Theo Scofield
tsco.dev | tk60wkc8u@mozmail.com

Education

University of British Columbia – Bachelor's in Computer Science May 2022

- **GPA**: 3.6
- Did coursework in the following: *Functional and logic programming, Artificial Intelligence, Computer Vision, Hardware and Operating Systems, Computability theory, Symbolic Logic, Vector Calculus, Statistics*

Projects

FoldList: A rust crate for a new data structure I came up with. The basic idea is that it's a list that also lets you perform "fold" operations in $O(\log n)$ time, similar to a segment tree, but also allowing the list to grow or shrink. All basic operations you can think of to manipulate a list of items are $O(\log n)$. This project took me about 10 months, mostly for two reasons. The first is that I initially wrote it in kotlin, but then rewrote it in rust to get a better idea of its performance when not everything needs to be an "Object", and therefore have its own allocation. The second reason is that my initial idea of the data structure was much more "flat", resembling a B-Tree with a large width, and with few children. The design of it was much more clever, and even had its own rebalancing mechanism, but I eventually found out that a simple binary tree was faster. This led to me basically rewriting the whole thing from scratch. The crate has been extensively tested, and is published on crates.io.

Te Test: An R package which calculates a particular statistical test. My version is $O(n)$, whereas naïvely implementing the version from the paper would be $O(n^2)$. I came up with the speedup myself (although I seem to be the only one to have implemented this particular test anyway, beside the paper's authors). The test is an instance of a "T-test," which solves the following problem: Suppose that you have two sets of samples of a number, from two different sources (distributions). How certain should you be that one distribution's average is greater than the other's average? Note that this is talking about the true averages of the distributions themselves, not just the averages of the samples that were taken from them (the samples may have misled you). This version of the test is unique because it works even when the two distributions have different standard deviations *and* when the samples have different sample sizes. (Link)

ProofHelper-v2: A command-line proof assistant I made for first-order logic. You can input some logical statements, as assumptions, and then write commands to prove additional statements. It only lets you perform correct deductions, and supports proofs by contradiction. It only supports statements about simple names, so it can't directly make statements about integers, but it can still describe them (i.e. it can't make statements about `succ(n)`, but it can make statements of the form "for all m, if m is the successor of n, then …" for some "is-sucessor-of" relation that you've defined). The proof system is from the textbook of a formal logic course I took, but it's a personal project, and not part of the course. (Link)

Extracurriculars/Volunteer Experience

- Member of Ballard High School Robotics Team 2014
- Participated in the PSCSTA Programming competition December 2015

Skills

- Strong experience in low-level strongly typed languages: C, Rust
- Experience in functional languages: OCaml, Haskell
- Background in math/logic/type systems/programming language design
- Other languages: Python, Javascript, Typescript, Prolog, Java, Kotlin