

Weighted Finite Automata with Failure Transitions: Algorithms and Applications

Cyril Allauzen

Google, New York

Editors: François Coste, Faissal Ouardi and Guillaume Rabusseau

Abstract

Weighted finite automata (WFA) are used in many applications including speech recognition, speech synthesis, machine translation, computational biology, image processing, and optical character recognition. Such applications often have strict time and memory requirements, so efficient representations and algorithms are paramount. We examine one useful technique, the use of failure transitions, to represent automata compactly. A failure transition is taken only when no immediate match to the input is possible at a given state. Automata with failure transitions, initially introduced for string matching problems, have found wider use including compactly representing language, pronunciation, transliteration and semantic models.

In this talk, we will address the extension of several weighted finite automata algorithms to automata with failure transitions (Φ -WFAs). Efficient algorithms to intersect two Φ -WFAs, to remove failure transitions, to trim, and to compute the shortest distance in a Φ -WFA will be presented.

We will demonstrate the application of some of these algorithms on two language modeling tasks: the distillation of arbitrary probabilistic models as weighted finite automata with failure transitions and the federated learning of n-gram language models. We will show the relevance of these methods to the privacy-preserving training of language models for virtual keyboard applications for mobile devices.

This talk covers work in collaboration with Michael Riley, Ananda Theertha Suresh, Brian Roark, Vlad Schogol, Mingqing Chen, Rajiv Mathews, Adeline Wong, and Françoise Beaufays.