Research Repository UCD

Title	Industrial Memories: Exploring the Findings of Government Inquiries with Neural Word Embedding and Machine Learning
Authors(s)	Leavy, Susan, Pine, Emilie, Keane, Mark T.
Publication date	2019-01-18
Publication information	Leavy, Susan, Emilie Pine, and Mark T. Keane. "Industrial Memories: Exploring the Findings of Government Inquiries with Neural Word Embedding and Machine Learning." ECML-PKDD, 2019.
Conference details	The European Conference on Machine Learning and Principles and Practice of Knowledge Discovery in Databases 2018 (ECML-PKDD 2018), Dublin, Ireland, 10-14 September 2018
Series	European Conference, ECML PKDD 2018, Dublin, Ireland, September 10–14, 2018, Proceedings, Part III
Publisher	ECML-PKDD
Item record/more information	http://hdl.handle.net/10197/10321
Publisher's statement	The final publication is available at www.springerlink.com
Publisher's version (DOI)	10.1007/978-3-030-10997-4_52

Downloaded 2023-04-07T10:52:29Z

The UCD community has made this article openly available. Please share how this access benefits you. Your story matters! (@ucd_oa)

© Some rights reserved. For more information

Industrial Memories: Exploring the Findings of Government Inquiries with Neural Word Embedding and Machine Learning

Susan Leavy, Emilie Pine, and Mark T Keane

University College Dublin, Dublin, Ireland
{susan.leavy,emilie.pine, mark.keane}@ucd.ie

Abstract. We present a text mining system to support the exploration of large volumes of text detailing the findings of government inquiries. Despite their historical significance and potential societal impact, key findings of inquiries are often hidden within lengthy documents and remain inaccessible to the general public. We transform the findings of the Irish government’s inquiry into industrial schools and through the use of word embedding, text classification and visualization, present an interactive web-based platform that enables the exploration of the text in new ways to uncover new historical insights.

Keywords: Word Embeddings · Text Classification · Visualization

1 Introduction

The Irish Government published one of the most important documents in the history of the state in 2009. It spanned over 1 million words and detailed the findings of a 9 year investigation into abuse and neglect in Irish industrial schools. The document however, remains largely unread. The structure and style of the report obscures significant information and prevents the kind of system-wide analysis of the industrial school regime which is crucial to identifying recurring patterns and preventing the re-occurrence of such atrocities.

The Industrial Memories project, publicly available online¹, presents a web-based interactive platform where the findings of CICA, commonly known as the Ryan Report, may be explored by researchers and the general public in new ways. The machine learning approach developed as part of this project addressed the challenge of learning within the context of limited volumes of training data. Techniques such as named entity recognition, word embedding and social network analysis are also used to extract information and unearth new insights from the report.

2 System Overview

The Industrial Memories project developed a web-based platform where the narrative form of the Ryan Report is deconstructed and key information extracted.

¹ <https://industrialmemories.ucd.ie>

In the Ryan Report most of the chapters deal separately with individual industrial schools. This segregation of information prevents a system-wide analysis of common trends and patterns across the industrial school system. To address this, a random forest classifier was employed to annotate the report. An interactive web based platform enables users to extract information and interact with the report using a comprehensive search tool (Fig. 1). Word embeddings were also used to identify actions pertaining to meetings and communication that occurred within the industrial school system facilitating an analysis of the dynamics of power and influence using social network analysis. Each individual named in the report was extracted using the Stanford NE system [1] which for the first time, allows the history of individuals to be traced over the entire industrial school system thus responding to a call for increased open data in government [2].

Fig. 1. Text Navigation Tool

2.1 Semantic Classification of Text

While some semantic categories or themes were conducive to extraction using a rule based approach, categories including the transfer events, descriptions of abuse and witness testimony were extracted using a random forest learning algorithm. The random forest algorithm was chosen because, as an ensemble learner that creates a ‘forest’ of decision trees by randomly sampling from the training set, it is suited to learning from smaller datasets [4]. The Ryan Report was represented in a relational database with annotated excerpts comprising 6,839 paragraphs (597,651 words). Each paragraph was considered a unit-of-analysis in the Report for the purposes of annotation.

2.2 Feature Extraction with Word Embedding

Feature selection was based on the compilation of domain-specific lexicons in order to address the constraint in this project concerning low volumes of available training data. As with many projects that are based in the humanities, compiling

training data is a manual process involving close reading of texts. To address this, domain-specific lexicons were compiled based on the entire dataset yielding high accuracy with low numbers of training examples.

Lexicons were generated using the word2vec algorithm from a number of seed terms [3]. Given that the word2vec algorithm performs best with larger datasets than the Ryan Report, five ensembles were generated and the top thirty synonyms for each seed term were extracted. Terms common across each ensemble were identified to generate domain specific synonyms. Features for the classification process were then derived from these lexicons.

Fig. 2. Collocation Network of Entities in the Ryan Report

2.3 Exploring Power, Influence and Migration within the System with Word Embedding and Network Analysis

Neural network analysis was used to generate lexicons describing interactions in the form of meetings, letters and telephone calls to be compiled and excerpts describing these communications extracted from the report. A social network based on these excerpts combined with entities named in the report was constructed along with a network detailing how entities were related in the narrative of the Ryan Report through collocation within paragraphs (Fig. 2). The sub-corpus that was generated by automatically extracting excerpts detailing transfers was examined using association rule analysis revealing the recurring patterns governing the movement of religious staff often in response to allegations of abuse (Fig. 3).

Fig. 3. System-Wide Analysis of the Transfer of Religious Throughout the Industrial School System

3 Conclusion

The Industrial Memories project presents a state-of-the-art text mining application to navigate the findings of the Irish Government’s Inquiry into Child Abuse at Industrial Schools. A demonstration video of the system is available online². Text analytic techniques were used to uncover new insights and system-wide patterns that are represented in an interactive web based platform accessible to the general public.

References

1. Finkel, J.R., Grenager, T., Manning, C.: Incorporating non-local information into information extraction systems by gibbs sampling. In: Proceedings of the 43rd annual meeting on association for computational linguistics. pp. 363–370. Association for Computational Linguistics (2005)
2. Kitchin, R.: The data revolution: Big data, open data, data infrastructures and their consequences. Sage (2014)
3. Mikolov, T., Sutskever, I., Chen, K., Corrado, G.S., Dean, J.: Distributed representations of words and phrases and their compositionality. In: Advances in neural information processing systems. pp. 3111–3119 (2013)
4. Polikar, R.: Ensemble based systems in decision making. IEEE Circuits and systems magazine **6**(3), 21–45 (2006)

² <https://youtu.be/cV1xzuJ0dv0>