

# The Center for Cyber Defenders

Expanding computer security knowledge



## Specializing Neural Network Language Models Using Bayesian Networks

Samuel Richter, Missouri University of Science and Technology, M.S. Comp Sci

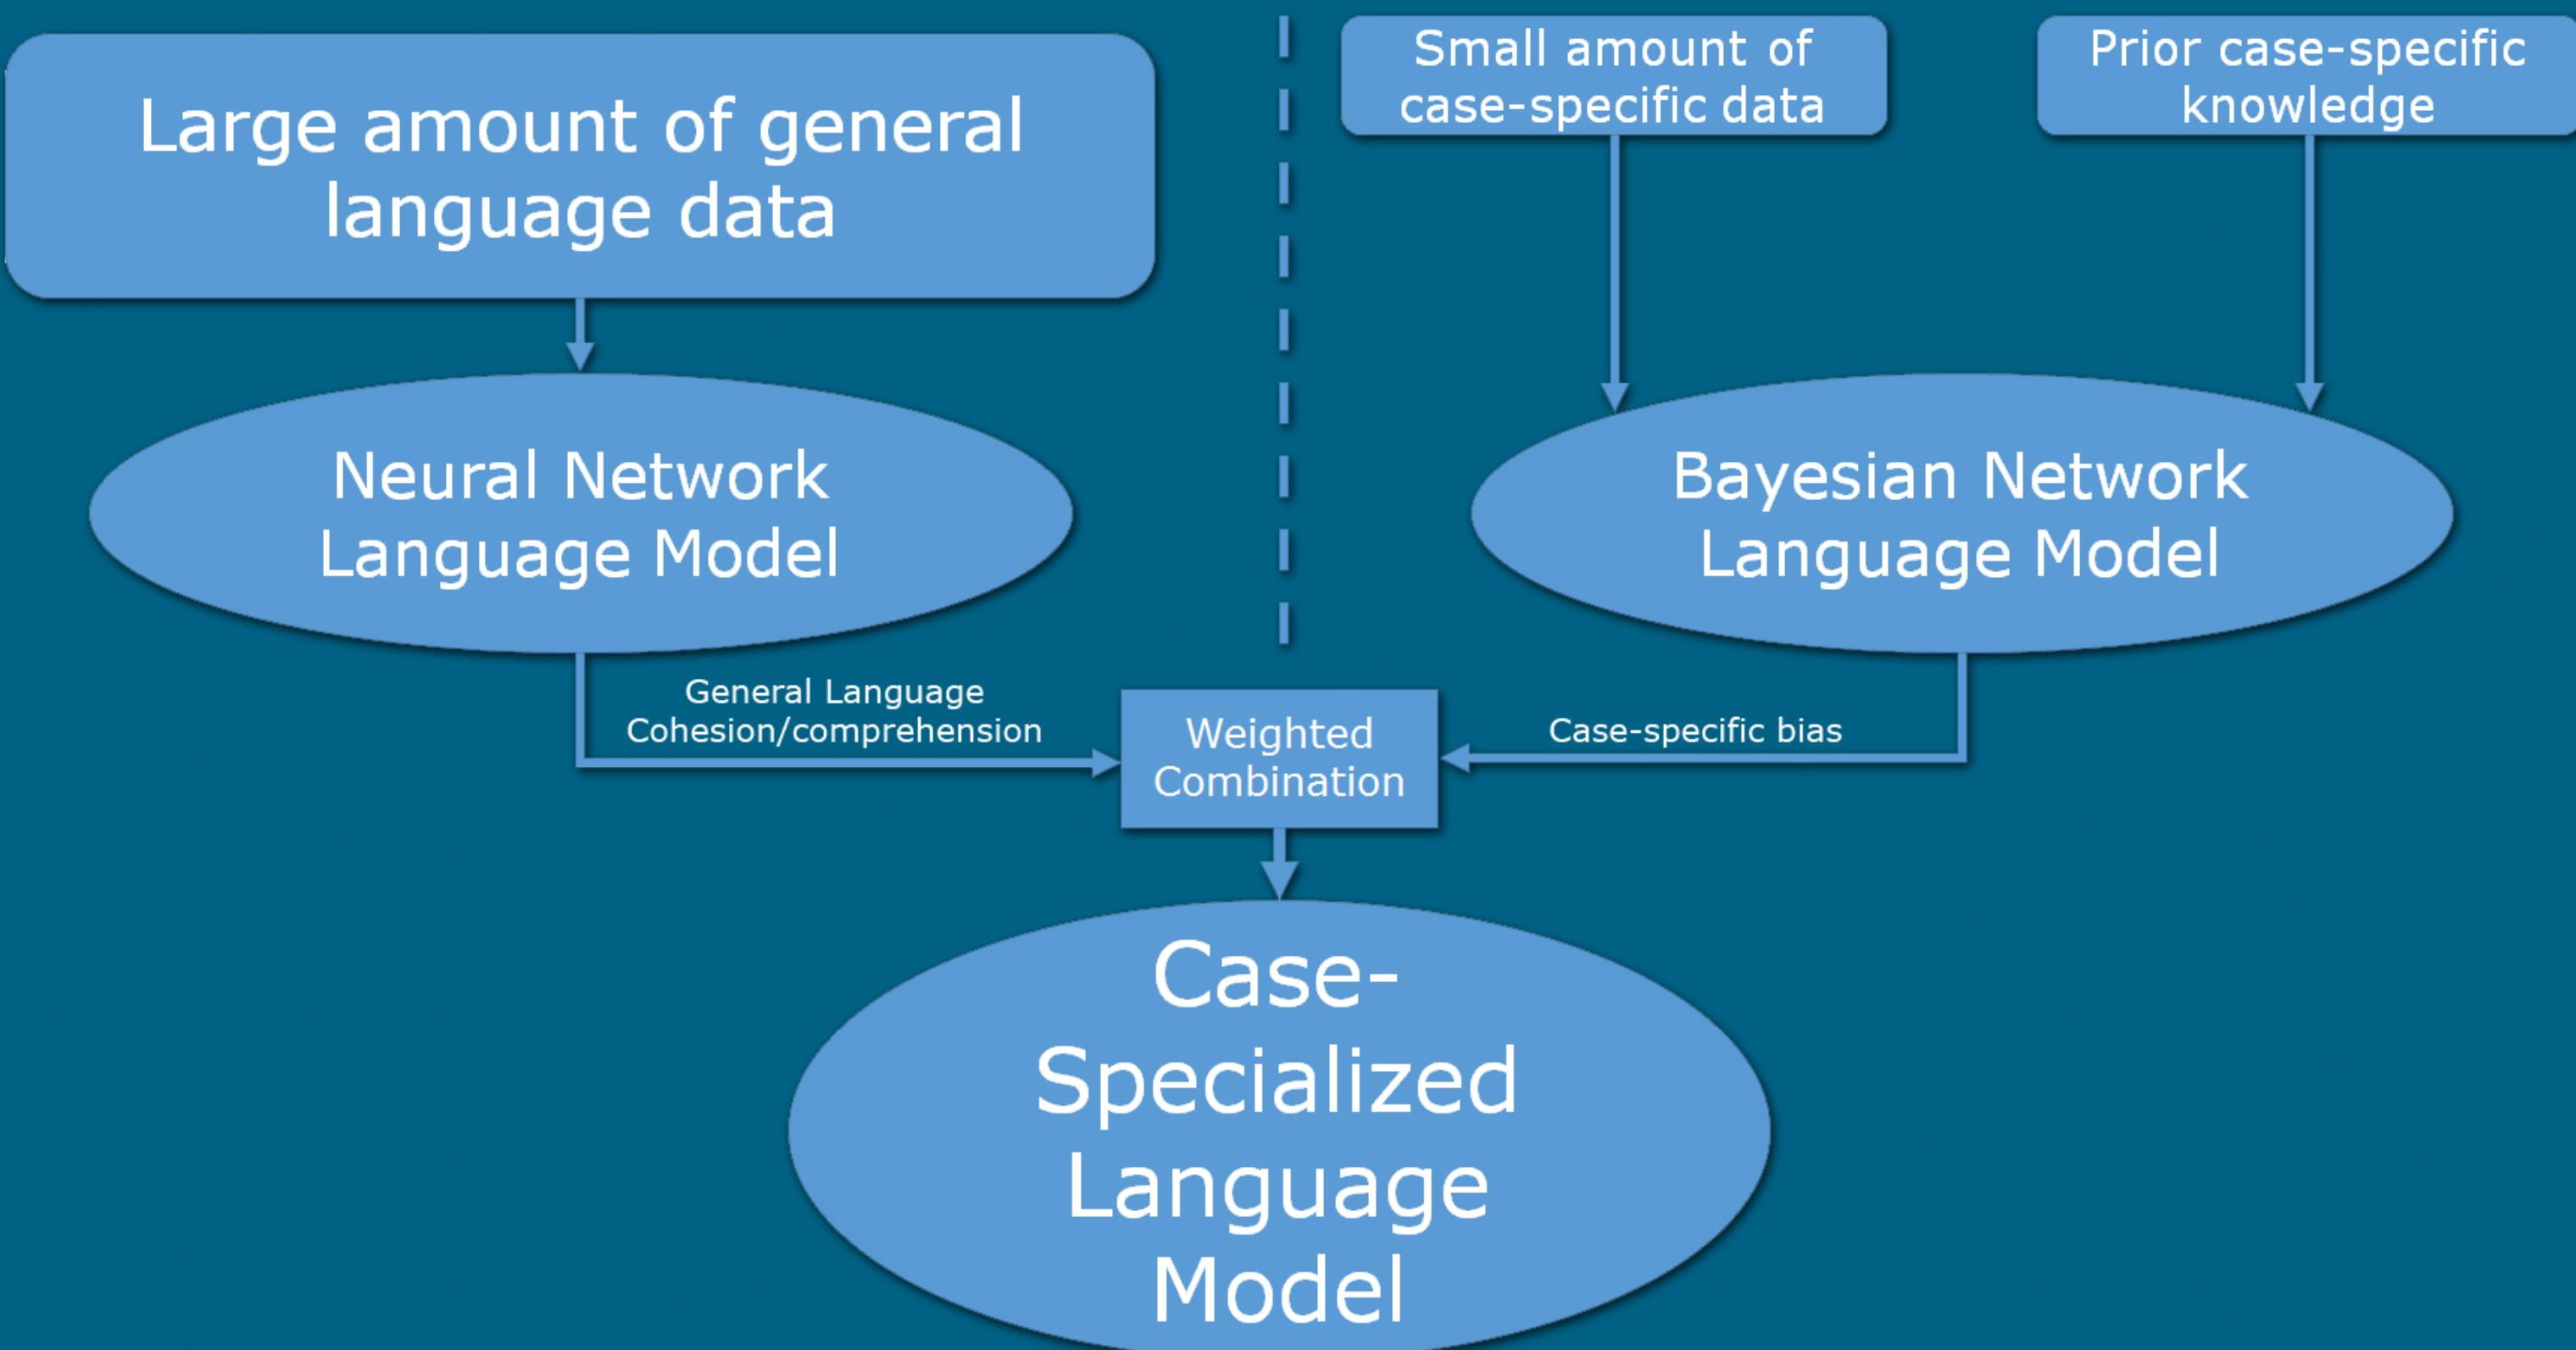
Project Mentor/PI: Hamilton Link, Org. 5874  
Project Advisor: Jacek Skryzalin, Org. 5853

### Problem Statement:

Neural networks can be used to build models of natural language, which measure the coherency and likelihood of written text. These models can be used, among other things, to aid in speech recognition, hand-writing transcription, and text generation. However, building these models requires significant amounts of data to train on, and while such data is available for language in general, data for more specialized cases, like dialects and particular reading levels, is more sparse. We can specialize a Neural Network Language model to a particular use case using Bayesian learning, which can produce structured models and leverage prior knowledge while requiring less data to train on.

### Approach:

We first train a regular Recurrent Neural Network (RNN) to produce a general model of the English language. We then train a Bayesian network on a more limited set of data from a more specialized language application, such as a particular dialect or reading level of English. We can then use the output of the neural network to measure coherency/correctness and the output of the Bayesian network to bias the language model towards words, phrases, and sentences that are more likely in the targeted use case. The outputs of both models are combined, with a bias factor on each model, to produce the final model output.



### Impact and Benefits:

This model can be used to improve the performance of any application which requires both a method of measuring language coherency and specialized knowledge of a particular case of the language. Potential example uses include generating article text for the Simple English Wikipedia, transcribing and translating dialects of a language, and measuring the amount of domain-specific knowledge in a document. Retooling the model for a different use case only requires retraining the Bayesian network with new case-specific data, since the neural network provides the more general model for the entire language and thus only requires training once.