T05W03 / Application of recent advances in Natural Language Processing for Public Policy Research

Topic : T05 / Policy Design, Evaluation, Policy Analysis

Chair : Edgar Salas Gironés (Delft University of Technology)

Second Chair : Dewulf Art (Wageningen University)

Third Chair : Pradeep Murukannaiah (Delft University of Technology)

CALL FOR PAPERS

Recent developments in Artificial Intelligence, and particularly in Natural Language Processing (NLP), allow public policy researchers to work with large collections of texts. Such developments offer ample opportunities for taking public policy research into new directions. However, the use of NLP in public policy has been limited, despite its adoption in adjacent disciplines (e.g. political science). We believe this has occurred due to the limited venues for researchers using NLP and due to the fast pace of development of NLP techniques.

For the aforementioned reasons, this workshop has two aims. First, to bring together public policy researchers using NLP for their research, and secondly to discuss innovative approaches in which NLP can be used to study public policy concepts and theories. In doing so, this workshop intends to contribute to accelerate the adoption in public policy research. The guiding questions of this workshop are: How can Natural Language Processing (NLP) techniques be used to research public policy concepts and theories? Which new research directions does NLP allow policy scholars to pursue?

This workshop particularly encourages the submission of empirical research using NLP and other approaches using text-as-data, as well as invites contributions linking such techniques with the advancement of public policy concepts and theories. The workshop format consists of paper presentations, followed by discussions. This workshop intends to organize, in the final session, an open discussion with the workshop participants on how we can structure the emerging field of NLP for public policy research and discuss ways to go forward.

ABSTRACT

Over the last decade, the use of Natural Language Processing (NLP) has gained momentum in the social sciences, allowing to work with large quantities of text. NLP offers new possibilities on how to conduct public policy research: for instance, by allowing new qualitative research strategies, collecting data using new methods, or automating activities like coding.

The adoption of NLP in public policy research remains limited, despite having gained prominence in related disciplines, e.g. political science. Moreover, it has been primarily used for exploratory research such as using topic modeling. While applications of NLP in public policy research have remained inductive and exploratory, theory-driven applications of these methods are what holds most promise. For instance, NLP has been used to identify public values (Pelaez et al., 2023), to retrieve context-specific values (Liscio et al., 2021), to extract arguments from opinions (van der Meer et al., 2022), to identify and extract argumentative structures (Lawrence & Reed, 2019), and to automate manual coding of texts (Zhou, 2018). We observe that such methods open exciting opportunities for policy scholars in studying concept and theories, such as agenda setting (Zahariadis, 2016), policy integration (Biesbroek et al., 2020), policy framing (Van Hulst & Yanow, 2014), or policy diffusion (Linder et al., 2020), to name a few examples.

This workshop invites applications of public policy research using NLP techniques. The aims are to bring together policy researchers using NLP and to discuss how can NLP be applied to study public policy concepts and theories. The workshop guiding questions are: How can Natural Language Processing (NLP) techniques be used to research public policy concepts and theories? Which new research directions does NLP allow policy scholars to pursue? It strongly encourages applications of empirical research using NLP and will consist of paper presentations followed by interventions of two discussants.

T05W03 / Application of recent advances in Natural Language Processing for Public Policy Research

Chair : Edgar Salas Gironés (Delft University of Technology)

Second Chair : Dewulf Art (Wageningen University)

Third Chair : Pradeep Murukannaiah (Delft University of Technology)

Session 1Introduction, NLP and policy cohesion

Wednesday, June 26th 10:15 to 12:00 (AULA 15)

Intergovernmental relations in a federalist system: policy alignment and student achievement differences in the state of São Paulo, Brazil

Filipe Recch (University of Pittsburgh)

In this study, I focus on Brazil’s federalist system, giving explicit attention to its influence on the provision of equitable education, particularly for its most underprivileged citizens who rely heavily on public education. Central to my research is a novel and, to the best of my knowledge, unique dataset that I develop. This dataset encompasses all long-term educational plans from Brazil's complex administrative system of 27 states and 5,568 municipalities, providing a comprehensive overview of the country's educational policy landscape.

Utilizing Optical Character Recognition (OCR), Natural Language Processing (NLP), and documental analysis, I construct an index to measure the alignment between São Paulo's state long-term plan and each of São Paulo’s 645 municipalities. This index is key to understanding how the alignment of these educational plans impacts student test scores and educational inequality. By applying Structural Topic Modeling (STM) in my NLP approach, I identify key themes and topics across these documents. Leveraging the prevalences from the STM model, I quantify the policy alignment between municipal and state educational policies, measuring the cosine similarity between each municipal plan and the São Paulo State Education Plan (Plano Estadual de Educação - PEE).

Integrating this analysis with data from the São Paulo State School Assessment System (SARESP), focusing on student test scores in mathematics and language from grades 5 to 9, and controlling for socioeconomic status and management quality, my study offers significant insights. My preliminary findings indicate that municipalities aligned with state policies exhibit smaller educational attainment differences between state and municipal school students. This finding suggests that policy alignment is more than just a bureaucratic ideal; it is a pivotal tool for addressing educational inequality, especially in a country where many depend on public education. Proponents of a market-centric education model, which centers on fostering competition to boost educational quality, must reckon with the consequences of policy misalignment and the oft-resulting intensification of educational inequalities.

This research extends beyond the borders of Brazil, as my primary focus on Brazil has allowed me to trace broader global themes: In nations where policy decisions produce parallel education systems—whether through charter schools in the US, academy or grammar schools in the UK, or private schools more universally—there is a magnified risk of escalating educational inequality, with cascading effects on job market opportunities.

Looking forward, my research sets the stage for a deeper investigation into the causal relationship between policy alignment and educational outcomes, not just in Brazil but globally. Future work should explore how changes in municipal plans impact student achievement differences over time and how political factors influence policy alignment and, subsequently, educational quality and equity. This expanded focus will contribute to a better understanding of the mechanisms through which policy alignment can enhance educational equity and effectiveness, informing policymakers and stakeholders in the field of education globally.

T05W03 / Application of recent advances in Natural Language Processing for Public Policy Research

Chair : Edgar Salas Gironés (Delft University of Technology)

Second Chair : Dewulf Art (Wageningen University)

Third Chair : Pradeep Murukannaiah (Delft University of Technology)

Session 2NLP and policy implementation

Wednesday, June 26th 14:00 to 17:15 (AULA 15)

An Ideational Policy Analysis of How the Universal Benefit Reform Emerged in Finland Utilizing Large Language Models

Pasi Moisio (Finnish Institute for Health and Welfare (THL))

Merita Mesiäislehto (Finnish Institute for Health and Welfare)

Johanna Peltoniemi (Finnish Institute for Health and Welfare)

Mika Pihlajamäki (Finnish institute fo health and welfare)

This study explores the evolution of the idea of Universal Benefit into a concrete policy reform proposal in Finland from 2007 to 2023. During the 2010s, various political parties began suggesting radical yet ambiguous plans to completely reform the rather unique universal social insurance system of Finland. The idea of a universal benefit emerged in the political debates of the 2010s as an answer to simplify social security and as a 'counterforce' to basic income initiatives. The parliamentary Social Security Committee was given the task to prepare a model for the simplification in 2000. Following the proposal of the Committee’s midterm report, the government initiated the Universal Benefit reform in 2023, aiming to consolidate social insurance benefits of the working-aged population into a Universal Basic benefit, aligning with existing principles and social risk categories of the social insurance system.

The conceptual framework is constructed around two hypotheses. The main hypothesis posits that there was a political struggle over the conditionality of residence-based social security of working-aged in Finnish social policy agenda in the 2010s. The idea of basic income emerging was the main challenger. The secondary hypothesis suggests that the need for simplification of social security is a political valence issue and the smallest common political denominator. However, parties used the idea of simplification either in favor of, against, or neutrally towards conditionality. These two hypotheses and concepts (conditionality and simplification) provide a conceptual framework for the analysis.

We utilize open-source large language models (LLMs) like Llama 70B on the Finnish IT Center for Science (CSC), and commercial models like the OpenAI GPT-4 and Google Bard, to conduct content, discourse, and actor-centric analysis with a large set of policy documents. Utilized language models are capable of Finnish text analysis, enabling advanced natural language processing of Finnish policy documents. We also compared results with LLMs using English only text by translating documents first to English, as hypothesis was that analysis would work better in English. The primary data comprises election manifestos, complemented by a comprehensive collection of publicly available policy documents from political parties, government programs, and high-level expert groups pertaining to social security between 2007 and 2023. The analysis conducted with the LLMs was hierarchical due to data-processing limitations.

We conclude that the societal disjunctions occurring in the 2010s, namely the dissolution of established party coalitions and the breakdown of tripartite general agreements, created opportunities for new political actors and interest groups to influence the social insurance policy agenda. We argue that this decade was marked by a political contest over the principles, direction, and control of the residence-based social insurance system. While the benefit structure was “technically” consolidated, the core principles of the residence-based social insurance system were preserved. The period also witnessed a significant shift in social policy direction, characterized by an intensified focus on the long-term sustainability of public finances and enhanced emphasis on employment as an essential component of social security. (We prefer to present this paper on-site.)

Unveiling the Nature of Flemish Ministerial Advisors: An Empirical Study through NLP Techniques (2000-2020)

Tom Bellens (Katholieke Universiteit Leuven)

This study employs advanced Natural Language Processing (NLP) techniques to analyze the role of Flemish ministerial advisors between 2000 and 2020. The focus is on the integration of NLP methodologies into the exploration of a unique dataset comprising career data of these advisors, alongside a comprehensive collection of parliamentary documents from the same period. Our approach begins with the implementation of mention detection algorithms on the parliamentary dataset to identify references to the advisors. This step is crucial for correlating the names in our career dataset with specific parliamentary discussions.

Subsequently, we utilize coreference resolution tools to extract pertinent paragraphs related to these mentions. This process ensures that we gather all relevant textual data associated with each advisor, enhancing the comprehensiveness of our analysis. The extracted text is then transformed into structured data through an encoder such as BERT. This representation facilitates a robust statistical analysis, allowing us to delve into the roles of ministerial advisors.

The study aims to discern whether these advisors function more as politicians within a bureaucratic framework or as bureaucrats in a political context. Additionally, it investigates the variation within this group and examines the applicability of established typologies to the empirical data. By leveraging NLP applications, this work represents a significant methodological advancement in the study of political actors, providing nuanced insights into the dynamics of ministerial advisorship. The paper will be presented on site at IWPP4 in Guadalajara.

T05W03 / Application of recent advances in Natural Language Processing for Public Policy Research

Chair : Edgar Salas Gironés (Delft University of Technology)

Second Chair : Dewulf Art (Wageningen University)

Third Chair : Pradeep Murukannaiah (Delft University of Technology)

Session 3NLP and grand challenges

Thursday, June 27th 09:00 to 12:15 (AULA 15)

(Virtual) Multilingual topic modelling of United Nations environmental initiatives lessons from collaborations on food, water and climate

Anne Sietsma (Wageningen University)

In a variety of domains, the United Nations (UN) asks countries and other actors to submit large volumes of documents. Natural Language Processing, especially topic modelling, can be used to gain insights into such datasets and their underlying politics. Here, we use topic models to analyse three different UN processes: the Global Stocktake under the Paris Agreement, the Food Systems Summit and the Water Action Agenda. We find marked similarities, despite the differing domains and underlying politics: submissions broadly represent the interests of the submitting party; moreover, many submissions discuss fairly general topics, while regionally dominant- and politically sensitive topics are either left out of the submissions (e.g. dietary change) or left out of official summaries (the Global Stocktake). We discuss the consequences for transparency efforts at the international level and share practical lessons on the topic modelling of multilingual datasets. BERTopic, in particular, can make use of recent advances in Large Language Models to create multilingual topic models without the need for translation. At present, however, it is also a more complex model to understand and fine-tune, so we discuss when it is appropriate to use such methods over more established alternatives, such as non-negative matrix factorization (NMF) and Latent Dirichlet Allocation (LDA). We place our findings in the context of tracking climate change impacts and adaptation and stress that multilingual data sources can and should be considered more to address structural data inequalities, especially given the continuing rapid advances in multilingual NLP.

Authors

Anne J. Sietsma, Wageningen University

Robbert Biesbroek, Wageningen University

Jeroen Candel, Wageningen University

Rick W. Groenendijk, University of Amsterdam

Kirsty Holstead, Wageningen University

Johanna Köhler, Wageningen University

Agenda setting for climate change adaptation in the Netherlands and the UK: an NLP analysis

Dewulf Art (Wageningen University)

Robbert Biesbroek (Wageningen University)

Climate change adaptation policy comes about in a dynamic playing field where policy, politics, traditional media, and social media interact. More insight into how attention for climate change adaptation develops over time in policy decisions, political debates, media stories, and on social media platforms is critical for a better understanding of when and how climate action happens or fails to happen. In this study we focus on the agenda setting process for climate change adaptation over the past decade in the Netherlands and the UK. Existing research shows that agenda-setting interrelationships exist between media and politics, both in the amount of attention for issues and in the way they are framed, but the size and direction of these relationships differs across (sub-)issues and contexts. To better understand how climate change adaptation

travels across different agendas, we track the discussion on climate change adaptation on (social) media agenda, the political agenda and the policy agenda. We analyse and compare these dynamics in two countries, the Netherlands and the UK, which differ in the presence and nature of climate change impacts and adaptation actions. We use datasets obtained from the Twitter API to map the social media agenda, and datasets from national news sources for the media agenda. Datasets of parliamentary debates and questions provide insight into the political agenda for climate change adapation, and policy documents on climate change adapation were downloaded from government websites to capture the policy agenda. We make use of data-driven methods from Natural Language Processing (NLP), including word vectorization and topic modeling, to map and compare climate change adaptation across the different agendas.

Inferring Sceptical Stances of Climate Change on Tweets from Multi-modal Representation Models

Nan Bai (Wageningen University)

Dewulf Art (Wageningen University)

Tamara Metze (Delft University of Technology)

Climate change has been a heated public-policy topic in the recent decades, triggering debates in public arenas including social media. A wide range of stakeholders, including scientists, news media, policymakers, and action groups, actively express their views on this politically controversial issue, trying to make their voices heard and have influence on the political agenda. Understanding the stances of people regarding climate change, i.e., whether they support or deny it, is crucial. Whereas several studies have used textual information to infer climate change stances with the help of Natural Language Processing (NLP), the visuals are equally important throughout the debates. Visuals can be complementary and/or conflicting with texts, making both modalities critically informative for transferring messages. Recent advances in machine learning have made it possible to combine textual and visual information as high-dimensional vectors. Multi-modal representation models with Contrastive Language-Image Pre-Training (CLIP) are shown effective in fusing and generalizing the semantic similarities of both modalities. The resulting multi-modal embeddings are consistently used in downstream tasks such as topic modelling, semantic search, few-shot learning, etc. However, the CLIP models were trained on a general corpus and cannot necessarily distinguish the nuances between supports and sceptics towards a particular topic (e.g., climate change) per se.

This paper explores the capability of CLIP-based representation models to infer the multi-modal tweet stances towards climate change. A Dutch Twitter dataset concerning Sea-Level Rise (Zeespiegelstijging) composed of images and texts was collected. A retweeting network among users was constructed, where community detection algorithms demonstrated two main clusters: users who support the mainstream view on anthropogenic climate change issues, and users who are skeptical about the existence or anthropogenic nature of climate change. The user-level pseudo-labels agreed upon by domain-experts were assigned to the level of posts, i.e., pairs of texts and [possibly missing] images, resulting in a partially-labelled multi-modal dataset with 220,494 samples. After filtering out the repetitive and ambiguous examples, 29,306 posts showed climate-change mainstream stances, and 20,010 demonstrated skeptical stances. CLIP-based models were consulted to extract the textual and visual embeddings. After modality fusion and domain-knowledge enhancing as a newly-proposed pipeline, the multi-modal embeddings were fed into an additional Multi-layer Perceptron (MLP) classifier to infer whether a multi-modal post supports or denies climate change. Overlapping visuals were found to exist in posts of both stances. The intermediate layers of the trained MLP provided transformed embeddings that are sufficiently aware of climate change stances, improving the accuracy of zero-shot classification roughly from 62% to 72% for different CLIP variants.

The proposed methodological innovations yielded both a multi-modal dataset about climate change stances and an embedding model to obtain climate-aware vector representations, both of which are rare in literature but provide potentials in data science and public policy research. It can be further extended to improve domain-specific multi-modal topic modelling and to assist multi-modal framing analysis, preparing for a more systematic understanding of discourse coalitions during public debates and policy controversies in the context of climate change and beyond.

T05W03 / Application of recent advances in Natural Language Processing for Public Policy Research

Chair : Edgar Salas Gironés (Delft University of Technology)

Second Chair : Dewulf Art (Wageningen University)

Third Chair : Pradeep Murukannaiah (Delft University of Technology)

Session 4NLP and science, technology, & innovation

Thursday, June 27th 14:00 to 17:15 (AULA 15)

(Virtual) Using natural language processing to assess institutionalization of techno-economic paradigms: an evaluation of European Commission reports on the bioeconomy

Matt Wilder (University of Toronto)

Large corpora of textual data contain valuable policy-relevant information that cannot be easily processed using conventional research methods (e.g., qualitative analysis and manual coding). Challenges are particularly difficult when documents contain technical information that is beyond researchers' subject matter expertise. Fortunately, recent advances in unsupervised machine learning have the potential to overcome these challenges. The proposed paper employs structural topic modelling to analyze European Commission reports on the bioeconomy published between 2005 and 2023. Substantively, the method is shown to be useful for ascertaining the institutionalization of techno-economic paradigms, which may be undermined by diminishing returns, policy layering, and policy drift. Evidence of all three of these factors is found in the model output, which conveys the evolution of topical content of European Commission reports on the bioeconomy over time.

The paper is organized in the format of a research note. The first section describes the problem of limited information processing capacity and implications for theory testing in the social sciences. The second section covers the data generation process, specifically as it relates to pre-processing PDF files of varying formats (e.g., wide format versus columned text). The third section covers optical character recognition (OCR) and automated methods for cleaning text. The fourth section describes the final stages of pre-processing, namely: chunking, stemming, and stop word removal. The fifth section delves into the intuition of topic modelling by comparing Latent Dirichlet Allocation (LDA) and Non-Negative Matrix Factorization (NMF). The sixth section discusses methods for estimating the number of topics to be included in the model. The seventh section reviews the model output and methods of theory-driven interpretation, as informed by the political economy literature on paradigm stability and change. The eighth section confirms the robustness of the findings to both NMF and LDA applications and adjustments to the number of topics called. The ninth section concludes with a discussion of applications of the approach to applied public policy, including the recent adoption of the technique by the Government of Canada to analyze the content of workplace accident reports.

Revealing semantics: evaluating the implementation and effectiveness of the EU Smart Specialisation Strategy

Deyu Li (Utrecht University)

The goal of the place-based Smart Specialisation Strategy (S3) policy in the European Union is to identify and develop new and more competitive activities in the region. It aims to set up priorities by leveraging existing strengths, identifying hidden diversification opportunities, and generating novel platforms upon which regions can build competitive advantage in new activities (Foray, 2015). S3 policies started to be implemented in 2014 for a period of 7 years (2014-2020), and were a precondition for regions to receive funding from the European Regional Development Fund. More than 200 S3 policies have been developed by EU regions. However, concerns have been raised about the implementation of S3 policies that it is not clear whether regions actually receive funding in the priorities they set up (Gianelle et al., 2020). It is also not clear how effective S3 policies are in terms of promoting the development of new technologies.

This paper aims to provide an quantitative evaluation of the implementation of the EU S3 policy. This paper

uses original text data from the description of the policy priorities identified in S3 documents and the description of projects funded by the European Regional Development Fund (ERDF) (Bachtroegler et al., 2020), and the description of technological classes in the International Patent Classifications. Applying the pre-trained sentence transformer model all-mpnet-base-v2 (Reimers and Gurevych, 2019), this paper first maps the semantic similarity between the description of ERDF funded projects in region to the description of priorities in the S3 policy of the same region to identify projects that focused on the priorities in the region's S3 policy. Second, we matches the description of the ERDF funded projects with the technological classes of the International Patent Classifications in the similar way. Doing so, this paper manages to quantify the amount of funding each region received in the prioritized area they set up, and the effectiveness of S3 policy regarding promoting the development of new technology specialisation of regions.

The methodology of this paper is in line with the scheme of the workshop applying natural language processing for public policy. The finding of this paper will also be relevant for the discussion of the effectiveness of place-based innovation policy.

Reference

Bachtroegler, J., Doussineau, M., Reschenhofer, P., 2020. Dataset of projects co-funded by the ERDF during the multi-annual financial framework 2014-2020. JRC Research Reports, JRC Research Reports.

Foray, D., 2015. Smart Specialisation: Opportunities and Challenges for Regional Innovation Policy. Routledge.

Gianelle, C., Guzzo, F., Mieszkowski, K., 2020. Smart Specialisation: what gets lost in translation from concept to practice? Regional Studies 54, 1377–1388. https://doi.org/10.1080/00343404.2019.1607970

Reimers, N., Gurevych, I., 2019. Sentence-BERT: Sentence Embeddings using Siamese BERT-Networks. https://doi.org/10.48550/arXiv.1908.10084

T05W03 / Application of recent advances in Natural Language Processing for Public Policy Research

Chair : Edgar Salas Gironés (Delft University of Technology)

Second Chair : Dewulf Art (Wageningen University)

Third Chair : Pradeep Murukannaiah (Delft University of Technology)

Session 5NLP and public participation

Friday, June 28th 09:00 to 12:15 (AULA 15)

What Voices are Heard in the Regulatory Policymaking Process? Key Point Analysis of Regulatory Comments using NLP

Edgar Salas Gironés (Delft University of Technology)

Pradeep Murukannaiah (Delft University of Technology)

Public participation is an important means of influencing policy changes. In principle, given an intended policy change, policymakers solicit public opinion on specific areas or aspects of the intended change (Bobbio, 2019). The participatory process broadens the views to consider in the design process, allowing to incorporate both ‘old’ and ‘new’ actors (e.g., organized interests and ordinary citizens).

As new tools of public participation are adopted, key questions emerge about which actors influence policy changes and to what extent. For instance, given a regulatory change, one may consider that an actor successfully influenced policy change if their views are included in the final draft of the policy. Addressing such questions through a qualitative research design may be infeasible, due to resource constrains (e.g., time and manpower). In this paper, we propose a novel technique to semi-automate this process, using a Natural Language Processing (NLP) method called Key Point Analysis (KPA). In short, KPA is a summarization task, which extracts important statements (e.g., arguments), called key points, from a corpus. Non-selected statements are later mapped to key points (Alshomary et al., 2021). Thus, KPA allows summarizations that consider both relevance (based on the number of statements associated to a key point) and quality (by selecting high-quality arguments as key points).

To test this approach, we use data from regulations.gov, which contains public comments submitted by individuals and organizations regarding a proposed regulatory change in the United States. We apply our approach to different regulatory changes (also known as dockets) in different policy domains. Each docket contains a proposed rule and a final rule, as well as public comments. In this presentation, we will describe our methodological approach—to apply KPA to policy texts, results, and forms of evaluation. Preliminary findings show that KPA can be used to analyze how the content of final regulation relates to comments. Additionally, KPA allows us to map comments to areas or aspects that policymakers were interested in opening for public participation. Although our contribution is based only on data from regulations.gov, a similar approach could be used with different data sources.

References

Alshomary, M., Gurcke, T., Syed, S., Heinisch, P., Spliethöver, M., Cimiano, P., Potthast, M., & Wachsmuth, H. (2021). Key Point Analysis via Contrastive Learning and Extractive Argument Summarization. In K. Al-Khatib, Y. Hou, & M. Stede (Eds.), Proceedings of the 8th Workshop on Argument Mining (pp. 184–189). Association for Computational Linguistics. https://doi.org/10.18653/v1/2021.argmining-1.19
Bobbio, L. (2019). Designing effective public participation. Policy and Society, 38(1), 41–57. https://doi.org/10.1080/14494035.2018.1511193
Howlett, M. (2019). Designing Public Policies: Principles and Instruments (2nd ed.). Routledge. https://doi.org/10.4324/9781315232003

The effect of large language modelling on government transparency: the case of regulations.gov

Alex Ingrams (Institute of Public Administration - Leiden University)
Pradeep Murukannaiah (Delft University of Technology)
Sarah Giest (Leiden University)
Edgar Salas Gironés (Delft University of Technology)

In 2018 the New York Attorney General, Letitia James, investigated the case of the Federal Communications Commission (FCC) proposed deregulation of net neutrality rules (Office of the New York State Attorney General, May 2021). During the heated political debate that ensued, a secondary talking point developed after it emerged that the majority of the millions of public comments submitted on the matter through the open regulation site regulations.gov were actually fake comments that had been generated by machine learning algorithms. In the wake of the controversy, the executive office responsible for regulations.gov, the General Services Administration introduced a series of new features to the public commenting platform regulations.gov that would detect and block mass comments campaigns. Individual agencies also took matters into their own hands. In order to prevent unwanted comments from entering into the public and internal review discussion, they adopted more aggressive approaches while cognizant of the constraints of the Administrative Procedures Act to determine which comments should be considered acceptable for consideration in the regulatory review process and, ultimately, for publication on the platform

This paper reports on the repercussions of this series of events for the transparency and accountability of the open regulations process. Prior scholar ship has widely considered the impact that algorithms such as large language models (LLMs) have on administrative rules and processes (e.g., Coglianese and Lehr, 2017; Ingrams, 202 3; Kowalski, Esteve, and Jankin Mikhaylov , 2020). However, we still understand very little about the success and impacts of the remedial approaches that public organizations have undertaken including their own use of Natural Language Processing and LLMs. In this paper, the authors examine the approaches of regulations.gov and individual agencies to tackle fake public posts. It also draws on evidence from 100 ‘high volume’ public comments dockets on regulations.gov. These dockets have thousands or millions of public comments and the majority of them withheld a certain proportion of the comments from publicity. The characteristics of the missing comments are considered and the processes involved in detecting and withholding spurious comments is reproduced using a variety of modeling techniques in order to understand what impact this has on the public comments process.

References:

Coglianese, C., & Lehr, D. (2019). Transparency and algorithmic governance. Administrative Transparency and algorithmic governance. *Administrative Law aw Review, 71*(1), 1--56.56.

Ingrams, A. (2023). Do public comments make a difference in open rulemaking? Insights from information management using machine learning and QCA analysis. information management using machine learning and QCA analysis. *Government Information Quarterly, 40*(1), 101778.

Kowalski, R., Esteve, M., & Jankin Mikhaylov, S. (2020). Improving public services by mining citizen feedback: An application of natural language processing. feedback: An application of natural language processing. *Public Administration, 98*(4), 1011--1026.

Office of the New York State Attorney General (May, 2021). *Attorney General James Issues Report Detailing Millions of Fake Comments, Revealing Secret Campaign to Influence FCC’s 2017 Repeal of Net Neutrality Rules*. https://ag.ny.gov/press-release/2021/attorney-general-james-issues-report-detailing-millions-fake-comments-revea