

CAN A STATISTICAL MODEL KNOW WHAT IT DOESN'T KNOW? AN APPROACH TO MODELING EPISTEMIC UNCERTAINTY IN RECURRENT NEURAL NETWORKS (RNNS)

SIMÃO GONÇALVES

ABSTRACT. Statistical models always produce an output, even for out-of-distribution inputs. But critical applications such as modeling medical prognosis through time, where physicians may act based on a model's inference, have a risk of negatively impacting patients' health. Therefore the decision making process of such applications would greatly benefit from a notion of uncertainty of the model.

Recurrent Neural Networks are a special case of basis regression functions and are used to model non-linear relationships $Y|X$ where X is a variable-sized sequence. Recent work by Yarin Gal et al. has bridged the RNN model with the variational posterior of a bayesian model. We propose a novel approach that leverages the proposed variational RNN to obtain an uncertainty that is correlated with model correctness.

(Simão Gonçalves) NOVA SCHOOL OF SCIENCES AND TECHNOLOGY, PORTUGAL
Email address: `sg.goncalves@campus.fct.unl.pt`