

Review of

Continuous Speech Recognition Using Support Vector Machines

by Ganapathiraju, Hamaker and Picone

Submitted to Computer Speech and Language

This paper is about the application of the powerful pattern classification approach of Support Vector Machines (SVM) to speech recognition. It has a basic introductory section, a lengthy section on the basics of SVMs with some examples, a section on adapting the classifier design to deliver probability estimates that can be integrated with HMMs and an experimental results section.

I do not think the paper is acceptable for publication in CSL for the following reasons:

- a) The novelty of material is low and is restricted to simple application of a widely known pattern classification method to the speech context. The important part of the process is to tweak the SVM outputs to be estimates of probabilities. This part of the paper is not a thorough investigation of different possible ways of achieving it. The experiments, too, are not that strong - a static vowels classification task and a isolated digits - task. Thus on the novelty front the paper doesn't score well either on the theoretical / algorithmic side, or on the scle of the applications side.
 - b) The balance in material is weak. Section 2 is a lengthy exposition of material already available widely. It may be argued that this is new to the Speech community, but I think speech researchers regularly read Machine Learning literature where they will find Section 2 in greater depth and thoroughly reviewed in tutorial articles.
 - c) Section 3, where the approaches to massaging SVM outputs to integrate with HMMs is not written well. Ad Hoc approaches of other researchers are not reviewed thoroughly, nor are the particular decisions of authors of this manuscript well justified; e.g. the equal variance assumption leading to Eqn (34) - why should having large amounts of data lead to equal variance posteriors?
-