
Misuse of linguistic evidence in a study of media bias

Geoffrey K. Pullum¹ and Brett Reynolds²

Abstract

Jackson (2024) presents what is claimed to be a “large-scale proof of historical bias against Palestine” in coverage by *The New York Times*, using computational linguistic methods. Fundamental errors in both linguistic analysis and computational methodology vitiate the study. The analysis rests on a profound misunderstanding of the grammatical notion of ‘passive voice’, and the quantitative results rest entirely on the failed grammatical analysis. Moreover, the computational methodology employs overly narrow keyword filters (not specified in the published paper), excludes relevant data, and lacks a necessary baseline for comparison. The alleged systematic bias remains conjectural. We remark in conclusion that if computational linguistic tools are to be used in media analysis, the linguistic analysis must be sound and coherent, and the computational analysis must be rigorous and consistent.

Keywords

Passive voice, media bias, Natural Language Processing (NLP), Israeli-Palestinian conflict, *New York Times*, linguistics

Introduction

The title claim of Jackson (2024b) is that “*The New York Times* distorts the Palestinian struggle.” Jackson claims to offer “a case study of anti-Palestinian bias in US news coverage of the first and second Palestinian intifidas,” based on “a methodologically novel, large-scale proof” of systematic anti-Palestinian bias in *The New York Times* (henceforth *NYT*). She seeks to ground establish her thesis through computational linguistic analysis of voice and sentiment, but her paper is fundamentally undermined

¹University of Edinburgh, ²Humber Polytechnic & University of Toronto
Email: ¹pullum@gmail.com, ²brett.reynolds@humber.ca

by a profound misunderstanding of basic English grammar – specifically the notion of the ‘passive voice’ – and by critical flaws in the computational methodology employed.

The question of media bias in coverage of the Israeli-Palestinian conflict certainly deserves attention. Multiple studies have examined how narrative framing, source selection, and terminology might favor particular perspectives in this contentious arena (Zelizer et al. 2002; Bazian 2015; Tenenboim-Weinblatt et al. 2016; Neureiter 2017; Pinker 2025). A rigorous computational analysis could in principle offer valuable insights into patterns that qualitative studies cannot easily capture. Unfortunately, Jackson’s paper fails at every level to provide such an analysis.

Jackson’s methodology hinges on the correct identification of active versus passive constructions, and making a connection between these identifications and statements about Israeli or Palestinian agency. Jackson fails dramatically on the first count, exhibiting a clear failure to identify passive clauses. The mistake she makes is amazingly common, especially among 20th-century writing advice books and stylistic critics, as (Pullum 2014) documents extensively. Jackson is just one of a long line of writers who fail to distinguish the grammatical property of being a passive from the vague semantic notion of being inexplicit about agency or responsibility. This basic linguistic error runs right through the foundation of her quantitative analysis, and invalidates it. In most cases such an error merely makes anecdotal critiques of writing style meaningless (typically by falsely accusing writers of using passives, as if that would be inherently bad). But in Jackson’s case they form the basis of what is supposed to be an empirical research contribution, reported in a peer-reviewed journal, supporting an implied moral judgment on the ethics of the entire journalistic staff of one of America’s premier newspapers. That is much more serious.

The way Jackson counts relevant instances introduces a different flaw that would be vitiating even if she correctly distinguished actives from passives. She relies on an overly narrow and arguably biased filter, counting sentences only when the grammatical subject matches a small predefined list of terms for “Israeli” or “Palestinian”. This ignores crucial context, such as the agent role in passive clauses (often the focus of bias analysis) and the use of related but unlisted terms to refer to the parties involved. It also explicitly excludes sentences deemed relevant to both sides. Furthermore, the paper provides no established baseline or counterfactual model for what unbiased reporting patterns would look like, making it impossible to meaningfully interpret the statistical patterns presented as evidence of bias.

The publication of this paper, despite its elementary linguistic errors and flawed methodology, and the subsequent refusal by the journal’s editors to take these flaws seriously when they were pointed out (see Kamm 2025), raises concerns about peer review and editorial standards. In this response we aim to correct the record by detailing the paper’s central failures and demonstrating why its conclusions about media bias in the *NYT*’s coverage are not supported by its evidence.

Errors in linguistic analysis: The passive voice

A core component of Jackson’s analysis (Jackson 2024b) rests upon the classification of sentences (more precisely, individual clauses) as active or passive. However, the paper exhibits a profound misunderstanding of this basic grammatical distinction, sufficient to completely invalidate the quantitative analysis based on it.

Conceptual misunderstanding of voice

It is worth noting first that the term ‘voice’, traditionally used to name the grammatical dimension on which *Russia invaded Ukraine* differs from *Ukraine was invaded by Russia*, is a singularly unhelpful one, because the active/passive distinction has nothing to do with voice in the phonetic sense (*her loud voice*) or the literary sense (*her distinctive authorial voice*). But there is a clear definition of what constitutes a passive clause in the grammatical sense, as we explain below.

Jackson sets out her approach to grammatical analysis as follows: “I identify whether actions by Palestinian and Israeli groups are being described in the active or passive voice. For every verb, I identified the perpetrator and recipient of the action (i.e. whether they were a Palestinian or Israeli group or individual)” (Jackson 2024b, 120).

The misconception is clear: Jackson takes voice to be a semantic category concerned with whether agents (“perpetrators”) are specified, rather than as the syntactic category familiar to grammarians. Her focus is not on the facts of syntactic structure, where there can be full agreement concerning what is or is not passive, but on a much more subtle semantic property to be inferred from understanding of content.

The extent of this misunderstanding is starkly illustrated by an example Jackson herself selected (from Kingsley 2021) to demonstrate bias:

More than 67 Palestinians, including 16 children, have died since the start of the conflict on Monday

Jackson cites this sentence as an instance of passive voice, on the grounds that it obscures the identity of the “perpetrators” responsible for the deaths. As Jackson has since acknowledged in a corrigendum (after Oliver Kamm brought the issue to the attention of the editors), it is not a passive at all: *have died* is the present perfect active form of the intransitive verb *die*. It exhibits none of the structural characteristics of a passive construction.

Passive clauses in English

Passive clauses in English are defined by their distinctive syntactic structure. The main verb is transitive (i.e., capable of taking a direct object) and is in its past participle form. Typically there is an additional verb, often the auxiliary verb BE, but sometimes GET, as in *He got killed*, or a verb such as *have*, as in *The government had him killed* (for further details see Huddleston and Pullum 2002, pp. 1427–1447 or Huddleston et al. 2022, pp. 362–371). To illustrate, these examples are both passive:

- (1) a. *Six civilians were killed in the shooting.* [Passive, no agent]

- b. *Six civilians were killed by the soldiers.* [Passive, with agent]

Notice that the inclusion or omission of a noun phrase denoting the agent is an orthogonal matter: passive constructions may omit agents, as in (1a), or include them in *by*-phrases, as in (1b). Moreover, many active constructions lack explicit agents:

- (2) a. *The stock market crashed.* [Active intransitive, no agent]
 b. *Many civilians died in the conflict.* [Active intransitive, no agent]
 c. *Fighting erupted near the border.* [Active intransitive, no agent]

As Pullum (2014) observes, passive voice is often incorrectly identified in popular discourse: it gets confused with the semantic property of not specifying the doer of an action. Jackson makes this familiar mistake.

Implications for methodological validity

Jackson's approach magnifies her conceptual error. Her problems extend far beyond misclassifying any specific example. Even if her computational tools correctly classified active and passive constructions (which is implausible, given Jackson's demonstrated inability to cite a relevant example), her interpretation of that result is fatally flawed in other ways. She repeatedly asserts that the type of clause she takes to be passive inherently "de-emphasizes or hides those perpetrating such negative action" (Jackson 2024b, 124). This ignores the fact that passives can, and often do, include explicit agent phrases, as in (1b).

Moreover, if Jackson believes that the omission of agents indicates bias, then her exclusive focus on passive constructions misses numerous active constructions that similarly omit agents. Intransitive verbs like *die*, *elapse*, *emerge*, *erupt*, *perish*, *suffer*, *undergo*, and hundreds of others naturally occur in active constructions that do not specify external causation. Equating passives with instances of obscuring agency is a gross mistake.

Jackson repeatedly asserts that the passive voice inherently "de-emphasizes or hides those perpetrating such negative action" (Jackson 2024b, 124). But this overlooks the crucial role of context in interpretation. Even when a passive construction lacks an explicit agent phrase (a *by*-phrase), the surrounding text or shared knowledge often makes the agent or cause perfectly clear. Consider the following examples drawn from the corpus data used by Jackson (Ferraro et al. 2018):

- (3) *The Israelis accuse Saadat of ordering the murder of the Israel's tourism minister, Rehavam Zeevi, who was gunned down in a Jerusalem hotel on Oct. 17.*
 (4) *The group said it was in retaliation for the assassination of Saadat's predecessor, Abu Ali Mustapha, who was killed in a rocket attack on his office.*

In (3), the main clause attributes the order for the assassination of Zeevi to Saadat, making the ultimate responsibility clear, even though the agent of the passive verb was gunned down is syntactically absent and potentially unknown. Adding a generic phrase like *by*

assassins would offer little clarification beyond what *gunned down* already conveys in this context; the significant information regarding responsibility is already provided.

Similarly, in (4), while the agent of *was killed* is omitted, the cause or means (in a rocket attack) is specified immediately following the verb. For readers familiar with the conflict context surrounding the targeted assassination of a figure like Mustapha, the likely agent responsible for such an attack is strongly suggested. Explicitly adding the agent (e.g., *in an Israeli rocket attack*) would specify this inference but doesn't dramatically change the core understanding already available from the description and context.

The supposed connection between voice and bias

Jackson nowhere establishes the plausibility of treating passive voice as a valid proxy for media bias. Her entire justification consists of a single sentence: "An analysis of voice is common in other studies of bias because the passive voice, which is often discouraged in formal writing, allows the writer to de-emphasize or entirely omit the perpetrator of an action in a sentence" (Jackson 2024b, 120). This perfunctory claim is her sole theoretical support for connecting grammatical voice to bias. She provides no substantive discussion of what distribution of active/passive constructions would characterize unbiased reporting, or why differences in voice usage across subjects would necessarily indicate bias rather than reflecting other factors such as prose style, reporting conventions, information presentation, or the nature of the events being covered.

It is true that use of passive clauses is often discouraged by writing advice sources, but in saying that they are giving very poor advice. Many of the advice sources that deprecate the passive are written by authors who cannot tell an active from a passive (Pullum 2014), and they nearly always ignore the fact that sometimes (depending on whether new information is being introduced) a passive is exactly right for expressing the intended meaning. In normal prose by excellent professional writers, between 10 to 20 percent of the transitive verbs are in passive clauses.

And as already noted, if the concern is tracking down a case of omission of agency (failing to name "perpetrators"), finding a passive clause is neither necessary nor sufficient.

Establishing the validity of a quantitative measure for an abstract concept like media bias demands more than face-value interpretation. As Carmines and Zeller (1979) note, criterion-related validity is often inapplicable to such concepts due to the lack of suitable external criteria, and content validity assessment remains subjective. Instead, validation primarily relies on establishing construct validity – demonstrating empirically that the measure behaves as expected within a network of theoretically related concepts (Carmines and Zeller 1979). This requires testing specific hypotheses about how voice patterns should correlate with other indicators or dimensions of bias. Jackson's paper offers no such theoretical framework or empirical tests, failing to establish the construct validity necessary to justify voice distribution as a measure of bias.

Without establishing this foundational connection between her proxy measure (voice distribution) and her object of study (media bias), Jackson's subsequent statistical

analysis becomes an exercise in counting grammatical constructions (often misidentified) that bear no demonstrated relationship to her research question.

Jackson's computational methodology

Beyond the fundamental linguistic errors concerning passive voice (Section), Jackson's computational methodology, as she describes and used it, suffers from critical flaws in transparency, implementation, and conceptualization that render its quantitative results unreliable as a measure of media bias. While the paper claims reliance on "state-of-the-art natural language processing toolkits" (Keras (Chollet et al. 2015); NLTK 3.6.2 (Bird et al. 2021); scikit-learn, (Pedregosa et al. 2011); spaCy (Honnibal and Montani 2021)) and high accuracy for its models (Jackson 2024b, Abstract), such tools require linguistically informed application and validation. Given the author's demonstrated misunderstanding of the core grammatical concept being investigated, her ability to correctly design the computational analysis, critically evaluate the tool's output, or perform the claimed "qualitative validation step" (Jackson 2024b, 116) is severely compromised.

The subject filters

A primary methodological flaw lies in the extremely narrow filter apparently used to determine subject relevance. The specific words used to identify "Israeli" or "Palestinian" subjects are not disclosed within the published paper, but can be found only by scrutinizing the Python scripts, hindering reproducibility and critical assessment. Examination of the author's publicly available code (Jackson 2024a) reveals that it classifies voice only for sentences where the grammatical subject (or its directly associated modifiers) explicitly matches one of the terms within the following short, predefined lists:

- Israeli relevance: *Israel, Israeli, Israelis, IDF, I.D.F.*
- Palestinian relevance: *Palestine, Palestinian, Palestinians, PLO, P.L.O., Fatah*

The paper provides no justification for the specific composition or limited size of these lists. This selection represents a critically consequential exercise of "researcher degrees of freedom" (Gelman 2018). Given the list's asymmetry and failure to cover key actors, particularly on the Palestinian side (discussed below), this undocumented choice severely undermines the foundation of the comparative analysis. It also fails to address the inherent trade-off between precision (ensuring captured subjects definitely belong to the target group) and recall (capturing a representative range of references). The lack of justification is particularly problematic given the complexity of referring expressions – a significant area of study in linguistics – and their especially nuanced nature in reporting on the two Intifadas, where terminology carries distinct political implications and where the groups involved change over time.

Furthermore, the lack of justification for the keyword lists is compounded by the paper's apparent treatment of the First Intifada (approx. 1987–1993) and the Second (approx. 2000–2005) as uniform periods for filtering via a single keyword set. These

conflicts differed significantly in their characteristics, primary actors, and international context. For instance, the prominence of groups like Hamas grew substantially by the time of the Second Intifada, yet key terms referring to it seem absent from the filter lists used, while terms perhaps more central to the First Intifada (like PLO) are included (Jackson 2024a). Applying a single, static, and narrow filter across these distinct historical periods without justification further undermines the methodology's ability to capture relevant data accurately and highlights its insensitivity to crucial contextual shifts.

The filtering also ignores the significant asymmetry between the conflict parties' representation in the lists (Jackson 2024a). While one side is primarily represented by a state actor and its military (terms mostly included), the lists fail to reflect the evolving Palestinian representation. The PLO (included) was key during the First Intifada, but the Palestinian Authority (PA), Hamas, and other major factions central to the Second Intifada (like Islamic Jihad, PFLP, and Al-Aqsa Martyrs Brigades) are all absent. Crucial terms related to status or geography (e.g., *refugee*, *militant*, *West Bank*, *Gaza*) are also missing.

This narrowness means sentences are excluded based merely on the lexical choice of the subject term, disproportionately impacting Palestinian representation. For instance, an active sentence like *Hamas fighters launched several; rockets* would be entirely excluded despite its relevance to voice usage concerning Palestinian actors. Similarly, a passive construction like *the checkpoint was attacked by militants from Gaza* would likely be ignored because *checkpoint* is not on a keyword list, discarding information about action and agency relevant to the conflict.

Furthermore, the filter discards sentences where the head noun of the subject itself is neutral (i.e., not on the keyword lists) even when the sentence clearly relates to the conflict parties through context, agents, or modifiers, as illustrated by examples drawn from the corpus used by Jackson (Ferraro et al. 2018):

- (5) a. ...*a leader of Al Aksa Martyrs Brigade killed... by an Israeli shell.*
- b. *That proposal was not only adopted by the Arab summit...*
- c. ...*The Saudi initiative has now been turned into an Arab initiative.*
- d. ...*the Arab leaders are being driven...*

In (5a5d), the passive clauses are discarded because the grammatical subjects (*leader*, *proposal*, *initiative*, *leaders*) are absent from the keyword lists, ignoring readily available contextual information linking them to the conflict parties.

Additionally, the logic of the script used explicitly requires exclusive relevance to one side based on the subject keyword, meaning subjects relating to both (e.g., *Israeli-Palestinian violence*) are systematically ignored. This arbitrary exclusion further distorts the dataset intended to represent reporting on the conflict.

The compounding effect of these filtering decisions is a dataset that is fundamentally unsuitable for valid comparative analysis. What Jackson presents as differences in voice distribution or sentiment scores between Israeli and Palestinian subjects likely reflects little more than artifacts of the filtering methodology interacting with natural linguistic

patterns of reference and the inherent asymmetries of the chosen keywords. This arbitrary and asymmetric filtering renders any subsequent quantitative comparison between the groups meaningless as a measure of reporting patterns, let alone as evidence for media bias.

Technical implementation concerns

Further doubts arise from the technical implementation suggested by the associated code (Jackson 2024a). The logic used to automatically detect voice appears overly simplistic, likely examining only the primary verb structure within the main clause of each sentence. Such an approach inherently ignores subordinate clauses, whether active or passive. For instance, in a structure like *The document, which was reviewed yesterday, contained errors*, the passive construction within the relative clause would likely be missed. Corpus studies suggest that the proportion of the passive may not differ significantly between main and subordinate clauses (e.g., Leong 2014), so this may not systematically bias analyses in a predictable way, but it discards a good deal of relevant data. Thus, Jackson's claim to have identified the perpetrator and recipient "for every verb," is far from accurate.

Missing counterfactual analysis

The study lacks any established baseline or counterfactual model for what patterns of voice distribution would characterize unbiased reporting on asymmetric conflicts. Without such a baseline, it is impossible to meaningfully interpret the statistical differences that Jackson presents as evidence of bias. The paper simply assumes that equal distribution of voice constructions across subjects would indicate neutrality – an assumption that ignores how different types of events, actors, and journalistic conventions might naturally produce different syntactic patterns regardless of bias.

Concluding remarks

Jackson's vaunted "large-scale proof" of systematic, disproportionate anti-Palestinian bias in the *NYT* is entirely unsupported by the evidence provided. It represents an unwarranted leap from flawed data to unsubstantiated conclusions about media bias.

Jackson does not even give a convincing argument for *why* a difference in passive subject counts, even if they were accurately measured, should necessarily indicate bias against the group more frequently denoted by the subject of a passive clause. Her paper disregards plausible alternative explanations for observed differences in voice usage; it neglects the issue of how events themselves shape reporting choices; it neglects how journalistic conventions influence voice selection; and it neglects how actual agency distribution affects syntactic patterns. Her analysis jumps directly from differential counts produced by a flawed methodology to a very strong claim about bias, ignoring other plausible linguistic and contextual factors associated with passive clauses (Huddleston and Pullum (2002, Ch. 16); Pullum (2014)). Although the paper includes a discussion of limitations, that section fails to critically interrogate the study's

foundational assumptions or the potential impact of its core methodology on the results. Instead of demonstrating a rigorous search for potential weaknesses in the study's main thrust, the discussion focuses on secondary or speculative points.

For instance, Jackson speculates whether the NLP tools themselves might carry an "anti-Arab, anti-Muslim bias" due to training on corpora with "Orientalist sentiments" (Jackson 2024b, 123). While algorithmic bias is a valid general concern, its prominence here – absent specific substantiation and serving primarily to deflect from unaddressed core flaws in linguistic understanding and methodological design – is revealing. Ultimately, the limitations section does not reflect a genuine attempt to identify or grapple with the most significant factors that could undermine the study's central claims about bias.

The study also fails to account for how source accessibility disparities and asymmetries in the conflict might naturally affect reporting patterns irrespective of bias. For example, Israeli government sources, benefiting from robust state communication infrastructure, generally have a greater capacity for regular engagement with international media through briefings and official statements, while Palestinian sources operate under significant infrastructural constraints and restrictions imposed by Israel (Abdel-Wahab 2025). This asymmetry is reflected in sourcing patterns, where outlets like the *NYT* have historically used Israeli officials a sources far more frequently than Palestinian officials (Viser 2003). Without addressing such practical realities of conflict reporting, differences in voice patterns cannot validly be interpreted as evidence of bias.

In short, Jackson's central claim is entirely unsupported by evidence. Again, this does not mean we dismiss the potential value of investigating how media language might obscure agency or downplay certain actors' roles – which was Jackson's goal. Studying such phenomena through linguistic analysis could be a valid and important endeavour in media research. But the project would require a sophisticated interdisciplinary approach. It would necessitate a careful combination of deep linguistic expertise (to correctly identify relevant structures beyond mishandled traditional grammatical labels), sound computational methodology (appropriately designed and validated for the specific linguistic task), and substantive domain knowledge (taking account of reporting conventions, source dynamics, and the specific context, such as the Israeli-Palestinian conflict). Furthermore, robust validation protocols, potentially involving consultation with subject matter experts, are essential to ensure that the analysis accurately captures the intended phenomena and its real-world significance. Effectively integrating these diverse skill sets – linguistics, computation, domain expertise, understanding of journalism, and statistical validation – is challenging, and the necessary combination is unfortunately scarce. It underscores the need for genuine interdisciplinary collaboration or consultation to successfully execute such complex, high-stakes research, ensuring that methods are appropriate and findings are reliable.

The publication of this paper, with its elementary linguistic errors and disastrously flawed methodology, raises concerns about the peer-review process. Our concerns were significantly amplified by the journal's editorial response when these fundamental issues were brought to their attention (see also Kamm 2025). The editors mandated only a narrow corrigendum addressing the misidentification of a single example sentence, while

declining to engage with or require correction of the paper's profound misunderstanding of passive voice or the systemic flaws identified in its computational methodology – issues invalidating its core quantitative claims. Having asked for this minimal correction, the editors then stated that they would not correspond further on the substantive critiques and indicated that if we felt more was needed, we should publish a critique. This disdain for normal procedures that ultimately lead to retraction of scientific papers raises serious questions about the adequacy of the journal's post-publication review process and its commitment to upholding basic standards of competence, rigour, and scientific ethics.

References

- Abdel-Wahab A (2025) Gaza's telecommunications: Occupied and destroyed. *Al-Shabaka* URL <https://al-shabaka.org/briefs/gazas-telecommunications-occupied-and-destroyed/>.
- Bazian H (2015) The Islamophobia industry and the demonization of Palestine: Implications for American studies. *American Quarterly* 67(4): 1057–1066. DOI:10.1353/aq.2015.0073.
- Bird S, Klein E and Loper E (2021) Natural Language Toolkit. <https://www.nltk.org/>. Version 3.6.2.
- Carmines EG and Zeller RA (1979) *Reliability and Validity Assessment, Quantitative Applications in the Social Sciences*, volume 17. Beverly Hills, CA: Sage Publications. DOI:10.4135/9781412985642.n2.
- Chollet F et al. (2015) Keras. <https://keras.io>.
- Ferraro F, Cubbedge D, Finlayson M, Guillaume T, Katz B, Konstantynova H, Luo L, Martin K, Mayfield J, Miller T, Peterson R, Pfeifer M, Porter T, Roncone A, Sil A, Strassel S, Thomas E, Traum D, Van Durme B, White P, Woods S, Yarmohammadi M and Zhang D (2018) Concretely annotated New York Times LDC2018T12. Web Download. URL <https://catalog.ldc.upenn.edu/LDC2018T12>.
- Gelman A (2018) Ethics in statistical practice and communication: Five recommendations. *Significance* 15(5): 40–43. DOI:10.1111/j.1740-9713.2018.01193.x.
- Honnibal M and Montani I (2021) spaCy: Industrial-strength natural language processing. <https://spacy.io>.
- Huddleston R and Pullum GK (2002) *The Cambridge Grammar of the English Language*. Cambridge: Cambridge University Press. DOI:10.1017/9781316423530.
- Huddleston R, Pullum GK and Reynolds B (2022) *A Student's Introduction to English Grammar*. Cambridge: Cambridge University Press. DOI:10.1017/9781009085748.
- Jackson HM (2024a) NYT_Content_Analysis. URL https://github.com/hollyjackson/NYT_Content_Analysis. GitHub repository.
- Jackson HM (2024b) *The New York Times* distorts the Palestinian struggle: A case study of anti-Palestinian bias in US news coverage of the First and Second Palestinian intifidas. *Media, War & Conflict* 17(1): 116–135. DOI:10.1177/17506352231178148.
- Kamm O (2025) Spurious research as tendentious propaganda. *Quillette* URL <https://quillette.com/2025/??> URL to come.
- Kingsley P (2021) As Gaza war escalates, new front opens in Israeli cities. *The New York Times* URL <https://www.nytimes.com/2021/05/12/world/middleeast/israel-gaza-violence.html>.
- Leong PA (2014) The passive voice in scientific writing. The current norm in science journals. *Journal of Science Communication* 13(01): A03. DOI:10.22323/2.13010203.
- Neureiter M (2017) Sources of media bias in coverage of the Israeli-Palestinian conflict: The 2010 Gaza flotilla raid in German, British, and US newspapers. *Israel Affairs* 23(1): 66–86. DOI: 10.1080/13537121.2016.1244381.
- Pedregosa F, Varoquaux G, Gramfort A, Michel V, Thirion B, Grisel O, Blondel M, Prettenhofer P, Weiss R, Dubourg V, Vanderplas J, Passos A, Cournapeau D, Brucher M, Perrot M and

Duchesnay (2011) Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research* 12: 2825–2830. URL <http://jmlr.org/papers/v12/pedregosa11a.html>.

Pinker EJ (2025) An analysis of the *New York Times* coverage of the war between Israel and Hamas. Technical Report 5104625, Yale School of Management. DOI:10.2139/ssrn.5104625.

Pullum GK (2014) Fear and loathing of the English passive. *Language & Communication* 37: 60–74. DOI:10.1016/j.langcom.2013.08.009.

Tenenboim-Weinblatt K, Hanitzsch T and Nagar R (2016) Beyond peace journalism: Reclassifying conflict narratives in the Israeli news media. *Journal of Peace Research* 53(2): 151–165. DOI: 10.1177/0022343315609091.

Viser M (2003) Attempted objectivity: An analysis of the *New York Times* and *Ha'aretz* and their portrayals of the Palestinian-Israeli conflict. *Harvard International Journal of Press/Politics* 8(4): 114–120. DOI:10.1177/1081180X03256999.

Zelizer B, Park D and Gudelunas D (2002) How bias shapes the news: Challenging the *New York Times*' status as a newspaper of record on the Middle East. *Journalism* 3(3): 283–307. DOI: 10.1177/146488490200300305.