

Good Uses for Crummy Knowledge Graphs

Keynote Abstract

Douglas W. Oard
University of Maryland
College Park, MD
USA
oard@umd.edu

ABSTRACT

In 1993, Ken Church and Ed Hovy suggested that before we ask how well some new technology meets the need we envision for it, we should pause and first reflect on the question of whether – now that we know something about what can be built – we are envisioning the right uses for what we have. They titled their paper “Good Applications for Crummy Machine Translation” [1]. At about that same time, information retrieval researchers obliged them by (generally without having read their paper) starting to work on cross-language information retrieval; arguably the best application for crummy machine translation ever invented. Now we have some crummy knowledge graphs – and this time we have read the Church and Hovy paper – so perhaps the time is right for us to ask whether we have yet envisioned good uses for crummy knowledge graphs. In this talk, I will seek to seed that discussion.

References

- [1] K. W. Church and E. H. Hovy. Good applications for crummy machine translation. *Machine Translation*, 8:239–258, 1993. URL <http://dx.doi.org/10.1007/BF00981759>.