

# Subtyping

Gabriele Keller

Liam O'Connor

Johannes Åman Pohjola

November 13, 2025

## 1 Subtyping

With type classes, the programmer can use the same overloaded function symbol `+` for addition of both floating point values and integer values, and the compiler or runtime will figure out which one to use. However, the following expression would still be rejected by the MinHs compiler:

$$1 + 1.75$$

This is because addition can be applied to two integers or two floats, but not a combination of both.<sup>1</sup> We explicitly have to convert the `Int` value to `Float` to add the two values.

C solves this “problem” using something called *integer promotion*: the basic types are ordered and if operations like `+` or `==` are applied to mixed operands, the one which is the lowest in the hierarchy is automatically cast to the higher type. This is quite convenient, but can easily lead to unexpected behaviour and subtle bugs, in particular with respect to signed/unsigned types.

The idea behind subtyping is similar to the approach in C. The key technical device is that that types can be partially ordered in a subtype relationship

$$\tau \leq \sigma$$

such that, whenever a value of some type  $\sigma$  is required, it is also fine to provide a value of type  $\tau$ , as long as  $\tau$  is a subtype of  $\sigma$ . For example, we could have the following subtype relationship:

$$\mathbf{Int} \leq \mathbf{Float} \leq \mathbf{Double}$$

With subtyping, it would then be ok to have

$$1 +_{\mathbf{Float}} 1.75$$

because even though the LHS is an `Int`, any subtype of `Float` is acceptable.

### 1.1 Coercion Interpretation

There are different ways to interpret the subtype relationship. One is to define  $\tau$  to be a subtype of  $\sigma$  if  $\tau \subseteq \sigma$ . After all, in mathematics, the integers are a subset of the rationals. However, this interpretation is quite restrictive for a programming language: `Int` is *not* a subset of `Float` since they have very different representations. However, there is an obvious *coercion* from `Int` to `Float`. For our study of subtyping, we will focus on this so-called *coercion interpretation* of subtyping:  $\tau$  is a subtype of  $\sigma$ , if there is a *sound*<sup>2</sup> coercion from values of type  $\tau$  to values of type  $\sigma$ .

As another example, consider a `Graph` and a `Tree`. Since trees are a special case of graphs, trees can be converted into graphs, and we can view the tree type as subtype of the graph type in the coercion interpretation.

<sup>1</sup>In Haskell, this expression by itself would be fine, as constants are also overloaded and 1 has type `Num a => a`. However, the compiler would also reject the addition of integer and float values, for example `(1 :: Int) + (1.7 :: Float)`.

<sup>2</sup>More about that later

## 1.2 Properties

For a subtyping relationship to be sound, it has to be reflexive, transitive, and antisymmetric (with respect to type isomorphism). In other words  $\leq$  is a *partial order*. This holds both for the subset interpretation and the coercion interpretation. For the subset interpretation, all three properties follow directly from the properties of the subset relation. In the coercion interpretation, reflexivity holds because the identity function is a coercion from  $\tau \rightarrow \tau$ . Transitivity holds since, given a coercion function from  $f : \tau_1 \rightarrow \tau_2$  and  $g : \tau_2 \rightarrow \tau_3$ , the composition of  $f$  and  $g$  result in a coercion function from  $\tau_1 \rightarrow \tau_3$ .

In the coercion relation, the antisymmetry of subtyping means that if we can coerce  $\tau$  to  $\rho$  and  $\rho$  back to  $\tau$ , then it must be the case that  $\tau \simeq \rho$ . This is only true if the coercion functions are *injective* — that is, we can map each element of the *domain* (input) of the function to a *unique* element of the *codomain* (output).

## 1.3 Coherency of Coercion

Coercion should be *coherent*. That is, if there are two ways to coerce a value, both coercions should yield the same result.

For example, let us assume we define *Int* to be a subtype of *Float*, and both to be subtypes of *String*, with coercion functions

```
intToFloat :: Int → Float
```

```
intToString :: Int → String
```

```
floatToString :: Float → String
```

On first sight, this looks reasonable. It is not coherent, however, because there are two coercion function from *Int* to *String*: the provided function *intToString*, but also *intToFloat* composed with *floatToString*. Unfortunately, applied to the number 3, for example, one would result in the string "3", the other in "3.0".

One reason why type promotion in C can be so tricky is precisely that it is not coherent in this way.

## 1.4 Variance

If we add subtyping to MinHS, one question that arises is how the subtyping relationship interacts with our type constructors. For example, if  $\text{Int} \leq \text{Float}$ , what about pairs, sums and function over these types? How do they relate to each other?

For pairs and sums, the answer is quite straight forward. Obviously, given a coercion function *intToFloat*, we can easily define coercion functions on pairs and sums:

```
p1 :: (Int × Int) → (Float × Float)
p1 (x, y) = (intToFloat x, intToFloat y)
```

```
p2 :: (Int × Float) → (Float × Float)
p2 (x, y) = (intToFloat x, y)
...
```

```
s1 :: (Int + Int) → (Float + Float)
s1 x = case x of
  InL v → intToFloat v
```

$\text{InR } v \rightarrow \text{intToFloat } v$

...

This means that, if two types  $\tau_1$  and  $\tau_2$  are subtypes of  $\sigma_1$  and  $\sigma_2$ , respectively, then the product type  $\tau_1 \times \tau_2$  is also a subtype of  $\sigma_1 \times \sigma_2$ . The same is true for sums. More formally, we have:

$$\frac{\tau_1 \leq \rho_1 \quad \tau_2 \leq \rho_2}{(\tau_1 \times \tau_2) \leq (\rho_1 \times \rho_2)}$$

$$\frac{\tau_1 \leq \rho_1 \quad \tau_2 \leq \rho_2}{(\tau_1 + \tau_2) \leq (\rho_1 + \rho_2)}$$

The following diagram also shows the subtype relationship between pair types, for **Int** and **Float**:

$$\begin{array}{ccc} \mathbf{Int} \times \mathbf{Int} & \Longrightarrow & \mathbf{Int} \times \mathbf{Float} \\ \Downarrow & & \Downarrow \\ \mathbf{Float} \times \mathbf{Int} & \Longrightarrow & \mathbf{Float} \times \mathbf{Float} \end{array}$$

Since product and sum types interact with subtyping in such a natural way, it is tempting to expect the same to work for all type constructors. Resist this temptation.

Consider function types. Is  $\mathbf{Int} \rightarrow \mathbf{Int}$  a subtype of  $\mathbf{Float} \rightarrow \mathbf{Int}$ ? That is, if a function of type  $\mathbf{Float} \rightarrow \mathbf{Int}$  is required, would it be ok to provide a function of type  $\mathbf{Int} \rightarrow \mathbf{Int}$  instead? Considering that the type **Int** is more restricted than the type **Float**, this means that a function which only works on the “smaller” type **Int** is also, in some sense, less powerful. Or, coming back to our second example, if we need a function to process any graph, then a function which only works on trees (and maybe relies on the fact that there are no cycles in a tree) is clearly not sufficient. We are also not able to define a coercion function in terms of our coercion function *intToFloat*:

$$c :: (\mathbf{Int} \rightarrow \mathbf{Float}) \rightarrow (\mathbf{Float} \rightarrow \mathbf{Float})$$

The other direction, however, is actually quite easy:

$$\begin{aligned} c' &:: (\mathbf{Float} \rightarrow \mathbf{Float}) \rightarrow (\mathbf{Int} \rightarrow \mathbf{Float}) \\ c' f &= \mathbf{let } g \ x = f \ (\text{intToFloat } x) \\ &\quad \mathbf{in } g \end{aligned}$$

Therefore, somewhat surprisingly, we have  $(\mathbf{Float} \rightarrow \mathbf{Float}) \leq (\mathbf{Int} \rightarrow \mathbf{Float})$ .

So, what about the result type of a function: is  $\mathbf{Int} \rightarrow \mathbf{Int}$  as subtype of  $\mathbf{Int} \rightarrow \mathbf{Float}$ , vice versa, or are these types not in a subtype relationship at all? If we need a function which returns a **Float** and get one that returns an **Int**, that’s ok since **Ints** can be coerced to **Floats**. Similarly, if we need a function which returns a graph, and we get a tree, that’s ok since a tree is a special case of a graph.

$$\begin{aligned} c'' &:: (\mathbf{Int} \rightarrow \mathbf{Int}) \rightarrow (\mathbf{Int} \rightarrow \mathbf{Float}) \\ c'' f &= \mathbf{let } g \ x = \text{intToFloat } (f \ x) \ \mathbf{in } g \end{aligned}$$

To summarise, the subtype relationship on functions over **Int** and **Float** is as follows (of course, you can substitute any type  $\tau$  for **Int**,  $\rho$  for **Float** here, as long as  $\tau \leq \rho$ ):

$$\begin{array}{ccc} \mathbf{Int} \rightarrow \mathbf{Int} & \Longrightarrow & \mathbf{Int} \rightarrow \mathbf{Float} \\ \Uparrow & & \Uparrow \\ \mathbf{Float} \rightarrow \mathbf{Int} & \Longrightarrow & \mathbf{Float} \rightarrow \mathbf{Float} \end{array}$$

The subtype propagation rule for function types expresses exactly the same relationship:

$$\frac{\tau_1 \leq \rho_1 \quad \tau_2 \leq \rho_2}{(\rho_1 \rightarrow \tau_2) \leq (\tau_1 \rightarrow \rho_2)}$$

Another example of a type which interacts with subtyping in a non-obvious manner are updateable arrays and reference types. To understand what is happening, let us have a look at Haskell-style updatable references. We have the following basic operations on this type:

```

newIORef :: a -> IO (IORef a)      — Returns an initialised reference
writeIORef :: a -> IORef a -> IO () — Updates the value of a reference
readIORef :: IORef a -> IO a      — Returns the current value

```

All other operations can be expressed in terms of these three operations.

The question now is: if  $\tau \leq \sigma$ , what is the subtype relationship between `IORef  $\tau$`  and `IORef  $\sigma$` ? To check whether, for example, `IORef Int`  $\leq$  `IORef Float`, think about what happens if we apply `writeIORef (1.5 :: Float)` to an `IORef Int` instead of an `IORef Float`. Clearly, this would not work, as the floating point value can't be stored in an `Int` reference. It would be okay the other way around: if we `writeIORef (1 :: Int)` apply to an `IORef Float`, it would be fine, since we could first coerce the value to a `Float` and store the result in the `Float` reference. This seems to suggest that `IORef Float`  $\leq$  `IORef Int`.

However, if we assume that `IORef Float`  $\leq$  `IORef Int`, we run into trouble with `readIORef`. If `readIORef` expects an `IORef Int`, and we give it an `IORef Float` instead, `readIORef` will return a floating point value which we cannot coerce to `Int`. In this case, the other direction would be fine: if we expect an `IORef Float`, we could apply it to an `IORef Int`, and then cast the resulting `Int` value to `Float`.

This means that  $\tau \leq \sigma$  implies no subtype relationship between `IORef Float` and `IORef Int` at all: when a mutable reference of a certain type is required, we cannot substitute the reference for a sub- or supertype: we must give exactly that type.

We encounter exactly the same situation with updatable arrays. In Java, the language allows subtyping for arrays, at the cost of dynamic checks since this violates type safety.

Type constructors like product and sum, which leave the subtype relationship intact, as called *covariant*. type constructors which reverse the relationship, like the function type in its first argument, are called *contravariant*. Type constructors like `IORef`, which do not imply a subtype relationship at all, are called *invariant*.