Modeling the effects of memory on human online sentence processing with particle filters

Roger Levy
Department of Linguistics
University of California, San Diego
rlevy@ling.ucsd.edu

Florencia Reali **Thomas L. Griffiths**
Department of Psychology
University of California, Berkeley
{floreali,tom_griffiths}@berkeley.edu

Abstract

Language comprehension in humans is significantly constrained by memory, yet rapid, highly incremental, and capable of utilizing a wide range of contextual information to resolve ambiguity and form expectations about future input. In contrast, most of the leading psycholinguistic models and fielded algorithms for natural language parsing are non-incremental, have run time superlinear in input length, and/or enforce structural locality constraints on probabilistic dependencies between events. We present a new limited-memory model of sentence comprehension which involves an adaptation of the particle filter, a sequential Monte Carlo method, to the problem of incremental parsing. We show that this model can reproduce classic results in online sentence comprehension, and that it naturally provides the first rational account of an outstanding problem in psycholinguistics, in which the preferred alternative in a syntactic ambiguity seems to grow more attractive over time even in the absence of strong disambiguating information.

1 Introduction

Nearly every sentence occurring in natural language can, given appropriate contexts, be interpreted in more than one way. The challenge of comprehending a sentence is identifying the intended intepretation from among these possibilities. More formally, each interpretation of a sentence $\boldsymbol{w}$ can be associated with a structural description T, and to comprehend a sentence is to infer T from $\boldsymbol{w}$ – *parsing* the sentence to reveal its underlying structure. From a probabilistic perspective, this requires computing the posterior distribution $P(T|\boldsymbol{w})$ or some property thereof, such as the description T with highest posterior probability. This probabilistic perspective has proven extremely valuable in developing both effective methods by which computers can process natural language [1, 2] and models of human language processing [3].

In real life, however, people receive nearly all linguistic input *incrementally*: sentences are spoken, and written sentences are by and large read, from beginning to end. There is considerable evidence that people also comprehend incrementally, making use of linguistic input moment by moment to resolve structural ambiguity and form expectations about future inputs [4, 5]. The incremental parsing problem can, roughly, be stated as the problem of computing the posterior distribution $P(T|w_{1\dots i})$ for a partial input $w_{1\dots i}$. To be somewhat more precise, incremental parsing involves constructing a distribution over partial structural descriptions of $w_{1\dots i}$ which implies the posterior $P(T|w_{1\dots i})$. A variety of "rational" models of online sentence processing [6, 7, 8, 9] take exactly this perspective, using the properties of $P(T|w_{1\dots i})$ or quantities derived from it to explain why people find some sentences more difficult to comprehend than others.

Despite their success in capturing a variety of psycholinguistic phenomena, existing rational models of online sentence processing leave open a number of questions, both theoretical and empirical. On the theoretical side, these models assume that humans are "ideal comprehenders" capable of computing $P(T|w_{1\dots i})$ despite its significant computational cost. This kind of idealization is common in rational models of cognition, but raises questions about how resource constraints might affect language processing. For structured probabilistic formalisms in widespread use in compu-