

## COMP 105 Assignment: Hindley-Milner Type Inference

Due **Monday**, November 21 at 11:59PM.

In this assignment you will implement Hindley-Milner type inference, which represents the current "best practice" for flexible static typing. The assignment has two purposes:

- To help you develop a deep understanding of type inference
- To help you continue to build your ML programming skills

### Setup

Make a clone of the book code:

```
git clone linux.cs.tufts.edu:/comp/105/build-prove-compare
```

The code you need is in `bare/nml/ml.sml`.

### Individual Problems

*Working on your own*, please solve Exercise 1 on page 528 and Exercise 2 on page 528 of Ramsey. These exercises explore some implications of type inference.

#### Problem Details

1. *Exploring the meaning of polymorphic types I*. Do Exercise 1 on page 528 of Ramsey.
2. *Exploring the meaning of polymorphic types II*. Do Exercise 2 on page 528 of Ramsey.

### Pair Problems

*Working with a partner*, please solve 18, 19 and 20 from pages 531–532 of Ramsey, and the two exercises S and T below.

#### Problem Details

**18. Implementing a constraint solver.** Do Exercise 18 on page 531 of Ramsey. This exercise is probably the most difficult part of the assignment. *Before proceeding with type inference, make sure your solver produces the correct result on our test cases and on your test cases.* You may also want to show your solver code to the course staff before proceeding to type inference.

**19. Implementing type inference.** Do Exercise 19 on page 532 of Ramsey.

**20. Adding primitives.** Do Exercise 20 on page 532 of Ramsey.

**S. Test cases for the solver.** Write three test cases for the constraint solver. At least two of these test cases should be constraints that have no solution. Assuming that *we provide a function* `constraintTest : con -> answer`, write your test cases as three successive calls to `constraintTest`. Do *not* define `constraintTest` yourself.

Here is a sample set of test cases:

```
val _ = constraintTest (TYVAR "a" ~ TYVAR "b")
val _ = constraintTest (CONAPP (TYCON "list", [TYVAR "a"]) ~ TYCON "int")
val _ = constraintTest (TYCON "bool" ~ TYCON "int")
```

Naturally, you will supply your own test cases, different from these.

## COMP 105 Assignment: nML Type Inference

**T. Test cases for type inference.** Write three test cases for type inference. At least two of these test cases should be for terms that fail to type check. Each test case should be a definition written in nML. Here is a sample set of test cases:

```
(val weird (lambda (x y z) (cons x y z)))  
(+ 1 #t)  
(lambda (x) (cons x x))
```

Naturally, you will supply your own test cases, different from these.

---

### Extra Credit

For extra credit, you may complete any of the following:

- Mutation, as in Exercise 23(a)(b) and possibly (c)
- Explicit types, as in Exercise 25
- Better error messages, as in Exercise 24(a)(b) and possibly (c)
- Tuples, as in Exercise 22

Of these exercises the most interesting are probably Mutation (easy) and Explicit types (not easy).

---

### What to submit: Individual Problems

You should submit two files:

- a README file containing
  - ◆ the names of the people with whom you collaborated
  - ◆ the number of hours you worked on the assignment.
- a file `meaning.nml` containing your code for Exercise 1 on page 528 and Exercise 2 on page 528. Your answers to Exercise 2 should appear in a comment.

### How to submit: Individual Problems

When you are ready, run `submit105-ml-inf-solo` to submit your work. Note you can run this script multiple times; we will grade the last submission.

### What to submit: Pair Problems

You should submit four files:

- a README file containing
  - ◆ the names of the people with whom you collaborated
  - ◆ the numbers of any extra credit problems you solved
  - ◆ the number of hours you worked on the assignment.
- `ml.sml`, implementing a complete interpreter for nano-ML which includes your answers to Exercises 18, 19, and 20.
- `stest.sml`, containing your answer to Exercise S
- `ttest.nml`, containing your answer to Exercise T

### How to submit: Pair Problems

When you are ready, run `submit105-ml-inf-pair` to submit your work. Note you can run this script multiple times; we will grade the last submission.

## COMP 105 Assignment: nML Type Inference

### Hints, guidelines, and test code

This is one assignment where it pays to run a lot of tests, of both good and bad definitions. *The most effective test of your algorithm is not that it properly assign types to correct terms, but that it reject ill-typed terms.* I have posted a [functional topological sort](#) that makes an interesting test case.

If you call your interpreter `ml .sml`, you can build a standalone version in `a.out` by running `mosmlc ml .sml` or a faster version in `ml` by running `mlton -output a.out ml .sml`. To help you with the solver, once you have implemented `solve`, the following code redefines `solve` into a version that checks itself for sanity (ie, idempotence). It is a good idea to check that the substitution returned by your solver is idempotent before using it in your type inferencer.

```
fun isStandard pairs =
  let fun distinct a' (a, tau) = a <> a' andalso not (member a' (freetyvars tau))
      fun good (prev', (a, tau)::next) =
          List.all (distinct a) prev' andalso List.all (distinct a) next
          andalso good ((a, tau)::prev', next)
      | good (_, []) = true
  in good ([], pairs)
  end

val solve =
  fn c => let val theta = solve c
          in if isStandard theta then theta
             else raise BugInTypeInference "non-standard substitution"
          end
```

In writing the type-inference code, you should refer to the typing rules of nml, which appear on pages 500-501 of Ramsey. With your solver in place, the type inference should be straightforward, with two exceptions: `let` and `letrec`. You can emulate the implementations for `val` and `val-rec`, but **you must split the constraint** into local and global portions. The splitting is covered in detail in the book, on pages 469–472, with particular focus in the sidebar on page 470.

### Testing

The course interpreter is located in `/comp/105/bin/nml`. If your interpreter can process the initial basis and infer correct types, you are doing OK.

The *real* test of your interpreter is that it should reject incorrect definitions. You should prepare a dozen or so definitions that should not type check, and make sure they don't. For example:

```
(val bad (lambda (x) (cons x x)))
(val bad (lambda (x) (cdr (pair x x))))
```

Pick your toughest three test cases to submit for Exercise T.

### Avoid common mistakes

Here some common mistakes:

- A common mistake is to create **too many fresh variables** or to fail to constrain your fresh variables.
- Another surprisingly common mistake is to include **redundant cases** in the code for inferring the type of a list literal. As is almost always true of functions that consume lists, it's sufficient to write one case for `NIL` and one case for `PAIR`.
- It's a common mistake to **define a new exception and not handle it**. If you define any new exceptions, make sure they are handled. It's not acceptable for your interpreter to crash with an unhandled exception just because some nano-ML code didn't type-check.
- It's a common mistake to omit the initial basis for testing and then to **forget to include an initial basis in the interpreter you submit**.

## COMP 105 Assignment: nML Type Inference

There are also some common assumptions which are mistaken:

- It is a mistake to assume that an element of a literal list always has a monomorphic type.
- It is a mistake to assume that `begin` is never empty.

### How your work will be evaluated

Your solutions are going to be evaluated automatically. We must be able to compile your solution in Moscow ML by typing, e.g.,

```
mosmlc ml.sml
```

If there are errors or warnings in this step, your work will earn No Credit for functional correctness.

We will focus most of our evaluation on your constraint solving and type inference.

	Exemplary	Satisfactory	Must improve
Form	<ul style="list-style-type: none"> <li>• The code has no offside violations.</li> <li>• <i>Or</i>, the code has just a couple of minor offside violations.</li> <li>• Indentation is consistent everywhere.</li> <li>• The submission has no bracket faults.</li> <li>• The submission has a few minor bracket faults.</li> <li>• <i>Or</i>, the submission has no bracketed names, but a few bracketed conditions or other faults.</li> </ul>	<ul style="list-style-type: none"> <li>• The code has several offside violations, but course staff can follow what's going on without difficulty.</li> <li>• In one or two places, code is not indented in the same way as structurally similar code elsewhere.</li> <li>• The submission has some redundant parentheses around function applications that are under infix operators (not checked by the bracketing tool)</li> <li>• <i>Or</i>, the submission contains a handful of bracketing faults.</li> <li>• <i>Or</i>, the submission contains more than a handful of bracketed names or conditions.</li> </ul>	<ul style="list-style-type: none"> <li>• Offside violations make it hard for course staff to follow the code.</li> <li>• The code is not indented consistently.</li> <li>• The submission contains more than a handful of parenthesized names as in <code>(x)</code></li> <li>• The submission contains more than a handful of parenthesized <code>if</code> conditions.</li> </ul>
Names	<ul style="list-style-type: none"> <li>• Type variables have names beginning with <code>t</code>; types have names beginning with <code>τ</code> or <code>tau</code>; constraints have names beginning with <code>c</code>; substitutions have names beginning with <code>theta</code>; lists of things have names that begin conventionally and end in <code>s</code>.</li> </ul>	<ul style="list-style-type: none"> <li>• Types, type variables, constraints, and substitutions mostly respect conventions, but there are some names like <code>x</code> or <code>l</code> that aren't part of the typical convention.</li> </ul>	<ul style="list-style-type: none"> <li>• Some names misuse standard conventions; for example, in some places, a type variable might have a name beginning with <code>τ</code>, leading a careless reader to confuse it with a type.</li> </ul>
Structure	<ul style="list-style-type: none"> <li>• The nine cases of simple type equality are handled by these five patterns: <code>TYVAR/any</code>, <code>any/TYVAR</code>, <code>CONAPP/CONAPP</code>, <code>TYCON/TYCON</code>, <code>other</code>.</li> <li>• The code for solving <math>\hat{I} \pm \hat{a} - \hat{a} \ \hat{1} \hat{4} - \hat{I}</math> has exactly three cases.</li> </ul>	<ul style="list-style-type: none"> <li>• The nine cases are handled by nine patterns: one for each pair of value constructors for <code>ty</code></li> <li>• The code for <math>\hat{I} \pm \hat{a} - \hat{a} \ \hat{1} \hat{4} - \hat{I}</math> has more than three cases, but the nontrivial cases all look different.</li> </ul>	<ul style="list-style-type: none"> <li>• The case analysis for a simple type equality does not have either of the two structures on the left.</li> <li>• The code for <math>\hat{I} \pm \hat{a} - \hat{a} \ \hat{1} \hat{4} - \hat{I}</math> has more than three cases, and different nontrivial cases share duplicate or near-duplicate code.</li> </ul>

## COMP 105 Assignment: nML Type Inference

<ul style="list-style-type: none"><li>• The constraint solver is implemented using an appropriate set of helper functions, each of which has a good name and a clear contract.</li><li>• Type inference for list literals has no redundant case analysis.</li><li>• Type inference for expressions has no redundant case analysis.</li><li>• In the code for type inference, course staff see how each part of the code is necessary to implement the algorithm correctly.</li><li>• Wherever possible appropriate, submission uses <code>map</code>, <code>filter</code>, <code>foldr</code>, and <code>exists</code>, either from <code>List</code> or from <code>ListPair</code></li></ul>	<ul style="list-style-type: none"><li>• The constraint solver is implemented using too many helper functions, but each one has a good name and a clear contract.</li><li>• The constraint solver is implemented using too few helper functions, and the course staff has some trouble understanding the solver.</li><li>• Type inference for list literals has one redundant case analysis.</li><li>• Type inference for expressions has one redundant case analysis.</li><li>• In some parts of the code for type inference, course staff see some code that they believe is more complex than is required by the typing rules.</li><li>• Submission sometimes uses a <code>fold</code> where <code>map</code>, <code>filter</code>, or <code>exists</code> could be used.</li></ul>	<ul style="list-style-type: none"><li>• Course staff cannot identify the role of helper functions; course staff can't identify contracts and can't infer contracts from names.</li><li>• Type inference for list literals has more than one redundant case analysis.</li><li>• Type inference for expressions has more than one redundant case analysis.</li><li>• Course staff believe that the code is significantly more complex than what is required to implement the typing rules.</li><li>• Submission includes one or more recursive functions that could have been written without recursion by using <code>map</code>, <code>filter</code>, <code>List.exists</code>, or a <code>ListPair</code> function.</li></ul>
---	---	--

Back to the [class home page](#)