

Topic Matters: How Linguistic Properties can Shape Reading Behaviour in Selective Exposure Studies

Thomas Krämer

thomas.kraemer@gesis.org

GESIS - Leibniz Institute for the Social Sciences
Cologne, Germany

Thomas Kosch

thomas.kosch@hu-berlin.de

HU Berlin
Berlin, Germany

Dagmar Kern

dagmar.kern@gesis.org

GESIS - Leibniz Institute for the Social Sciences
Cologne, Germany

Daniel Hienert

daniel.hienert@gesis.org

GESIS - Leibniz Institute for the Social Sciences
Cologne, Germany

Abstract

Research on selective exposure frequently relies on eye tracking to study reading behaviour, often assuming that texts across different controversial topics are comparable once basic controls are applied. This assumption is problematic if topic-dependent linguistic properties systematically shape how users read and allocate attention. We therefore examine whether such properties relate to differences in reading behaviour in selective exposure contexts. We analyse linguistic features and eye-tracking data from a laboratory study in which 68 participants searched for and read news articles on climate change and migration policy. Our results reveal systematic differences in both textual characteristics and reading behaviour across topics. These findings identify an important methodological confound in selective exposure research and highlight the need to account for topic-specific linguistic properties when interpreting eye-tracking measures and designing systems intended to mitigate biased information consumption.

CCS Concepts

• **Human-centered computing** → *Empirical studies in HCI*; *Laboratory experiments*; HCI design and evaluation methods.

Keywords

Eye tracking, Reading behaviour, News consumption, Selective Exposure

1 Introduction

Public discourse in online information environments is increasingly shaped by controversial and polarizing topics. As users are confronted with large volumes of potentially conflicting information, concerns have grown about how individuals select and attend to information. One prominent explanation for these patterns of information consumption is selective exposure [27], understood as a behavioural manifestation of the broader phenomenon of confirmation bias [23], both of which have long been central topics in political science and are discussed as key mechanisms contributing to echo chambers and filter bubbles (e.g., [10, 14, 22]). In recent years, these concepts have also garnered increasing

attention in Human-Computer Interaction (HCI), particularly in the form of cognitive biases [3]. HCI research extends this discourse by investigating how interface design and system features can either exacerbate or mitigate selective exposure of information [16].

A substantial body of HCI work explores design interventions aimed at reducing selective exposure, for example, through source position indicators [18], source expertise indicators [19], stance and credibility labels [8], or more broadly through the development of bias-aware systems [1]. A fundamental prerequisite for such systems is the reliable identification of selective or biased user behaviour. To this end, researchers employ a variety of behavioural and physiological measures, including interaction logs [9, 14], eye-tracking data [29, 30], and other physiological signals such as electrodermal activity, and brain hemodynamic responses [1], as well as electroencephalography (EEG) [15].

Recent advances in eye-tracking technology, such as the availability of webcam-based approaches and growing possibilities for real-time analysis using Artificial Intelligence ¹, or web-based use of professional eye-trackers outside the lab [26], lead to an increased interest in detecting selective exposure through eye-movement behaviour. Eye movements provide indicators of cognitive processing during reading. For instance, longer fixation durations are commonly associated with higher processing load [13]. Fixations, saccades, and derived reading measures vary with individual differences [28], as well as with lexical properties and comprehension demands [12, 24, 25].

Empirical studies on selective exposure, with and without eye tracking, frequently rely on controlled experimental settings, often using mock-up websites to simulate news feeds (e.g., Twitter/X in [2] or Facebook in [29]), partisan online news sites [9], or curated sets of arguments that participants are asked to read and evaluate [18]. The primary advantage of these mock-ups lies in the researchers' ability to control textual properties, such as text length [2, 17] or readability (e.g., Flesch Reading Ease scores; [1]). Some studies additionally annotate texts with respect to stance [7] or evaluative polarity (positive, negative, neutral, as in [30]). Alongside these controlled approaches, more naturalistic studies have examined selective exposure in free web searches [30] or news search tasks [15], without the ability to control the selected and read text.



This work is licensed under a Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International License.

¹<https://sdk.eyedid.ai/>

Across both controlled and naturalistic studies, researchers typically investigate selective exposure and confirmation bias in the context of controversial or highly debated topics, including abortion rights [2], political progressivism [1], the death penalty [18, 19], vegetarianism [2, 18], same-sex marriage [9], migration [15, 29], and climate change [1, 15]. While different topics are often used to test the robustness or generalizability of selective exposure effects, the texts associated with these topics are typically treated as functionally equivalent once basic controls (e.g., length or readability) are applied. At the same time, prior research has demonstrated that individual characteristics such as topic interest [1, 29], topic involvement [17], low-effortful thinking [2], and strong political beliefs [2] shape selective exposure behaviour. However, comparatively little attention has been paid to how topic-specific textual properties themselves might influence selective exposure factors.

In this paper, we analyse textual properties and reading behaviour from a user study on selective exposure using two controversial topics: climate change and migration policy. We observe systematic differences in both linguistic features and eye-movement patterns across the two topics.

2 User Study and Dataset

We analysed data from a laboratory user study conducted in the context of selective exposure to address the following research question: How do linguistic properties of texts on different controversial topics relate to differences in reading behaviour in selective exposure contexts?

In the following, we describe the aspects of the user study relevant to answering this research question. The user study included 68 participants (native German speakers, 42 female, 24 male, and 2 others, age: $M=27$, $SD=7.9$) and was designed to simulate a real-world scenario of searching for news online. Each participant completed two eight-minute search tasks² in a counterbalanced order. The search engine³ indexed a corpus of 192,685 German-language news articles crawled from mainstream outlets via the Fundus news crawler⁴. Participants completed one task on migration policy and one task on climate change. These topics were selected because they frequently trigger public controversy and polarized discourse [20]. After the search and reading phase, participants rated a subset of sentences for familiarity and agreement. These data are not analysed in the present paper. The study protocol received institutional ethics approval.

The study was conducted in a laboratory setting, and eye movements were recorded using a Tobii Pro Spectrum eye tracker at 300 Hz. We used EyeLiveMetrics [11] to compute eye-tracking metrics on the sentences read by participants, including fixation metrics (e.g., Fixation Duration, Fixation Count), saccade metrics (e.g., Saccades Count, Rightward Saccades Length), and reading metrics (e.g., Regression-Path Duration, Re-Reading Duration). As EyeLiveMetrics provides word-level metrics, we aggregated these measures to the sentence level. To account for variation in

sentence length, all metrics (except Rightward Saccades Average Length, Rightward Saccades Average Amplitude and Rightward Saccades Scan Speed) were normalized by the number of words. To analyse the relationship between linguistic properties and reading behaviour, we restricted our analysis to the sentences participants read. We further applied the following filtering steps: Only the first visit to a page was included. Subsequent visits after navigating away were excluded. Sentences were included only if their total fixation duration was at least 100 ms, which is considered the minimum time required for basic lexical or semantic processing of one or two words of a sentence [13, 25]. After filtering, our dataset included 10,030 sentences for migration policy and 10,387 for climate change. Sentences could appear multiple times in the dataset when read by different participants.

For analysing the textual properties, we used the ELFEN (Efficient Linguistic Feature Extraction for Natural Language Datasets) framework⁵[21]. ELFEN has been developed and used for the classification of human-generated vs. LLM-generated content[6]. To enable an unbiased comparison of linguistic properties across topics, we considered only unique sentences for this analysis. This resulted in 4,095 sentences for migration policy (average sentence length 16.9 tokens, $SD = 9.6$) and 4,381 unique sentences for climate change (average sentence length 17.3 tokens, $SD = 10.4$)⁶. The average sentence length did not differ significantly between topics.

3 Results and Discussion

3.1 Effects of Topic on Stimulus Properties

In the text-centric analysis, we analysed 208 different linguistic metrics for all read sentences computed by the ELFEN framework. The metrics are based on the entire sentence text, not only the fixated parts. The metrics were normalized by the number of tokens (e.g., words, punctuation, or white space) in a sentence. Some metrics are expected to be highly correlated, for example, average word length (`avg_word_length`) and the number of long words (`n_long_words`). We therefore first computed a Spearman's rank correlation matrix, then clustered metrics with a correlation threshold of > 0.8 , and selected one representative metric per cluster with the highest variance, reducing the number to 148 metrics and limiting the over-representation of closely related metrics. We computed a Linear Model⁷ for each z-scaled linguistic metric across topics, allowing the regression coefficients to be interpreted and compared as standardized effect sizes. We control the False Discovery Rate (FDR) by applying Benjamini-Hochberg correction. We found a significant difference ($p < .001$) between the two topics, migration policy and climate change, for 26 linguistic metrics.

In Figure 1, only those metrics are visualized that reached a minimum standardized effect size of $|\beta| \geq 0.1$ to filter out negligible variations that have $p_{adj} \leq 0.001$. The labels refer to the acronyms and definitions from Universal Dependencies⁸ as also referred to in the ELFEN Framework. The analysis reveals linguistic divergence

²Task: Your train is delayed, and you now have some unexpected free time. You decide to catch up on the current debate about [migration policy/climate change]. Use the news search and enter any search terms you like.

³Based on <https://www.elastic.co/elasticsearch>

⁴<https://github.com/flairNLP/fundus/tree/master/src/fundus/publishers/de>.

⁵Details see <https://elfen.readthedocs.io/en/latest/tutorials.html#ratio-normalize>

⁶Tokenisation via spaCy with the German News transformer model https://spacy.io/models/de#de_dep_news_trf

⁷R package `lm()`

⁸<https://universaldependencies.org/docs/>

across 19 metrics between the two topics. The circle size represents the variance explained by the topic (R^2), and the position indicates the magnitude of the difference.

The metrics with the highest effect measured for climate change are the number of hapax legomena (# GLOBAL TOKEN HAPAX LEGOMENA, $\beta = +0.15$), indicating tokens that occur only once in our sentence corpus, and the number of words with few synonyms (# LOW SYNSETS, $\beta = +0.12$), indicating more specialized vocabulary. For migration policy, the metrics with the highest effects are the number of verbs in past tense (# VERB TENSE PAST, $\beta = -0.12$) and pronouns with male genus (# PRON GENDER MASC, $\beta = -0.12$).

While the topics are statistically distinguishable across many linguistic dimensions (26 metrics), the low R^2 values suggest that topic-specific vocabulary and grammar reflect minor differences, and hence the results should be interpreted with caution. For instance, climate change exhibited a slightly more technical and informational profile, defined by high noun density, a diverse and specialized vocabulary (hapax legomena), and higher Age of Acquisition (AOA) scores, while migration policy appears as a more narrative style, utilizing significantly more personal pronouns and past-tense verbs, indicating a focus on human actors and historical events.

3.2 Effects of Topic on Reading Behaviour

Except Rightward Saccades Average Length, Rightward Saccades Average Amplitude and Rightward Saccades Scan Speed, eye-tracking measures were word-normalized to capture visual processing effort per unit of text rather than total sentence length, thereby aligning the gaze analysis with the sentence-level linguistic analysis. We modelled sentence-level gaze metrics using Generalized Linear Mixed Models (GLMM) to account for the hierarchical data structure (sentences nested within participants) and unbalanced trial numbers. For each gaze metric, we fit a model with the factor topic, with migration policy as reference category, and random intercepts for participant and sentence.

Table 1 shows differences of gaze metrics between the two topics. For the First-Pass Fixation Duration the -4.8% difference ($p = .043$) suggests that migration-related vocabulary required slightly more effort during early lexical processing. Regression metrics in the first-pass reading of a sentence, such as Re-Reading Duration, Regression Path Duration, and Selective Regression Path Duration, showed the most substantial effects, with higher durations for migration policy (9.6% to 13.6% increase, $p < .01$). This suggests that participants return significantly longer to previous parts of the sentence during the first-pass reading of the topic migration policy. Higher Regression-Path Duration indicates difficulties in integrating words when they are fixated [5, p. 349], while Re-Reading Duration indicates general comprehension difficulties [5, p. 363] or provides information about revision and reanalysis [4]. For Look-back Time, the 4.5% change ($p = .034$) for the topic migration policy indicates that readers were significantly more likely to return to previous sentences after reading further in the text to resolve, for example, ambiguities.

Table 1: Differences between eye-tracking metrics with respect to the reference category migration policy to climate change in percent. Negative values correspond to lower values for sentences read in the topic climate change compared to migration policy. All metrics are normalized by word per sentence, if not otherwise stated.

Metric	% Δ	p	Sig	SE	R_m^2	R_c^2
1st Pass Regression-Path Duration	-11.9	<.001	***	.034	.002	.407
1st Pass Re-Reading Duration	-13.6	<.01	**	.048	.002	.396
1st Pass SRPD	-9.6	<.01	**	.034	.001	.429
Look Back Time ($\geq$ 2nd Pass)	-4.5	.034	*	.022	.000	.440
1st Pass First Fixation Duration	-4.8	.043	*	.024	.000	.543
Fixation Duration	-3.2	.077	n.s.	.019	.000	.482
Fixation Count	-3.1	.083	n.s.	.018	.000	.485
Saccades Count	3.4	.263	n.s.	.030	.000	.456
Rightward Saccades ($\emptyset$ amplitude)	-2.3	.278	n.s.	.021	.000	.396
Rightward Saccades Count	3.1	.312	n.s.	.030	.000	.453
Rightward Scan Speed	1.1	.391	n.s.	.013	.000	.294
1st Pass Duration	-1.1	.633	n.s.	.023	.000	.398
1st Pass Regression	-0.7	.719	n.s.	.021	.000	.192
Rightward Saccades Length	-.6	.854	n.s.	.031	.000	.154
Rightward Saccades Length ($\emptyset$ in px)	0.0	.999	n.s.	.021	.000	.373

3.3 Relationship of Reading Behaviour and Linguistic Sentence Properties

To better understand the relationship between linguistic properties of the text and gaze behaviour, we computed the correlation matrix of the 19 linguistic metrics identified in Section 3.1 and the 5 significant gaze metrics from Section 3.2. Due to the characteristic right-skew of gaze metrics (confirmed via Q-Q plots and Shapiro-Wilk tests), Spearman's rank correlations were employed to provide a robust measure of the relationship between linguistic features and gaze metrics.

Again, given the high dimensionality of the data, we controlled the False Discovery Rate (FDR) while maintaining sufficient statistical power to detect meaningful associations that might be suppressed by more conservative family-wise error rate corrections (e.g., Bonferroni). In Figure 2 only those associations are shown that remained significant after Benjamini-Hochberg correction, with a significance threshold set to $\alpha = 0.001$.

We observed several correlations between linguistic and gaze metrics. First-Pass First Fixation Duration is affected similarly across both topics: Determiners and pronouns (# DET Number Sing, # DET Gender Neut, and # PRON Gender Masc/Fem) show the strongest negative correlations ($\rho \approx -0.05$ to -0.10), indicating that higher proportions of these features are associated with shorter first-fixation durations on a word. A similar pattern is observed for infinitives and past tense verbs. First-Pass First Fixation Duration is surprisingly also reduced with a higher density of words that occur only once (Global Hapax Legomena) or twice (Dislegomena) in the sentence corpus. Because Hapax Legomena are defined relative to our specific stimulus corpus rather than a general language corpus, even words that are frequent in the language can function as hapaxes if they occur only once in our dataset. Only a higher density of nouns and of words with few synonyms consistently increased First-Pass First Fixation Duration for both topics.




Figure 1: Linguistic metrics: effects with at least 0.1 standard deviation difference and maximum 0.001 p_{adj} (FDR-corrected significance level). The circle sizes refer to R^2 . With migration policy as reference category in the linear model, $\beta > 0$ means the feature ratio is higher in Climate Change sentences, $\beta < 0$ means the feature ratio is higher in Migration Policy sentences (in standardized units).

Our results show topic-dependent associations between Regression-Path Duration and Selective Regression-Path Duration and noun-related features. Noun density and dative proper-noun density have a negative correlation in migration policy but a positive one in climate change. Pronoun density reverses sign (positive in migration policy and negative in climate change) for Regression-Path Duration. One possible explanation is that in migration discourse, nouns might refer to concrete actors, whereas in climate change, nouns might be more abstract, technical concepts that are harder to capture.

The positive correlations between pronouns and verbs and Re-Reading Duration suggest that these features increase late-stage integration effort. Although identified quickly during the First-Pass First Fixation, they appear more difficult to incorporate into the broader sentence structure. Look-back Time patterns are consistent across both topics: Low synonym density increases Look-back time⁹, while Hapax (Dis)Legomata reduces it.

4 Conclusion

In this paper, we analysed textual properties and reading behaviour from a selective exposure user study comparing news articles on two controversial topics: climate change and migration policy. By examining linguistic characteristics of the texts alongside participants' eye-movement patterns, we identified systematic differences in both textual features and reading behaviour across topics. While the observed correlations between textual features and reading measures are relatively small, their statistical significance across thousands of observations suggests that linguistic features reliably

capture aspects of processing effort during natural reading. In the context of the complex cognitive processes involved in reading, even modest effects can account for consistent variance in behavioural measures.

Taken together, these findings indicate that topic-dependent textual structure can meaningfully influence reading behaviour and, consequently, the interpretation of selective exposure measures based on eye-tracking data. Our results highlight the importance of accounting for topic-specific textual characteristics in selective exposure research and in the design of systems intended to mitigate biased information consumption, rather than assuming equivalence across content domains. Furthermore, in the context of bias-aware systems, user interfaces might take these findings into account and dynamically adapt, e.g., via glossary overlays for technical content to support specialized vocabulary, or with context anchors, such as pronoun-referent highlighting in narrative content.

Acknowledgments

For this publication, AI was used in accordance with the AI guidelines of our institution and the ACM. During the preparation of this work, ChatGPT v5 and Gemini v3 were used for grammar and style editing and for minor coding tasks. All content was reviewed and edited by the authors, who take full responsibility for the accuracy and final publication of this work.

This work is funded by the German Research Foundation (DFG) project "Overcome Selective Exposure in Web Search by Considering Eye Movements and Physiological Signals" (Project-ID 525041402). Furthermore, this work is supported by the German Research Foundation (DFG), CRC 1404: "FONDA: Foundations

⁹thresholds for low synonym density: 2 or less synonyms




Figure 2: Linguistic metrics with influence on eye-movement behaviour. Features shown are those that reached statistical significance (Benjamini-Hochberg-corrected $p \leq .001$) in either or both the Migration Policy and Climate Change sentences.

of Workflows for Large-Scale Scientific Data Analysis” (Project-ID 414984028).

References

- [1] Nattapat Boonprakong, Xiuge Chen, Catherine Davey, Benjamin Tag, and Tilman Dingler. 2023. Bias-Aware Systems: Exploring Indicators for the Occurrences of Cognitive Biases When Facing Different Opinions. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*. Association for Computing Machinery, New York, NY, USA, 1–19.
- [2] Nattapat Boonprakong, Saumya Pareek, Benjamin Tag, Jorge Goncalves, and Tilman Dingler. 2025. Assessing Susceptibility Factors of Confirmation Bias in News Feed Reading. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems (CHI '25)*. Association for Computing Machinery, New York, NY, USA, Article 418, 19 pages. doi:10.1145/3706598.3713873
- [3] Nattapat Boonprakong, Benjamin Tag, Jorge Goncalves, and Tilman Dingler. 2025. How Do HCI Researchers Study Cognitive Biases? A Scoping Review. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems (CHI '25)*. Association for Computing Machinery, New York, NY, USA, Article 473, 20 pages. doi:10.1145/3706598.3713450
- [4] Charles Clifton, Fernanda Ferreira, John M. Henderson, Albrecht W. Inhoff, Simon P. Liversedge, Erik D. Reichle, and Elizabeth R. Schotter. 2016. Eye movements in reading and information processing: Keith Rayner’s 40year legacy. *Journal of Memory and Language* 86 (2016), 1–19. doi:10.1016/j.jml.2015.07.004
- [5] Charles Clifton, Adrian Staub, and Keith Rayner. 2007. Chapter 15 - Eye movements in reading words and sentences. In *Eye Movements*, Roger P.G. Van Gompel, Martin H. Fischer, Wayne S. Murray, and Robin L. Hill (Eds.). Elsevier, Oxford, 341–371. doi:10.1016/B978-008044980-7/50017-3
- [6] Esra Dönmez, Maximilian Maurer, Gabriella Lapesa, and Agnieszka Falenska. 2025. AI Argues Differently: Distinct Argumentative and Linguistic Patterns of LLMs in Persuasive Contexts. In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*, Christos Christodoulopoulos, Tanmoy Chakraborty, Carolyn Rose, and Violet Peng (Eds.). Association for Computational Linguistics, Suzhou, China, 34595–34626. doi:10.18653/v1/2025.emnlp-main.1755
- [7] Tim Draws, Oana Inel, Nava Tintarev, Christian Baden, and Benjamin Timmermans. 2022. Comprehensive Viewpoint Representations for a Deeper Understanding of User Interactions With Debated Topics. In *Proceedings of the 2022 Conference on Human Information Interaction and Retrieval (Regensburg, Germany) (CHIIR '22)*. Association for Computing Machinery, New York, NY, USA, 135–145. doi:10.1145/3498366.3505812
- [8] Mingkun Gao, Ziang Xiao, Karrie Karahalios, and Wai-Tat Fu. 2018. To Label or Not to Label: The Effect of Stance and Credibility Labels on Readers’ Selection and Perception of News Articles. *Proc. ACM Hum.-Comput. Interact.* 2, CSCW, Article 55 (Nov. 2018), 16 pages. doi:10.1145/3274324
- [9] R. Kelly Garrett. 2009. Echo chambers online?: Politically motivated selective exposure among Internet news users. *Journal of Computer-Mediated Communication* 14, 2 (01 2009), 265–285. doi:10.1111/j.1083-6101.2009.01440.x
- [10] R. Kelly Garrett. 2009. Politically Motivated Reinforcement Seeking: Reframing the Selective Exposure Debate. *Journal of Communication* 59, 4 (12 2009), 676–699. doi:10.1111/j.1460-2466.2009.01452.x
- [11] Daniel Hienert, Heiko Schmidt, Thomas Krämer, and Dagmar Kern. 2024. EyeLiveMetrics: Real-time Analysis of Online Reading with Eye Tracking. In *Proceedings of the 2024 Symposium on Eye Tracking Research and Applications (Glasgow, United Kingdom) (ETRA '24)*. Association for Computing Machinery, New York, NY, USA, Article 81, 7 pages. doi:10.1145/3649902.3656495

- [12] Halszka Jarodzka and Saskia Brand-Gruwel. 2017. Tracking the reading eye: Towards a model of real-world reading. 193–201 pages.
- [13] Marcel A Just and Patricia A Carpenter. 1980. A theory of reading: from eye fixations to comprehension. *Psychological review* 87, 4 (1980), 329.
- [14] Silvia Knobloch-Westerwick and Jingbo Meng. 2011. Reinforcement of the Political Self Through Selective Exposure to Political Messages. *Journal of Communication* 61, 2 (04 2011), 349–368. doi:10.1111/j.1460-2466.2011.01543.x
- [15] Thomas Krämer, Daniel Hienert, Francesco Chiossi, Thomas Kosch, and Dagmar Kern. 2025. Escaping the Filter Bubble: Evaluating Electroencephalographic Theta Band Synchronization as Indicator for Selective Exposure in Online News Reading. In *Proceedings of the Extended Abstracts of the CHI Conference on Human Factors in Computing Systems (CHI EA '25)*. Association for Computing Machinery, New York, NY, USA, Article 228, 7 pages. doi:10.1145/3706599.3720078
- [16] Thomas Krämer, Francesco Chiossi, and Thomas Kosch. 2025. Evaluating Eye Tracking and Electroencephalography as Indicators for Selective Exposure During Online News Reading. In *Proceedings of the CHI 2025 Workshop on Human-AI Interaction for Augmented Reasoning*. ACM, Yokohama, Japan. doi:10.48550/arXiv.2503.22018 Workshop at the 2025 ACM CHI Conference on Human Factors in Computing Systems.
- [17] Q. Vera Liao and Wai-Tat Fu. 2013. Beyond the filter bubble: interactive effects of perceived threat and topic involvement on selective exposure to information. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (Paris, France) (CHI '13)*. Association for Computing Machinery, New York, NY, USA, 2359–2368. doi:10.1145/2470654.2481326
- [18] Q. Vera Liao and Wai-Tat Fu. 2014. Can you hear me now? mitigating the echo chamber effect by source position indicators. In *Proceedings of the 17th ACM Conference on Computer Supported Cooperative Work & Social Computing (Baltimore, Maryland, USA) (CSCW '14)*. Association for Computing Machinery, New York, NY, USA, 184–196. doi:10.1145/2531602.2531711
- [19] Q. Vera Liao and Wai-Tat Fu. 2014. Expert voices in echo chambers: effects of source expertise indicators on exposure to diverse opinions. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (Toronto, Ontario, Canada) (CHI '14)*. Association for Computing Machinery, New York, NY, USA, 2745–2754. doi:10.1145/2556288.2557240
- [20] Steffen Mau, Thomas Lux, and Linus Westheuser. 2023. *Triggerpunkte*. Suhrkamp Verlag, Berlin, Germany.
- [21] Maximilian Maurer. 2026. ELFEN: A Python Package for Efficient Linguistic Feature Extraction for Natural Language Datasets. In *Proceedings of the 19th Conference of the European Chapter of the Association for Computational Linguistics: System Demonstrations*. Association for Computational Linguistics, Rabat, Morocco.
- [22] Judith Möller. 2021. Filter bubbles and digital echo chambers 1. In *The Routledge companion to media disinformation and populism*. Routledge, London, 92–100.
- [23] Raymond S Nickerson. 1998. Confirmation bias: A ubiquitous phenomenon in many guises. *Review of general psychology* 2, 2 (1998), 175–220.
- [24] Keith Rayner. 1998. Eye movements in reading and information processing: 20 years of research. *Psychological bulletin* 124, 3 (1998), 372.
- [25] Erik D Reichle, Alexander Pollatsek, and Keith Rayner. 2006. E–Z Reader: A cognitive-control, serial-attention model of eye-movement behavior during reading. *Cognitive systems research* 7, 1 (2006), 4–22.
- [26] Simon Röhrli, Florian Hauser, Timur Ezer, Lisa Grabinger, and Prof. Dr. Jürgen Mottok. 2025. WebGazeTrack: A Web-based Eye Tracking Tool. In *Proceedings of the 6th European Conference on Software Engineering Education (ECSEE '25)*. Association for Computing Machinery, New York, NY, USA, 125–134. doi:10.1145/3723010.3723024
- [27] David O Sears and Jonathan L Freedman. 1967. Selective exposure to information: A critical review. *Public opinion quarterly* 31, 2 (1967), 194–213.
- [28] Li Shi and Jacek Gwizdka. 2025. The Effects of Confirmation Bias and Readability on Relevance Assessment: An Eye-Tracking Study. In *Information Systems and Neuroscience*, Fred D. Davis, René Riedl, Jan vom Brocke, Pierre-Majorique Léger, Adriane B. Randolph, and Gernot R. Müller-Putz (Eds.). Springer Nature Switzerland, Cham, 137–146.
- [29] Michael Sülflow, Svenja Schäfer, and Stephan Winter. 2019. Selective attention in the news feed: An eye-tracking study on the perception and selection of political news posts on Facebook. *New Media & Society* 21, 1 (2019), 168–190. doi:10.1177/1461444818791520
- [30] Arne Freya Zillich and Lars Guenther. 2021. Selective exposure to information on the internet: measuring cognitive dissonance and selective exposure with eye-tracking. *International Journal of Communication* 15 (2021), 20.