

Type Variables in Patterns

Richard A. Eisenberg
Bryn Mawr College
Bryn Mawr, PA, USA
rae@cs.brynmawr.edu

Joachim Breitner
University of Pennsylvania
Philadelphia, PA, USA
joachim@cis.upenn.edu

Simon Peyton Jones
Microsoft Research
Cambridge, UK
simonpj@microsoft.com

Abstract

For many years, GHC has implemented an extension to Haskell that allows type variables to be bound in type signatures and patterns, and to scope over terms. This extension was never properly specified. We rectify that oversight here. With the formal specification in hand, the otherwise-labyrinthine path toward a design for binding type variables in patterns becomes blindingly clear. We thus extend *ScopedTypeVariables* to bind type variables explicitly, obviating the *Proxy* workaround to the dustbin of history.

CCS Concepts • Software and its engineering → Patterns; Functional languages; Data types and structures;

Keywords Patterns, type variables, polymorphism, Haskell

ACM Reference Format:

Richard A. Eisenberg, Joachim Breitner, and Simon Peyton Jones. 2018. Type Variables in Patterns. In *Proceedings of the 11th ACM SIGPLAN International Haskell Symposium (Haskell '18)*, September 27–28, 2018, St. Louis, MO, USA. ACM, New York, NY, USA, 12 pages. <https://doi.org/10.1145/3242744.3242753>

1 Introduction

Haskell allows the programmer to write a *type signature* for a definition or expression, both as machine-checked documentation, and to resolve ambiguity (Section 2.1). For example,

```
prefix :: a → [[a]] → [[a]]  
prefix x yss = map xcons yss  
where xcons ys = x : ys
```

Sadly, it is not always possible to write such a type signature. For example, to give a signature for *xcons* we might try:

```
prefix :: a → [[a]] → [[a]]  
prefix x yss = map xcons yss  
where xcons :: [a] → [a]  
xcons ys = x : ys
```

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

Haskell '18, September 27–28, 2018, St. Louis, MO, USA

© 2018 Association for Computing Machinery.

ACM ISBN 978-1-4503-5835-4/18/09...\$15.00

<https://doi.org/10.1145/3242744.3242753>

But Haskell98’s scoping rules specify that the *a* in the signature for *xcons* is *locally* quantified, thus: $xcons :: \forall a. [a] \rightarrow [a]$. That is not what we want! We want the *a* in the signature for *xcons* to mean “the universally quantified type variable for *prefix*”, and *Haskell98 provides no way to do that*.

The inability to supply a type signature for *xcons* might seem merely inconvenient, but it is just the tip of an iceberg. Haskell uses type inference to infer types, and that is a wonderful thing. However, insisting on *complete* type inference—that is, the ability to infer types for any well-typed program with no help from the programmer—places serious limits on the expressiveness of the type system. GHC’s version of Haskell has gone far beyond these limits, and fundamentally relies on programmer-supplied annotations to guide type inference. As some examples, see the work of Peyton Jones et al. [2007], Vytiniotis et al. [2011], or Eisenberg et al. [2016].

So the challenge we address is this: *it should be possible for the programmer to write an explicit type signature for any sub-term of the program*. To do so, some type signatures must refer to a type that is already in the static environment, so we need a way to *name* such types. The obvious way to address this challenge is by providing language support for *lexically scoped type variables*. GHC has long supported scoped type variables: the *ScopedTypeVariables* extension is very popular, and 29% of Haskell packages on Hackage use it. But it has never been formally specified! Moreover, as we shall see, it is in any case inadequate to the task. In this paper we fix both problems, making the following contributions:

- In the days of Haskell98, scoped type variables were seldom crucial. Through a series of examples we show that, as Haskell’s type system has grown more sophisticated, the need for scoped type variables has become acute (Section 2), while GHC’s existing support for them has become more visibly inadequate (Section 3).
- To fix these inadequacies, we describe *visible type application in patterns*, a natural extension to GHC’s existing visible type applications from terms to patterns (Section 4).
- We give the first formal specification of scoped type variables for Haskell, formalizing the folklore, and providing a firm foundation for both design and implementation (Section 5).
- As part of this specification, we offer a new and simpler typing judgment for GADT pattern matching (Section 5.3), which treats uniformly the universal and existential variables of a data constructor.

```

show  :: Show a => a -> String
read  :: Read a => String -> a
(+)   :: [a] -> [a] -> [a]
concat :: [[a]] -> [a]

```

Figure 1. Types of standard functions

2 Motivation and Background

2.1 The Need for Type Annotations

One of the magical properties of ML-family languages, including Haskell, is that type inference allows us to write many programs with no type annotations whatsoever. In practice, however, Haskell programs contain many user-written type signatures, for two main reasons.

First, the type of a function can be extremely helpful as documentation, with the advantage that it is *machine-checked* documentation. Almost all programmers regard it as good practice to provide a signature for every top-level function. Indeed, GHC has a warning flag, `-Wmissing-signatures`, which enforces this convention.

Second, as Haskell's type system becomes increasingly expressive, complete type inference becomes intractable, and the type system necessarily relies on programmer-supplied type annotations. Here are some examples:

- *Type-class ambiguity* is present even in Haskell 98. Consider¹:

```

normalize :: String -> String
normalize s = show (read s)

```

This function parses a string to a value of some type, and then turns that value back into a string. But nothing in the code specifies that type, so the programmer must disambiguate. One way to do so is to provide a type signature that specifies the result type of the `read s` call, thus:

```
normalize s = show (read s :: Int)
```

- *Polymorphic recursion*. In ML, recursive calls to a function must be at monomorphic type, but Haskell has always supported polymorphic recursion, provided the function has a type signature. For example:

```

data T a = Leaf a | Node (T [a]) (T [a])
leaves :: T a -> [a]
leaves (Leaf x)    = [x]
leaves (Node t1 t2) = concat (leaves t1 + leaves t2)

```

- *Higher-rank types* [Peyton Jones et al. 2007]. Consider

```

f :: (∀a. [a] -> [a]) -> ([Char], [Bool])
f g = (g "Hello", g [True, False])

```

Here the type of `g` is polymorphic, so it can be applied to lists of different type. The type signature is essential to specify the type of the argument `g`; without it, `f` will be rejected.

- *Generalized algebraic data types* [Schrijvers et al. 2009]. The popular *GADTs* extension to GHC allows pattern matching to refine the type information available in the right hand side of an equation. Here is an example:

```

data G a where
  MkInt  :: G Int
  MkFun  :: G (Int -> Int)

```

When we learn that a value `g :: G a` is actually the constructor `MkInt`, then we simultaneously learn that `a` really is `Int`. GHC can use this fact during type checking the right-hand side of a function, like this:

```

matchG :: G a -> a
matchG MkInt  = 5
matchG MkFun  = (+10)

```

Again, however, `matchG` will only type-check if it is given a signature; see Schrijvers et al. [2009] for details.

- *Ambiguous types*. Consider

```

type family F a
ambig :: Typeable a => F a -> Int
test  :: Char -> Int
test x = ambig x

```

In `test` GHC must decide at what type to call `ambig`; that is, what type should instantiate the `a` in `ambig's` type. Any choice `a = τ` must ensure that `F τ ~ Char` but, because `F` might not be injective, that does not tell us what `a` should be. A type signature is not enough to resolve this case; we need a different form of type annotation, namely *visible type application* (Section 2.3).

There is a general pattern here: as the type system becomes more expressive, the type inference needs more guidance. Moreover, that guidance is extremely informative to the programmer, as well as to the compiler.

2.2 Support for Scoped Type Variables

Given the increasing importance of type annotations, a good principle is this: it should be possible to specify, via a type signature, the type of any sub-expression or any let-binding. Alas, as shown in the introduction, Haskell98 supports only *closed* type signatures, so there are useful type signatures that we simply cannot write.

The key deficiency in Haskell98 is that it provides no way to bring type variables into scope. GHC has recognized this need for many years, and the *ScopedTypeVariables* extension offers two ways to bring a type variable into scope:

- *Binding in a declaration signature*. Since 2004 GHC allows you to write

¹ Figure 1 gives the types of standard functions, such as `read` and `show`.

```

prefix :: ∀a. a → [[a]] → [[a]]
prefix x yss = map xcons yss
  where xcons :: [a] → [a]
        xcons ys = x : ys

```

The explicit “ \forall ” brings a into scope in the rest of the type signature (of course), but it *also* brings a into scope in the body of the named function, `prefix` in this case. The rule is a bit strange, because the definition of `prefix` is not syntactically “under” the \forall , and indeed the signature can be written far away from the actual binding. But in practice the rule works well, and we take it as-is for the purposes of this paper.

- *Pattern signatures*: binding a type variable in a pattern. For even longer, since 1998, GHC has allowed you to write this:

```

prefix (x :: b) yss = map xcons yss
  where xcons :: [b] → [b]
        xcons ys = x : ys

```

Here, we bind the type variable b in the pattern $(x :: b)$, and this binding scopes over the body of the binding. We describe pattern signatures in much more detail in Section 3.

2.3 Visible Type Application

The `TypeApplications` extension provides a relatively new form of type annotation: explicit type applications, first described by Eisenberg et al. [2016]. The idea is that an argument of the form `@ty` specifies a *type* argument. This can often be used more elegantly than a type signature. For example, a hypothetical unit-test for the function `isJust`: *Maybe a* \rightarrow *Bool*,

```

testIsJust1 = isJust (Just (2018 :: Int))    == True
testIsJust2 = isJust (Nothing :: Maybe Int) == False

```

can equivalently be written more elegantly using explicit type annotations

```

testIsJust1 = isJust (Just @Int 2018) == True
testIsJust2 = isJust (Nothing @Int) == False

```

Visible type application solves the awkward case of *ambig* in Section 2.1: we can disambiguate the call with a type argument. For example:

```

type family F a
type instance F Bool = Char
ambig :: Typeable a => F a -> Int
test :: Char -> Int
test x = ambig @Bool x

```

Here we specify that `ambig` should be called at `Bool`, and that is enough to type-check the program.

It is natural to wonder whether we can extend visible type application to patterns, just as we extended type signatures to

patterns. Doing so is the main language extension suggested in this paper: Section 4.

3 Pattern Signatures and Their Shortcomings

We see above that `ScopedTypeVariables` enables the user to bind type variables in patterns, by providing a *pattern signature*, that is, a type signature in a pattern. We explore pattern signatures and their shortcomings in this section.

3.1 The Binding Structure of a Pattern Signature

A pattern signature may *bind* a type variable, but it may also *mention* a type variable that is already in scope. For example, we may write

```

prefix (x :: a) yss = map xcons yss
  where xcons (ys :: [a]) = x : ys

```

Here, the pattern signature $(x :: a)$ binds a (as well as x), but the pattern signature $(ys :: [a])$ simply mentions a (which is already in scope), as well as binding ys . The rule is this: a use of a type variable p in a pattern signature is an *occurrence* of p if p is already in scope; but *binds* p if p is not already in scope.

It is entirely possible to have many different type variables in scope, all of which are aliases for the same type. For example:

```

prefix :: ∀a. a → [[a]] → [[a]]
prefix (x :: b) (yss :: [[c]]) = map xcons yss
  where xcons (ys :: [d]) = x : ys

```

Here a , b , c , and d are all in scope in the body of `xcons`, and are all aliases for the same type.

The current implementation of `ScopedTypeVariables` allows such lexically-scoped type variables to stand only for other *type variables*, and not for arbitrary *types*, a point we return to in Section 3.5.

3.2 Pattern Signatures Are Useful

Pattern signatures have merit even if there are no type variables around. Consider this Haskell program:

```

main = do x ← readLn
       if null x then putStrLn "Empty"
       else putStrLn "Not empty"

```

The types of the program are ambiguous: Clearly, x is some type that has a `Read` instance, and because it is passed to `null`, is a list², but the compiler needs to know the precise type, and rejects the program.

To fix this in Haskell98, the programmer has two options:

²Or, with a recent version of the standard library, it is something with a `Foldable` instance.

- They can wrap the call to `readLn` in a type annotation:

```
main = do x ← (readLn :: IO [Int])
        if null x then putStrLn "Empty!"
        else print x
```

but this is infelicitous because there is no question that `readLn` is in the *IO*, and with larger types this can get very verbose.

- They can wrap an *occurrence* of `x` in a type annotation:

```
main = do x ← readLn
        if null (x :: [Int]) then putStrLn "Empty!"
        else print x
```

but again this is unsatisfying, because it feels too late.

Both variants are essentially work-arounds for the natural way of specifying the type of `x`, namely at its binding site:

```
main = do (x :: [Int]) ← readLn
        if null x then putStrLn "Empty!"
        else print x
```

which is precisely what the *PatternSignatures* language extension provides³—the ability to write a type annotation in a pattern.

Users, especially beginners, who have to track down a confusing type error in their code, can now exhaustively type annotate not just their terms, but also their patterns, until they have cornered the bug.

3.3 Pattern Signatures Are Essential

Pattern signatures become more crucial when we consider existential types. The *ExistentialQuantification* extension allows users to bind *existential* variables in their data constructors. These are type variables whose values are “stored” by a constructor (but not really, because types are erased) and made available during pattern matching. Here are two examples:

data Ticker where

```
MkTicker :: ∀a. a → (a → a) → (a → Int) → Ticker
```

data Showable where

```
MkShowable :: ∀a. Show a ⇒ a → Showable
```

A *Ticker* contains an object of some type (but we do not know *what* type), along with an update function of type $a \rightarrow a$ and a way to convert an *a* into an *Int*. The *Showable* type packs a value of an arbitrary type that has a *Show* instance along with its *Show* dictionary. Here are some functions that operate on these types:

- Updates a ticker, returning whether or not the ticker
- has reached a limit

```
tick :: Ticker → Int → (Ticker, Bool)
```

```
tick (MkTicker val upd toInt) limit
    = (newTicker, toInt newVal ≥ limit)
  where newVal    = upd val
        newTicker = MkTicker newVal upd toInt
```

```
showAll :: [Showable] → String
showAll [] = ""
showAll (MkShowable x : ss) = show x # showAll ss
```

We see that the *tick* function can unpack the existential in the *Ticker* value and operate on the value of type *a* without ever knowing what *a* is. Similarly, the *showAll* function works with data of type *a* knowing only that *a* has a *Show* instance. (The *Show a* constraint is brought into scope by the pattern match.)

However, existentials can never escape, forbidding the following function:

```
jailbreak (MkTicker val _ _) = val
```

What should the type of *jailbreak* be? There is no answer to this question ($\text{jailbreak} :: \text{Ticker} \rightarrow a$ is clearly too polymorphic), and so GHC rejects this definition, correctly stating that type variable ‘a’ would escape its scope.

We naturally wish to *name* the existential type variable sometimes. For example, suppose we wanted to give *newVal* (in the *where* clause of *tick*) a type signature. Saying $\text{newVal} :: a$ would hardly do, because there is not yet a connection between the name *a* and the type unpacked from *MkTicker*. We have to do this:

```
tick (MkTicker (val :: b) upd toInt) limit = ...
  where newVal :: b
        newVal = upd val
```

The $\text{val} :: b$ in the pattern binds the type variable *b*, so we can refer back to it later.

3.4 Pattern Signatures Are Clumsy

Pattern signatures can be clumsy to use, when the type variable is buried deep inside an ornate type. Here is a contrived example:

data Elab where

```
MkElab :: Show a ⇒ [Maybe (Tree (a, Int))] → Elab
```

```
getE :: Elab → Int
```

```
getE (MkElab (xs :: [Maybe (Tree (a, Int))])) = ...a ...
```

To bring *a* into scope in *f*’s right-hand side, we have to repeat the *MkElab*’s elaborate argument type.

More seriously, it may be impossible, rather than merely clumsy, to bind the variable we need. Consider the following GADT:

data GM a where

```
MkMaybe :: GM (Maybe b)
```

```
matchGM :: a → GM a → Bool
```

```
matchGM x MkMaybe = isJust x
```

³Modern GHC actually folds *PatternSignatures* into *ScopedTypeVariables*, giving both extensions the same meaning. However, it is expositionally cleaner to separate the two, as we do throughout this paper.

This definition works just fine: GHC learns that x 's type is *Maybe* b (for some existential b) and the call to *isJust* is well typed. But what if we want to bind b in this definition? Annoyingly, *MkMaybe* has no argument to which we can apply a pattern signature. Nor does it work to wrap a pattern signature around the *outside* of the match, thus:

```
matchGM :: a → GM a → Bool
matchGM x (MkMaybe :: GM (Maybe b)) = isJust @b x
```

This definition is rejected. The problem is that the type annotation on the *MkMaybe* pattern is checked *before* the pattern itself is matched against⁴. Before matching the *MkMaybe*, we do not yet know that a is really *Maybe* b . Nor can we put the type annotation on x , as that, too, occurs before the *MkMaybe* pattern has been matched. A possible solution is this monstrosity:

```
matchGM :: a → GM a → Bool
matchGM x gm@MkMaybe = case gm of
  (_ :: GM (Maybe b)) → isJust @b x
```

This is grotesque. Of course, in this case, we can simply flip the argument order to *matchGM*, but we should not be forced to change argument order just because of clumsy syntax. We must do better, and we will in Section 4.

3.5 Pattern Signatures Resist Refactoring

In Section 3.1, we explained that a scoped type variable may refer only to another type variable. This means that the definition

```
prefix :: a → [[a]] → [[a]]
prefix (x :: b) yss = map xcons yss
  where xcons ys = x : ys
```

is accepted, because b stands for the type variable in the type of *prefix*. But suppose, for example, we specialize the type signature of *prefix* without changing its definition:

```
prefix :: Int → [[Int]] → [[Int]]
prefix (x :: b) yss = map xcons yss
  where xcons ys = x : ys
```

Now this definition is rejected with the error message “Couldn't match expected type b with actual type *Int*”, because b would have to stand for *Int*.

Since the design of *ScopedTypeVariables*, GHC has evolved, and with the advent of type equalities, the restriction itself becomes confusing. Should GHC accept the following definition?

```
prefix :: a ~ Int ⇒ Int → [[a]] → [[a]]
prefix (x :: b) yss = map xcons yss
  where xcons ys = x : ys
```

⁴This ordering arises because any type variable in a pattern signature is bound within the pattern.

Is b an alias for a (legal) or for *Int* (illegal)? Since a and *Int* are equal, the question does not really make sense. We therefore propose to simply drop this restriction (Section 4.3).

3.6 Pattern Signatures Are Inadequate

We end our growing list of infelicities with a case in which there is no way whatsoever to bind the type variable, short of changing the data type definition.

Type families [Chakravarty et al. 2005; Eisenberg et al. 2014] allow users to write type-level functions and encode type-level computation. For example, we might write this:

type family F a **where**

```
F Int = Bool
F Char = Double
F Float = Double
```

Naturally, we can use a type family to define the type of an argument to an existential data constructor:

data TF **where**

```
MkTF :: ∀a. Typeable a ⇒ F a → TF
```

The *MkTF* constructor stores a value of type (F a); it also stores a dictionary for a *Typeable* a constraint [Peyton Jones et al. 2016]—that is, we can use a runtime type test to discover the choice for a . We would thus like to write the following function:

```
toDouble :: TF → Double
toDouble (MkTF x) -- We expect x :: F a
  | Just HRefl ← isType @Int = if x then 1.0 else -1.0
  | Just HRefl ← isType @Char = x
  | Just HRefl ← isType @Float = x
  | otherwise = 0.0
  where
    isType :: ∀ty. Typeable ty ⇒ Maybe (a ≈: ty)
    isType = eqTypeRep (typeRep @a) (typeRep @ty)
```

The specifics of this function are not important here (see [Peyton Jones et al. 2016]). For our present purposes, the crucial point is this: the existentially-bound type a is mentioned in both the definition of *isType* and its type signature—but *there is no way to bring a into scope*. We might try using a pattern signature at the binding of x , thus:

```
toDouble (MkTF (x :: F a)) = ...
```

but that does not quite work. The problem is that F is not injective. The a in that pattern type annotation need not be the same one packed into the existential type variable by *MkTF*, and GHC rightly considers such an a to be ambiguous⁵.

The only workaround available today is to change *MkTF* to take a proxy argument:

⁵If F were an injective type family, we could label it as such to fix the problem [Stolarek et al. 2015]. But here we assume that it is not.

```
data Proxy a = Proxy -- in GHC's Data.Proxy
data TF where
```

```
MkTF :: ∀a. Typeable a ⇒ Proxy a → F a → TF
toDouble (MkTF (_ :: Proxy a) x) = ...
```

The *Proxy* type stores no runtime information (at runtime, it is isomorphic to `()`), but the type *Proxy a* carries the all-important type variable *a*. All datatypes are injective, so we can use this proxy argument to bind the type variable *a* in a way that we could not do previously.

As with many other examples, this is once again unsatisfying: it is a shame that we have to modify the data constructor declaration just to deal with type variable binding.

3.7 Conclusion

In this section we have seen that pattern signatures allow us to bring into scope the existentially-bound type variables of a data constructor, but that doing so can be clumsy, and occasionally impossible. We need something better.

4 Visible Type Application in Patterns

Consider again *Elab* from Section 3.4:

```
data Elab where
  MkElab :: Show a ⇒ [Maybe (Tree (a, Int))] → Elab
```

and suppose we want to build a value of type *Elab* containing an empty list. We cannot write just `MkElab []` because that is ambiguous: we must fix the type at which *MkElab* is called so that the compiler can pick the right *Show* dictionary. We can use a type signature, but it is clumsy, just as the pattern signature was clumsy in Section 3.4:

```
MkElab ([] :: [Maybe (Tree (Bool, Int))])
```

It is much nicer to use visible type application and write `MkElab @Bool []`. So it is natural to ask whether we could do the same in patterns, like this:

```
getE :: Elab → Int
getE (MkElab @a xs) = ...a ...
```

Here, we bind *a* directly⁶, as a type-argument pattern all by itself, rather than indirectly via a pattern signature.

We call this *visible type application in patterns*, a dual of visible type application in the same way that a pattern signature is a dual of type signatures. This section describes visible type application in patterns informally, while the next formalizes it.

⁶One might reasonably wonder how we can steal `@` in a pattern in this way. After all, Haskellers can also write, e.g., `f list@(x : xs) = ...` to alias *list* to the pattern `(x : xs)`. The new syntax is not actually ambiguous, however: an *as-pattern* always has a variable on its left, while our new form is always headed by a data constructor with all type patterns preceding value-level patterns.

This feature was first requested more than two years ago.⁷ Furthermore, binding type variables like this is useful for more than just disambiguation, as we will shortly see.

4.1 Examples

Visible type application in patterns immediately fixes the other problems of pattern signatures identified above. For example, in the GADT example of Section 3.4 we can write

```
matchGM :: a → GM a → Bool
matchGM x (MkMaybe @b) = isJust @b x
```

and for the type-family example of Section 3.6 we write

```
toDouble (MkTF @a x) = ...
```

4.2 Universal and Existential Variables

Visible type applications in patterns can be used for *all* the type arguments of a data constructor, whether existential or universal. As an example of the latter we may write

```
main = do (Just @Int x) ← readMaybe 'fmap' getLine
         putStrLn "Input was " + show x
```

as an alternative to

```
main = do (Just (x :: Int)) ← readMaybe 'fmap' getLine
         putStrLn$ "Input was " + show x
```

Visible type application in patterns considers the type of data constructor, *exactly as written by the user*. For example

```
data G a b where
  G1 :: ∀b. Char → G Int b
  G2 :: ∀p q b. p → q → b → G (p, q) b
  G3 :: ∀p q a b. (a ~ (p, q)) ⇒ p → q → b → G a b

f :: G a Bool → Int
f (G1 @Bool y)           = ord y
f (G2 @p @q @x y z)     = 0
f (G3 @p @q @a @Bool x y z) = 1
```

In this definition

- *G1* has one type argument.
- *G2* has three type arguments, but we have chosen to match only the first two.
- *G3* is morally identical to *G2*, because of the equality, but it is written with four type arguments, and visible type application in patterns follows that specification.

4.3 Type Aliases

In Section 3.5 we have seen that GHC currently restricts type variables to refer to type variables, but that this does not have to be the case. Similar questions arise in our function *f* above. Could we write this for *G3*?

```
f (G3 @p @q @(p, q) @b x y z) = 1
```

⁷<https://ghc.haskell.org/trac/ghc/ticket/11350>

Term variables	$\ni x, y, z, f, g, h$
Internal type vars	$\ni a, b$
User type vars	$\ni c$
Data constructors	$\ni K$
Atoms	$v ::= K \mid x$
Expressions	$e ::= v \mid \lambda x. e \mid e_1 e_2 \mid \mathbf{case} e \mathbf{of} \{ \overline{p \rightarrow e} \}$ $\mid \mathbf{let} x :: \tau = e_1 \mathbf{in} e_2$
Patterns	$p ::= x \mid K \overline{p}$
Polytypes	$\sigma ::= \forall \overline{a}. \tau$
Monotypes	$\tau, v ::= tv \mid T \overline{\tau} \mid \dots$
Type variables	$tv ::= a \mid c$
Type env	$\Gamma ::= \epsilon \mid \Gamma, v : \sigma$
Substitutions	$\theta ::= \overline{tv \mapsto \tau}$
Types of data cons	$\Gamma_t = \overline{K : \forall \overline{a}. \overline{\sigma} \rightarrow T \overline{a}}$

Figure 2. The initial grammar

Instead of a we have written (p, q) , which is equal to a . And instead of $Bool$ we have written b , thereby binding b to $Bool$.

Given the ubiquity of equalities, it no longer seems to make sense to restrict what a scoped type variable can stand for, so we propose simply to drop the restriction. Doing so simplifies the specification and the implementation of both pattern signatures and visible type application in patterns. A GHC proposal by one of the authors [Breitner 2018] is underway. Relaxing the requirement also allows the user to use type variables as “local type synonyms” that stand for possibly long types:

```
processMap :: Map Int (Maybe (Complex Type)) → ...
processMap (m :: Map key value) = ...
```

5 Formal Specification

We give the first formal specification of a number of extensions to Haskell related to pattern matching and the scoping of type variables: annotations in patterns, scoped type variables, and type application syntax in patterns. This section builds up these specification step by step, starting with a specification of the language without these features.

Our specification does not cover **let**-bindings and declaration type signatures. We focus instead on pattern signatures, which is where our new contribution lies. The scoping of forall-bound type variables from declaration type signatures would be straightforward to add.

5.1 The Baseline

We begin with a reduced model of Haskell98 terms, which knows nothing yet about scoped type variables nor type

equalities (i.e., no GADTs). We also removed type class constraints. The syntax is given in Fig. 2, and the typing rules are in Fig. 3. In the typing rules, we use a convention where an over-bar indicates a list, optionally with a superscript index to indicate the iterator. Iterators are additionally annotated with length bounds, where appropriate.

The grammar includes separate metavariables for internal type variables a and user type variables c . The former are type variables as propagated by the compiler, while the latter are type variables the user has written. It is as if internal type variables a are spelled with characters unavailable in source Haskell. This distinction becomes important in Section 5.5. The language also includes only annotated **let**-bindings; no **let**-generalization here. (The “generalization” you might spot in Rule LET is simply quantifying over the variables the user has lexically written in the type signature.) This keeps our treatment simple and avoids the challenges of type inference. Allowing full **let**-generalization and un-annotated **lets** changes none of the conclusions presented here.

The judgment $\Gamma \vdash e : \tau$ indicates that the term e has type τ in the context Γ , where Γ is a list of term variables and their (possibly polymorphic) types. Data constructors are globally fixed in an initial top-level context Γ_t ; it is assumed that any context Γ contains the global Γ_t binding data constructors.

The type-checking of possibly nested patterns, as they occur in a case statement, is offloaded to the judgment $\Gamma \vdash_p p : \sigma \Rightarrow \Gamma'$, which checks that p is a pattern for a value of type σ and possibly binds new term variables, which are added to Γ and returned in the extended environment Γ' . The auxiliary judgment $\Gamma \vdash_p^* \overline{p_i : \sigma_i^i} \Rightarrow \Gamma'$ straightforwardly threads the environment through a list of such pattern typings.

These rules should be unsurprising, but provide a baseline from which to build.

5.2 Support for GADTs

Now we extend this language with support for GADTs, with their existential type variables and equality constraints. See Fig. 4. The term syntax is unchanged, but polytypes now can mention constraints, which can either be empty (and elided from this text), an equality between two monotypes, or a conjunction of constraints. We leave the possibility open for additional constraints, as indicated by the ellipsis.

The environment Γ is extended with two new forms. First, we track the scope of type variables by adding $a : *$ to Γ ⁸. Second, we add constraints Q to Γ , to indicate which constraints (bound by a GADT pattern match) are in scope. Conversely, constraints are proved by the $\Gamma \Vdash Q$ entailment relation. As type inference and entailment is not the subject of this paper, we leave this relation abstract. The concrete instantiation of this judgment by, e.g., that of Vytiniotis et al.

⁸Haskell supports higher kinds, but we elide that here for simplicity, and assume that all type variables have kind $*$.

$$\begin{array}{c}
\boxed{\Gamma \vdash e : \tau} \quad \text{Expression typing} \\
\frac{\Gamma, x : \tau_1 \vdash e : \tau_2}{\Gamma \vdash \lambda x. e : \tau_1 \rightarrow \tau_2} \text{ ABS} \qquad \frac{\Gamma \vdash e_1 : \tau_1 \rightarrow \tau_2 \quad \Gamma \vdash e_2 : \tau_2}{\Gamma \vdash e_1 e_2 : \tau_2} \text{ APP} \qquad \frac{\Gamma \vdash e : v \quad \bar{c} = \text{ftv}(v)}{\Gamma, x : \forall \bar{c}. v \vdash e_2 : \tau} \text{ LET} \\
\frac{(\nu : \forall \bar{t}v. v) \in \Gamma}{\Gamma \vdash \nu : [\bar{t}v \mapsto \bar{\tau}]v} \text{ VARCON} \qquad \frac{\Gamma \vdash e : v \quad \overline{\Gamma \vdash_p p_i : v \Rightarrow \Gamma'_i} \quad \overline{\Gamma'_i \vdash e_i : \tau^i}}{\Gamma \vdash \text{case } e \text{ of } \{\bar{p}_i \rightarrow e_i^i\} : \tau} \text{ CASE} \\
\boxed{\Gamma \vdash_p p : \sigma \Rightarrow \Gamma'} \quad \text{Pattern typing} \\
\frac{}{\Gamma \vdash_p x : \sigma \Rightarrow \Gamma, x : \sigma} \text{ PATVAR} \qquad \frac{(\text{K} : \forall \bar{a}_i^i. \bar{\sigma}_k^k \rightarrow \text{T} \bar{a}_i^i) \in \Gamma \quad \overline{\Gamma \vdash_p^* p_k : [\bar{a}_i \mapsto \bar{\tau}_i^i] \sigma_k} \Rightarrow \Gamma'}{\Gamma \vdash_p \text{K} \bar{p}_k^k : \text{T} \bar{\tau}_i^i \Rightarrow \Gamma'} \text{ PATCON98} \\
\boxed{\Gamma \vdash_p^* \bar{p}_i : \bar{\sigma}_i^i \Rightarrow \Gamma'} \quad \text{Pattern sequence typing} \\
\frac{\overline{\Gamma_{i-1} \vdash_p p_i : \sigma_i \Rightarrow \Gamma_i}^{i \in 1..n}}{\Gamma_0 \vdash_p^* \bar{p}_i : \bar{\sigma}_i^{i \in 1..n} \Rightarrow \Gamma_n} \text{ PATSEQ}
\end{array}$$

Figure 3. Typing of Haskell98 patterns

$$\begin{array}{l}
\text{Updates to grammar:} \\
\text{Constraints} \quad Q ::= \epsilon \mid Q_1 \wedge Q_2 \mid \tau_1 \sim \tau_2 \mid \dots \\
\text{Polytypes} \quad \sigma ::= \forall \bar{a}. Q \Rightarrow \tau \\
\text{Type env} \quad \Gamma ::= \epsilon \mid \Gamma, \nu : \sigma \mid \Gamma, a : * \mid \Gamma, Q \\
\text{Types of data cons} \quad \Gamma_t = \text{K} : \forall \bar{a}. Q \Rightarrow \bar{\sigma} \rightarrow \text{T} \bar{\tau} \\
\text{Constraint entailment} \quad \Gamma \Vdash Q \\
\boxed{\Gamma \vdash e : \tau} \quad \text{Expression typing} \\
\frac{(\nu : \forall \bar{a}. Q \Rightarrow v) \in \Gamma \quad \Gamma \vdash e : \tau_1 \quad \Gamma \Vdash [\bar{a} \mapsto \bar{\tau}]Q}{\Gamma \vdash \nu : [\bar{a} \mapsto \bar{\tau}]v} \text{ VARCONQ} \qquad \frac{\Gamma \Vdash \tau_1 \sim \tau_2}{\Gamma \vdash e : \tau_2} \text{ EQ} \\
\frac{\overline{\Gamma \vdash_p p_i : v \Rightarrow \Gamma'_i} \quad \overline{\Gamma'_i \vdash e_i : \tau^i}}{\Gamma \vdash e : v \quad \text{ftv}(\tau) \subseteq \text{dom}(\Gamma)} \text{ CASETV} \\
\frac{\Gamma, \bar{a} : * \vdash e : [\bar{c} \mapsto \bar{a}]v \quad \bar{c} = \text{ftv}(v) \quad \bar{a} \# \text{dom}(\Gamma) \quad \Gamma, x : \forall \bar{c}. v \vdash e_2 : \tau}{\Gamma \vdash \text{let } x :: v = e \text{ in } e_2 : \tau} \text{ LETTV} \\
\boxed{\Gamma \vdash_p p : \sigma \Rightarrow \Gamma'} \quad \text{Pattern typing} \\
\frac{(\text{K} : \forall \bar{a}. Q \Rightarrow \bar{\sigma}_i^i \rightarrow \text{T} \bar{\nu}_j^j) \in \Gamma \quad \bar{a} \# \text{dom}(\Gamma) \quad \Gamma, \bar{a} : *, \bar{\nu}_j \sim \bar{\tau}_j^j, Q \vdash_p^* \bar{p}_i : \bar{\sigma}_i^i \Rightarrow \Gamma'}{\Gamma \vdash_p \text{K} \bar{p}_i^i : \text{T} \bar{\tau}_j^j \Rightarrow \Gamma'} \text{ PATCON}
\end{array}$$

Figure 4. Adding support for GADTs

[2011] would be appropriate in an implementation of this type system.

Support for GADTs can be seen in the new form of data constructor types, listed in Fig. 4. Note that the arguments to T in the return type are no longer confined to be \bar{a} , the quantified type variables; instead they can be arbitrary monotypes. In addition, a constructor can include a constraint Q.

When a data constructor is used in an *expression*, then the type equalities must be satisfied in the current environment, as expressed by the new premise of Rule VARCONQ in Fig. 4. We see also that the type equalities in the environment can be used for implicit coercion, as expressed in the Rule EQ.

When *pattern-matching* a data constructor, Rule PATCON brings the type variables \bar{a} into scope, by extending Γ . We require that these bound variables are fresh with respect to other variables in scope, a requirement we can satisfy by α -renaming if necessary. We also add the type equalities that we have learned to the environment—that is, the equivalence between the $\bar{\nu}_j^j$ from the data constructor’s type and the $\bar{\tau}_j^j$ from the pattern type.

Finally, we update Rule CASETV to prevent skolem escape and Rule LETTV to track the internal variables brought into scope. Note that these variables are *internal* only—the user cannot write them in a program.

At this point, our type system is comparable in expressiveness to the specification given by Vytiniotis et al. [2011]. A notable difference is that we explicitly handle nested patterns. This is important, as in the presence of GADTs, the precise formulation of how nested pattern are type-checked matters. For example, consider:

data $G\ a\ \text{where}$
 $G1 :: G\ \text{Bool}$
 $G2 :: G\ a$
 $f :: (G\ a, a, G\ a) \rightarrow \text{Bool}$
 $f\ (G1, \text{True}, _) = \text{False}$
 $f\ (_, \text{True}, G1) = \text{False}$

Here the first equation for f is fine, but the second is not, because the pattern `True` cannot match against an argument of type a until *after* the constructor $G1$ has been matched—and matching in Haskell is left-to-right.

5.3 Treating Universals and Existentials Uniformly

A technical contribution of this paper is that Rule PATCON is simpler and more uniform than the one usually given [e.g. by Vytiniotis et al. 2011], in that *it does not distinguish the universal and existential type variables of the data constructor*. Instead, *all* the type variables are freshly bound, with the equalities $\bar{v}_j \sim \bar{\tau}_j^j$ linking them to the context. In particular, these equalities take the place of the substitution written in the previous Rule CON98.

However, there is a worry: pattern-matching involving GADTs lacks principal types, and hence usually requires a type signature (see Section 2.1). If we treat vanilla, non-GADT Haskell98 data types in the same way as GADTs, do we lose type inference for ordinary Haskell98 definitions? Specifically, Vytiniotis et al. [2011, Section 5.6.1] describe how assumed local constraints can interfere with type inference, essentially by making certain unification variables “untouchable” (that is, unavailable for unification). That section also describes how to make more unification variables touchable in the non-GADT case, when the constraints entail no equalities. But *our typing rule introduces equalities even in the non-GADT case*, so this mechanism fails for us.

Let us investigate Rule PATCON specialized to the case of an ordinary, non-GADT constructor, which binds no context Q and does not constrain its result type arguments:

$$\frac{(\mathbb{K} : \forall \bar{a}_j^j. \bar{\sigma}_i^i \rightarrow \mathbb{T} \bar{a}_j^j) \in \Gamma \quad \bar{a} \# \text{dom}(\Gamma) \quad \Gamma, \bar{a}_j : *^j, \bar{a}_j \sim \bar{\tau}_j^j \vdash_p^* \bar{p}_i : \bar{\sigma}_i^i \Rightarrow \Gamma'}{\Gamma \vdash_p \mathbb{K} \bar{p}_i^i : \mathbb{T} \bar{\tau}_j^j \Rightarrow \Gamma'} \text{PATCON98}'$$

We see that all of its assumed equality constraints take the form $a_j \sim \tau_j$, where a_j is freshly bound. We can view such equalities not as true assumed equalities (which lead to the type inference problems for GADTs), but instead as a form of local **let**-binding: the context simply gives us the definition of these type variables. In this interpretation, it is critical that the type variable in the equality assumption is freshly bound—that is, we are not referring to a type variable from a larger scope. Viewing the equalities in Γ' as **let**-like, it is sensible to extend the ad-hoc extension of Vytiniotis et al. [2011] to include such forms. Indeed, doing so is an independently-useful improvement to type inference, and

$$\text{Patterns} \quad p ::= \dots \mid p :: \sigma$$

$$\frac{\text{ftv}(\sigma') = \emptyset \quad \Gamma \Vdash \sigma \leq \sigma' \quad \Gamma \vdash_p p : \sigma' \Rightarrow \Gamma'}{\Gamma \vdash_p (p :: \sigma') : \sigma \Rightarrow \Gamma'} \text{PATSIG}$$

Figure 5. Syntax and typing rule for pattern signatures

$$\text{Expressions} \quad e ::= \dots \mid \text{let } x :: \sigma = e_1 \text{ in } e_2$$

$$\frac{\sigma = \forall \bar{c}. Q \Rightarrow v \quad \text{ftv}(\sigma) \subseteq \text{dom}(\Gamma) \quad \Gamma, \bar{c} : *, Q \vdash e : v \quad \Gamma, x : \sigma \vdash e_2 : \tau}{\Gamma \vdash \text{let } x :: \sigma = e \text{ in } e_2 : \tau} \text{LETFORALL}$$

$$\frac{\bar{c} = \text{ftv}(\sigma') \setminus \text{dom}(\Gamma) \quad \Gamma' = \Gamma, \bar{c} : *, \bar{c} \sim \bar{\tau} \quad \text{isInternalTypeVar}(\tau) \quad \Gamma' \Vdash \sigma \leq \sigma' \quad \Gamma' \vdash_p p : \sigma' \Rightarrow \Gamma''}{\Gamma \vdash_p (p :: \sigma') : \sigma \Rightarrow \Gamma''} \text{PATSIGTV}$$

Figure 6. Typing with scoped type variables

GHC has already adopted it, in response to a request⁹ from one of this paper’s authors. Thus, despite the addition of equalities in Rule PATCON, we do not have a negative effect on type inference.

5.4 Closed Pattern Signatures

Our next step is to formalize *PatternSignatures*, which allows the user to annotate patterns with type signatures, but for now we will only handle *closed* pattern signatures. We simply add one new typing rule PATSIG, shown in Fig. 5. Note that the user is allowed to give the pattern a more specific type, as in this example, which requires *RankNTypes*:

$f :: (\forall a. a \rightarrow a) \rightarrow \text{Int}$
 $f\ (x :: \text{Int} \rightarrow \text{Int}) = x\ 42$

The typing rule expresses this through the premise $\Gamma \Vdash \sigma \leq \sigma'$, an appeal to the subtype relation on polytypes. This subtype relationship checks that the expected type of the pattern σ is more general than the annotated type σ' . Note that this relationship is *backwards* from the usual expected/actual relationship in typing because patterns are in a negative position. The subtleties of polytype subtyping are well explored in the literature¹⁰ and need not derail our exploration here. However, note that Rule PATSIG checks pattern p with the annotated type σ' , not the more general σ —after all, the user has asked us to use σ' .

5.5 Scoped Type Variables

Next, we add support for two features that can bring type variables into scope: *open* pattern signatures and **let** with an explicit \forall . The grammar now allows a polytype σ as the

⁹<https://ghc.haskell.org/trac/ghc/ticket/15009>

¹⁰GHC’s current implementation of subtyping is described by Eisenberg et al. [2016], who also cite other relevant publications on the subject.

annotation to a **let**-bound identifier. We additionally replace PATSIG with PATSIGTV and add LETFORALL in Fig. 6.

The LETFORALL rule allows programmers to bring variables \bar{c} into scope when an explicit \forall is mentioned in the source. Note that this rule does not do any implicit lexical generalization: echoing GHC’s behavior, if the user writes a \forall , *all* new variables to be used in the type signature must be bound explicitly.

In Rule PATSIGTV, the last two premises are identical to those of its predecessor PATSIG. The first premise extracts the type variables \bar{c} that are free in the user-written type signature σ' , but not already in scope in Γ . The “not already in scope” part reflects the discussion of Section 3.1.

But what if Γ contains a binding, introduced by rule PATCON, for a type variable that just happens to have the same name as one of the \bar{c} in a user-written signature? After all, the names of the type variables in PATCON are arbitrary internal names; they just need to be fresh. Our solution is simple: we take advantage of the difference between internal type variables and external ones. The user cannot accidentally capture an internal variable.

The second (top-right) premise of Rule PATSIGTV is the most unusual. It brings the variables \bar{c} into scope, but then *also* assumes that each variable c equals some other type τ ; the following premise asserts that each τ is, in fact, just an internal type variable b . Strikingly, the $\bar{\tau}$ are mentioned nowhere else in the rule. This setup essentially says that the \bar{c} are merely a renaming of existing in-scope internal variables. In practice, the $\bar{\tau}$ are chosen in order to make the subtyping relationship $\Gamma' \Vdash \sigma \leq \sigma'$ hold; GHC checks this subtyping relationship, unifying the \bar{c} with internal variables \bar{b} as necessary. Because the subtyping relationship is checked with respect to a context that contains the $\bar{c} \sim \bar{\tau}$ equalities, the $\bar{\tau}$ do not need to be explicitly mentioned again in the rule. For example, consider

```
data ExIntegral where -- packs an Integral value
```

```
  MkEx :: ∀a. Integral a ⇒ a → ExIntegral
```

```
getInt :: ExIntegral → Integer
```

```
getInt (MkEx (x :: c)) = toInteger @c x
```

The pattern match on *MkEx* brings an internal existential variable a into scope, via the PATCON rule. Recall that the user cannot type the name of such a variable. Instead, the user annotates the pattern x with the user-written type variable c . This annotation triggers Rule PATSIGTV, which must find a type τ such that $a : *$, $c : *$, $c \sim \tau \Vdash a \leq c$. The answer is that we must choose τ to be equal to the variable a , and the rule succeeds. We have thus renamed the internal variable a to become the user-written variable c and can successfully use c in the pattern’s right-hand side.

Contrast that behavior with this (failing) example:

```
notAVar :: Int → Int
```

```
notAVar (x :: c) = x
```

Patterns

$$p ::= x \mid K @ \tau \bar{p} \mid p :: \sigma$$

$$(K : \forall \bar{a}_j^j. Q \Rightarrow \bar{\sigma}_k^k \rightarrow T \bar{v}_i^i) \in \Gamma$$

$$\bar{c}_l^l = \text{ftv}(\bar{\tau}_j^j) \setminus \text{dom}(\Gamma) \quad \bar{a}_j^j \# \text{dom}(\Gamma)$$

$$\Gamma' = \Gamma, \bar{a}_j : *^j, \bar{c}_l : *^l, \bar{c}_l \sim \bar{\tau}_l^l, \bar{v}_i \sim \bar{\tau}_i^i, Q$$

$$\Gamma' \Vdash \bar{\tau}_j^j \sim \bar{a}_j^j \quad \text{isInternalTypeVar}(\bar{\tau}_l^l)$$

$$\Gamma' \Vdash_p^* \bar{p}_k : \bar{\sigma}_k^k \Rightarrow \Gamma''$$

$$\Gamma \vdash_p K @ \tau_j^j \bar{p}_k^k : T \bar{\tau}_i^i \Rightarrow \Gamma'' \quad \text{PATCONTYAPP}$$

Figure 7. Typing of type applications in patterns

Here, we are trying to bind a user-written type variable c to *Int*. GHC rejects this function, saying that c does not match with *Int*. In terms of Rule SIGPATTV, there exists no τ such that $c : *$, $c \sim \tau \Vdash \text{Int} \leq c$ and *isInternalTypeVar*(τ) holds.

There is a free design choice embodied in Rule SIGPATTV: the rule asserts that each c must be a renaming of a type *variable*. Instead, we could simply drop the *isInternalTypeVar*(τ), allowing each type variable to rename a *type*. Nothing else in the system would have to change. Indeed, understanding this very fact is one of the primary motivators for writing this specification in the first place.

5.6 Type Applications in Patterns

Having nailed down the status quo, it is now easy to specify what it should mean to use type applications in patterns. This version supports type applications only in *constructor* patterns; we study pattern synonyms in the appendix of the extended version¹¹. The syntax and new typing rule are shown in Fig. 7. Rule PATCONTYAPP looks scary, but it just integrates the concepts seen in Rule PATSIGTV into Rule PATCON. We have kept all the iteration indices to help the reader match up which lists are expected to have the same size.

Let us look at each premise separately:

- Once again, the type variables \bar{c}_l^l are those that occur in the explicit type patterns but are not yet in scope. These are treated like type variables in a pattern signature: they are brought into scope here, each as a short-hand for some type $\bar{\tau}_l^l$, as long as that type is an internal type variable.
- The environment Γ is extended to Γ' and contains now the (internal) type variables \bar{a}_j^j , the user-written scoped type variables \bar{c}_l^l , the type equations equating each c_l to its internal type variable b_l , the GADT equalities $\bar{v}_i \sim \bar{\tau}_i^i$, and the constraint Q captured by K .
- The type patterns are checked against the types they match against. In contrast to pattern signatures, we use type equality here (\sim), not the subtyping relation

¹¹<https://arxiv.org/abs/1806.03476>

(\leq): no types involved can be polytypes, and so the subtyping relation degenerates to type equality.

As written here, the rule requires a type application for each type variable (note that the $\overline{\tau}_j^j$ use the same indices as the quantified type variables \overline{a}_j^j in K 's type). However, we can weaken this requirement simply by dropping some τ_j^j 's from both the conclusion and the relevant premises.

Just as in Rule SIGPATV:

- The $\overline{\tau}_1^l$ are mentioned nowhere else in the rule; instead, they are fixed such that the equality constraints for the $\overline{\tau}_j^j$ are entailed by Γ' .
- The rule requires that each user-written type variable stands for an internal *variable*, but we can once again simply drop the $isInternalTypeVar(\tau_j^j)$ premise to relax this restriction.

5.7 Type Safety

At this point, after developing a set of inference rules defining a type system, one would normally prove that the language is type safe. We do not do so here. Not only would defining an operational semantics and writing out a proof distract us from our main point (the precise description and specification of the use of type variables in patterns), but it is also largely unnecessary. Let us assume that GHC/Haskell, without our new extension, is type safe. (See, e.g., [Sulzmann et al. \[2007\]](#) for a related proof.) If we compare PATCON to our new PATCONTYAPP, we see that the difference is only the new type variables brought into scope. Yet this same rule insists that these type variables are equal to existing types. In other words, the type variables are merely abbreviations or renamings of other types. Furthermore, the changes have no effect on operational behavior: the changes are all at compile-time. There appears to be no way that introducing such variables can cause a type system to lose safety—everything we have done here amounts only to syntactic convenience¹², thus obviating the need for a full-blown proof.

5.8 Conclusion

Through the incremental building of rules, we can see precisely how the new feature of explicit binding sites for type variables fits into the existing typing framework. We have also explored two further extensions¹³:

- Allowing type application in patterns headed by *pattern synonym* [[Pickering et al. 2016](#)]. Our framework extends well in this new context, offering no surprises.
- Incorporating explicit binding sites for type variables in the patterns of a λ -expression. This is slightly subtler

(though the end result adds only one, simple typing rule), but is relegated to the appendix because it requires reasoning about bidirectional type checking. Bringing all the necessary context into scope would take us too far afield here.

6 Alternative Approaches

6.1 Universals vs. Existentials

Type theorists are wont to separate quantified type variables in data constructors into two camps: *universals* and *existentials*. Here is a contrived but simple example:

data *UnivEx a where*

MkUE :: $\forall a b. a \rightarrow b \rightarrow \text{UnivEx } a$

matchUE :: $\forall a. \text{UnivEx } a \rightarrow \dots$

matchUE (*MkUE* x y) = ...

In the constructor *MkUE*, the variable a is universal (it is fixed by the return type *UnivEx a*) while b is existential (it is not fixed by the result type). When we match on *MkUE* in *matchUE*, we might want to bind b , as it is first brought into scope by the match. However, we never need to match a , as it is already in scope from *matchUE*'s type signature.

An alternative design for type applications in patterns is to allow matching only existentials in pattern matches, thus:

matchUE :: $\forall a. \text{UnivEx } a \rightarrow \dots$

matchUE (*MkUE* $@b$ x y) = ...

Indeed, this forms the main payload of the original GHC proposal for binding type variables [[Suarez 2017](#)]. This design is attractive because the bindings would be concise: only those variables that need to be bound would be available. However, there are two distinct drawbacks:

Universals and existentials are hard to differentiate.

Given the complexity of Haskell, the line between these two is blurry. Clearly, a is universal in the constructor *MkUE* above. But what if its type were *MkUE* :: $\forall a b. a \rightarrow b \rightarrow \text{UnivEx } (Id\ a)$, where *Id* is a type synonym? An injective type family? If we add $a \sim b$ to the constraints of *MkUE*, then b is also fixed by the result type—does that make it a universal?

The question of whether the value of a type variable is fixed by the return type depends on how smart the compiler is, and any specification would have to draw an arbitrary line. In the end, this would leave our users just very confused.

Universals can be instantiated in expressions. When using a data constructor in an *expression*, the caller is free to instantiate both universals and existentials. Indeed, universals and existentials are utterly indistinguishable in expressions. That means that one might write *MkUE @Int @Bool 5 True* in an expression. If we could match against only existentials in patterns, though, we would write a pattern *MkUE @b x y*, remembering to skip the universal a . This would both be confusing to users and weaken the ergonomics of patterns,

¹²In the case that a variable is ambiguous, such as the example in Section 3.6, our new features indeed change what is possible. However, this should be seen more as an infelicity of the way the previous binding structure worked than a new feature we are introducing.

¹³In the appendix of the extended version at <https://arxiv.org/abs/1806.03476>

whose chief virtue is that deconstructing a datatype resembles closely the syntax of constructing one.

We thus prefer not to differentiate universals and existentials in this way.

6.2 The Type-Lambda Approach

A plausible alternative approach to adding scoped type variables is to take a hint from System F, the explicitly-typed polymorphic lambda calculus [Girard 1990]. In System F, a *type lambda*, written “ Λ ”, binds a type variable, just as a term lambda, written “ λ ”, binds a term variable. For example:

$$\begin{aligned} id &: \forall \alpha. \alpha \rightarrow \alpha \\ id &= \Lambda \alpha. \lambda x : \alpha. x \end{aligned}$$

A term $\Lambda \alpha. e$ has type $\forall \alpha. \tau$, for some type τ , just as a term $\lambda x. e$ has type $\tau_1 \rightarrow \tau_2$. Hence, a very natural idea is to bind a source-language type variable with a source-language type lambda. This “the type-lambda approach” is the one adopted by SML 97 [Milner et al. 1997]. In SML one can write:

```
fun 'a prefix (x : 'a) yss =
  let fun xcons (ys : 'a list) = x :: ys in
  map xcons yss
```

Here, “’a” following the keyword `fun` is the binding site of an (optional) type parameter of `prefix`; it scopes over the patterns of the definition and its right hand side.

Just as Haskell has implicit quantification in type signatures, SML allows the programmer to introduce *implicit* type lambdas. This definition is elaborated into the previous one:

```
fun prefix (x : 'a) yss =
  let fun xcons (ys : 'a list) = x :: ys in
  map xcons yss
```

The language definition gives somewhat intricate rules to explain how to place the implicit lambdas. For example:

```
fun f x = ....(fun (y:'a) => y)....
```

Where is the type lambda that binds the type variable ‘a’? In SML one cannot answer that question without knowing both what the “....” is, and the context for the definition `fun f`. Roughly speaking, the type lambda for an implicitly-scoped type variable ‘a’ is placed on the innermost function definition that encloses all the free occurrences of ‘a’. The rule [Milner et al. 1997] is only one informal, albeit carefully worded, paragraph; the formal typing rules assume that a pre-processing pass has inserted an explicit binding for every type variable that is implicitly bound by the above rule.

The type-lambda approach explicitly connects *lexical scoping* and *quantification*. In contrast, our approach presented here decouples the two, by treating a lexically scoped type variable merely as an alias for a type (or type variable). The

appendix, included in the extended version of this paper, additionally has the details of an extension of this work to include explicit binding of type variables in λ -expressions, providing a similar experience to what we see in SML above.

Acknowledgments

This material is based upon work supported by the National Science Foundation under Grant No. 1319880, Grant No. 1521539, and Grant No. 1704041.

References

- Joachim Breitner. 2018. Allow ScopedTypeVariables to refer to types. GHC proposal. <https://github.com/ghc-proposals/ghc-proposals/pull/128>
- Manuel M. T. Chakravarty, Gabriele Keller, Simon L. Peyton Jones, and Simon Marlow. 2005. Associated types with class. In *POPL*. ACM, 1–13. <https://doi.org/10.1145/1040305.1040306>
- Richard A. Eisenberg, Dimitrios Vytiniotis, Simon L. Peyton Jones, and Stephanie Weirich. 2014. Closed type families with overlapping equations. In *POPL*. ACM, 671–684. <https://doi.org/10.1145/2535838.2535856>
- Richard A. Eisenberg, Stephanie Weirich, and Hamidhasan G. Ahmed. 2016. Visible Type Application. In *ESOP (LNCS)*, Vol. 9632. Springer, 229–254. https://doi.org/10.1007/978-3-662-49498-1_10
Extended version at <https://cs.brynmawr.edu/~rae/papers/2016/type-app/visible-type-app-extended.pdf>.
- J-Y Girard. 1990. The System F of variable types: fifteen years later. In *Logical Foundations of Functional Programming*, G Huet (Ed.). Addison-Wesley.
- Robin Milner, Mads Tofte, Robert Harper, and David MacQueen. 1997. *The Definition of Standard ML (Revised)*. MIT Press, Cambridge, Massachusetts. xii+114 pages.
- Simon L. Peyton Jones, Dimitrios Vytiniotis, Stephanie Weirich, and Mark Shields. 2007. Practical type inference for arbitrary-rank types. *Journal of Functional Programming* 17, 1 (2007), 1–82. <https://doi.org/10.1017/S0956796806006034>
- Simon L. Peyton Jones, Stephanie Weirich, Richard A. Eisenberg, and Dimitrios Vytiniotis. 2016. A Reflection on Types. In *A List of Successes That Can Change the World - Essays Dedicated to Philip Wadler on the Occasion of His 60th Birthday (LNCS)*, Vol. 9600. Springer, 292–317. https://doi.org/10.1007/978-3-319-30936-1_16
- Matthew Pickering, Gergo Érdi, Simon L. Peyton Jones, and Richard A. Eisenberg. 2016. Pattern synonyms. In *Haskell Symposium*. ACM, 80–91. <https://doi.org/10.1145/2976002.2976013>
- Tom Schrijvers, Simon L. Peyton Jones, Martin Sulzmann, and Dimitrios Vytiniotis. 2009. Complete and decidable type inference for GADTs. In *ICFP*. ACM, 341–352. <https://doi.org/10.1145/1596550.1596599>
- Jan Stolarek, Simon L. Peyton Jones, and Richard A. Eisenberg. 2015. Injective type families for Haskell. In *Haskell Symposium*. ACM, 118–128. <https://doi.org/10.1145/2804302.2804314>
- Emmanuel Suarez. 2017. Binding existential type variables. GHC proposal. <https://github.com/ghc-proposals/ghc-proposals/pull/96>
- Martin Sulzmann, Manuel M. T. Chakravarty, Simon L. Peyton Jones, and Kevin Donnelly. 2007. System F with type equality coercions. In *TLDI*. ACM, 53–66. <https://doi.org/10.1145/1190315.1190324>
- Dimitrios Vytiniotis, Simon L. Peyton Jones, Tom Schrijvers, and Martin Sulzmann. 2011. OutsideIn(X) Modular type inference with local assumptions. *Journal of Functional Programming* 21, 4-5 (2011), 333–412. <https://doi.org/10.1017/S0956796811000098>