

Homework 7

Due Sunday, Apr 26 by 11:59pm

Handout 20
CSCI 334: Spring 2020

Turn-In Instructions

For text response questions, provide a \LaTeX source file and pre-built PDF. For example, for question 3, provide both “q3.tex” and “q3.pdf” files. For full credit, your \LaTeX file should build properly. Be sure to `git add` all necessary files (e.g., images) if your \LaTeX depends on it.

For each coding question in this assignment, create a project directory. For example, the source directory for question 1 should be in a folder called “q1”. You should be able to `cd` into this directory and then run the program by typing the command “`dotnet run`”. Each program should be split into two pieces: a “Program.fs” file that contains the `main` method and associated program-startup routines (like argument parsers), and another “Library.fs” file that contains the function(s) of interest in the question. All library code should be in a module named “CS334”. Be sure to provide usage output (defined in `main`) for all programs that require arguments. For full credit, your program should both build and run correctly.

Turn in your work using the Github repository assigned to you. The name of the Github repository will have the form `cs334hw7_<your user name>`. For example, my repository would be `cs334hw7_dbarowy`. You should have received an invite to commit in the repository in your email. If you did not receive this email, please contact me right away!

A good way to check that you submitted your assignment correctly is to `git clone` your repository to a new folder and then try building/running everything.

Pair Programming and Honor Code: You may optionally collaborate with one other person on the submission for this assignment. If you work with a partner, choose one of the two partner repositories for your submission. In the other repository, be sure to leave a `collaborators.txt` file that states who you worked with and the name of the repository where the code may be found. If you would like me to pair you with a partner, please let me know.

This assignment is due on Sunday, Apr 26 by 11:59pm.

Reading

1. **(Required)** “Parser combinators” from course packet.
2. **(Required)** “What is Logo?” from https://el.media.mit.edu/logo-foundation/what_is_logo/index.html, particularly the section “A Logo Primer.”
3. **(Optional)** “Appendix B: Original SML Specification” from course packet.

Problems

Q1. (30 points) Parsing with Combinators

(a) Given the following grammar in Backus-Naur Form,

```
<expr> ::= <var>
        | <abs>
        | <app>
<var>   ::=  $\alpha \in \{a \dots z\}$ 
<abs>   ::= (L<var>.<expr>)
<app>   ::= (<expr><expr>)
```

and the algebraic data type,

```
type Expr =  
  | Variable of char  
  | Abstraction of char * Expr  
  | Application of Expr * Expr
```

provide an implementation for the parser function

```
parse(s: string) : Expr option
```

In other words, given a `string s` representing a valid expression, `parse` should return `Some` abstract syntax tree (the `Expr`) of the expression. When given a `string s` that is not a valid expression, `parse` should return `None`.

You may use any of the combinator functions defined in the assigned reading on parser combinators in your solution. You may find the `pletter`, `pchar`, `pstr`, `pseq`, `pbetween`, `pleft`, `peof`, `<|>`, and `|>>` combinators to be the most useful.

Note that because your expression parser will be recursive, we need to do a little bookkeeping to keep F# happy. The first parser that appears your implementation should be written like:

```
let expr, exprImpl = recparser()
```

`recparser` defines two things: a declaration for a parser called `expr` and an implementation for that same parser called `exprImpl`. Later, once you have defined all of the parsers that `expr` depends on, write:

```
exprImpl := (* your expr parser implementation here *)
```

Note: use of `recparser` is admittedly a little bit of a hack. We use it because F# does not like us to use functions before we define them, which means that F# gets unhappy when it notices that we define recursive functions in terms of other recursive functions. `recparser` is one way around this problem. You should only need to use `recparser` once in this problem. To be *crystal clear*, your code should probably have at least the following definitions in it:

```
let expr, exprImpl = recparser()  
let variable : Parser<Expr> = (* variable parser implementation *)  
let abstraction : Parser<Expr> = (* abstraction parser implementation *)  
let application : Parser<Expr> = (* application parser implementation *)  
exprImpl := (* expr parser implementation *)
```

- (b) Provide a function `prettyprint(e: Expr) : string` that turns an abstract syntax tree into a string. For example, the program fragment

```
let asto = parse "((Lx.x)(Lx.y))"  
match asto with  
| Some ast -> printfn "%A" (prettyprint ast)  
| None      -> printfn "Invalid program."
```

should print the string

```
Application(Abstraction(Variable(x), Variable(x)), Abstraction(Variable(x), Variable(y)))
```

- (c) Be sure to document all of your functions using XML comment blocks.
- (d) (Optional) For a fun challenge, provide a second implementation (copy your implementation above to a new directory called `bonus` and modify it) that does one or more of the following:

- i. Accepts arbitrary amounts of whitespace between elements.
- ii. Accepts λ in addition to L for abstractions.
- iii. Allows the user to omit parens when the lambda calculus' rules of precedence and associativity allow the expression to be parsed unambiguously.

The last item is a bit challenging, but it is quite satisfying to produce a parser that can read in expression just as they are written in the course packet. If you're looking for a challenge, give it a try!

The project directory for this question should be called "q1". You should be able to run your program on the command line by typing, for example, "`dotnet run "((Lx.x)(Lx.y))"`" and output like the kind shown above should be printed to the screen.

Q2. (10 points) Five 5's (Bonus)

Consider the well-formed arithmetic expressions using the numeral 5. These are expressions formed by taking the integer literal 5, the four arithmetic operators +, -, *, and /, parentheses. Examples are 5, 5 + 5, and (5 + 5) * (5 - 5 / 5). Such expressions correspond to binary trees in which the internal nodes are operators and every leaf is a 5. Write an F# program that answers each of the following questions:

- (a) What is the smallest positive integer that cannot be computed by an expression involving exactly five 5's?
- (b) What is the largest prime number that can be computed by an expression involving exactly five 5's?
- (c) Generate an expression that evaluates to that prime number.

You should start by defining an algebraic data type (i.e., a **type**) to represent arithmetic expressions involving the number 5 and the arithmetic operations, as well as an **eval** function that evaluates them. This question involves only integer arithmetic, so be sure to use **int** data. You should then write a function that generates all possible expressions that contain a given number of 5's.

Answering the questions will involve an exhaustive search for all numbers that can be computed with a fixed number of 5's, so good programming techniques are important. Avoid unnecessary operations. You will also have to address the possibility of division by zero inside **eval** using an exception handler.

Be sure to document your code using (* F# comment blocks *).

The project directory for this question should be called "q2". You should be able to run your program on the command line by typing, for example, "`dotnet run`" and the program's output should clearly indicate the answers to the questions above.

Q3. (60 points) Project Proposal

For this question, you will propose your final project: a programming language of your own design. You are strongly encouraged to work with a partner on this assignment, however you will be permitted to work by yourself if you feel up to the extra challenge.

At this stage, your proposal is still non-binding, however, you should proceed as if this is the language that you want to implement as your final project.

Many programming language specifications begin as informal proposals, and this is the template that we will follow in this class. With each stage of your project, you will revisit your document, making it clearer and more precise as you work. It will be a "living document." By the end of this class, your final specification should include a formal syntax and an informal, but precise, description of your language's semantics. It should clearly document your software artifact, which will be an interpreter for the language.

Structure of the Proposal

Version 1.0 of your specification should include the following sections. Please explicitly include these sections. The purpose of this document is to *convince yourself* that such a language implementation is possible. If you aren't convinced, add more detail to convince yourself!

- (a) Introduction
2+ paragraphs. What problem does your language solve? Why does this problem need its own programming language?
- (b) Design Principles
1+ paragraphs. Languages can solve problems in many ways. What are the guiding aesthetic or technical principles that underpin its design?
- (c) Examples
3+ examples. Keeping in mind that your syntax is informal, sketch out 3 or more sample programs in your language.
- (d) Language Concepts
1+ paragraphs. What are the core concepts a user needs to understand in order to write programs? Think in terms of both “primitives” and “combining forms” (think back to Hughes’ use of this term and also our use of it in the language of parser combinators). What are the key ideas and how are they combined?
- (e) Syntax
As much as is needed. Sketch out the syntax of the language. For now, this can be an English description of the key syntactical elements and how they fit together. We will eventually transform this into a formal syntax section written in Backus-Naur Form.
- (f) Semantics
5+ paragraphs. How is your program interpreted? This need not be formal yet, however, you should demonstrate that you’ve thought about how your program will be represented and evaluated on a computer. It should answer the following questions.
 - i. What are the primitive kinds of values in your system? For example, a primitive might be a number, a string, a shape, a sound, and so on.
 - ii. What are the “actions” or compositional elements of your language? In other words, how are values used? For example, your system might have “commands” like “move pen” or compositional elements like “sequence of notes.”
 - iii. How is your program represented? If it helps you to think about this using ML algebraic data types, please use them. Otherwise, rough sketches such as class hierarchy drawings are fine.
 - iv. How do elements “fit together” to represent programs as abstract syntax? For the three example programs you gave earlier, provide sample abstract syntax trees.
 - v. How is your program evaluated? In particular,
 - A. Do programs in your language read any input?
 - B. What is the effect (output) of evaluating a program?
 - C. Evaluation is usually conceived of as a depth-first, post-order traversal of an AST. Describe how such a traversal yields the effect you just described and *provide illustrations for clarity*. Demonstrate evaluation for at least one of your example programs.

Goals

Your language need not be (and I discourage you from trying to build) a Turing-complete programming language. Instead, focus on solving a small class of problems. In other words, design a *domain-specific programming language*.

Your project must *eventually* achieve all of the following objectives:

- (a) It should have a grammar capable of generating either an infinite or practically-infinite number of possible programs.
- (b) It should have a parser that recognizes a grammatically-correct program, outputting the corresponding abstract syntax tree.

- (c) It should have an evaluator capable of “running” any valid program.
- (d) The language should do useful computational work.

If it is convenient to solve your problem by reducing to or extending a lambda calculus interpreter, you may propose such a solution. However, in most cases, it will likely be easier to design a purpose-built interpreter.

Sample Specifications

The following documents, described above in the readings section, are useful to see how others structure their language specifications. Read these as necessary in order to get a feel for your proposal. Note that Standard ML is a large, extensively-developed general-purpose language and therefore has long and very detailed specification. Logo has short and imprecise specification as it is a special-purpose language.

Your specification should be as long as is necessary and no longer. Writing good technical documentation is a real art, and we can only scratch the surface in this class. In general, try to be concise and precise, but not so concise or precise that you sacrifice understandability. The Logo documentation, especially the section titled “A Logo Primer” is an example of lucid, but informal, technical writing that I think captures the ideas of the language beautifully, and without any formal specification.

How to Organize

The project directory for this question should be called “q3”. You must use \LaTeX for your specification.