

## Preface

**Catherine Havasi**

Massachusetts Institute of Technology

**Doug Lenat**

Cycorp

**Benjamin Van Durme**

Johns Hopkins University

When we are confronted with unexpected situations, we deal with them by falling back on our general knowledge or making analogies to other things we know. When software applications fail, on the other hand, they often do so in brittle and unfriendly ways. The sheer amount of commonsense knowledge one would need to represent makes it challenging to acquire, to represent, to reason efficiently with, and to harness in applications.

This is, ultimately, the bottleneck to strong AI, and so it has remained one of the central topics of research interest for 50 years, from John McCarthy, Patrick Hayes, and colleagues grappling with representation and reasoning, to Doug Lenat, Push Singh, and Lenhart Schubert conducting large scale engineering projects to construct collections

of background knowledge and special-purpose reasoners to support general inference. Recent advances in text mining, crowdsourcing, and professional knowledge engineering efforts have finally led to commonsense knowledge bases of sufficient breadth and depth for practical applications.

A growing number of research projects now seek to use these knowledge collections in a wide variety of applications, including computer vision, speech processing, robotics, dialogue and text understanding, and apply them to real-world tasks such as healthcare and finance, where brittleness is unacceptable. At the same time, new application domains are giving fresh insights into desiderata for common sense reasoners and guidance for knowledge collection efforts.