

Textual use-cases have been traditionally used at the design stage of the software development process to describe software functionality from the user's perspective. Because use-cases typically rely on natural language, they cannot be directly subject to formal verification. Another important artefact is the domain model, a high-level overview of the most important concepts in the problem space. A domain model is usually not constructed en bloc, yet it undergoes refinement starting from the first prototype elicited from text. This thesis covers two closely related topics - formal verification of use-cases and elicitation of a domain model from text. The former is a method (called FOAM) that features simple user-definable annotations inserted into a use-case to make it suitable for verification. A model-checking tool is employed to verify temporal invariants associated with the annotations while still keeping the use-cases understandable for non-experts. The latter is a method (titled Prediction Framework) that features an in-depth linguistic analysis of text and a sequence of statistical classifiers (log-linear Maximum Entropy models) to predict the domain model.