

Abstract

Over the years we have collected a large collection of Haskell programs written by students in a first-year functional programming course using the Helium compiler. The mining of such a collection is not trivial, especially since the programming was done *in vivo*, and hence largely outside our control. We have developed a library in Haskell, called NEON, for computing characteristics of this collection of programs and presenting the results visually. These computations range from simple kinds of analyses like computing the average length of a program per student to determining how long it takes for a programmer to resolve a type error.