

What Makes Inspections Work? Understanding How and Why Different Inspection Methods Impact Effectiveness and Cost

Adam Porter*, Harvey Siy
Computer Science Department
University of Maryland
College Park, Maryland 20742
{aporter,harvey}@cs.umd.edu

Lawrence G. Votta, Jr.
Software Production Research Department
AT&T Bell Laboratories
Naperville, Illinois 60566
votta@research.att.com

Abstract

For two decades, software inspections have proven effective for detecting defects in software. We have reviewed the different ways software inspections are done, created a taxonomy of inspection methods, and examined claims about the cost-effectiveness of different methods.

We detect a disturbing pattern in the evaluation of inspection methods. Although there is near universal agreement on the effectiveness of software inspection, their economics are uncertain. Our examination of several empirical studies leads us to conclude that the benefits of inspections are often overstated and the costs (especially for large software developments) are understated. Furthermore, some of the most influential studies establishing these costs and benefits are 20 years old now, which leads us to question their relevance to today's software development processes.

Extensive work is needed to determine exactly how, why, and when software inspections work, and whether some defect detection techniques might be more cost-effective than others. In this tutorial we ask some questions about measuring effectiveness of software inspections and determining how much they really cost when their effect on the rest of the development process is considered.

*This work is supported in part by a National Science Foundation Faculty Early Career Development Award CCR-9501354. Mr. Siy was also partly supported by AT&T's Summer Employment Program