7     println!("{}", cp.x); // prints 10.5
8 }
```

Enumerations

One of the more powerful kinds of types in Rust are enumerations

```
1  enum IpAddressType {  
2      Ipv4,  
3      Ipv6,  
4  }
```

- An enumeration (listing) of different *variants*
- Each variant is an alternative value of the enum, you pick a single value to create an instance

```
1  fn main() {  
2      let ip_type = IpAddressType::Ipv4;  
3  }
```


Enumerations

But enums get more powerful, because each variant can have associated data with it

```
1 enum IPAddress {
2     Ipv4(u8, u8, u8, u8),
3     Ipv6(u16, u16, u16, u16, u16, u16, u16, u16),
4 }
```

- This way, the associated data and the variant are bound together
- Impossible to create an ipv6 address while only giving a 32 bits integer


```
1 fn main() {
2     let ipv4_home = IPAddress::Ipv4(127, 0, 0, 1);
3     let ipv6_home = IPAddress::Ipv6(0, 0, 0, 0, 0, 0, 0, 1);
4 }
```

- Note: an enum always is as large as the largest variant

IPAddress::Ipv4(127, 0, 0, 1)



IPAddress::Ipv6(0, 0, 0, 0, 0, 0, 0, 1)



Pattern matching

To extract data from enums we can use pattern matching using the `if let [pattern] = [value]` statement

```
1 fn accept_ipv4(ip: IpAddress) {
2     if let IpAddress::Ipv4(a, b, _, _) = ip {
3         println!("Accepted, first octet is {} and second is {}", a, b);
4     }
5 }
```

- `a` and `b` introduce local variables within the body of the if that contain the values of those fields
- The underscore (`_`) can be used to accept any value

Match

But pattern matching is very powerful if combined with the match statement

```
1 fn accept_home(ip: IpAddress) {
2   match ip {
3     IpAddress::Ipv4(127, 0, 0, 1) => {
4       println!("You are home!");
5     },
6     IpAddress::Ipv6(0, 0, 0, 0, 0, 0, 0, 1) => {
7       println!("You are in your new home!");
8     },
9     _ => {
10      println!("You are not home");
11    },
12  }
```

- Every part of the match is called an arm
- A match is exhaustive, which means that all values must be handled by one of the match arms
- You can use a catch-all `_` arm to catch any remaining cases if there are any left

Match as an expression

The match statement can even be used as an expression

```
1 fn get_first_byte(ip: IpAddress) {
2     let first_byte = match ip {
3         IpAddress::Ipv4(a, _, _, _) => a,
4         IpAddress::Ipv6(a, _, _, _, _, _, _, _) => a / 256 as u8,
5     };
6     println!("The first byte was: {}", first_byte);
7 }
```

- The match arms can return a value, but their types have to match
- Note how here we do not need a catch all ``_`` arm because all cases have already been handled by the two arms

Generics

Enums become even more powerful if we introduce a little of generics

```
1 struct PointFloat(f64, f64);
2 struct PointInt(i64, i64);
```

We are repeating ourselves here, what if we could write a datastructure for both of these cases?

```
1 struct Point<T>(T, T);
2
3 fn main() {
4     let float_point: Point<f64> = Point(10.0, 10.0);
5     let int_point: Point<i64> = Point(10, 10);
6 }
```

Generics are much more powerful, but this is all we need for now

Option

A quick look into the basic enums available in the standard library

- Rust does not have null, but you can still define variables that optionally do not have a value
- For this you can use the `Option<T>` enum

```
1  enum Option<T> {
2      Some(T),
3      None,
4  }
5
6  fn main() {
7      let some_int = Option::Some(42);
8      let no_string: Option<String> = Option::None;
9  }
```

Option

A quick look into the basic enums available in the standard library

- Rust does not have null, but you can still define variables that optionally do not have a value
- For this you can use the `Option<T>` enum

```
1  enum Option<T> {
2      Some(T),
3      None,
4  }
5
6  fn main() {
7      let some_int = Some(42);
8      let no_string: Option<String> = None;
9  }
```

Error handling

What would we do when there is an error?

```
1 fn divide(x: i64, y: i64) -> i64 {
2     if y == 0 {
3         // what to do now?
4     } else {
5         x / y
6     }
7 }
```


Error handling

What would we do when there is an error?

```
1 fn divide(x: i64, y: i64) -> i64 {
2     if y == 0 {
3         panic!("Cannot divide by zero");
4     } else {
5         x / y
6     }
7 }
```

- A panic in Rust is the most basic way to handle errors
- A panic error is an all or nothing kind of error
- A panic will immediately stop running the current thread/program and instead immediately work to shut it down, using one of two methods:
 - Unwinding: going up through the stack and making sure that each value is cleaned up
 - Aborting: ignore everything and immediately exit the thread/program
- Only use panic in small programs if normal error handling would also exit the program
- Avoid using panic in library code or other reusable components

Error handling

What would we do when there is an error? We could try and use the option enum instead of panicking

```
1 fn divide(x: i64, y: i64) -> Option<i64> {  
2     if y == 0 {  
3         None  
4     } else {  
5         Some(x / y)  
6     }  
7 }
```

Result

Another really powerful enum is the result, which is even more useful if we think about error handling

```
1  enum Result<T, E> {
2      Ok(T),
3      Err(E),
4  }
5
6  enum DivideError {
7      DivisionByZero,
8      CannotDivideOne,
9  }
10
11 fn divide(x: i64, y: i64) -> Result<i64, DivideError> {
12     if x == 1 {
13         Err(DivideError::CannotDivideOne)
14     } else if y == 0 {
15         Err(DivideError::DivisionByZero)
16     } else {
17         Ok(x / y)
18     }
19 }
```

Handling results

Now that we have a function that returns a result we have to think about how we handle that error at the call-site

```
1 fn div_zero_fails() {
2     match divide(10, 0) {
3         Ok(div) => println!("{}", div),
4         Err(e) => panic!("Could not divide by zero"),
5     }
6 }
```

- We made the signature of the `divide` function explicit in how it can fail
- The user of the function can now decide what to do, even if it is panicking
- Note: just as with `Option` we never have to use `Result::Ok` and `Result::Err` because they have been made available globally

Handling results

Especially when writing initial prototyping code you will often find yourself wanting to write error handling code later, Rust has a useful utility function to help you for both `Option` and `Result`:

```
1 fn div_zero_fails() {
2     let div = divide(10, 0).unwrap();
3     println!("{}", div);
4 }
```

- `unwrap` checks if the `Result/Option` is `Ok(x)` or `Some(x)` respectively and then return that `x`, otherwise it will panic your program with an error message
- Having `unwraps` all over the place is generally considered a bad practice
- Sometimes you can ensure that an error won't occur, in such cases `unwrap` can be a good solution

Handling results

Especially when writing initial prototyping code you will often find yourself wanting to write error handling code later, Rust has a useful utility function to help you for both `Option` and `Result`:

```
1 fn div_zero_fails() {
2     let div = divide(10, 0).unwrap_or(-1);
3     println!("{}", div);
4 }
```

Besides `unwrap`, there are some other useful utility functions

- `unwrap_or(val)`: If there is an error, use the value given to `unwrap_or` instead
- `unwrap_or_default()`: Use the default value for that type if there is an error
- `expect(msg)`: Same as `unwrap`, but instead pass a custom error message
- `unwrap_or_else(fn)`: Same as `unwrap_or`, but instead call a function that generates a value in case of an error

Result and the `?` operator

Results are so common that there is a special operator associated with them, the `?` operator

```
1 fn can_fail() -> Result<i64, Error> {
2     let intermediate_result = match divide(10, 0) {
3         Ok(ir) => ir,
4         Err(e) => return Err(e);
5     };
6
7     match divide(intermediate_result, 0) {
8         Ok(sec) => Ok(sec * 2),
9         Err(e) => Err(e),
10    }
11 }
```

Look how this function changes if we use the `?` operator

```
1 fn can_fail() -> Result<i64, Error> {
2     let intermediate_result = divide(10, 0)?;
3     Ok(divide(intermediate_result, 0)? * 2)
4 }
```

Result and the `?` operator

```
1 fn can_fail() -> Result<i64, Error> {  
2     let intermediate_result = divide(10, 0)?;  
3     Ok(divide(intermediate_result, 0)? * 2)  
4 }
```

- The `?` operator does an implicit match, if there is an error, that error is then immediately returned and the function returns early
- If the result is `Ok()` then the value is extracted and we can continue right away

Intermission: Impl blocks

In the past few slides we saw a syntax which wasn't explained before:

```
1 fn main() {  
2     let x = Some(42);  
3     let unwrapped = x.unwrap();  
4     println!("{}", unwrapped);  
5 }
```

- The syntax `x.y()` looks similar to how we accessed a field in a struct
- We can define functions on our types using impl blocks
- Impl blocks can be defined on any type, not just structs (with some limitations)

Intermission: Impl blocks

```
1  enum IPAddress {
2      Ipv4(u8, u8, u8, u8),
3      Ipv6(u16, u16, u16, u16, u16, u16, u16, u16),
4  }
5
6  impl IPAddress {
7      fn as_u32(&self) -> Option<u32> {
8          match self {
9              IPAddress::Ipv4(a, b, c, d) => a << 24 + b << 16 + c << 8 + d
10             _ => None, _
11         }
12     }
13 }
14
15 fn main() {
16     let addr = IPAddress::Ipv4(127, 0, 0, 1);
17     println!("{:?}", addr.as_u32());
18 }
```

Intermission: Impl blocks, self and Self

- The `self` parameter defines how the method can be used.
- The `Self` type is a shorthand for the type on which the current implementation is specified.

```
1  struct Foo(i32);
2
3  impl Foo {
4      fn consume(self) -> Self {
5          Self(self.0 + 1)
6      }
7
8      fn borrow(&self) -> &i32 {
9          &self.0
10     }
11
12     fn borrow_mut(&mut self) -> &mut i32 {
13         &mut self.0
14     }
15
16     fn new() -> Self {
17         Self(0)
18     }
19 }
```

Intermission: Impl blocks, the self parameter

The self parameter is called the *receiver*.

- The self parameter is always the first and it always has the type on which it was defined
- We never specify the type of the self parameter
- We can optionally prepend `&` or `&mut` to self to indicate that we take a value by reference
- Absence of a self parameter means that the function is an associated function instead

```
1 fn main () {
2     let mut f = Foo::new();
3     println!("{}", f.borrow());
4     *f.borrow_mut() = 10;
5     let g = f.consume();
6     println!("{}", g.borrow());
7 }
```

Vec: storing more of the same

The vector is an array that can grow

- Compare this to the array we previously saw, which has a fixed size

```
1 fn main() {
2     let arr = [1, 2];
3     println!("{:?}", arr);
4
5     let mut nums = Vec::new();
6     nums.push(1);
7     nums.push(2);
8     println!("{:?}", nums);
9 }
```

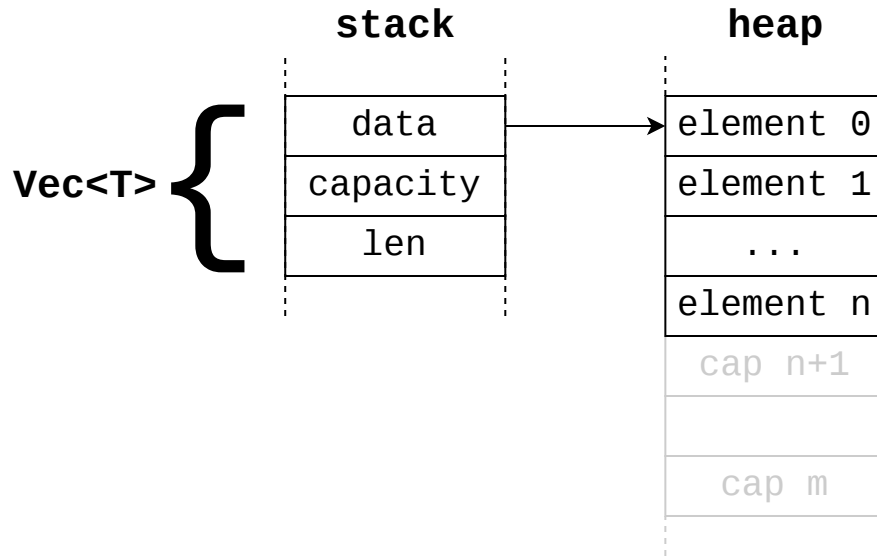
Vec

Vec is such a common type that there is an easy way to initialize it with values that looks similar to arrays

```
1 fn main() {  
2     let mut nums = vec![1, 2];  
3     nums.push(3);  
4     println!("{:?}", nums);  
5 }
```

Vec: memory layout

How can a vector grow? Things on the stack need to be of a fixed size




Put it in a box

That pointer from the stack to the heap, how do we create such a thing?

- Boxing something is the way to create data that is stored on the heap
- A box uniquely owns that data, there is no one else that also owns the same data
- Even if the type inside the box is `Copy`, the box itself is not, move semantics apply to a box.

```
1 fn main() {  
2     // put an integer on the heap  
3     let boxed_int = Box::new(10);  
4 }
```



Boxing

There are several reasons to box a variable on the heap

- When something is too large to move around
- We need something that is sized dynamically
- For writing recursive datastructures

```
1 struct Node {  
2     data: Vec<u8>,  
3     parent: Node,  
4 }
```

Boxing

There are several reasons to box a variable on the heap

- When something is too large to move around
- We need something that is sized dynamically
- For writing recursive datastructures

```
1 struct Node {  
2     data: Vec<u8>,  
3     parent: Box<Node>,  
4 }
```

Vectors and arrays

What if we wanted to write a sum function, we could define one for arrays of a specific size:

```
1 fn sum(data: &[i64; 10]) -> i64 {  
2     let mut total = 0;  
3     for val in data {  
4         total += val;  
5     }  
6     total  
7 }
```

Vectors and arrays

Or one for just vectors:

```
1 fn sum(data: &Vec<i64>) -> i64 {  
2     let mut total = 0;  
3     for val in data {  
4         total += val;  
5     }  
6     total  
7 }
```

Slices

But what if we want something to work on arrays of any size? Or what if we want to support summing up only parts of a vector?

- A slice is a dynamically sized view into a contiguous sequence
- Contiguous: elements are laid out in memory such that they are evenly spaced
- Dynamically sized: the size of the slice is not stored in the type, but is determined at runtime
- View: a slice is never an owned datastructure
- Slices are typed as `[T]`, where `T` is the type of the elements in the slice

Slices

```
1 fn sum(data: [i64]) -> i64 {
2     let mut total = 0;
3     for val in data {
4         total += val;
5     }
6     total
7 }
8
9 fn main() {
10    let data = vec![10, 11, 12, 13, 14];
11    println!("{}", sum(data));
12 }
```

```
1     Compiling playground v0.0.1 (/playground)
2 error[E0277]: the size for values of type `[i64]` cannot be known at compilation time
3 --> src/main.rs:1:8
4 |
5 1 | fn sum(data: [i64]) -> i64 {
6 |     ^^^^^ doesn't have a size known at compile-time
7 |
8 = help: the trait `Sized` is not implemented for `[i64]`
9 help: function arguments must have a statically known size, borrowed types always have a known size
```


Slices

```
1 fn sum(data: &i64) -> i64 {
2     let mut total = 0;
3     for val in data {
4         total += val;
5     }
6     total
7 }
8
9 fn main() {
10    let data = vec![10, 11, 12, 13, 14];
11    println!("{}", sum(&data));
12 }
```

```
1    Compiling playground v0.0.1 (/playground)
2    Finished dev [unoptimized + debuginfo] target(s) in 0.89s
3    Running `target/debug/playground`
4    60
```

Slices

- `[T]` is an incomplete type: we need to know how many `T` there are
- Types that have a known compile time size implement the `Sized` trait, raw slices do not implement it
- Slices must always be behind a reference type, i.e. `&[T]` and `&mut [T]` (but also `Box<[T]>` etc)
- The length of the slice is always stored together with the reference



Creating slices

Because we cannot create slices out of thin air, they have to be located somewhere. There are three possible ways to create slices:

- Using a borrow
 - We can borrow from arrays and vectors to create a slice of their entire contents
- Using ranges
 - We can use ranges to create a slice from parts of a vector or array
- Using a literal (for immutable slices only)
 - We can have memory statically available from our compiled binary

Creating slices

Using a borrow

```
1 fn sum(data: &[i32]) -> i32 { /* ... */ }
2
3 fn main() {
4     let v = vec![1, 2, 3, 4, 5, 6];
5     let total = sum(&v);
6     println!("{}", total);
7 }
```

Creating slices

Using ranges

```
1 fn sum(data: &[i32]) -> i32 { /* ... */ }
2
3 fn main() {
4     let v = vec![0, 1, 2, 3, 4, 5, 6];
5     let all = sum(&v[..]);
6     let except_first = sum(&v[1..]);
7     let except_last = sum(&v[..5]);
8     let except_ends = sum(&v[1..5]);
9 }
```

- The range `start..end` contains all values `x` with `start <= x < end`.
- Note: you can also use ranges on their own, for example in a for loop:

```
1 fn main() {
2     for i in 0..10 {
3         println!("{}", i);
4     }
5 }
```

Creating slices

From a literal

```
1 fn sum(data: &[i32]) -> i32 { /* ... */ }
2
3 fn get_v_arr() -> &'static [i32] {
4     &[0, 1, 2, 3, 4, 5, 6]
5 }
6
7 fn get_v_vec() -> &'static [i32] {
8     &vec![0, 1, 2, 3, 4, 5, 6]
9 }
10
11 fn main() {
12     let all = sum(get_v_arr());
13     let all_vec = sum(get_v_vec());
14 }
```

- Interestingly `get_v_arr` works, even though the literal looks like it would only exist temporarily
- Literals actually exist during the entire lifetime of the program
- `&'static` here is used to indicate that this slice will exist the entire lifetime of the program

Strings

We have already seen the `String` type being used before, but let's dive a little deeper

- Strings are used to represent text
- In Rust they are always valid UTF-8
- Their data is stored on the heap
- A `String` is almost the same as `Vec<u8>` with extra checks to prevent creating invalid text

Strings

Let's take a look at some strings

```
1 fn main() {
2     let s = String::from("Hello world\nSee you!");
3     println!("{:?}", s.split_once(" "));
4     println!("{}", s.len());
5     println!("{:?}", s.starts_with("Hello"));
6     println!("{}", s.to_uppercase());
7     for line in s.lines() {
8         println!("{}", line);
9     }
10 }
```

String literals

We have already seen string literals being used while constructing a string. The string literal is what arrays are to vectors

```
1 fn main() {  
2     let s1 = "Hello world";  
3     let s2 = String::from("Hello world");  
4 }
```

String literals

We have already seen string literals being used while constructing a string. The string literal is what arrays are to vectors

```
1 fn main() {  
2     let s1: &'static str = "Hello world";  
3     let s2: String = String::from("Hello world");  
4 }
```

- `s1` is actually a slice, a string slice

String literals

We have already seen string literals being used while constructing a string. The string literal is what arrays are to vectors

```
1 fn main() {  
2     let s1: &str = "Hello world";  
3     let s2: String = String::from("Hello world");  
4 }
```

- `s1` is actually a slice, a string slice

str - the string slice

It should be possible to have a reference to part of a string. But what is it?

- Not `[u8]`: not every sequence of bytes is valid UTF-8
- Not `[char]`: we could not create a slice from a string since it is stored as UTF-8 encoded bytes
- We introduce a new special kind of slice: `str`
- For string slices we do not use brackets!

str, String, array, Vec

Static	Dynamic	Borrowed
<code>[T; N]</code>	<code>Vec<T></code>	<code>&[T]</code>
-	<code>String</code>	<code>&str</code>

- There is no static variant of `str`
- This would only be useful if we wanted strings of an exact length
- But just like we had the static slice literals, we can use `&'static str` literals for that instead!

String or str

When do we use String and when do we use str?

```
1 fn string_len(data: &String) -> usize {  
2     data.len()  
3 }
```

String or str

When do we use String and when do we use str?

```
1 fn string_len(data: &str) -> usize {  
2     data.len()  
3 }
```

- Prefer `&str` over `String` whenever possible
- If you need to mutate a string you might try `&mut str`, but you cannot change a slice's length
- Use `String` or `&mut String` if you need to fully mutate the string

Summary

- Rust uses ownership and borrowing to give memory safety without a garbage collector
- Rust has structs and enums to structure your data
- Use `panic!`, `Result` and `Option` for handling errors and missing values
- Define methods and associated functions with `impl` blocks
- Use `Vec<T>` for growable array storage
- Use `Box<T>` to put something on the heap
- Use slices whenever possible instead of owned `Vec` and `String` types

Exercises

- We'll be doing the A2 exercises, see <https://101-rs.tweede.golf>
- To keep in contact we will use Discord: <https://discord.gg/pzv92cAZ>
- Join one of the voice channels and ask us to join you in the `#lab-sessions` channel when you need help!
- Don't hesitate to ask when you get stuck!

