

DOCTORAL CANDIDATE:

DEGREE: Philosophiae Doctor
FACULTY: The Faculty of Mathematics and Natural Sciences
DEPARTMENT: Department of Informatics
AREA OF EXPERTISE: Recommender Systems for Software Engineering
SUPERVISORS: Leon Moonen, Razieh Behjati and Magne Jørgensen
DATE OF DISPUTATION: 1st of September 2017

DISSERTATION TITLE: *Improving History-Based Change Recommendation Systems for Software Evolution*

The software that we depend on every day is constantly changed. These changes are necessary to comply with shifting user requirements, keeping a competitive advantage, adapting to changes in other software, and to fix the ever-present bugs. It is crucial that the impact of these changes is well understood, as failure to do so may well lead to additional bugs being introduced, which may directly affect the longevity and success of the software.

In order to understand the impact of a change, the parts of the software which is affected by the change must be uncovered. This thesis explores through five papers how the change history of software can be leveraged to address this challenge. From a change history, it is possible to identify the files and methods that typically change together. Furthermore, these change patterns can successfully be used to derive association rules which can predict the impact of future changes.

As with all recommendation systems, a danger is always that parts of a recommendation are wrong or missing (false positives and negatives). This thesis presents steps towards reducing such issues for change recommendation systems utilizing change patterns and association rule mining. The contributions with respect to each paper are as follows: (A) A targeted association rule mining algorithm for filtering out irrelevant association rules (B) An approach for aggregating association rules, increasing overall recommendation precision (C) A study of parameters for association rule mining in this context (D) A study of the effects of change history length and age on recommendation precision (E) And lastly, an approach for predicting whether a change recommendation is relevant. These contributions help drive forward a field which has long been dominated by approaches based on static and dynamic analysis. By building on top of version control systems, our approach has considerably less overhead and is inherently language agnostic.