

Right-wing Mnemonics

Phillip Stenmann Baun¹, Kristoffer Nielbo^{2,3}

¹Department of Global Studies, Aarhus University, 8000 Aarhus C, Denmark

²Center for Humanities Computing Aarhus, Aarhus University, 8000 Aarhus C, Denmark

³Interacting Minds Centre, Aarhus University, 8000 Aarhus C, Denmark

Abstract

This paper presents a natural language processing technique for studying memory on the far-right political discussion forum /po1/ on 4chan.org. Memory and the use of history play a pivotal role on the far-right for temporally structuring beliefs about social life and order. However, due in part to methodological limitations, there is a lack of knowledge regarding the specific historical entities that make up the far-right memory culture and wider historiography. To better grasp the structure of far-right memory, this paper opts for a data-intensive methodology, using machine learning on a data set of approximately 66 million posts from /po1/ from 2020. 19,821 random posts were manually annotated, according to the presence of historical entities. After evaluating interrater reliability, data were used to train a naïve Bayes text classifier to learn the lexical features of so-called “posts of memory” (POMs). After parameter tuning, the model extracted from the dataset a total of 1.083.471 POMs with a precision score of 98.43%. It is argued that this technique provides a novel way to automate the identification of historical entities within the far-right authored text, of benefit for the fields of memory studies and far-right studies, two fields that have traditionally relied on more qualitative close-reading approaches. By investigating the mnemonic features of the /po1/ posts during steps in the methodological pipeline, the paper contributes important insights into the challenges of identifying and classifying lexical features in hyper-vernacular digital spaces like 4chan, where communication is highly defined by intertextuality, semantic ambiguity, and cacography.

Keywords

4chan, Text classification, Far-right memory, Media and memory, Right-wing extremism

Introduction

Memory plays a pivotal role for far-right groups as it does in most processes of collective identity formation. Previous work on the memory practices of far-right groups and actors has emphasized, how elements from the past are strategically energized within this ideological environment as symbolic and cultural building blocks for forming identity, strengthening group affiliations, systematizing ideological strands, and directing contemporary political objectives [2, 3, 5, 8, 10, 16, 22, 25, 32, 38, 40, 41]. Collectively imagined beliefs about social life and order, borne out of today’s far-right mnemonic practices, can function as key analytical entryways to understanding those concomitant ideals that structure their particular worldview.

CHR 2022: Computational Humanities Research Conference

✉ psb@cas.au.dk (P. S. Baun); kln@cas.au.dk (K. Nielbo)

🌐 <https://psbaun.github.io/> (P. S. Baun); <https://knielbo.github.io/> (K. Nielbo)

🆔 0000-0002-2433-8212 (P. S. Baun); 0000-0002-5116-5070 (K. Nielbo)



© 2020 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

CEUR Workshop Proceedings (CEUR-WS.org)

No study has to our knowledge mapped the far-right memory-scape in its totality. This is due in part to the limited ability of qualitative methods to make generalizable claims about the wider far-right milieu. Instead, such methodological approaches are more often employed for the purpose of ethnographic studies of select subsections or case studies of the far-right, such as the memory practices of single actors or groups [25, 38]. As a result, there is limited knowledge about the general historical elements that bind together this overarching far-right collective memory.

Large-scale data-driven approaches offer a partial solution to this issue of generalizability. More specifically, statistical techniques for supervised and unsupervised learning make it possible to analyze the lexical qualities of memory on the far-right in ways normally unfeasible by traditional close-observational studies.

To reconstruct the memory-scape of the far-right, this paper presents an observational study that combines systematic annotation and text classification to the ‘politically incorrect’-board (colloquially known as /po1/) of 4chan.org, a social media site and chat forum well known for harboring far-right views and content [6, 9, 17, 26, 36, 37, 39].

Methods

Data

Posts from 4chan’s /po1/-board were randomly extracted from a dataset of threads from 2020 using 4chan’s API. In total 19,821 random posts were manually annotated by one of the authors in *Doccano* [27] for the presence of memory, sorting all posts into two non-overlapping groups: posts of memory (*POM*) and posts not of memory (*non-POM*). Conceptually, memory on /po1/ is defined as *reference to historical entities in posts* targeting those narrative categories in which a disorganized past becomes meaningful in discourse. Without assuming a priori the precise contents of such entities, they are lexically and semantically speaking related to historical signifiers such as the names of past individuals, events, periods, etc. In order to delineate between entities conceptually understood by users as belonging to either the past or present, the concept only considers entities related to before the year 2000.

With this approach out of the 19,821 posts, 1,236 meet the criteria as being *POMs*, with the other 18,585 posts consequently being classified as *non-POMs*. In other words, 6.24% of posts contained references to a historical entity before the year 2000.

Validation

To test annotator reliability, three independent raters were recruited and tasked with annotating a subset of the 19,821 posts. After being instructed in the coding procedure, they were each provided with a dataset of 500 posts consisting of 50/50 randomly selected posts from each category (*POM/non-POM*). Cohen’s κ was calculated for the original annotations and each rater independently: Rater 1 $\kappa = .94$, rater 2 $\kappa = .92$, and rater 3 $\kappa = 0.82$. Landis and Koch’s benchmark for interpreting kappa values suggests that a κ of more than 0.81 is considered ‘almost perfect’ [24]. Fleiss’ κ – recommended for multiple raters determining among nominal

categories [13, 15] – were also calculated for all raters $\kappa = 0.86$. Statistics like Scott’s π and Krippendorff’s α coefficient yielded similar results [23, 33]. We take the high level of annotation agreement of *POMs* as strong evidence of the original annotations’ reliability.

Text Normalization and Representation

Before model training, a series of text normalization transformations were applied to the posts, specifically lemmatization, case-folding, and removal of punctuation. The posts were subsequently structured using a vector space model of lexical features, specifically the empirical probability of words and sequential combinations (i.e., n-grams). The n-gram range was treated as a parameter value together with minimum and maximum document feature frequency.

Parameter Optimization & Model Training

For model training, data were balanced with under-sampling, that is, by randomly sampling 1, 236 *non-POMs* without replacement, corresponding to the total number of *POMs*, and removing the rest of the majority category, resulting in a balanced dataset of 50% posts from each category. As opposed to over-sampling (i.e., multiplying data from the minority category to match the level of the majority category), research suggests that under-sampling provides better sensitivity at the cost of specificity [1]. Given that we are interested in accurately detecting true positives of *POMs*, under-sampling was the more sensible approach. Following a similar logic, precision was chosen as the performance measure, since precision is defined as the number of true positives (ground truth *POMs*), divided by the number of predicted positives (model classified *POMs*, either correctly or incorrectly), thereby giving a sense of the exactness of how well the model is able to find actual *POMs*, whilst limiting the number of false positives.

The multinomial Naive Bayes algorithm was chosen as a classification technique due to its computational efficiency and level of explainability [35, 31]. The algorithm is based on Bayes’ Theorem and computes the posterior probability for the target and non-target class (*POM* vs. *non-POM*) for each document using the Bayes rule and assigns the document to a class based on the maximum posterior probability, or formally, the probability of a document d belonging in class c , $P(c | d)$ is:

$$P(c | d) \propto P(c) \prod_{i=1}^m P(t_i | c) \quad (1)$$

and the class of a document d is then computed as:

$$c_{MAP} = \arg \max_{c \in \{c_1, c_2\}} P(c | d) \quad (2)$$

Model parameters were optimized using a train/test split ratio of 75/25 with 149, 760 candidate parameter values for nine parameters using a five-fold cross-validation method, totaling 748, 800 fits. When comparing an unoptimized model without default parameters to the optimized model, precision increased from 63.49% to 76.19%. By shifting the decision threshold value, making the model more discriminant in only classifying posts with a predicted probability over 0.9 as a *POM*, precision increased to 98.43%.

A note on the interpretation of parameter optimization, the optimal n-gram range was [1, 5], suggesting that historical entities can be expressed through a multitude of complex word sequences. The minimum document frequency of a feature was determined as two documents (i.e., the model filters out words that appear in less than two posts). The fact that the optimal minimum document frequency is not higher, suggests that even very rare words contribute to model learning, seeing as they might represent obscure historical entities. Conversely, optimal maximum document frequency was determined to be 30%, a somewhat extensive upper boundary, suggesting that the corpus is riddled with many common words that have little significance in terms of learning.

Results

To understand the classifier's behavior, we looked at the specific n-grams that were most predictive of the two categories. By counting the number of times each word occurs in all posts divided by the number of posts in each category, we calculated the percentage of times a word has appeared in one or the other category. In other words, by dividing the number of times a word appears in the *POM* category by the number of times it appears in the *non-POM* category, we are able to calculate a 'memory ratio' for each word and rank them according to this scale. To avoid division by zero in the instances where a word only occurs in one category, a pseudo-count of 1 was added to all word counts, similar to the Laplace smoothing technique used to regularize the naive Bayes algorithm to avoid a probability estimate of zero when a feature value does not occur in a given category. Looking at the top and bottom of this sorted list of n-grams in table 1, we see what words most strongly correlate with each category i.e., what features are most and least 'memory-like'.

The results indicate that the model correctly picks up on key historical entities that to a human interpreter is understandable as having mnemonic significance, for example, distinct historical persons, events, or places like 'Hitler', 'Rome', 'Holocaust', 'Weimar' and 'WW2', but also more general historical or temporal terms such as 'history', 'century', 'ancient', 'civilization', and 'ancestor'. There are also less historically signifying terms like 'catholic', 'Germany', or 'Europe', implying that these have been used in mnemonic context across multiple posts for the model to start associating these terms with the *POM* category. Such re-contextualization of generally nondescript terms is also interesting in the cases of 'build' and 'steal', which indicate that /pol/ users ascribe specific sentiments to their mnemonic discussions, expressed through not only historical entities themselves, but also with descriptive modifiers such as verbs.

Conversely, the features in table 1 with the lowest memory ratio expectedly signify little mnemonic substance or contain less contextual meaning (such as words like 'she', 'her', 'anyway', 'no no', and 'ah'). The low scores of words like 'Biden', 'virus', 'test', and 'death' also indicate that discussions about memory do not overlap with discussions about contemporary events in 2020, such as the election of Joe Biden and the COVID-19 pandemic. Terminology specific to 4chan like 'tfw' (that face when), 'kek' (synonymous to LOL; laughing out loud), 'image', 'post', as well as 'rare' (referring to the 'rare flag' meme about users from small or unfamiliar countries), are also unrelated to the *POM* category, suggesting that discussions about memory are a distinct subtheme on /pol/, branched off from more general 4chan topics.

Table 1Ratio for top and bottom 20 words most and least related to the *POM* category.

Memory words	Memory ratio	Non-memory words	Memory ratio
hitler	13.75	biden	0.24
history	6.87	virus	0.31
rome	6.21	twf	0.32
century	5.61	kek	0.32
catholic	4.83	she	0.34
pyramid	4.63	her	0.36
holocaust	4.55	death	0.38
ancient	4.50	test	0.38
germany	4.35	make you	0.40
soviet	4.26	anyway	0.40
weimar	4.18	image	0.40
civilization	4.17	love	0.40
ancestor	4.09	no no	0.42
build	4.05	fit	0.43
ww2	3.88	post	0.43
the pyramid	3.80	lel	0.43
steal	3.76	ah	0.44
nazi	3.69	rare	0.45
europe	3.67	community	0.45
slave	3.63	google	0.46

Looking at the predicted probability of *POMs* provided by the model, many posts with high predicted probability are also generally lengthier and mention several historical entities. For example, one post with a word length of 133 (decidedly higher than the corpus average of 44 number of words per post), and with a 99.23% probability of being a *POM*, repeats the word ‘USSR’ four times as well as consisting of other significant historical entities, such as ‘Soviet Union’ and the names of former Russian leaders. In contrast, scanning through posts with low *POM* probability, these are much shorter (posts with only 10% *POM* probability contain on average 23 words), and use many unspecific words.

Discussion

Examples from the data set can provide further insights into the model’s learning behavior. Consider a post from the dataset that reads: ‘enjoy your bat meat, you fucking barbarians.’ While the term ‘barbarian’ was originally used in ancient Greece to refer to non-Greek peoples based on cultural-linguistic differences, and would therefore adhere to the target class criteria, the context within which it is used here – involving the zoonotic origin of the coronavirus – is detracted from its original historical context. This is despite the remnant of its historically rhetorical function still undergirding its usage in the posts (i.e., to ‘uncivilize’ someone by stereotyping them as barbarian). The semantic ambiguity – that there is not a direct correlation between the use of a term and that term always represents a historical entity – necessarily

complicates the machine learning procedure, because the model relies on the assumption that there exists an unambiguous lexical distinction between the *POM* and *non-POM* category's textual content. Semantic ambiguity at a lexical level tends to be the rule rather than the exception in natural language, and hence an inevitable condition of any machine learning process dealing with unstructured text from a real-life environment. Some terms will inherently fluctuate in their potential for expressing memory.

Depending on the circumstance, the composition of entities can be more or less metonymically representative of that particular conceptual category to which they are assumed to belong. By exploring the lexical features of an abstract category such as memory, we are not only taking the necessary, precautionary steps of transparently revealing the data that goes into the machine learning model but are also shedding light on the complex ways that memory is expressed through language in decidedly ambiguous ways. This can be demonstrated from an example in the dataset: 'we wuz vikangz.' Briefly put, the post is a satirical rehashing of the meme colloquially known as 'we wuz kangz', which was originally directed towards a type of Afrocentric memory concerning the disputed and anachronistic claim that ancient Egypt was a black civilization. Consequently, the rehashing is now being used to satirize the pretense of people claiming to have Viking ancestry. While certainly interesting on its own as a case for how memory can be embedded in multilayered intertextual contexts, the post also directs our attention towards a specific characteristic of the data. That is, historical entities can be represented in text by lexical symbols that may be synonymous in their meaning but which are orthographically heterogeneous in their written format. The words 'Vikings', and 'Vikangz' both connote the same historical entity, but since they are spelled differently it blocks the machine learning model from disambiguating their semiotic synonymy.

The presence of homonyms also needs to be taken into consideration as these also distort the informational quality of the data. As an example, the word 'Rome' was often used to refer to the historical Mediterranean civilization but is also used to denote the capital of present-day Italy. Likewise, 'Romans' would generally refer to the people of ancient Rome, but would also figure in discussions of biblical scripture, referring to the Letter to the Romans, part of the Pauline epistles. While rare, these homonymic cases exemplify some of the linguistic features of the data that need to be considered when constructing the machine learning model.

The subtle variations of meaning and semantics explored in the dataset present difficulty for accurately extracting memory on /po1/. Nevertheless, by highlighting these cases, the article has identified some of the preliminary linguistic characteristics that undergird discourses about memory on /po1/, which ultimately become important to consider when building and evaluating machine learning models for 4chan-specific and/or history-related classification problems. Consequently, a major contribution of this study lies in its impelling of a still rather nascent and sparse field of research that employs various natural language processing techniques in the study of memory in the digital realm specifically, e.g. [42, 34, 14, 19, 12, 18].

It also speaks directly to the difficulty of operationalizing a socially situated memory concept in contrast to more individual-oriented studies attempted by neurobiologists and cognitive psychologists, most recently pointed out by [29] and [28]. Similar critiques of the supposed conceptual unclarity of "memory", regarding the concept's supposed over-extension and semantic overloading (for example in the metaphorical misuse of psychological terms such as "trauma" in supra-individual social contexts), resulting in redundant and unsophisticated uses

of the concept as a rhetorical signal, rather than as a clearly defined analytical tool, have also been pointed out by [20, 4, 11, 30, 21] and [7]. Recognizing this criticism within the field of memory studies, this article has attempted to offer a more systematized and reproducible way of identifying and understanding memory, specifically in the context of memory's lexical manifestation.

However, we do not argue to have presented anything approaching a *sui generis* methodology, completely separate and distinct from previous endeavors. If anything, the detailing of many of the elements of note involved in applying machine learning, from evaluating statistical biases to the exploration of linguistic variations, shows how a project like this requires an inclusive, interdisciplinary outlook capable of combining and balancing the potential from multiple perspectives: including both theory and method, the close reading spotlight and the distant reading bird's eye view, human awareness and computer industriousness.

There are also certain limitations involved with this method. As has been pointed out previously, there is a general disconnect on both a practical and theoretical level between the lexical data that the models were trained on and those conceptual definitions that structured the qualitative interpretation of that data. While "historical entities" as a concept, allegorically symbolizing objects of some abstract and intangible past, may be comprehensible on a theoretical level as the basic constituents of mnemonic communication, there is no guarantee that such a concept, once operationalized, is bijectional translated into exactly replicable lexical items making up a conversation on /po1/. In other words, the word tokens that are the fundamental components of the machine learning models are dimensionally, linguistically, and conceptually speaking different from the theoretical definition of memory, even when such inferential parallelism is assumed in the study's design. The polysemantic and cacography of natural language thus impedes machine learning, for example when it has to determine the mnemonic significance of ambivalent features, or when it needs to learn from quirky spelling and terminology.

Moreover, given the multifarious representations that conceptually fall under the category of historical entities, there is also a certain bias related to the initial, manually coded dataset. Even though the dataset was constructed from a random sampling of /po1/ posts, there are quite likely numerous historical entities that did not make it into this relatively small subset of the general /po1/ conversation, even if they might have appeared relatively frequently and would have been important to the board's collective memory. Such biases would then be replicated in the model's algorithm, resulting in favoritism of historical entities that statistically might describe very well the core memory on /po1/, represented by historical entities and other "memory lingo" so common as to appear distinctly in the random sampling. However, it would be blind to the "edges" of this collective memory, to the speckling of historical entities that the model never got a chance to see, and which were subsequently not included in the extraction process and therefore most likely wouldn't feature as some of the top words in the final topic modeling. This is, of course, but an unavoidable condition of necessarily decomplexifying the exceedingly amorphous cultural concept of memory into quantifiable entities suitable for study. With any perspective striving for the macroscopic, there is bound to be a corresponding loss in detail.

References

- [1] S. Alsaif and A. Hidri. “Impact of Data Balancing During Training for Best Predictions”. In: *Informatica* 45 (2021). DOI: 10.31449/inf.v45i2.3479.
- [2] A. J. Bauer. “The Alternate Historiography of the Alt-Right: Conservative Historical Subjectivity from the Tea Party to Trump.” In: *Far-Right Revisionism and the End of History: Alt/Histories*. Ed. by L. D. Valencia-García. Routledge, 2020, pp. 120–137.
- [3] P. S. Baun. “Memory and far-right historiography: The case of the Christchurch shooter”. In: *Memory Studies* 15.4 (2022), pp. 650–665. DOI: 10.1177/17506980211044701. URL: <http://journals.sagepub.com/doi/abs/10.1177/17506980211044701>.
- [4] D. Berliner. “The Abuses of Memory: Reflections on the Memory Boom in Anthropology”. In: *Anthropological Quarterly* 78.1 (2005), pp. 197–211.
- [5] H. Betz. “Revisiting Lepanto: the political mobilization against Islam in contemporary Western Europe.” In: *Patterns of Prejudice* 43.3-4 (2009), pp. 313–334.
- [6] K. Bezio. “Ctrl-Alt-Del: gamergate as a precursor to the rise of the alt-right.” In: *Leadership* 14.5 (2018), pp. 556–566.
- [7] A. Confino. “Collective Memory and Cultural History: Problems of Method”. In: *The American Historical Review* 102.5 (1997), pp. 1386–1404.
- [8] J. Dozier. “Hate Groups and Greco-Roman Antiquity Online: To Rehabilitate or Reconsider?” In: *Far-Right Revisionism and the End of History: Alt/Histories*. Ed. by L. D. Valencia-García. Routledge, 2020, pp. 251–296.
- [9] B. Elley. ““The rebirth of the West begins with you!”—Self-improvement as radicalisation on 4chan”. In: *Humanities and Social Sciences Communications* 8.1 (2021), p. 67. DOI: 10.1057/s41599-021-00732-x. URL: <https://doi.org/10.1057/s41599-021-00732-x>.
- [10] A. B. Elliott. “Internet medievalism and the White Middle Ages”. In: *History Compass* 16.3 (2018), e12441. DOI: <https://doi.org/10.1111/hic3.12441>. URL: <https://compass.online.library.wiley.com/doi/abs/10.1111/hic3.12441>.
- [11] J. Fabian. “Remembering the Other: Knowledge and Recognition in the Exploration of Central Africa”. In: *Critical Inquiry* 26 (1999), pp. 49–69.
- [12] M. Ferron and P. Massa. “Beyond the encyclopedia: Collective memories in Wikipedia”. In: *Memory Studies* 7 (2013), pp. 22–45. DOI: 10.1177/1750698013490590.
- [13] J. L. Fleiss. “Measuring nominal scale agreement among many raters”. In: *Psychological Bulletin* 76.5 (1971), pp. 378–382. DOI: 10.1037/h0031619.
- [14] R. García-Gavilanes, A. Mollgaard, M. Tsvetkova, and T. Yasseri. “The memory remains: Understanding collective memory in the digital age”. In: *Science Advances* 3.4 (2017), e1602368. DOI: [doi: 10.1126/sciadv.1602368](https://doi.org/10.1126/sciadv.1602368). URL: <https://www.science.org/doi/abs/10.1126/sciadv.1602368>.
- [15] N. Gisev, J. S. Bell, and T. F. Chen. “Interrater agreement and interrater reliability: key concepts, approaches, and applications”. In: *Res Social Adm Pharm* 9.3 (2013), pp. 330–8. DOI: 10.1016/j.sapharm.2012.04.004.

- [16] R. Griffin. "Fixing Solutions: Fascist Temporalities as Remedies for Liquid Modernity". In: *Journal of Modern European History* 13.1 (2015), pp. 5–23. DOI: 10.17104/1611-8944\2015\1\5. URL: <https://journals.sagepub.com/doi/abs/10.17104/1611-8944%5C%5F2015%5C%5F1%5C%5F5>.
- [17] G. Hawley. *The Alt-Right: What Everyone Needs to Know*. New York: Oxford University Press, 2019.
- [18] A. Jatowt, D. Kawai, and K. Tanaka. *Digital History Meets Wikipedia: Analyzing Historical Persons in Wikipedia*. Conference Paper. 2016. DOI: 10.1145/2910896.2910911. URL: <https://doi.org/10.1145/2910896.2910911>.
- [19] N. Kanhabua, T. N. Nguyen, and C. Niederée. "What triggers human remembering of events? A large-scale analysis of catalysts for collective memory in Wikipedia". In: *IEEE/ACM Joint Conference on Digital Libraries* (2014), pp. 341–350.
- [20] W. Kansteiner. "Finding Meaning in Memory: A Methodological Critique of Collective Memory Studies". In: *History and Theory* 41.2 (2002), pp. 179–197.
- [21] K. L. Klein. "On the Emergence of Memory in Historical Discourse". In: *Representations* 69 (2000), pp. 127–150.
- [22] C. Kølvrå. "Embodying 'the Nordic race': imaginaries of Viking heritage in the online communications of the Nordic Resistance Movement". In: *Patterns of Prejudice* 53.3 (2019), pp. 270–284. DOI: 10.1080/0031322x.2019.1592304. URL: <https://doi.org/10.1080/0031322X.2019.1592304>.
- [23] K. Krippendorff. *Content analysis: An introduction to its methodology*. 3rd edition. Thousand Oaks, CA: Sage, 2013.
- [24] J. R. Landis and G. G. Koch. "The Measurement of Observer Agreement for Categorical Data". In: *Biometrics* 33.1 (1977), pp. 159–174. DOI: 10.2307/2529310. URL: <http://www.jstor.org/stable/2529310>.
- [25] C. Miller-Idris. *The extreme gone mainstream: commercialization and far right youth culture in Germany*. Princeton: Princeton University Press, 2017.
- [26] A. Nagle. *Kill All Normies: Online Culture Wars from 4chan and Tumblr to Trump and the Alt-Right*. Zero Books, 2017.
- [27] H. Nakayama, T. Kubo, J. Kamura, Y. Taniguchi, and X. Liang. *doccano: Text Annotation Tool for Human*. 2018. URL: <https://github.com/doccano/doccano>.
- [28] J. Olick. *The Politics of Regret: On Collective Memory and Historical Responsibility*. Routledge, 2007.
- [29] J. K. Olick, Vinitzsky-Seroussi, V., and D. Levy. "Introduction". In: *The Collective Memory Reader*. Ed. by D. L. Jeffrey K. Olick Vered Vinitzky-Seroussi. Oxford University Press, 2011, pp. 3–62.
- [30] G. J. R. "Memory and Identity: The History of a Relationship." In: *Commemorations*. Ed. by J. R. Gillis. Princeton: Princeton University Press, 1994, pp. 3–24.

- [31] S. Raschka. “Naive Bayes and Text Classification I - Introduction and Theory”. In: *arXiv:1410.5329 [cs]* (2017). URL: <http://arxiv.org/abs/1410.5329>.
- [32] I. Richards. “A Philosophical and Historical Analysis of “Generation Identity”: Fascism, Online Media, and the European New Right”. In: *Terrorism and Political Violence* 34.1 (2022), pp. 28–47. DOI: 10.1080/09546553.2019.1662403. URL: <https://doi.org/10.1080/09546553.2019.1662403>.
- [33] W. A. Scott. “Reliability of Content Analysis: The Case of Nominal Scale Coding”. In: *The Public Opinion Quarterly* 19.3 (1955), pp. 321–325. URL: <http://www.jstor.org/stable/2746450>.
- [34] Y. Sumikawa, A. Jatowt, and M. Düring. “Digital History meets Microblogging: Analyzing Collective Memories in Twitter”. In: *Proceedings of the 18th ACM/IEEE on Joint Conference on Digital Libraries* (2018).
- [35] S. L. Ting, W. H. Ip, and A. H. C. Tsang. “Is Naïve Bayes a Good Classifier for Document Classification?” In: *International Journal of Software Engineering and Its Applications* 5.3 (2011), p. 11.
- [36] M. Tuters. “Esoteric Fascism Online: 4chan and the Kali Yuga.Far-Right Revisionism and the End of History.” In: *Far-Right Revisionism and the End of History: Alt/Histories*. Ed. by L. D. Valencia-García. Routledge, 2020, pp. 287–303.
- [37] M. Tuters and S. Hagen. “(((They))) rule: Memetic antagonism and nebulous othering on 4chan”. In: *New Media & Society* 22.12 (2019), pp. 2218–2237. DOI: 10.1177/1461444819888746. URL: <https://doi.org/10.1177/1461444819888746>.
- [38] L. D. Valencia-García. *Far-Right Revisionism and the End of History: Alt/Histories*. Routledge., 2020.
- [39] M. Wendling. *Alt-Right: From 4chan to the White House*. London: Pluto Press, 2018.
- [40] R. Wodak and B. Forchtner. “Embattled Vienna 1683/2010: right-wing populism, collective memory and the fictionalisation of politics”. In: *Visual Communication* 13.2 (2014), pp. 231–255. DOI: 10.1177/1470357213516720. URL: <https://doi.org/10.1177/1470357213516720>.
- [41] D. Wollenberg. “The new knighthood: Terrorism and the medieval.” In: *Postmedieval* 5.1 (2014), pp. 21–33.
- [42] C.-m. A. Yeung and A. Jatowt. *Studying how the past is remembered: towards computational history through large scale text mining*. Conference Paper. 2011. DOI: 10.1145/2063576.2063755. URL: <https://doi.org/10.1145/2063576.2063755>.