Refining Predicates for Relation Extraction through Thesaurus Integration - Abstract

Ewan Hannaford[1], Youcef Benkhedda[2], Marc Alexander[1], Goran Nenadic[2] and Riza Batista-Navarro[2,*]

[1]*School of Critical Studies, University of Glasgow, Glasgow, G11 6EW, UK*
[2]*Department of Computer Science, University of Manchester, Oxford Road, Manchester M13 9PL, UK*

Using natural language processing (NLP) approaches combined with humanities expertise, the *Our Heritage, Our Stories* project seeks to connect community archives from across the UK — archives run for and by local communities that contain unique datasets of community-generated digital content (CGDC) — with materials in The National Archives (TNA) of the UK [1]. Foundational to this research is the refinement of NLP methods for their application to CGDC, in order to extract entities appearing in materials (people, places, dates, organisations, etc) and relations between them, and link these across collections [2].

A key step in this process is the definition of predicates (i.e., relation types) that form the basis of the automated extraction of relationships between named entities [3, 4]. As there may exist many different linguistic formulations of the same relationship, normalisation of these heterogeneous, yet synonymous, forms to a prototypical relation is often necessary to adequately capture similar relationships appearing in diverse materials. This ensures that materials expressing the same meaning in different ways are recognised as doing so, and that connections can be drawn across materials as a result. Wikidata, a canonical knowledge base for linked data approaches [5], contains an extensive list of relation types, which researchers can map relations appearing in their materials to as a means of relation normalisation. In the *Our Heritage, Our Stories* project, a set of ~30 key relations relevant to CGDC were identified, which covered the core relationships between entities typically presented in community archive materials.

However, these relations, as prototypical expressions, do not, and were not intended to, comprehensively cover the diverse ways in which such relations can be expressed in CGDC. As a result, in order to prevent relations that appeared in CGDC but not in lists of canonical relations going unrecognised and unrepresented, a broader range of expressions was required to capture the full range of linguistic manifestations of relationships between entities in materials. Using the *Historical Thesaurus of English* — the world's largest thesaurus of English [6] — the project team identified and enriched the set of predicates selected from Wikidata with synonymous terms from similar semantic fields, subsequently integrating these expanded relation types into annotation and processing.

*Corresponding author.

✉ ewan.hannaford@glasgow.ac.uk (E. Hannaford); youcef.benkhedda@manchester.ac.uk (Y. Benkhedda); marc.alexander@glasgow.ac.uk (M. Alexander); gnenadic@manchester.ac.uk (G. Nenadic); riza.batista@manchester.ac.uk (R. Batista-Navarro)

This talk discusses this work, explaining how key relations were selected, how synonymous terms for these relations were identified from the *Historical Thesaurus*, and how these were integrated into the project's NLP approaches to improve the automated interpretation of relationships appearing in CGDC. In doing so, it delineates how linguistic thesauri may be used as a means of refining relation extraction and linking, enabling more comprehensive capture of relations whilst maintaining normalisation necessary for linked data approaches. Consequently, it proposes a hybrid approach to relation extraction, demonstrating how NLP methods can be supported by manual integration of linguistic expertise and resources.

References

[1] E. Hannaford, M. Alexander, L. Hughes, R. Lewis, Our Heritage, Our Stories: Democratising the UK National Collection, in: Proceedings of the Digital Humanities Conference 2023 (DH 2023), Graz, Austria, 2023, p. 565.

[2] Y. Benkhedda, A. Skapars, V. Schlegel, G. Nenadic, R. Batista-Navarro, Enriching the Metadata of Community-Generated Digital Content through Entity Linking: An Evaluative Comparison of State-of-the-Art Models, in: Y. Bizzoni, S. Degaetano-Ortlieb, A. Kazantseva, S. Szpakowicz (Eds.), Proceedings of the 8th Joint SIGHUM Workshop on Computational Linguistics for Cultural Heritage, Social Sciences, Humanities and Literature (LaTeCH-CLfL 2024), Association for Computational Linguistics, St. Julians, Malta, 2024, pp. 213–220. URL: https://aclanthology.org/2024.latechclfl-1.20.

[3] Y. Lee, J. Son, M. Song, BertSRC: transformer-based semantic relation classification, BMC Medical Informatics and Decision Making 22 (2022) 234.

[4] A. Plum, T. Ranasinghe, S. Jones, C. Orasan, R. Mitkov, Biographical: A Semi-Supervised Relation Extraction Dataset, in: Proceedings of the 45th International ACM SIGIR Conference on Research and Development in Information Retrieval, 2022, pp. 3121–3130.

[5] D. Vrandečić, L. Pintscher, M. Krötzsch, Wikidata: The Making Of, in: Companion Proceedings of the ACM Web Conference 2023, WWW '23 Companion, Association for Computing Machinery, New York, NY, USA, 2023, p. 615–624. URL: https://doi.org/10.1145/3543873.3585579. doi:10.1145/3543873.3585579.

[6] C. Kay, M. Alexander, F. Dallachy, J. Roberts, M. Samuels, I. Wotherspoon, The Historical Thesaurus of English (2nd edn., version 5.0), University of Glasgow, 2024. Available online: https://ht.ac.uk.