

A BETTER WORDSCORES: SCALING TEXT WITH THE CLASS AFFINITY MODEL

BY PATRICK O. PERRY[†] AND KENNETH BENOIT^{*,‡}

Oscar Health[†] and London School of Economics and Political Science[‡]

Abstract

Probabilistic methods for classifying text form a rich tradition in machine learning and natural language processing. For many important problems, however, class prediction is uninteresting because the class is known, and instead the focus shifts to estimating latent quantities related to the text, such as affect or ideology. We focus on one such problem of interest, estimating the ideological positions of 55 Irish legislators in the 1991 *Dáil* confidence vote, a challenge brought by opposition party leaders against the then-governing *Fianna Fáil* party in response to corruption scandals. In this application, we clearly observe support or opposition from the known positions of party leaders, but have only information from speeches from which to estimate the relative degree of support from other legislators. To solve this scaling problem and others like it, we develop a text modeling framework that allows actors to take latent positions on a “gray” spectrum between “black” and “white” polar opposites. We are able to validate results from this model by measuring the influences exhibited by individual words, and we are able to quantify the uncertainty in the scaling estimates by using a sentence-level block bootstrap. Applying our method to the *Dáil* debate, we are able to scale the legislators between extreme pro-government and pro-opposition in a way that reveals nuances in their speeches not captured by their votes or party affiliations.

*Prepared for presentation at the 2022 COMPTTEXT Workshop, Dublin, 6–7 May. This is a work in progress and while we would love you to cite it, please wait a few months for a more definitive (and hopefully published) version. This research was supported by the European Research Council grant ERC-2011-StG 283794-QUANTESS. We thank Jouni Kuha for comments.