

Foreword

Context is always problematic in science - we would like our theories to be as general as possible and thus independent of context. For example, a formula like Newton's second law, $F = ma$, seems quite general, but implicitly presupposes the context of an unchanging, freely moving mass with constant acceleration at speeds much less than that of light. In fact, much of the difficulty of learning formal science involves understanding which formulas are appropriate in a given problem context.

Context is even more vexing in natural language. The main-stream treatment of formal semantics of language is based on the principle of strong compositionality, which states that the meaning of a (possibly complex) expression is totally determined by its form and is independent of context. As Manfred Krifka says in the online MIT Encyclopedia of Cognitive Science, "In its strict version, this claim is clearly wrong". The controversy over contextuality in language involves some of the deepest questions about the mind.

Context is also problematic in the myriad practical applications of computational linguistics including search engines, data mining, speech recognition, etc. Statistical techniques rely on massive amounts of data and implicitly assume that all of the data is from the same general context. There is now increasing recognition that context-independent techniques are approaching their limits and more efforts, including conferences, are focusing on the role of context on language theory and practice. Robert Porzel's work on Contextual Computing goes back several years, anticipating the current trends. This monograph, based on his doctoral dissertation, is a significant and timely contribution to the field. The book addresses a problem of great and increasing technical and practical importance - the role of context in natural language processing (NLP). Although language processing programs, such as search engines, play a huge role in society, essentially all of them deal only with language form, not meaning. It would be easy to add automated dictionaries, but this would not help much - the meaning of language is heavily dependent on the context of use. The first issue addressed is how to capture context in language processing. Following current practice, Porzel formalizes contextual knowledge as ontology data bases and includes some insights on building and evaluating such ontologies. The monograph explores the use of context in NLP through a range of

demonstrations and experiments. These range from small studies to a huge international project, SmartKom, in which Porzel himself played an important role. He considers the role of context in three important tasks: Automatic Speech Recognition, Semantic Interpretation, and Pragmatic Interpretation.

The book is organized into five chapters. The first chapter is a short introduction to the subject and presents the major aims and contribution of the work. In Chapter 2, Porzel describes the state of the art, covering previous definitions and approaches to contextual computing for NLP as well as the methodological background that is necessary for measuring and evaluating the performance of NLP systems.

In Chapter 3, different speech hypotheses that can result from a noisy input or a non-perfect speech-recognizer are mapped to concept-sets that correspond to sub-graphs of the ontology. Porzel then shows that evaluating the coherence of the sub-graph can be used for scoring the speech hypotheses and can lead to a significant gain in overall recognition quality. Then the graphs are extended by also using concepts from the surrounding discourse of an utterance, leading to even better results. Two further experiments extend this contextual approach further for word sense disambiguation tasks and the extraction of semantic relations.

Chapter 4 discusses integrating user and situation knowledge into dialog systems. The main challenge here is to resolve pragmatic ambiguities and find the intention of the user in a given utterance. Porzel illustrates the problem with a variety of examples and also presents real-world examples where these ambiguities occur. He suggests two different approaches. The first uses models based on empirical data, e.g. Bayesian networks, in order to resolve pragmatic ambiguities. The second uses explicit formal models of pragmatic information in an ontology. Such pragmatic patterns can then be used together with context data in order to help determine how a dialog system should react when multiple solutions are possible.

Using his previous training in experimental psychology, Porzel provides detailed empirical studies to support his computational context models. For most problems, a gold standard has been produced by human annotators and the performance of the proposed method has been compared to simpler baseline methods. The final chapter is the usual summary and suggestions for future work. Overall, the book represents a novel and insightful investigation into the potential of contextual information processing in NLP.

Berkeley, September 2010

Jerome Feldman



<http://www.springer.com/978-3-642-17395-0>

Contextual Computing
Models and Applications
Porzel, R.
2011, XI, 180 p., Hardcover
ISBN: 978-3-642-17395-0