

Can websites reveal the extent and degree to which a business's values reflect national policy? A text embeddings approach

Alexander Hogan¹, Stephanie Cussans Moran² , Kevin Hogan³, Beth Barker⁴, Richard Woodall⁵

¹AI Research and Design, Etic Lab LLP, Wales, ²Etic Lab LLP, Wales, ³Etic Lab LLP, Wales, ⁴Etic Lab LLP, Wales, ⁵Etic Lab LLP, Wales.

How to cite: Hogan, A.; Moran S.C.; Hogan K., Barker B., Woodall R. 2024. Can websites reveal the extent and degree to which a business's values reflect national policy? A text embeddings approach. In: 6th International Conference on Advanced Research Methods and Analytics (CARMA 2024). Valencia, 26-28 June 2024. <https://doi.org/10.4995/CARMA2024.2024.17801>

Abstract

This paper demonstrates a novel application of text embeddings and vectorisation to website texts to measure the prevalence and proliferation of values associated with the Well-being of Future Generations (Wales) Act 2015 amongst Welsh businesses. This was achieved by training a model to recognise topics of particular importance in texts describing the WFGA, and then using this model to analyse the text hosted on Welsh business's websites. The companies studied comprised a structured national sample of the Welsh economy. Our findings indicate that the method detects meaningful differences between regions, business sector and participants in a government innovation support program.

Keywords: *Econometrics; sociometrics; policy evaluation; big data analytics; text vectorisation; text embeddings.*

1. Introduction

Recent years have posed a series of pressing challenges for measuring business performance. Heightened social, political and economic volatility (including the pandemic, inflation and energy price shocks) have subjected economies and businesses to an accelerated rate of change, meaning that traditional survey-based business intelligence has lagged behind real-time rapid change. At the same time, an increased digitalisation of businesses driven by the pandemic has seen companies' rapid adoption of remote working practices and increased online engagement. The increased frequency of digital activity creates an opportunity to develop robust digital business intelligence metrics that can be collected at scale in near real-time on an ongoing basis.

This paper presents a new approach to measuring the prevalence and proliferation of a particular set of values amongst a specific community of businesses. It describes how we trained a model

to measure the text on business websites against a unique piece of national legislation, the Well-being of Future Generations (Wales) Act 2015, using a sentence vectorisation method. We applied this to text scraped from a structured sample of company websites across Wales, representative by business type and unitary authority, in order to measure the extent and degree to which they represent the values put forward in the Act.

The question we aim to answer with this case study is whether our metric is capable of meaningfully discriminating across a sample. We measured this by looking at variations across unitary authority or by business sector (represented by SIC code), and variations registering the impact of a Welsh government business support program, within the sample. When tested against a structured sample of 5,431 Welsh companies, the results showed meaningful differences in the prevalence and contextual salience of values associated with the WFGA across regions, business type and participants in a government innovation support program.

Wales and the Well-being of Future Generations (Wales) Act 2015, herein the WFGA, represent an interesting test case for this technology. Firstly, Wales and the WFGA offer a well-defined, discrete entity; the Act perfectly overlaps with geographic area and the businesses we are measuring, as it applies to businesses across the whole of Wales. Moreover, the compact size of Wales makes it practical to build a representative sample of businesses by Unitary Authority (UA) that reflects the distribution of businesses in each UA by business sector (SIC code). Secondly, the WFGA requires government departments, businesses and other organisations to factor the health, environmental, sustainable, equitable and economic prospects of future generations into their decision making. That is to say, the WFGA requires a public response from government and businesses, that demonstrates they subscribe to the values and are trying to act on them. This can be measured more easily than, for example, carbon emissions or volume of waste recycled. Finally, it offers a case study of a regional and rural economy that differs from highly digitised and capitalised urban cosmopolitan regions, hence providing an interesting test for this methodology and whether it can produce meaningful results across urban and rural geographies with differing levels of digitisation (Wales Digital Maturity Survey 2023).

2. The WFGA

The WFGA is a unique piece of Welsh legislation, passed in 2015, which creates a legal obligation for:

“public bodies in Wales to think about the long-term impact of their decisions, to work better with people, communities and each other, and to prevent persistent problems such as poverty, health inequalities and climate change” (Future Generations Commissioner for Wales).

The aims of the act, expressed through seven “wellbeing goals”, are: a prosperous Wales, a resilient Wales, a more equal Wales, a healthier Wales, a Wales of cohesive communities, Wales

of vibrant culture and thriving Welsh language, and a globally responsible Wales. There is an increasing prevalence of schemes that utilise government training and procurement processes to incentivise businesses to sign up to social projects such as net zero, community wealth-building and social value; the WFGA represents a unified, strategic attempt to achieve this on a national scale with a distinctive set of goals and values.

The Act creates responsibilities for 48 public bodies across Wales, with a commissioner to oversee them. It requires them to think and act for the long term, to integrate their activities with other public bodies, to involve citizens and communities in planning and delivery, to collaborate with communities and other public bodies, and to focus on preventing public ills. Businesses are expected to deliver prosperity and growth in ways that protect the environment, create good jobs and generate wealth for Welsh communities. Many of the deliverables associated with Act are concerned with the relationship between government and business, including training, business support and procurement. In other words, part of the work of the Act is to communicate its values to Welsh businesses.

3. Methodology

We found and created two sets of website text samples: (1) a reference set of texts relating to the WFGA and (2) the text from a structured sample of 5,431 Welsh companies' website. Text embeddings were used to measure how closely different pieces of text resemble each other, making a comparison between the language used and frequency and saliency of key terms. The reference set of texts comprised 26,489 words that explain and represent the values and principles underpinning the legislation, drawn from the extant Welsh Government literature about the act and its implementation. We deployed text vectorisation to create a model for measuring the semantic similarity of large samples of unstructured text. This was subjected to a normalising process, by which we were also able to find co-incident topics that arise in the wild alongside the stated values of the WFGA. This repository was used to train the model.

We then created a sample of website texts to test against the model. We maintain a structured sample of 5,431 Welsh companies selected to represent the distribution of businesses by SIC code in each of Wales' 22 Unitary authorities, providing a highly representative picture of the make-up of the Welsh economy that has previously been used for a national survey (Hogan et al., 2023). We crawled these websites to a depth of 0 (only the front page), to retrieve January 2024 text. The model analysed the prevalence and salience of topics related to the WFGA, which would become a Welsh Future Generations (WFG) Score for each company. Text was extracted only from HTML found on the websites; any text found between the script and styling tags was included. At the time of collecting, 272 of the 5,431 in the sample were inaccessible.

We generate a WFG score by converting the target text and the reference text to vectors using Google's "Universal Sentence Encoder". The text of each as a whole was encoded, rather than

an aggregate of the sentence level. We then find the cosine distance between the vectors, giving us a score from -1 to 1 where -1 is as different from the reference texts as it's possible to be, and 1 is exactly the same as the reference texts. This method is based on two influential papers on Word2Vec (Mikolov et al 2013) and Universal Sentence Encoder (Cer et al. 2018). The application of these methods for semantic text representation has been recently reviewed (Worth 2023), and others have developed it for similar purposes of textual semantic comparison (Agarwala et al. 2021; Jangabylova, A. et al. 2022). There are precedents for using it to measure shifting values in public discourse that reflect policy change (Rodman 2020). The end product of our process is a metric, a number between 0 and 1, for each company that reflects the degree of similarity between the material produced by an individual organisation and the content of the reference WFGA-encoded set. In other words, this offers a measure of the extent and degree to which the content of a business's website reflects or embodies the span of issues represented in discussions of the legislation.

Our proposition is that a higher WFG score indicates that this business has engaged with the values expressed in the WFGA legislation, and has made an attempt to embody these values in its public communications. We hypothesise that the presence of text on a website that reflects the language and values of the Act may be used as a proxy measure for compliance with the Act. We believe this leap is possible because (a) the Welsh government is very clear about its intention, and clear that they want businesses to very closely replicate the Act in the language they use; (b) they are very prescriptive about what businesses should say; and (c) the Act and these values are very well documented by the Welsh government, with a large set of reference texts to train a model on. This leads us to have confidence that when we find a close text match on a business's website, it can be interpreted as an expression of the Act's values.

4. Results

The results below indicate that our method is capable of discriminating meaningful differences in the prevalence and contextual salience of values associated with the WFGA across regions, business sector and participants in a government innovation support program. To the extent that the WFGA represents an attempt to intervene in the values and priorities shaping the Welsh economy, the WFG scores of Welsh businesses provide a good overall indication of the Act's success in achieving its goals.

4.1. The WFG Score results against the national sample

The WFG Score for the national sample reflects a normal distribution, with most scores clustered between 0.3 and 0.4 (shown in Figure 1, below). This provides reasonable evidence that the metric is measuring a real variation in the digital presence of the target companies. It is to be expected that that distribution is shifted to the left, because a) it is highly unlikely that any

website would yield a perfect score, and b) the legislation is still relatively new, so it is still early in the process of educating the business community.




Figure 1. Distribution of WFG Scores for our national sample of Welsh companies.

4.2. The WFG Score results by Unitary Authority

The score for each Unitary Authority (UA) is averaged, to accommodate the fact that some UAs have many more businesses than others. Results (shown in Figure 2, below) indicate that the WFGA metric successfully distinguishes between UAs in Wales, with UAs representing the urban centres of Cardiff and the South scoring highest. This is to be expected, due to high business density, high levels of digitalisation and a high percentage of businesses selling to the public sector given their proximity to the seat of government.

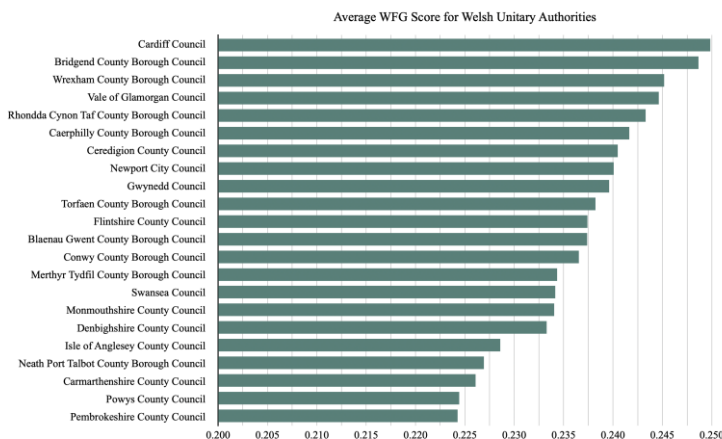


Figure 2. Average WFG Scores for Welsh Unitary Authorities.

4.2. The WFG Score results by business type

Our SIC group results (shown in Figure 3, below) also suggest that the WFGA metric successfully distinguishes between business sectors, with Business and IT (tending to be concentrated in the South and other urban areas) and Construction (tending to be engaged with government procurement and subject to strict criteria) scoring higher, while Retail/Wholesale Trade and Accommodation and Food Services are lower, as our graph below shows.




Figure 3. Average WFG Score by SIC Group.

We speculate that this difference may in part be due to the ways in which different kinds of business are required to comply with the Act. There are obvious places in which Construction businesses, for example, must comply, such as having environment policies on their websites, which is not the case for restaurants or holiday accommodation rentals. The highly digitalised kinds of services offered in the Business and IT sector trade significantly on credibility generated through their own digital properties, which is not true for many Accommodation and Food Services, who often use proprietary platforms such as AirBnB, Booking.com and Deliveroo to generate business. They do not therefore regularly update their websites, and so also miss the incidental and unconscious updates other sectors make regularly, that reflect more widespread changes in cultural values and attitudes promoted by the Act. This is substantiated by these sectors' lower Digital Growth and Digital Maturity recorded in the Wales Digital Maturity Survey (Welsh Digital Maturity Survey 2023).

4.3. The WFG Score results for a government-funded innovation support programme

Our results for a sample of 340 companies that took part in a government-funded innovation support programme (shown in Figure 5, below) are shifted to the right, meaning they show systematically higher scores than the national sample. This suggests that direct contact with the government tends to lead to higher WFG scores, whether through selection or training.

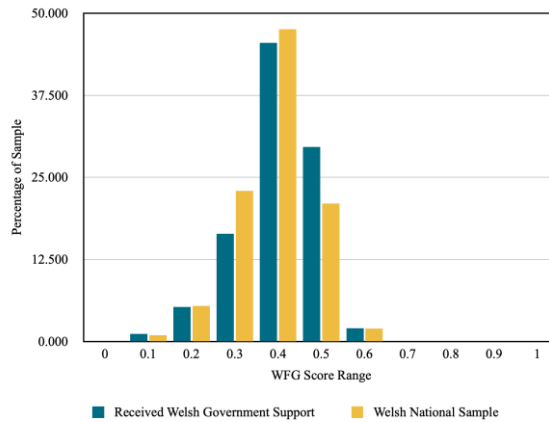


Figure 4. WFG Score comparisons.

4.4. English Counterfactual WFG Scores

We ran the WFG metric on an English counterfactual sample, where we matched English companies to the Welsh sample by sector (SIC code), company age, size (by accounts), and rural or urban (from their ONS postcode classification). As the graph below shows (Figure 5), the counterfactual sample performs considerably less well than the Welsh sample, as expected.




Figure 5. WFG Score of English counterfactual against the Welsh sample

5. Discussion

The results demonstrate the validity of the method, suggesting it is possible to use WebAI and word vectorisation to measure the prevalence of a set of ideas or values among a community of businesses, to measure differences between business type and geographic location, and to measure the impact of targeted government support programmes. We argue that it is possible to

make a meaningful comparison in this case between the reference texts and the target texts because of the volume of documentation and its prescriptiveness. Possible future applications of this include monitoring the effects of the act over time as well as monitoring and evaluating the results of specific interventions.

We have applied a similar method to developing other metrics, such as our Sense of Place metric, developed for a funding evaluation programme (Turner et al. 2024). We are currently working on deploying a DEI (diversity, equity and inclusion) metric to evaluate the DEI values of nursing colleges in the US; given the recent legal challenges and rapid change around DEI in US higher education (Murray et al. 2023) this is a timely task, offering the ability to monitor changing responses as they occur.

5.1. Limitations and drawbacks

This method relies on businesses having a website and the web domain being discoverable and validated. It works well for SMEs and large corporate entities, but may miss sole traders and micro businesses that do not have websites, that rely on word of mouth or social media platforms to do business.

5.2. Where next

Our next steps are to test the results against convergent validities, including testing against an English counterfactual sample and historic website data. We have a method for collecting website data for our sample over using the Wayback machine, and we know the date the legislation took effect (2015), so it should be possible to find out if there is a noticeable change in average scores before and after the Act's implementation.

Having established the premise for measuring the proliferation of values by regulatory bodies in the context of the WFGA, which represents a fairly well-defined sandbox, we would like to try applying it to other policy areas such as DEI, financial regulation and EU AI regulation.

References

- Agarwala, S. et al. (2021). Detecting Semantic Similarity Of Documents Using Natural Language Processing. *Procedia Computer Science*, Vol 189, 128-135, ISSN 1877-0509, doi: 10.1016/j.procs.2021.05.076.
- Cer, D. et al. (2018). Universal Sentence Encoder. arXiv:1803.11175v2 [cs.CL], <https://doi.org/10.48550/arXiv.1803.11175>.
- Future Generations Commissioner for Wales. Retrieved from <https://www.futuregenerations.wales/about-us/future-generations-act>.
- Jangabylova, A. et al. (2022). Greedy Texts Similarity Mapping. *Computation*, 10, 200. <https://doi.org/10.3390/computation10110200>.

- Mikolov, T. et al. (2013) Efficient Estimation of Word Representations in Vector Space. arXiv:1301.3781v3 [cs.CL], <https://doi.org/10.48550/arXiv.1301.3781>.
- Murray, T.A. et al. (2023). Anti-DEI legislation targeting colleges and universities: Its potential impacts on nursing education and the pursuit of health equity. *Nursing Outlook* 71 (2023), <https://doi.org/10.1016/j.outlook.2023.101994>.
- Rodman, E. (2020). A Timely Intervention: Tracking the Changing Meanings of Political Concepts with Word Vectors. *Political Analysis*, 28(1), 87–111. doi:10.1017/pan.2019.23.
- Turner, D., Ailes, O., Thomas, M., Hogan, A. "Welsh Government SMART Suite Evaluation". Cardiff, Welsh Government, Wavehill Limited. Interim 2024.
- United Nations (2021). Our Common Agenda: UN policy brief. Retrieved from <https://www.un.org/sites/un2.un.org/files/our-common-agenda-policy-briefs-a-quick-summary.pdf>
- Welsh digital maturity survey 2023. Retrieved from <https://welshdigitalmaturitysurvey2023.etiqlab.co.uk/>.
- Welsh Government (2022). Well-being of Wales: 2021. Retrieved from <https://www.gov.wales/wellbeing-wales-2021-html>.
- Worth, P. (2023) Word Embeddings and Semantic Spaces in Natural Language Processing. *International Journal of Intelligence Science*, **13**, 1-21. doi: 10.4236/ijis.2023.131001.