

Haskell

from the very beginning

John Whittington

COHERENT PRESS

Cambridge

Published in the United Kingdom by Coherent Press, Cambridge

© Coherent Press 2019

This publication is in copyright. Subject to statutory
exception no reproduction of any part may take place
without the written permission of Coherent Press.

First published October 2019
Reprinted with corrections May 2020

A catalogue record for this book is available from the British Library

ISBN 978-0-9576711-3-3 Paperback

by the same author

PDF Explained (O'Reilly, 2012)

OCaml from the Very Beginning (Coherent, 2013)

More OCaml: Methods, Algorithms & Diversions (Coherent, 2014)

A Machine Made this Book: Ten Sketches of Computer Science (Coherent, 2016)

Contents

Getting Ready	ix
1 Starting Off	1
2 Names and Functions	9
Using Scripts	21
3 Case by Case	23
4 Making Lists	31
Two Ways of Thinking	41
5 Sorting Things	43
6 Functions upon Functions upon Functions	51
7 When Things Go Wrong	59
8 Looking Things Up	67
9 More with Functions	73
10 New Kinds of Data	79
11 Growing Trees	87
12 The Other Numbers	95
13 Being Lazy	105
14 In and Out	111
15 Building Bigger Programs	127
16 The Standard Prelude and Base	135
Answers to Questions	143
Hints for Questions	189
Coping with Errors	197
Index	201

Preface

This book is based on the Author's experience of teaching programming to students in the University of Cambridge supervisions system. In particular, working with students for the first-year undergraduate course "Foundations of Computer Science", lectured for many years by Lawrence C. Paulson.

An interesting aspect of supervising students from a wide range of backgrounds – some with no previous experience at all, taking Computer Science as an additional subject within the Cambridge Natural Sciences curriculum, and some with a great deal of programming experience already – is the level playing field which the functional family of languages (like Haskell) provide. Sometimes, those students with least prior programming experience perform the best.

I have tried to write a book which has no prerequisites – and with which any intelligent undergraduate ought to be able to cope, whilst trying to be concise enough that someone coming from another language might not be too annoyed by the tone.

One caveat: most things in life are small and elegant, or large and unwieldy. Haskell, as practised, is in the unusual position of being large and elegant. This may be the first Haskell book you read, but it will probably not be the last.

Special note to those who have already written programs

When I was a boy, our class was using a word processor for the first time. I wanted a title for my story, so I typed it on the first line and then, placing the cursor at the beginning, held down the space bar until the title was roughly in the middle. My friend taught me how to use the centring function, but it seemed more complicated to me, and I stuck with the familiar way – after all, it worked. Later on, of course, when I had more confidence and experience, I realized he had been right.

When starting a language which is fundamentally different from those you have seen before, it can be difficult to see the advantages, and to try to think of every concept in terms of the old language. I would urge you to consider the possibility that, at the moment, you might be the boy holding down the space bar.

Acknowledgments

The book was improved markedly by the comments of the technical reviewer, Stephen Dolan. Helpful additional review was provided by Michał Gajda, Jim Stuttard, Stuart Kurtz, Krystal Maughan, David Feuer, Gregory Popovitch, Steven Leiva, Charles Parker, and Dirk Markert. Any remaining errors are due solely to the author.

Getting Ready

This book is about teaching the computer to do new things by writing *computer programs*. Just as there are different languages for humans to speak to one another, there are different *programming languages* for humans to speak to computers.

We are going to be using a programming language called **Haskell**. A Haskell system might already be on your computer, or you may have to find it on the internet and install it yourself. We will be using the Glasgow Haskell system. You will know that you have it working when you see something like this:

```
GHCi, version 8.6.1: http://www.haskell.org/ghc/  :? for help
Prelude>
```

Please make sure the version number is at least 8. Haskell is waiting for us to type something. Try typing `1` `space` `+` `space` `2` followed by the `Enter` key. You should see this:

```
GHCi:
Prelude> 1 + 2
3
Prelude>
```

Haskell tells us the result of the calculation. You may use the left and right arrow keys on the keyboard to correct mistakes and the up and down arrow keys to look through a history of previous inputs. To leave Haskell, give the `:quit` command, again followed by `Enter`:

```
GHCi:
Prelude> :quit
Leaving GHCi.
```

You should find yourself back where you were before. If you make a mistake when typing, you can use the arrow keys on your keyboard to edit the text. To abandon typing, and ask Haskell to forget what you have already typed, enter `Ctrl-C` (hold down the `Ctrl` key and tap the `c` key). This will allow you to start again.

We are ready to begin.

Chapter 1

Starting Off

We will cover a fair amount of material in this chapter and its questions, since we will need a solid base on which to build. You should read this with a computer running Haskell in front of you.

Consider first the mathematical expression $1 + 2 \times 3$. What is the result? How did you work it out? We might show the process like this:

$$\begin{aligned} & 1 + 2 \times 3 \\ \Rightarrow & 1 + 6 \\ \Rightarrow & 7 \end{aligned}$$

How did we know to multiply 2 by 3 first, instead of adding 1 and 2? How did we know when to stop? Let us underline the part of the expression which is dealt with at each step:

$$\begin{aligned} & 1 + \underline{2 \times 3} \\ \Rightarrow & \underline{1 + 6} \\ \Rightarrow & 7 \end{aligned}$$

We chose which part of the expression to deal with each time using the familiar mathematical rules. We stopped when the expression could not be processed any further.

Computer programs in Haskell are just like these expressions. In order to give you an answer, the computer needs to know all the rules you know about how to process the expression correctly. In fact, $1 + 2 \times 3$ is a valid Haskell expression as well as a valid mathematical one, but we must write `*` instead of `×`, since there is no `×` key on the keyboard:

```
GHCi:
Prelude> 1 + 2 * 3
7
```

Here, `PreLude>` is Haskell prompting us to write an expression, and `1 + 2 * 3` is what we typed (the Enter key tells Haskell we have finished our expression). We'll see what `PreLude` means later. Haskell responds with the answer 7.

Let us look at our example expression some more. There are two *operators*: `+` and `×`. There are three *operands*: 1, 2, and 3. When we wrote the expression down, and when we typed it into Haskell, we put spaces between the operators and operands for readability. How does Haskell process it? First, the text we wrote must be split up into its basic parts: 1, +, 2, *, and 3. Haskell then looks at the order and sort of the

operators and operands, and decides how to parenthesize the expression: $(1 + (2 \times 3))$. Now, processing the expression just requires doing one step at a time, until there is nothing more which can be done:

$$\begin{aligned} & (1 + (2 \times 3)) \\ \Rightarrow & \underline{(1 + 6)} \\ \Rightarrow & 7 \end{aligned}$$

Haskell knows that $+$ refers not to 1 and 2 but to 1 and the result of 2×3 , and parenthesizes the expression appropriately. We say the \times operator has *higher precedence* than the $+$ operator. An *expression* is any valid Haskell program. To produce an answer, Haskell *evaluates* the expression, yielding a special sort of expression, a *value*. In our previous example, $1 + 2 \times 3$, $1 + 6$, and 7 were all expressions, but only 7 was a value. Here are some mathematical operators on numbers:

Operator	Description
$a + b$	addition
$a - b$	subtract b from a
$a * b$	multiplication

The $*$ operator has higher precedence than the $+$ and $-$ operators. For any operator \oplus above, the expression $a \oplus b \oplus c$ is equivalent to $(a \oplus b) \oplus c$ rather than $a \oplus (b \oplus c)$ (we say the operators are *left associative*). Negative numbers are written with $-$ before them, and if we use them next to an operator we may need parentheses too:

```
GHCi:
Prelude> 5 * (-2)
-10
```

Of course, there are many more things than just numbers. Sometimes, instead of numbers, we would like to talk about truth: either something is true or it is not. For this we use *boolean values*, named after the English mathematician George Boole (1815–1864) who pioneered their use. There are just two boolean things:

```
True
False
```

How can we use these? One way is to use one of the *comparison operators*, which are used for comparing values to one another:

```
GHCi:
Prelude> 99 > 100
False
Prelude> 4 + 3 + 2 + 1 == 10
True
```

Here are the comparison operators:

Operator	Description
$a == b$	true if a and b are equal
$a < b$	true if a is less than b
$a <= b$	true if a is less than or equal to b
$a > b$	true if a is more than b
$a >= b$	true if a is more than or equal to b
$a /= b$	true if a is not equal to b

Notice that if we try to use operators with things for which they are not intended, Haskell will not accept the program at all:

GHCi:

```
Prelude> 1 > True
```

```
<interactive>:2:1: error:
```

- No instance for (Num Bool) arising from the literal '1'
 - In the first argument of '(>)', namely '1'
- In the expression: 1 > True
In an equation for 'it': it = 1 > True

Do not expect to understand the details of this error message for the moment. We shall return to them later on. You can find more information about error messages in Haskell in the appendix “Coping with Errors” on page 197.

There are two operators for combining boolean values (for instance, those resulting from using the comparison operators). The expression $a \ \&\& \ b$ evaluates to `True` only if expressions a and b both evaluate to `True`. The expression $a \ || \ b$ evaluates to `True` if a evaluates to `True` or b evaluates to `True`, or both do. In each case, the expression a will be tested first – the second may not need to be tested at all. The `&&` operator (pronounced “and”) is of higher precedence than the `||` operator (pronounced “or”), so $a \ \&\& \ b \ || \ c$ is the same as $(a \ \&\& \ b) \ || \ c$.

We shall also be using *characters*, such as ‘a’ or ‘?’. We write these in single quotation marks:

GHCi:

```
Prelude> 'c'
'c'
```

So far we have looked only at operators like `+`, `==` and `&&` which look like familiar mathematical ones. But many constructs in programming languages look a little different. For example, to choose a course of evaluation based on some test, we use the **if ... then ... else** construct:

GHCi:

```
Prelude> if 100 > 99 then 0 else 1
0
```

The expression between **if** and **then** (in our example `100 > 99`) must evaluate to either `True` or `False`, and the expression to choose if true and the expression to choose if false must be the same sort of thing as one another – here they are both numbers. The whole expression will then evaluate to that sort of thing too, because either the **then** part or the **else** part is chosen to be the result of evaluating the whole expression:

```
if 100 > 99 then 0 else 1
```

Diagram illustrating the types of the components in the conditional expression:

- `100 > 99` is a **boolean**.
- `0` is a **number**.
- `1` is a **number**.
- The entire expression `if 100 > 99 then 0 else 1` is a **number**.

We have covered a lot in this chapter, but we need all these basic tools before we can write interesting programs. Make sure you work through the questions on paper, on the computer, or both, before moving on. Hints and answers are at the back of the book.

Questions

1. What sorts of thing do the following expressions represent and what do they evaluate to, and why?

```
17
1 + 2 * 3 + 4
400 > 200
1 /= 1
True || False
True && False
if True then False else True
'%'
```

2. These expressions are not valid Haskell. In each case, why? Can you correct them?

```
1 + -1
A == a
false || true
if 'A' > 'a' then True
'a' + 'b'
```

3. A programmer writes $1+2 * 3+4$. What does this evaluate to? What advice would you give them?
4. Haskell has a remainder operator, which finds the remainder of dividing one number by another. It is written ``rem``. Consider the evaluations of the expressions $1 + 2 \text{ `rem` } 3$, $(1 + 2) \text{ `rem` } 3$, and $1 + (2 \text{ `rem` } 3)$. What can you conclude about the `+` and ``rem`` operators?
5. Why not just use, for example, the number 0 to represent falsity and the number 1 for truth? Why have a separate `True` and `False` at all?
6. What is the effect of the comparison operators like `<` and `>` on alphabetic characters? For example, what does `'p' < 'q'` evaluate to? What about `'A' < 'a'`? What is the effect of the comparison operators on the booleans `True` and `False`?

So Far

1 Numbers ... -3 -2 -1 0 1 2 3... Booleans True and False. Characters like 'X' and '!'.
1

Mathematical operators + - * which take two numbers and give another.

Operators == < <= > >= /= which compare two values and evaluate to either True or False.

The “conditional” construct **if** *expression1* **then** *expression2* **else** *expression3*, where *expression1* evaluates to something boolean and *expression2* and *expression3* evaluate to the same sort of thing as one another.

The boolean operators && and || which allow us to build compound boolean expressions. The remainder operator `rem`.

Chapter 2

Names and Functions

So far we have built only tiny toy programs. To build bigger ones, we need to be able to name things so as to refer to them later. We also need to write expressions whose result depends upon one or more other things. Before, if we wished to use a sub-expression twice or more in a single expression, we had to type it multiple times:

```
GHCi:
Prelude> 200 * 200 * 200
8000000
```

Instead, we can define our own name to stand for the expression, and then use the name as we please:

```
GHCi:
Prelude> x = 200
Prelude> x * x * x
8000000
```

To write this all in a single expression, we can use the **let** ... = ... **in** ... construct:

```
GHCi:
Prelude> let x = 200 in x * x * x
8000000
Prelude> let a = 500 in (let b = a * a in a + b)
250500
Prelude> let a = 500 in let b = a * a in a + b
250500
```

We can use the special name **it** for the value resulting from the most recently evaluated expression, which can be useful when we forget to name something whilst experimenting:

```
GHCi:
Prelude> 200 * 200
40000
Prelude> it * 200
8000000
```

```
Prelude> it * 200
1600000000
```

The `it` name is not a part of the Haskell language – it is just a shortcut to make experimenting easier.

In Chapter 1, we talked about how values could be different “sorts of things” such as numbers and booleans and characters, but in fact Haskell knows about this idea – these “sorts of things” are called *types*, and every value and indeed every expression has a type. For example, the type of `False` is **Bool**. We can ask Haskell to tell us the type of a value or expression by using the `:type` command:

```
GHCi:
Prelude> :type False
False :: Bool
Prelude> :type False && True
False && True :: Bool
Prelude> :type 'x'
'x' :: Char
```

Note that commands like `:type` are not part of the Haskell language, and so cannot form part of expressions. We can read `False && True :: Bool` as “The expression `False && True` has type **Bool**”. An expression always has same type as the value it will evaluate to. There is a further complication, which we shall only explain in detail later, but which we must confront on its surface now:

```
GHCi:
Prelude> :type 50
50 :: Num a => a
```

We might expect the type of `50` to be something like **Number** but it is the rather more cryptic **Num a => a**. You can read this as “if `a` is one of the types of number, then `50` can have type `a`”. In Haskell, integers and other numbers are sorts of **Num**. For now, we will not worry too much about types, just making sure we can read them without being scared. The purpose is to allow, for example, the expression `50` to do the job of an integer and a real number, as and when required. For example:

```
GHCi:
Prelude> 50 + 0.6
50.6
Prelude> 50 + 6
56
```

In the first line, `50` is playing the part of a real number, not an integer, because we are adding it to another real number. In the second, it plays the part of an integer, which is why the result is `56` rather than `56.0`.

The letter `a` in the type is, of course, arbitrary. The types **Num a => a** and **Num b => b** and **Num frank => frank** are interchangeable. In fact, Haskell does not always use `a` first. On the author’s machine our example reads:


```
GHCi:
Prelude> :type 50
50 :: Num p => p
```

However, we shall always use the letters `a`, `b` etc. Let us move on now to consider *functions*, whose value depends upon some input (we call this input an *argument* – we will be using the word “input” later in the book to mean something different):


```
GHCi:
Prelude> cube x = x * x * x
Prelude> cube 200
8000000
```

We chose `cube` for the name of the function and `x` for the name of its argument. If we ask for its type, Haskell will reply by telling us that its type is `Num a => a -> a`. This means it is a function which takes a number as its argument, and, when given that argument, evaluates to the same sort of number. To use the function, we just write its name followed by a suitable argument. In our example, we calculated 200^3 by giving the `cube` function `200` as its argument.

The `cube` function has type `Num a => a -> a`, we gave it a number `200`, and so the result is another number. Thus, the type of the expression `cube 200` is `Num a => a` (remember that the type of any expression is the type of the thing it will evaluate to, and `cube 200` evaluates to `8000000`, a number of type `Num a => a`). In diagram form:



It might be easier to see what is going on if we imagine missing out the part to the left of the \Rightarrow symbol in each type:



If we try an argument of the wrong type, the program will be rejected:

```
GHCi:
Prelude>cube False

<interactive> error:
• No instance for (Num Bool) arising from a use of 'cube'
• In the expression: cube False
  In an equation for 'it': it = cube False
```

You can learn more about how to understand such messages in “Coping with Errors” on page 197. Here is a function which determines if a number is negative:

```
GHCi:
Prelude> neg x = if x < 0 then True else False
Prelude> neg (-30)
True
```

But, of course, this is equivalent to just writing

```
GHCi:
Prelude> neg x = x < 0
Prelude> neg (-30)
True
```

because $x < 0$ will evaluate to the appropriate boolean value on its own – True if $x < 0$ and False otherwise. What is the type of `neg`?

```
GHCi:
Prelude> neg x = x < 0
Prelude> :type neg
neg :: (Num a, Ord a) => a -> Bool
```

We can read this as “The argument to our function can have type `a` if `a` is an one of the class of types **Num** and also one of the class of types **Ord**. The result of the function is of type **Bool**”. The class of types, or *typeclass* **Ord** is for things which can be ordered – in other words, ones on which we can use `<` and other comparison operators. A type which is one of a class of types is called an *instance* of that class. Here is another function, this time of type **Char** \rightarrow **Bool**. It determines if a given character is a vowel or not:

```
GHCi:
Prelude> isVowel c = c == 'a' || c == 'e' || c == 'i' || c == 'o' || c == 'u'
Prelude> :type isVowel
isVowel :: Char -> Bool
Prelude> isVowel 'x'
False
```

The line is getting a little long. We can type a function (or any expression) over multiple lines by preceding it with `{` and following it with `}`, pressing the Enter key between lines as usual. Haskell knows that we are finished when we type `}` followed by the Enter key. Notice also that we press space a few times so that the second line appeared a little to the right of the first. This is known as *indentation*.

```
GHCi:
Prelude> {:
Prelude| isVowel c =
Prelude|   c == 'a' || c == 'e' || c == 'i' || c == 'o' || c == 'u'
Prelude| :}
```

The start of the second line must be to the right of the name of the function: Haskell is particular about this. There can be more than one argument to a function. For example, here is a function which checks if two numbers add up to ten:

```
GHCi:
Prelude> {:
Prelude| addToTen a b =
Prelude|   a + b == 10
Prelude| :}
Prelude> addToTen 6 4
True
```

We use the function in the same way as before, but writing two numbers this time, one for each argument the function expects. The type is **(Eq a, Num a)** \Rightarrow `a` \rightarrow `a` \rightarrow **Bool** because the arguments are both numbers, and both capable of being tested for equality (hence **Eq**) and the result is a boolean.

```
GHCi:
Prelude> {:
```

```

Prelude| addToTen a b =
Prelude|   a + b == 10
Prelude| :}
Prelude> :type addToTen
addToTen :: (Eq a, Num a) => a -> a -> Bool

```

Note that **Eq** and **Ord** are different. Not everything which can be tested for equality with `==` can be put in order with `<` and similar operators.

A *recursive function* is one which uses itself. Consider calculating the factorial of a given number – the factorial of 4 (written 4! in mathematics), for example, is $4 \times 3 \times 2 \times 1$. Here is a recursive function to calculate the factorial. Note that it uses itself in its own definition.

```

GHCi:
Prelude> :{
Prelude| factorial n =
Prelude|   if n == 1 then 1 else n * factorial (n - 1)
Prelude| :}
Prelude> :type factorial
factorial :: (Eq a, Num a) => a -> a
Prelude> factorial 4
24

```

How does the evaluation of `factorial 4` proceed?

$$\begin{aligned}
 & \text{factorial } 4 \\
 \implies & 4 * \text{factorial } (4 - 1) \\
 \implies & 4 * (3 * \text{factorial } (3 - 1)) \\
 \implies & 4 * (3 * (2 * \text{factorial } (2 - 1))) \\
 \implies & 4 * (3 * (2 * 1)) \\
 \implies & 4 * (3 * 2) \\
 \implies & 4 * 6 \\
 \implies & 24
 \end{aligned}$$

For the first three steps, the **else** part of the **if** (or *conditional expression*) is chosen, because the argument `a` is greater than one. When the argument is equal to 1, we do not use `factorial` again, but just evaluate to 1. The expression built up of all the multiplications is then evaluated until a value is reached: this is the result of the whole evaluation. It is sometimes possible for a recursive function never to finish – what if we try to evaluate `factorial (-1)`?

$$\begin{aligned}
 & \text{factorial } (-1) \\
 \implies & -1 * \text{factorial } (-1 - 1) \\
 \implies & -1 * (-2 * \text{factorial } (-2 - 1)) \\
 \implies & -1 * (-2 * (-3 * \text{factorial } (-3 - 1))) \\
 & \vdots \\
 & \vdots
 \end{aligned}$$

The expression keeps expanding, and the recursion keeps going. You can interrupt this infinitely-long process by typing `Ctrl-C` on your keyboard (it may take a little while to work):

```
GHCi:
Prelude> factorial (-1)
^CInterrupted.
```

This is an example of a problem Haskell cannot find by merely looking at the program text – it can only be uncovered during the process of evaluation. Later in the book, we will see how to prevent people who are using our functions from making such mistakes.

One of the oldest methods for solving a problem (or *algorithm*) still in common use is Euclid's algorithm for calculating the greatest common divisor of two numbers (that is, given two positive integers a and b , finding the biggest positive integer c such that neither a/c nor b/c have a remainder). Euclid was a Greek mathematician who lived about three centuries before Christ. Euclid's algorithm is simple to write as a function with two arguments:

```
GHCi:
Prelude> :{
Prelude| gcd' a b =
Prelude|   if b == 0 then a else gcd' b (rem a b)
Prelude| :}
Prelude> gcd' 64000 3456
128
```

The function built-in function `rem` finds the remainder of dividing a by b . If we like, we can surround the function `rem` in backticks as ``rem`` (we have already seen this in Question 4 of the previous chapter). This allows us to put its two arguments either side, making it an operator like `+` and `||`:

```
GHCi:
Prelude> :{
Prelude| gcd' a b =
Prelude|   if b == 0 then a else gcd' b (a `rem` b)
Prelude| :}
```

Here is the evaluation:

```

gcd' 64000 3456
⇒ gcd' 3456 (64000 `rem` 3456)
⇒ gcd' 1792 (3456 `rem` 1792)
⇒ gcd' 1664 (1792 `rem` 1664)
⇒ gcd' 128 (1664 `rem` 128)
⇒ 128
```

Why did we call our function `gcd'` instead of `gcd`? Because Haskell has a built in function `gcd`, and we should not reuse the name. Later on, when we load our programs from files, Haskell will in fact not let us reuse the name. This is another way in which Haskell is being rather careful, to prevent us being tripped up when writing larger programs.

Finally, here is a simple function on boolean values. In the previous chapter, we looked at the `&&` and `||` operators which are built in to Haskell. The other important boolean operator is the `not` function, which returns the boolean complement (opposite) of its argument – `True` if the argument is `False`, and

vice versa. This is again built in, but it is easy enough to define ourselves, as a function of type **Bool** → **Bool**.

```
GHCi:
Prelude> :{
Prelude| not' x =
Prelude|   if x then False else True
Prelude| :}
Prelude> :type not'
not' :: Bool -> Bool
Prelude> not' True
False
```

Almost every program we write will involve functions such as these, and many larger ones too. In fact, languages like Haskell are often called *functional languages*.

A more formal look at types

Most readers will wish to skip this section, and the extra questions which relate to it, and not worry too much about types, coming back to it after a few more chapters have been worked through. However, for those who refuse to take things on trust without understanding them, it is perhaps best to tackle it now.

Every expression in Haskell has a *type*, which indicates what sort of thing it will eventually evaluate to. Simple types include **Bool** and **Char**. For example, the expression `False || True` has the type **Bool** because, when evaluated, it will result in a boolean value. So a type represents a collection of values. For example, the **Bool** type has two values: `True` and `False`, but the **Char** type has many more.

The purpose of types is to make sure that no part of the program receives something it was not expecting, and for which it cannot sensibly do anything. For example, the addition operator `+` being asked to add a number to a boolean. This avoids, at a stroke, a huge class of possible program misbehaviours, or bugs. Haskell can do this automatically, by working out the types of everything in the program and making sure they all fit together, and that no function can possibly receive an argument of the wrong type. This is called *type inference*, because the types are inferred (worked out) by Haskell.

When we ask Haskell what the type of `42` is, we get the surprising answer **Num a** \Rightarrow `a`, rather than something simple like **Number**. The letters `a`, `b`, `c`... are *type variables* standing for types. A *typeclass* like **Num** is a collection of types. So, a typeclass is a collection of types, each of which is a collection of values. A type with a \Rightarrow symbol in it has a left-hand and right-hand part. The left-hand part says which typeclasses one or more of the type variables on the right-hand side must belong to. So if `42` has the type **Num a** \Rightarrow `a` we may say "Given that the type variable `a` represents a type which is an instance of the typeclass **Num**, `42` can have type `a`". Remember our example where a number was used as both an integer and a real number, even though it was written the same. Of course, many types do not have a \Rightarrow symbol, which means either they are very specific, like **Bool**, or very generic, like `a`, which represents any type at all.

We have also introduced functions, which have types like $a \rightarrow b$. For example, if `a` is **Char** and `b` is **Bool**, we may have the type **Char** \rightarrow **Bool**. Of course, functions may have a left-hand part too. For example, the function which adds two numbers may have the type **Num a** \Rightarrow `a` \rightarrow `a` \rightarrow `a`. That is to say, the function will add any two things both of a type `a` which is an instance of the typeclass **Num**, and the result is a number of the same type.

So this is what is rather confusing to us about the type **Num a** \Rightarrow `a`: it is actually rather harder to understand for the beginner than the function types in the previous paragraph, and yet it represents what we expect to be a simple concept: the number. All will be explained in Chapter 12.

We can have more than one constraint on a single type variable, or constraints on multiple type variables. They are each called *class constraints*, and the whole left hand part is sometimes called the *context*. For example, the type **(Num a, Eq b)** \Rightarrow `a` \rightarrow `b` \rightarrow `a` is the type of a function of two arguments, the first of which must be of some type from typeclass **Num** and the second of some type from typeclass **Eq**.

Further complicating matters, sometimes every type of a certain typeclass is by definition also part of one or more other ones. In the case of the typeclasses we have seen so far, every type in the typeclass **Ord** is also in the typeclass **Eq**. What this means is that if we list the constraint **Ord** we need not also list **Eq**.

Questions

1. Write a function which multiplies a given number by ten. What is its type?
2. Write a function which returns `True` if both of its arguments are non-zero, and `False` otherwise. What is the type of your function?
3. Write a recursive function `sum'` which, given a number n , calculates the sum $1 + 2 + 3 + \dots + n$. What is its type?
4. Write a function `power x n` which raises x to the power n . Give its type.
5. Write a function `isConsonant` which, given a lower-case character in the range `'a'... 'z'`, determines if it is a consonant.
6. What is the result of the expression `let x = 1 in let x = 2 in x + x`?
7. Can you suggest a way of preventing the non-termination of the `factorial` function in the case of a zero or negative argument?

For those who are confident with types and typeclasses. To be attempted in the first instance without the computer.

8. Here are some expressions and function definitions and some types. Pair them up.

<code>1</code>	<code>Ord a ⇒ a → a → Bool</code>
<code>1 + 2</code>	<code>(Ord a, Num a) ⇒ a → a → Bool</code>
<code>f x y = x < y</code>	<code>Num a ⇒ a → b → c → a</code>
<code>g x y = x < y + 2</code>	<code>Num a ⇒ a</code>
<code>h x y = 0</code>	<code>Num a ⇒ b → c → a</code>
<code>i x y z = x + 10</code>	<code>Num a ⇒ a</code>

9. Infer (work out) types for the following expressions or function definitions.

`46 * 10` `2 > 1`
`f x = x + x` `g x y z = x + 1 < y`
`i a b c = b`

10. Why are the following expressions or function definitions not accepted by Haskell?

`True + False`
`6 + '6'`
`f x y z = (x < y) < (z + 1)`

11. Which of the following types are equivalent to one another and which are different? Which are not valid types?

`Num a ⇒ b` `Num t1 ⇒ t1`
`Num b ⇒ b → a` `Num a ⇒ a → b`
`(Ord a, Num a) ⇒ a → a` `Num a ⇒ a → a`
`(Num a, Ord a) ⇒ a → a` `Num a ⇒ a`

12. These types are correct, but have some constraints which are not required. Remove them.

`(Eq a, Ord a) ⇒ a → b → a`
`(Ord a, Eq a, Eq b) ⇒ b → b → a`

So Far

1 Numbers ... -3 -2 -1 0 1 2 3... Booleans True and False. Characters like 'X' and '!'.
Mathematical operators + - * which take two numbers and give another.

Operators == < <= > >= /= which compare two values and evaluate to either True or False.

Operators == < <= > >= /= which compare two values and evaluate to either True or False.

The “conditional” construct **if** *expression1* **then** *expression2* **else** *expression3*, where *expression1* evaluates to something boolean and *expression2* and *expression3* evaluate to the same sort of thing as one another.

The boolean operators && and || which allow us to build compound boolean expressions. The remainder operator `rem`.

2 Assigning a name to an expression using the *name = expression* construct. Building compound expressions using **let** *name1 = expression1* **in** **let** *name2 = expression2* **in** ...

Functions, introduced by *name argument1 argument2 ... = expression*. These have type $a \rightarrow b$, $a \rightarrow b \rightarrow c$ etc. for some types a, b, c etc. Recursive functions. Turning a two-argument function into an operator with backticks like `rem`.

The types **Bool** and **Char**. The typeclasses **Num**, **Ord**, and **Eq**. A function from values of type a to type b with a in typeclass **Eq** and b in typeclass **Ord** would have type **(Eq a, Ord b) \Rightarrow $a \rightarrow b$** .

The special value `it`. The command `:type` and the use of Ctrl-C to interrupt a computation.

Using Scripts

From now on, instead of showing the actual Haskell session...

```
GHCi:
Prelude> :{
Prelude| factorial n =
Prelude|   if n == 1 then 1 else n * factorial (n - 1)
Prelude| :}
```

... we will usually just show the program in a box, together with its type:

```
factorial :: (Eq a, Num a) => a -> a

factorial n =
  if n == 1 then 1 else n * factorial (n - 1)
```

In fact, this is just how Haskell programs are normally written, in a text file with the `.hs` (haskell script) extension, rather than typed directly into Haskell. We can include the type in our `.hs` file, or leave it out and let Haskell infer it.

We can use the `:load` and `:reload` commands to access the program from Haskell. Assuming we have a file `Script.hs` which looks like the contents of the box above, we can use it like this:

```
GHCi:
Prelude> :load Script.hs
[1 of 1] Compiling Main                ( Script.hs, interpreted )
Ok, one module loaded.
*Main> factorial 24
620448401733239439360000
```

When we have made a change to the file `Script.hs` in our text editor (and saved the file), we can reload the new one:

```
GHCi:
*Main> :reload
[1 of 1] Compiling Main                ( Script.hs, interpreted )
Ok, one module loaded.
```


Chapter 3

Case by Case

In the previous chapter, we used the conditional expression `if ... then ... else` to define functions whose results depend on their arguments. For some of them we had to nest the conditional expressions one inside another. Programs like this are not terribly easy to read, and expand quickly in size and complexity as the number of cases increases.

Haskell has a nicer way of expressing choices – *pattern matching*. For example, recall our factorial function:

```
factorial :: (Eq a, Num a) => a -> a
factorial n =
  if n == 1 then 1 else n * factorial (n - 1)
```

We can rewrite this using pattern matching:

```
factorial :: (Eq a, Num a) => a -> a
factorial 1 = 1
factorial n = n * factorial (n - 1)
```

We can read this as “See if the argument matches the pattern `1`. If it does, just return `1`. If not, see if it matches the pattern `n`. If it does, the result is `n * factorial (n - 1)`.” Patterns like `n` are special – they match anything and give it a name. Remember our `isVowel` function from the previous chapter?

```
isVowel :: Char -> Bool
isVowel c =
  c == 'a' || c == 'e' || c == 'i' || c == 'o' || c == 'u'
```


Here is how to write it using pattern matching:

```
isVowel :: Char → Bool

isVowel 'a' = True
isVowel 'e' = True
isVowel 'i' = True
isVowel 'o' = True
isVowel 'u' = True
isVowel _  = False
```

The special pattern `_` matches anything. If we miss out one or more cases – for example leaving out the final case, Haskell can warn us:

```
<interactive> warning: [-Wincomplete-patterns]
  Pattern match(es) are non-exhaustive
  In an equation for 'isVowel':
    Patterns not matched:
      p where p is not one of {'u', 'o', 'i', 'e', 'a'}
```

To enable this behaviour, you must start Haskell by writing `ghci -Wincomplete-patterns` instead of just `ghci`. Writing `ghci -Wall` enables all warnings. Haskell does not reject the program outright, because there may be legitimate reasons to miss cases out, but for now we will make sure all our pattern matches are exhaustive. Finally, let us rewrite Euclid's Algorithm from the previous chapter:

```
gcd' :: Integral a ⇒ a → a → a

gcd' a b =
  if b == 0 then a else gcd' b (a `rem` b)
```

Now in pattern matching style:

```
gcd' :: Integral a ⇒ a → a → a

gcd' a 0 = a
gcd' a b = gcd' b (a `rem` b)
```

We use pattern matching whenever it is easier to read and understand than `if ... then ... else` expressions.

What about this `Integral` typeclass? We did not try `:type` on the `gcd'` function in the last chapter, so we did not see this. A type of number which is an `Integral` has an additional property to one which is merely a `Num`, which is that whole-number division and remainder operations work on it. Since every thing which is an `Integral` is also a `Num`, we do not see `(Num a, Integral a)`, but just `Integral a` in the type.

Sometimes we need more than just a pattern to decide which case to choose in a pattern match. For example, in `gcd'` above, we only needed to distinguish between `0` and any other value of `b`. Consider, though, the function to determine the sign of a number, producing `-1` for all numbers less than zero, `0` for just the number zero, and `1` for all numbers above zero:

```

sign :: (Ord a, Num a, Num b) => a -> b

sign x =
  if x < 0 then -1 else if x > 0 then 1 else 0

```

We cannot rewrite this using a pattern match with three cases. Haskell has a facility called *guarded equations* to help us (each line in our pattern matched functions can also be called an equation). A *guard* is an extra check to decide if a case of a pattern match is taken based upon some condition, for example $x < 0$. Here is our sign function written using guarded equations:

```

sign :: (Ord a, Num a, Num b) => a -> b

sign x | x < 0      = -1
       | x > 0      = 1
       | otherwise = 0

```

There is no need to line up the equals signs vertically, but we do so to make it easier to read. The cases are considered one after another, just like when using pattern matching, and the first case which matches the guard is taken. The **otherwise** guard matches anything, so it comes last. We use an **otherwise** case to make sure every possibility is handled. We can read the `|` symbol as “when”. A function can be defined using multiple equations, each of which has multiple guarded parts.

The layout rule

We have mentioned indentation, noting that Haskell is particular about it. Indeed, programs will not be accepted unless they are properly indented:

```

GHCi:
Prelude> {:
Prelude| sign x =
Prelude| if x < 0 then -1 else if x > 0 then 1 else 0
Prelude| :}

```

```

<interactive> error:
  parse error (possibly incorrect indentation or mismatched brackets)

```

Haskell is telling us that it cannot work out what we mean. Since the **if ... then ... else ...** expression is part of the `sign` function, it must be indented further than the beginning of the whole `sign` expression. This applies at all times – even when the start of the whole expression is itself indented. In the case of **if ... then ... else ...** itself, it is in fact permitted not to indent:

```

GHCi:
Prelude> {:
Prelude| if 1 < 0
Prelude| then 2
Prelude| else 3
Prelude| :}
3

```

However, we shall often do so, when it is easier to read:

```
GHCi:
Prelude> :{
Prelude| if 1 < 0
Prelude|   then 2
Prelude|   else 3
Prelude| :}
3
```

Consider again our sign function:

```
sign :: (Ord a, Num a, Num b) => a -> b

sign x | x < 0    = -1
      | x > 0    = 1
      | otherwise = 0
```

We have already mentioned that lining up the equals signs is not necessary. However, we must always indent the cases. Here, we start the cases on the next line:

```
Prelude| sign x
Prelude|   | x < 0 = -1
Prelude|   | x > 0 = 1
Prelude|   | otherwise = 0
Prelude| :}
```

The layout rule is not complicated, but it can be frustrating to the beginner, especially when the error message is not clear.

Questions

1. Rewrite the `not'` function from the previous chapter in pattern matching style.
2. Use pattern matching to write a recursive function `sumMatch` which, given a positive integer n , returns the sum of all the integers from 1 to n .
3. Use pattern matching to write a function which, given two numbers x and n , computes x^n .
4. For each of the previous three questions, comment on whether you think it is easier to read the function with or without pattern matching. How might you expect this to change if the functions were much larger? Write each using guarded equations too.
5. Use guarded equations to write a function which categorises characters into three kinds: kind 0 for the lowercase letters `a...z`, kind 1 for the uppercase letters `a...z`, and kind 2 for everything else.
6. Experiment with the layout of the function definitions in this and the previous chapter. Which kinds of layout are allowed by Haskell? Which of the allowed layouts are aesthetically pleasing, or easy to read? Do any of your layouts make the program harder to change?

So Far

1 Numbers ... -3 -2 -1 0 1 2 3... Booleans `True` and `False`. Characters like `'X'` and `'!'`.

Mathematical operators `+` `-` `*` which take two numbers and give another.

Operators `==` `<` `<=` `>` `>=` `/=` which compare two values and evaluate to either `True` or `False`.

The “conditional” construct `if expression1 then expression2 else expression3`, where `expression1` evaluates to something boolean and `expression2` and `expression3` evaluate to the same sort of thing as one another.

The boolean operators `&&` and `||` which allow us to build compound boolean expressions. The remainder operator ``rem``.

2 Assigning a name to an expression using the `name = expression` construct. Building compound expressions using `let name1 = expression1 in let name2 = expression2 in ...`

Functions, introduced by `name argument1 argument2 ... = expression`. These have type `a → b`, `a → b → c` etc. for some types `a`, `b`, `c` etc. Recursive functions. Turning a two-argument function into an operator with backticks like ``rem``.

The types `Bool` and `Char`. The typeclasses `Num`, `Ord`, and `Eq`. A function from values of type `a` to type `b` with `a` in typeclass `Eq` and `b` in typeclass `Ord` would have type `(Eq a, Ord b) ⇒ a → b`.

The special value `it`. The command `:type` and the use of `Ctrl-C` to interrupt a computation.

3 Matching patterns using `f pattern1 = expression1 ↔ f pattern2 = expression2` etc. ... The expressions `expression1`, `expression2` etc. must have the same type as one another. Writing functions using guarded equations like `f x | guard = expression ↔ | guard2 = expression2 | otherwise ...`. The typeclass `Integral`.

Chapter 4

Making Lists

A *list* is a collection of elements. Here is a list of three numbers:

```
[1, 2, 3]
```

We write a list between square brackets [and], separating the elements with commas. The list above has type `Num a => [a]`, because it is a list of elements each of type `Num a => a`. All elements of the list must have the same type. The elements in the list are ordered (in other words, `[1, 2, 3]` and `[2, 3, 1]` are not the same list).

The first element is called the *head*, and the rest are collectively called the *tail*. In our example, the head is the number 1 and the tail is the list `[2, 3]`. So you can see that the tail has the same type as the whole list. Here is a list with no elements (called “the empty list” or sometimes “the nil list”):

```
[]
```

It has neither a head nor a tail. Here is a list with just a single element:

```
[5]
```

Its head is the number 5 and its tail is the empty list `[]`. So every non-empty list has both a head and a tail. Lists may contain elements of any type: numbers, booleans, functions, even other lists. For example, here is a list containing elements of type `Bool`:

```
[False, True, False] :: [Bool]
```

Haskell defines two operators for lists. The `:` operator (pronounced “cons”) is used to add a single element to the front of an existing list:

```
False : [True, False]
=      [False, True, False]
```

The `cons` operation is completed in a constant amount of time, regardless of the length of the list. The `++` operator (pronounced “append” or “concatenate”) is used to combine two lists together:

```
[1, 2] ++ [3, 4, 5]
=>      [1, 2, 3, 4, 5]
```


This takes time proportional to the length of the list on the left hand side of the ++ operator (that is, a list of length 100 will take roughly twice as long as one of length 50). We will see why soon.

Now, how do we write functions using lists? We can use pattern matching as usual, with some new types of pattern. For example, here is a function which tells us if a list is empty:

<pre>isNil :: [a] → Bool isNil [] = True isNil _ = False</pre>	<p style="text-align: right;"><i>the list is empty</i> <i>it has at least one element</i></p>
--	---

The argument has type [a] because this function does not inspect the individual elements of the list, it just checks if the list is empty. And so, this function can operate over any type of list. Functions like this are known as *polymorphic*. We can also use : in our patterns, this time using it to deconstruct rather than construct the list:

<pre>length' :: Num b ⇒ [a] → b length' [] = 0 length' (x:xs) = 1 + length' xs</pre>	<p style="text-align: right;"><i>the list has zero elements (the "base case")</i> <i>x is the head, xs the tail</i></p>
--	---

The traditional name xs for the tail is pronounced *exes*. If two types are represented by the same letter they must have the same type. If they are not, they may have the same type, but do not have to. For example, in length', type a might happen to be in type class Num also, but it does not have to be. Here is how the evaluation might proceed:

```
length' [5, 5, 5]
⇒ 1 + length' [5, 5]
⇒ 1 + (1 + length' [5])
⇒ 1 + (1 + (1 + length' []))      base case
⇒ 1 + (1 + (1 + 0))
⇒* 3                               (⇒* means we are not showing all the steps)
```

This works by recursion over the list, then addition of all the resultant 1s. It takes time proportional to the length of the list. Can you see why? Since x is not used in the expression 1 + length' xs, this function is also polymorphic. Indeed we can replace x in the pattern with the special pattern _ since there is no use giving a name to something we are not going to refer to:

<pre>length' :: Num b ⇒ [a] → b length' [] = 0 length' (_:xs) = 1 + length' xs</pre>
--

A very similar function can be used to add a list of numbers:

```

sumElts :: Num a => [a] -> a

sumElts [] = 0
sumElts (x:xs) = x + sumElts xs

```

the sum of no elements is zero
otherwise, add the head to the sum of the tail

However, since we are actually using the individual list elements (by adding them up), this function is not as general – it operates over lists of type `Num a => [a]` only. Functions can, of course, return lists too. Here is a function to return the list consisting of the first, third, fifth and so on elements in a list:

```

oddElements :: [a] -> [a]

oddElements [] = []
oddElements [x] = [x]
oddElements (x:_:xs) = x : oddElements xs

```

the list has zero elements
the list has one element
the list has more than one element

Consider the evaluation of `oddElements [2, 4, 2, 4, 2]`:

```

oddElements [2, 4, 2, 4, 2]
=> 2 : oddElements [2, 4, 2]
=> 2 : 2 : oddElements [2]
=> 2 : 2 : [2]
=>* [2, 2, 2]

```

You might notice that the first two cases in the pattern match return exactly their argument. By reversing the order, we can reduce this function to just two cases:

```

oddElements :: [a] -> [a]

oddElements (x:_:xs) = x : oddElements xs
oddElements l = l

```

there is something to skip over
there is nothing to skip over

We have seen how to use the `++` (append) operator to concatenate two lists:

```

[1, 2] ++ [3, 4, 5]
=> [1, 2, 3, 4, 5]

```

How might we implement `++` ourselves, if it were not provided? Consider a function `append xs ys`. If the first list is the empty list, the answer is simply `ys`. But what if the first list is not empty? Then it has a head `x` and a tail `xs`. So we can start our result list with the head, and the rest of the result is just `append xs ys`.

```

append :: [a] -> [a] -> [a]

append [] ys = ys
append (x:xs) ys = x : append xs ys

```

Consider the evaluation of `append [1, 2, 3] [4, 5, 6]`:

```

      append [1, 2, 3] [4, 5, 6]
⇒ 1 : append [2, 3] [4, 5, 6]
⇒ 1 : 2 : append [3] [4, 5, 6]
⇒ 1 : 2 : 3 : append [] [4, 5, 6]
⇒ 1 : 2 : 3 : [4, 5, 6]
⇒* [1, 2, 3, 4, 5, 6]

```

This takes time proportional to the length of the first list – the second list need not be processed at all. What about reversing a list? For example, we want `reverse' [1, 2, 3, 4]` to evaluate to `[4, 3, 2, 1]`. One simple way is to reverse the tail of the list, and append the list just containing the head to the end of it:

```

reverse' :: [a] → [a]

reverse' [] = []
reverse' (x:xs) = reverse' xs ++ [x]

```

Here is how the evaluation proceeds:

```

      reverse' [1, 2, 3, 4]
⇒ reverse' [2, 3, 4] ++ [1]
⇒ reverse' [3, 4] ++ [2] ++ [1]
⇒ reverse' [4] ++ [3] ++ [2] ++ [1]
⇒ reverse' [] ++ [4] ++ [3] ++ [2] ++ [1]
⇒ [ ] ++ [4] ++ [3] ++ [2] ++ [1]
⇒* [4, 3, 2, 1]

```

This is a simple definition, but not very efficient – can you see why?

Two more useful functions for processing lists are `take'` and `drop'` which, given a number and a list, either take or drop that many elements from the list:

```

take' :: (Eq a, Num a) ⇒ a → [b] → [b]
drop' :: (Eq a, Num a) ⇒ a → [b] → [b]

take' 0 _ = []
take' n (x:xs) = x : take' (n - 1) xs

drop' 0 l = l
drop' n (_:xs) = drop' (n - 1) xs

```

For example, here is the evaluation for `take' 2 [2, 4, 6, 8, 10]`:

```

                take' 2 [2, 4, 6, 8, 10]
    ⇒          2 : take' 1 [4, 6, 8, 10]
    ⇒          2 : 4 : take' 0 [6, 8, 10]
    ⇒          2 : 4 : []
    ⇒*        [2, 4]

```

And for `drop' 2 [2, 4, 6, 8, 10]`:

```

                drop' 2 [2, 4, 6, 8, 10]
    ⇒          drop' 1 [4, 6, 8, 10]
    ⇒          drop' 0 [6, 8, 10]
    ⇒          [6, 8, 10]

```

Note that these functions contain incomplete pattern matches. The function fails if the arguments are not sensible – that is, when we are asked to take or drop more elements than are in the argument list. Later on, we will see how to deal with that problem. Note also that for any sensible value of n , including zero, `take' n l` and `drop' n l` split the list into two parts with no gap. So `drop'` and `take'` often appear in pairs.

Lists can contain anything, so long as all elements are of the same type. So, of course, a list can contain lists. Here is a list of lists of numbers:

```
[[1], [2, 3], [4, 5, 6]] :: Num a => [[a]]
```

Each element of this list is of type `Num a => [a]`. Within values of this type, it is important to distinguish the list of lists containing no elements

```
[] :: [a]
```

from the list of lists containing one element which is the empty list:

```
[[[]]] :: [[a]]
```

Other ways to build lists

Haskell provides a convenient shorthand for building useful lists of numbers by giving a starting point, and ending point, and perhaps one other. For example, we can write `[1 .. 10]` to define the list `[1, 2, 3, 4, 5, 6, 7, 8, 9, 10]`:

```

GHCi:
Prelude> [1 .. 10]
[1,2,3,4,5,6,7,9,10]
Prelude> [1, 3 .. 10]
[1,3,5,7,9]

```

We can use what is called a *list comprehension* to select and process elements from one list into another. For example, to calculate the list of squares of the numbers from 1 to 10:

```
GHCi:
Prelude> nums = [1 .. 10]
Prelude> [x * x | x <- nums]
[1,4,9,16,25,36,49,64,81,100]
```

We can read this as “Each element of the list will be calculated using the expression $x * x$, where x is drawn from the list `nums`”. More than one name can be assigned. For example, to produce a times table:

```
GHCi:
Prelude> [x * y | x <- [1 .. 5], y <- [1 .. 4]]
[1,2,3,4,2,4,6,8,3,6,9,12,4,8,12,16,5,10,15,20]
```

You can see that, in the case of multiple names, each later one cycles more quickly than the earlier one, so we have $x = 1, y = 1$ followed by $x = 1, y = 2$ rather than $x = 2, y = 1$. We can also specify *guards* to select and process only those items in the initial list which meet certain criteria. For example, to find all the even squares:

```
GHCi:
Prelude> [x * x | x <- [1 .. 20], x `rem` 2 == 0]
[4,16,36,64,100,144,196,256,324,400]
```

When we write a function involving the `[..]` construct, the type is different – it has an additional type constraint:

```
f : (Num a) => a -> [a]
g : (Num a, Enum a) => a -> [a]

f x = [1, x]
g x = [1 .. x]
```

The typeclass **Enum** includes such things as can be enumerated – that is to say, given one, we can find the next one. The integers, for example, are enumerable.

Lists of characters

A list of characters is special, and is called a *string*. It is printed between double quotation marks “like this”. The types **String** and **[Char]** are interchangeable – it really is a list of characters, so we can use ordinary list functions:

```
GHCi:
Prelude> reverse' "stressed"
"desserts"
```

We will always write **String** instead of **[Char]** in our types for consistency.