



Research Note
RN/14/05

The Plastic Surgery Hypothesis

16/03/2014

Earl Barr, University College London

Yuriy Brun, University of Massachusetts

Prem Devanbu, University of California Davis

Mark Harman, University College London

Federica Sarro, University College London

Abstract

Recent work on Genetic Programming based approaches to automatic program patching have relied upon the insight that the content of new code can often be assembled out of fragments of code that already exist in the code base. This insight has been dubbed the plastic surgery hypothesis; successful, well-known automatic repair tools such as GenProg and ClearView rest on this hypothesis, but it has never been validated. We formalize and validate the plastic surgery hypothesis and empirically measure the extent to which raw material for changes actually already exists in projects. In this paper, we mount a large-scale study of several large Java projects, and examine a history of 15,723 commits to determine the extent to which these commits are graftable, viz., reconstituted from existing code, and find an encouraging degree of graftability, surprisingly independent of commit size and type of commit. With a view to investigating the difficulty of finding these grafts, we study the fecundity of possible sources of such grafts: the immediately previous version, prior history, and other projects. We also examine contiguity or chunking of these grafts, and the degree to which grafts can be found in the same file. Our results overall are quite promising and are indicative of a optimistic future for automatic program patching methods which search for raw material within already extant code in the project being patched.