

# COMP105 Assignment: Operational Semantics

Due Friday, February 1 at 5:59PM, **six hours before midnight**.

The purpose of this assignment is to help you develop rudimentary skills with operational semantics, inference rules, and syntactic proof technique. You will use these skills heavily throughout the first two-thirds of the course, and you will use them again later if you ever want to keep up with the latest new ideas in programming languages or if you want to go on to advanced study.

Some of the essential skills are

- Understanding what judgment forms mean, as developed by reading and writing them
- Understanding what constitutes a valid syntactic proof, known as a *derivation*
- Understanding how a valid derivation in the operational semantics relates to the successful evaluation of an expression
- Proving facts about families of programs by reasoning about proofs, a technique known as *metatheory*
- Using operational semantics to express language features and language-design ideas
- Connecting operational semantics with informal English explanations of language features
- Connecting operational semantics with code in compilers or interpreters

Few of these skills can be mastered in a single assignment. When you've completed the assignment, I hope you will feel confident of your knowledge of exactly the way judgment forms, inference rules, and derivations are written. On the other skills, you'll have made a start.

## Part A: Adding Local Variables to the Interpreter (Work with a partner, 25 percent)

This exercise will help you understand how language changes can be realized in C code. You will do Exercise 26 from page 58 of Ramsey's book. **We recommend that you solve this problem with a partner, but this solution must be kept separate from your other solutions. Your programming partner, if any, must not see your other work.**

- Get your copy of the code from the book by running

```
git clone linux.cs.tufts.edu:/comp/105/book-code
```

or if that doesn't work, from a lab or linux machine, try

```
git clone /comp/105/book-code
```

You can find the source code from Chapter 2 in subdirectory `bare/impcore` or `commented/impcore`. The `bare` version, which we recommend, contains just the C code from the book, with simple comments identifying page numbers. The `commented` version, which you may

use if you like, includes part of the book text as commentary.

- We provide new versions of `all.h`, `ast.c`, `toplevel-code.c` and `parse.c` that handle local variables. These version are found in subdirectory `bare/impcore-with-locals`. There are not many changes; to see what is different, try running

```
diff -r bare/impcore bare/impcore-with-locals
```

You may wish to try the `-u` or `-y` options with `diff`

In the directory `bare/impcore-with-locals`, you can build an interpreter by typing `make`, but when you run the interpreter, it will halt with an assertion failure. You'll need to change the interpreter to add local variables:

- In `impcore.c`, you will have to modify the functions in the initial basis to use the new syntax.
  - In `eval.c`, you will have to modify the evaluator to give the right semantics to local variables. Local variables that have the same name as a formal parameter should hide that formal parameter, as in C.
  - You also have the right to modify other files as you see fit.
- Create a file called `README` in this directory (your `impcore-with-locals` directory). Use this file to describe your solution to this problem.

## Part B: Language structure and operational semantics (Individual work, 75 percent)

These are exercises intended to help you become fluent with operational semantics. **Do not share your solutions with any programming partners.** We encourage you to discuss ideas, but **nobody else may see your rules, your derivations, or your code.**

If you have difficulty, find a TA, who can help you work a couple of similar problems.

For these exercises you will turn in two files: `theory.pdf` and `14.imp`. For file `theory.pdf`, you will probably find it easiest to write your answers on paper and [scan it](#). Please see the [note about how to organize your answers](#).

- Do Exercises 13 and 14 on page 56 of Ramsey's book. The purpose of these exercises is to give you a feel for the kinds of choices language designers can make. Please include your answer to exercise 13 as part of `theory.pdf`. Please put your answer to exercise 14 in file `14.imp`.
- Do Exercise 10 on page 55 of Ramsey's book. The purpose of the exercise is to help you develop your understanding of derivations, so be sure to make your derivation *complete* and *formal*. You can write out a derivation with `BEGIN ... => 3 ...` in its conclusion, or if you prefer, you can write a sequence of judgments, and say for each one what rules and what previous judgments justify that judgment.

Please include your answer as part of file `theory.pdf`.

- Do Exercise 11 on page 55 of Ramsey's book. Now that you know how to *write* a derivation, in this exercise you start *reasoning* about derivations.

Please include your answer as part of file `theory.pdf`.

- Do Exercise 20 on pages 56–57 of Ramsey's book. In this exercise you raise your game again, reasoning about the *set* of all *valid* derivations. When you have got your thinking to this level, you can see how language designers use operational semantics to show nontrivial properties of their languages—and how these properties can guide implementors.

Please include your answer in file `theory.pdf`.

Metatheoretic proofs are probably unfamiliar, so you may want to look at some [sample cases](#) we have provided to help you. Also, to relieve some of the tedium (which is very common in programming-language proofs), we suggest that you allow your proof for the `AddApply` case to stand in for all other cases involving primitive operators. We also suggest that you simplify by leaving out the global environment `xi`.

## Organizing the answers to Part B

To help us read your answers to Part B, we need for you to organize them carefully:

- The answer to each question must **start on a new page**.
- The answers **must appear in order**: Exercises 10, 11, 13, and finally 20.

## Submitting

Before submitting your code, **test it**. We do not provide any tests; you must write your own.

- To submit part B, which you will have done by yourself, change into the appropriate directory and run `submit105-opsem-solo` to submit your work. In addition to files `14.imp` and `theory.pdf`, please also include a file called `README`. Use your `README` file to
  - Tell us *how long it took you to complete the entire assignment* (parts A and B)
  - Tell us how well you think you did on each of the four evaluation dimensions for operational semantics below: Semantics, Rules, Derivations, and Metatheory.
  - Tell us anything else you think it useful for us to know.

(If you wish to use PDF, then please submit `README.pdf` instead of `README`.)

- To submit part A, which you will have done with a partner, change into `bare/impcore-with-locals` and run `submit105-opsem-pair` to submit your work.

## How your work will be evaluated

### Adding local variables to Impcore

	Exemplary	Satisfactory	Must improve
Locals	<ul style="list-style-type: none"><li>• Change to interpreter appears motivated either by changing the semantics as little as possible or by changing the code as little as possible.</li></ul>	<ul style="list-style-type: none"><li>• Course staff believe they can see motivation for changes to interpreter, but more</li></ul>	<ul style="list-style-type: none"><li>• Course staff cannot understand what ideas were used to change the interpreter.</li><li>• Local variables for</li></ul>

	<ul style="list-style-type: none"> <li>Local variables for Impcore pass simple tests.</li> </ul>	<p>changes were made than necessary.</p> <ul style="list-style-type: none"> <li>Local variables for Impcore pass some tests.</li> </ul>	<p>Impcore pass few or no tests.</p>
Form	<ul style="list-style-type: none"> <li>Code taken from the book is either totally unchanged or is changed in a meaningful way.</li> <li>All code, including code from the book, fits in 80 columns.</li> <li>All code, except possibly machine-altered code from the book, respects the <a href="#">offside rule</a></li> <li>Indentation is consistent everywhere.</li> <li>No code is commented out.</li> <li>Solution file contains no distracting test cases or print statements.</li> </ul>	<ul style="list-style-type: none"> <li>One or two lines are wider than 80 columns.</li> <li>The code contains one or two violations of the <a href="#">offside rule</a></li> <li>In one or two places, code is not indented in the same way as structurally similar code elsewhere.</li> <li>Solution file may contain clearly marked test <i>functions</i>, but they are never executed. It's easy to read the code without having to look at the test functions.</li> </ul>	<ul style="list-style-type: none"> <li>Code taken from the book has been changed cosmetically (e.g., by changing line breaks or indentation, or by changing names) without changing the code's function.</li> <li>Three or more lines are wider than 80 columns.</li> <li>The code contains three or more violations of the <a href="#">offside rule</a></li> <li>The code is not indented consistently.</li> <li>Solution file contains code that has been commented out.</li> <li>Solution file contains test cases that are run when loaded.</li> <li>When loaded, solution file prints test results.</li> </ul>
Naming	<ul style="list-style-type: none"> <li>Where the code implements math, the names of each variable in the code is either the same as what's in the math (e.g., <code>rho</code> for <math>\rho</math>), or is an English equivalent for what the code stands for (e.g., <code>parameters</code> or <code>parms</code> for <math>\rho</math>).</li> </ul>	<ul style="list-style-type: none"> <li>Where the code implements math, the names don't help the course staff figure out how the code corresponds to the math.</li> </ul>	<ul style="list-style-type: none"> <li>Where the code implements math, the course staff cannot figure out how the code corresponds to the math.</li> </ul>
Structure	<ul style="list-style-type: none"> <li>The code is so clear that course staff can instantly tell whether it is correct or incorrect.</li> <li>There's only as much code as is needed to do the job.</li> <li>The code contains no redundant case analysis.</li> </ul>	<ul style="list-style-type: none"> <li>Course staff have to work to tell whether the code is correct or incorrect.</li> <li>There's somewhat more code than is needed to do the job.</li> <li>The code contains a little redundant case analysis.</li> </ul>	<ul style="list-style-type: none"> <li>From reading the code, course staff cannot tell whether it is correct or incorrect.</li> <li>From reading the code, course staff cannot easily tell what it is doing.</li> <li>There's about twice as much code as is needed to do the job.</li> </ul>

- A significant fraction of the case analyses in the code, maybe a third, are redundant.

## Operational semantics

	Exemplary	Satisfactory	Must improve
Semantics	<ul style="list-style-type: none"> <li>• The program which is supposed to behave differently in Awk, Icon, and Impcore semantics behaves exactly as specified with each semantics.</li> </ul>	<ul style="list-style-type: none"> <li>• The program which is supposed to behave differently in Awk, Icon, and Impcore semantics behaves almost exactly as specified with each semantics.</li> </ul>	<ul style="list-style-type: none"> <li>• The program which is supposed to behave differently in Awk, Icon, and Impcore semantics gets one or more semantics wrong.</li> <li>• The program which is supposed to behave differently in Awk, Icon, and Impcore semantics looks like it is probably correct, but it does not meet the specification: either running the code does not print, or it prints two or more times.</li> </ul>
Rules	<ul style="list-style-type: none"> <li>• Every inference rule has a single conclusion which is a judgment form of the operational semantics.</li> <li>• In every inference rule, every premise is either a judgment form of the operational semantics or a simple mathematical predicate such as equality or set membership.</li> <li>• In every inference rule, if two states, two environments, or two of any other thing <i>must</i> be the same, then they are notated using a <i>single</i> metavariable that appears in multiple places. (Example: <math>\rho</math> or <math>\sigma</math>)</li> </ul>	<ul style="list-style-type: none"> <li>• In every inference rule, two states, two environments, or two of any other thing <i>must</i> be the same, yet they are notated using different metavariables. However, the inference rule includes a premise that these metavariables are equal. (Example: <math>\rho_1 = \rho_2</math>)</li> <li>• A new language design has a few too many new or changes a few too many existing rules.</li> <li>• Or, a new language design is missing a few rules that are needed, or it doesn't change a few existing rules</li> </ul>	<ul style="list-style-type: none"> <li>• Notation that is presented as an inference rule has more than one judgment form or other predicate below the line.</li> <li>• Inference rules contain notation, either above or below the line, that resembles a judgment form but is not actually a judgment form.</li> <li>• In every inference rule, two states, two environments, or two of any other thing <i>must</i> be the same, yet they are notated using different metavariables, and nothing in the rule forces these metavariables to be equal. (Example: <math>\rho</math> and <math>\rho'</math> are both used, yet they must be identical.)</li> <li>• In some inference rule, two states, two environments, or two other things <i>may</i> be different, but they are notated using a single metavariable. (Example: using <math>\rho</math> everywhere, but in some places, <math>\rho'</math> is needed.)</li> </ul>

	<ul style="list-style-type: none"> <li>• In every inference rule, if two states, two environments, or two of any other thing <i>may</i> be different, then they are notated using different metavariables. (Example: <math>\rho</math> and <math>\rho'</math>)</li> <li>• New language designs use or change just enough rules to do the job.</li> <li>• Inference rules use one judgment form per syntactic category.</li> </ul>	<p>that need to be changed.</p>	<ul style="list-style-type: none"> <li>• In a new language design, the number of new or changed rules is a lot different from what is needed.</li> <li>• Inference rules contain a mix of judgment forms even when describing the semantics of a single syntactic category.</li> </ul>
<p>Derivations</p>	<ul style="list-style-type: none"> <li>• In every derivation, every utterance is either a judgment form of the operational semantics or a simple mathematical predicate such as equality or set membership.</li> <li>• In every derivation, every judgement follows from instantiating a rule from the operational semantics. (Instantiating means substituting for meta variables.) The judgement appears below a horizontal line, and above <i>that</i> line is one derivation of each premise.</li> <li>• In every derivation, equal environments are notated equally. In a derivation, <math>\rho</math> and <math>\rho'</math> <i>must</i> refer to different environments.</li> <li>• Every derivation takes the form of a tree. The root of the tree,</li> </ul>	<ul style="list-style-type: none"> <li>• In one or more derivations, there are a few horizontal lines that appear to be instances of inference rules, but the instantiations are not valid. (Example: rule requires two environments to be the same, but in the derivation they are different.)</li> <li>• In a derivation, the semantics requires new bindings to be added to some environments, and the derivation contains environments extended with the right new bindings, but not in exactly the right places.</li> </ul>	<ul style="list-style-type: none"> <li>• In one or more derivations, there are horizontal lines that the course staff is unable to relate to any inference rule.</li> <li>• In one or more derivations, there are many horizontal lines that appear to be instances of inference rules, but the instantiations are not valid.</li> <li>• A derivation is called for, but course staff cannot identify the tree structure of the judgments forming the derivation.</li> <li>• In a derivation, the semantics requires new bindings to be added to some environments, and the derivation does not contain any environments extended with new bindings, but the new bindings in the derivation are not the bindings required by the semantics. (Example: the semantics calls for a binding of <i>answer</i> to 42, but instead <i>answer</i> is bound to 0.)</li> <li>• In a derivation, the semantics requires new bindings to be added to some environments, but the derivation does not contain any environments extended with new bindings.</li> </ul>

	<p>which is written at the bottom, is the judgment that is derived (proved).</p> <ul style="list-style-type: none"> <li>• In every derivation, new bindings are added to an environment exactly as and when required by the semantics.</li> </ul>		
Metatheory	<ul style="list-style-type: none"> <li>• Metatheoretic proofs operate by structural induction on derivations, and derivations are named.</li> <li>• Metatheoretic proofs classify derivations by case analysis over the final rule in each derivation. The case analysis includes every possible derivation, and cases with similar proofs are grouped together.</li> </ul>	<ul style="list-style-type: none"> <li>• Metatheoretic proofs operate by structural induction on derivations, but derivations and subderivations are not named, so course staff may not be certain of what's being claimed.</li> <li>• Metatheoretic proofs classify derivations by case analysis over the final rule in each derivation. The case analysis includes every possible derivation, but the grouping of the cases does not bring together cases with similar proofs.</li> </ul>	<ul style="list-style-type: none"> <li>• Metatheoretic proofs don't use structural induction on derivations (<b>serious fault</b>).</li> <li>• Metatheoretic proofs have incomplete case analyses of derivations.</li> <li>• Metatheoretic proofs are missing many cases (<b>serious fault</b>).</li> <li>• Course staff cannot figure out how metatheoretic proof is broken down by cases (<b>serious fault</b>).</li> </ul>