Network Science (2026), **14**, e20, 1–25
doi:10.1017/nws.2026.10040

CAMBRIDGE UNIVERSITY PRESS

RESEARCH ARTICLE

Beyond interaction networks: Practice mapping as a step towards the detection of destructive polarization in social networks

Axel Bruns, Katharina Esau, Kateryna Kasianenko, Ehsan Dehghan, Tariq Dos Santos Choucair, Vish Padinjaredath Suresh and Donald Johannessen

Digital Media Research Centre, Queensland University of Technology, Australia
Corresponding author: Axel Bruns; Email: a.bruns@qut.edu.au

Abstract

The study of polarization in social networks is complicated by conceptual confusion: key forms of polarization include issue, ideological, affective, interactional, interpretive, and perceived polarization, but consensus remains elusive on which of these concepts are most important, how they are interrelated, and whether they should be considered problematic. In this article, we show how the communicative practices of the various actors in a given public debate can be analyzed to reveal the networked structures of discursive alignments and divergences and assess the degree of alliance or opposition between them, and how this analysis in turn provides a basis for the identification of polarization and its potential diagnosis as destructive. We address the challenge of assessing polarization by introducing the new analytical approach of practice mapping. Conceptually supported by practice theory, practice mapping advances beyond conventional social network analysis by no longer grounding its network analysis only in the direct network ties made by participants in social networks, but conceptualizing each participant's broader patterns of action and interaction within the social network (including their engagement with other participants, but also their information sourcing, language choices, and other qualities of their participation) as that participant's unique communicative practice—and then constructing and mapping the network of similarities between these practices for all participants. This serves as a key step towards identifying and analyzing patterns of polarization in public discourse. To illustrate this approach, we apply it to public posts on Facebook about the Voice to Parliament referendum in Australia. Our practice mapping was able not only to clearly identify the deeply opposed Yes and No campaigns, but also to show substantial distinctions between the highly centralized and unified No camp and the loose and uncoordinated discursive alliances supporting the Yes camp. Using this case study, we demonstrate the utility of practice mapping in assessing polarization in social networks.

Keywords: polarization; network analysis; practice mapping; social media; Facebook; Australia; Voice to Parliament

1. Introduction

Public debate about contentious political and societal issues is inherently marked by disagreement. Such disagreements, on any given issue, can manifest as issue polarization, which can also be productive for democratic discourse. However, as McCoy and Somer (2019) have argued, polarization becomes problematic or *pernicious* when political parties develop mutually reinforcing social identities that structure how actors interpret information, evaluate what is legitimate, and engage with political opponents. In such contexts, disagreement no longer genuinely revolves around issues or policy positions, but instead hardens into antagonistic camps that undermine democratic norms and conventions. In this article, we show how the communicative practices of the

© The Author(s), 2026. Published by Cambridge University Press. This is an Open Access article, distributed under the terms of the Creative Commons Attribution licence (https://creativecommons.org/licenses/by/4.0/), which permits unrestricted re-use, distribution and reproduction, provided the original article is properly cited.

https://doi.org/10.1017/nws.2026.10040 Published online by Cambridge University Press

various actors in a given public debate can be analyzed to reveal the networked structures of discursive alignments and divergences and assess the degree of alliance or opposition between them, through the application of the novel practice mapping approach, and how this analysis in turn provides a basis for the identification of polarization and its potential diagnosis as destructive.

This article is positioned at the intersection of three fields of interest. Conceptually, it builds on our current understanding of partisanship and polarization in online public debate, which distinguishes several key types of polarization and emphasizes the need to pay close attention especially also to the signs that indicate when the agonism and antagonism of everyday debate and disagreement turns into problematic, pernicious, or *destructive* polarization (Esau et al., 2025; McCoy & Somer, 2019). Recent work in media and communication studies outlined the potential symptoms of such destructive polarization that can be observed within communicative practices (Esau et al., 2025a; Esau, 2025). Methodologically, the article develops an approach that builds on network analysis and visualization principles but extends them beyond the mapping of direct interactions and towards the identification of networks determined by shared practices: *practice mapping* (Bruns et al., 2025). Empirically, finally, it illustrates these conceptual and methodological contributions by applying practice mapping to a particularly notable and notorious case of polarization in national public debate in Australia: the campaigns surrounding the country's 2023 constitutional referendum to establish an Indigenous Voice to Parliament (Johannessen, 2025). This shows how the practice mapping approach serves as a key step towards the detection and analysis of the symptoms of destructive polarization. We apply further mixed-methods analysis to the results of the practice mapping, to demonstrate the validity of these results and show how they can be operationalized to further investigate the case study. Throughout this process, we understand social networks as constituted by their participants' complete range of communication practices, including explicit *inter*actions with others as well as broader communicative *actions* that address only an implicit, imagined audience within the discursive space.

The following sections work through these conceptual, methodological, and empirical considerations sequentially. We begin by reviewing the current state of the literature on polarization, and introduce the symptoms of destructive polarization. Since much of the current debate identifies social media as a particularly important locus of polarized and polarizing debate, we then outline the state of the art and limitations of current social network analysis approaches as applied to online public debate, and contextualize these especially against the substantial recent shifts in social media platform preferences and data accessibility that have necessitated the development of new data access, processing, and analysis approaches like practice mapping. Finally, we introduce the Voice to Parliament case study, and present a preliminary analysis of public debates surrounding the referendum as they were conducted on public pages and groups on Facebook; this both serves to demonstrate practice mapping in action, and in doing so points to the utility of this approach as a step towards investigating the symptoms of destructive polarization that may be present in the debate. We conclude by connecting these conceptual, methodological, and empirical strands of our discussion, and offer suggestions for further development.

2. Polarization: Which types, and to what degree?

In everyday public communication as well as in the scholarly literature, the idea of "polarization" is currently experiencing a substantial resurgence (e.g., Benson, 2023; Kreiss & McGregor, 2023); this is due not least to the clear evidence of widespread partisanship and division which have taken hold (and in some cases have been deliberately encouraged) in public debate across many nations (McCoy et al., 2018). Both in its everyday and its scholarly uses, however, "polarization" often remains remarkably amorphous and ill-defined as a concept (as noted, e.g., by Esau et al., 2025a; Kubin & von Sikorski, 2021): it is often used as a mere synonym for division and disagreement between two or more political groups.

https://doi.org/10.1017/nws.2026.10040 Published online by Cambridge University Press

This is in spite of the fact that polarization research has by now produced a confounding array of diverging polarization types. Amongst these, key concepts include issue, ideological, affective, interactional, interpretive, and perceived polarization (Esau et al., 2025a). *Issue polarization* refers to the divergence, increasing distance, or extremity of individual or collective positions on a specific issue (Leifeld & Brandenberger, 2019; Mason, 2015): in the case study we address in the empirical part of this article, for instance, the central issue is whether or not Australia's Indigenous peoples should be formally recognized in the country's constitution, and provided with a distinct institutional mechanism for making representations to the legislative: the Voice to Parliament. Other, related issues that may also be debated in this context include dedicated spending on Indigenous health and welfare programs, interventions to reduce the high rates of Indigenous incarceration, or the incorporation of Indigenous cultural elements into Australian government events.

Such polarization on specific issues can intersect with and merge into broader *ideological polarization*: here, divisions extend beyond distinct issues and are based on the overall political philosophies held by individuals and groups (Fiorina & Abrams, 2008; Levendusky, 2009). Support for or opposition to government spending on Indigenous welfare may align with broader views on the appropriate extent of the welfare state, for instance; these in turn may abstract further into positions that could be loosely (if reductively) described as "left-wing" or "right-wing." Issue and ideological polarization are therefore closely interrelated: political alignment between actors across multiple issues may result in a "partisan sorting" that produces increasingly clearly demarcated ideological camps, while such clearly identified ideological differences are also likely to affect the positioning of actors on new issues as they emerge.

Individually and collectively, this ideological positioning therefore also facilitates the development of a common identity amongst those who share the same ideology; this identity is political, in the first place, but may also extend into social and other demographic aspects, and the term "identity polarization" is therefore also used on occasion to describe the emergence of such differences between groups. Since self-identity—not least also when it extends beyond the narrowly political—always also carries strong personal meaning, such in-group membership therefore also aligns closely with a third type of polarization: *affective polarization* (Iyengar et al., 2012). This is commonly measured in survey research through "warm–cold" feeling thermometers toward opposing camps, and related social-distance items (Almond & Verba, 1963; Iyengar et al., 2012), and in media and communication research through the assessment of emotions and toxicity directed at out-groups (Lerman et al., 2024; Esau et al., 2025b). Here, the emphasis is on the emotive component of belonging to a (political, social, societal) camp—much as in cultural fandom, the stronger one feels one's connection to that camp, the greater one's opposition to opposing camps is likely to be (Barnes, 2022).

All three types of polarization introduced so far predominantly describe an individual's or group's internal attitudes: their specific rational or affective positioning towards an issue, ideology, or the in- and out-groups representing them. This specific positioning is often assessed against the overall range of positions held across the general population. Additionally, however, this positioning will usually also affect their external actions: this may manifest for instance in *interactional polarization*: based on their rational and/or affective stance towards others, individuals and groups may choose to preferentially engage with those whom they regard as sympathetic to their views, and deliberately ignore those whom they envision as antagonists; or they may choose to engage in differing ways with these groups—in a social media context, for instance, respectively by amplifying the messages of allied discourse participants and by attacking those of discursive opponents (Yarchi et al., 2021). Such interactional polarization, especially where it is expressed in public communication, produces the discursive signals that have usually been harnessed for empirical analyses of polarization in public debate, especially through network analysis.

A similar form of polarized action is *interpretive polarization*, which might also be considered as a variant of interactional polarization that targets the communicative contributions of other

https://doi.org/10.1017/nws.2026.10040 Published online by Cambridge University Press

participants, rather than those participants themselves (Kligler-Vilenchik et al., 2020). Here, as a result of their own issue, ideological, and/or affective positioning, individuals and groups choose to interpret the information and communication they encounter from others—including from news outlets, politicians and parties, activists, experts, and ordinary participants—through a lens that is shaped by their own views. Consequently, the very same news report or social media post might be interpreted in diametrically different ways by individuals who are positioned on opposite sides of an issue; such interpretive polarization may in turn be operationalized in empirical analysis if it is expressed publicly by the individual or group in their own public responses to the information, especially also including their content sharing practices in social media contexts. In this case, it may also be closely intertwined with interactional polarization, however: while interactional polarization is predominantly evident from *whom* an individual chooses to engage with, interpretive polarization is expressed more in *how* they do so.

Finally, then, any such public statements about other discursive agonists and antagonists are also likely to point to the individual's or group's *perceived polarization*: their understanding of how polarized the broader community is on the specific divisive issue or broader disputed matters (Lelkes, 2016). Such perceived polarization may be evident for instance in how they describe the views of others as they understand them, and in how aggressively or dismissively they address and attack these views. This perception of polarization will often diverge from *actual polarization* on a given issue, as it may be measured by opinion polls and similar assessments of policy and political views across society—and indeed political rhetoric and propaganda often seek to artificially inflate perceptions of polarization within the general public, in order to create stronger issue, ideological, and affective alignment between citizens and political parties.

In addition to these six key *types* of polarization—issue, ideological, affective, interactional, interpretive, and perceived—scholars often also distinguish between the different *levels* of collective that are said to be polarized in this way. On the one hand, polarization may occur between individuals or between groups (and we have discussed both those opportunities in this section, since individual and group polarization are often interrelated and may form a feedback loop). On the other hand, polarization may be more pronounced at the level of political elites (i.e., amongst parliamentarians, office holders, candidates, and other political professionals) or at the level of the masses (i.e., the broader population, with feedback effects between the two levels also possible; Hetherington, 2009). For our purposes, which include applying these concepts to the empirical study of communicative polarization in online debate, the distinction between or combined study of elites and masses can be addressed through the selection of datasets for analysis, while our practice mapping approach connects individual and group polarization by aggregating individual communicative patterns into overall network analyses that enable the detection of discursive alliances and antagonisms between broader groups of actors, as well as the identification of the key actors driving these processes.

Having identified the six types of polarization, and the possible levels of individuality or collectivity at which they might occur, however, we must also address a further challenge: assessing the extent, strength, or depth, or in short the *degree* of polarization. As several scholars have noted, it is illusory to imagine a complete absence of polarization in any political discussion; indeed, it would be impossible to have any political debates in the absence of polarization, since the very existence of diverging views on a given issue or topic implies at least a mild degree of polarization (cf. McCoy & Somer, 2019). From this perspective, then, *mild* polarization is not only an unavoidable feature of public debate, but is in fact even beneficial and exerts a constructive influence, as it provides the mechanisms that enable participants and observers to more clearly articulate, consolidate, and identify the various possible positions competing for majority support.

We note in this context also that—contrary to views which are overly influenced by the unusual, extreme, and dysfunctional reduction of virtually all political discussions to a "Republicans vs. Democrats" paradigm that is common to the United States, and almost exclusively *only* to the United States—this means that such ordinary polarization almost always describes a *multi-polar*

https://doi.org/10.1017/nws.2026.10040 Published online by Cambridge University Press

contest between possible positions (Kekkonen & Ylä-Anttila, 2021): few issues and topics in public debate have two and only two possible solutions, unless they are deliberately reduced to such a bipolar choice by extreme ideological and affective polarization, or—as in our case study—presented to voters in the form of a Yes/No constitutional referendum question. *A priori*, then, we should always expect the existence of a multitude of somewhat interrelated positions in public debate, which may however be streamlined, through partisan sorting and an increase in the degree of polarization, into a starker choice between a small number of—and eventually, potentially only two—mainstream options (Fiorina, 2017).

Such a reduction of available options indicates that polarization is turning from mild and constructive to pernicious and destructive: the oversimplification of policy-making to a binary, winner-takes-all choice between inherently opposed views severely undermines any potential for meaningful collaboration and compromise between the different sides of a debate. As the US case shows ever more dramatically, intractable policy gridlock, deep societal division, and a descent into illiberalism and autocracy are potential consequences. This slippage from constructive to destructive polarization is unlikely to be characterized by any one distinct tipping point, however: rather, as Esau et al. (2025a) have noted, there are a number of possible symptoms that indicate increased degrees of polarization and therefore serve as warning signs of a descent into destructive polarization:

- *Breakdown of communication*: the outright refusal of one side in a polarized debate to engage with others, or at least the refusal to do so *constructively* without personal attacks;
- *Discrediting and dismissing of information*: the rejection of information because of its content or source, potentially for fabricated reasons, regardless of the facticity of the information;
- *Erasure of complexities*: the reduction of complex policy and other considerations to simplistic, often binary, solutions, not least also through the ceaseless repetition of vapid talking points and slogans;
- *Exacerbated attention to and space for extreme voices*: the presence of and attention to fringe actors and views that seek to exploit the debate to gain a foothold in public communication and divide society;
- *Exclusion through emotions*: the deliberate use of emotive language and argumentation to swamp public debate and inflame sentiment against specific (often minority and vulnerable) groups.

This list of symptoms may not be exhaustive: others may be identified over time and added to this list. They are also likely to occur in different combinations depending on the issue context and communicative environment for the debate (Esau et al., 2025a: 1506). It is important to note that it is not necessary for all of these symptoms to be present simultaneously in order to classify a given public debate as destructively polarized; indeed, each symptom should be understood as representing a sliding scale from "absent" to "severe," and it is in fact the assessment of one or more of these symptoms as being present to a severe degree that would indicate the presence of destructive polarization in a particular communicative context. Such presence may also be asymmetrical, of course, if individuals and groups on one side of a debate act destructively while their opponents still seek constructive engagement; it is therefore often more useful to first examine the polarized or polarizing stances employed by the various communicative actors *within* a debate than to immediately seek to assess the polarization *of* the debate itself.

In order to conduct such an assessment for any public debate, then, it becomes necessary to identify the various individuals and groups participating in the debate, and to determine how the various discursive alliances and antagonisms that structure the debate arise from communicative practices of the participating groups. This groundwork subsequently enables us to

https://doi.org/10.1017/nws.2026.10040 Published online by Cambridge University Press

examine whether such structures and their cleavages are formed around shared positions on specific issues, broader commonalities in ideological views, and/or affective antipathies between in- and out-groups, and to assess the extent to which these discourse participants and their specific communicative practices exhibit the symptoms of destructive polarization that Esau et al. (2025a) have outlined. To do so, we introduce practice mapping as a powerful new framework for identifying discourse actors and the practices they share. To be clear: practice mapping in itself does not determine whether or to what degree destructive polarization is present in a given context—but it provides the framework to identify the various individual and collective actors within a discourse based on their sustained patterns of action, and in doing so enables the examination of their distinct discursive practices for symptoms of destructive polarization. We demonstrate this process in our empirical analysis below, but move first to making the case for practice mapping.

3. Practice mapping: Beyond interactions

From the hyperlink networks of the early blogosphere (Adamic & Glance, 2005) to present-day social media network mapping, the analysis and visualization of online public communication networks (for the detection of polarization as well as for many other purposes) has predominantly focussed on *interaction* networks. The nature of such interactions has evolved over time and depends inherently on the affordances of the platforms being studied: amongst blogs, interactions took the form of hyperlinks and link pingbacks, for instance; on Twitter, they included @mentions, retweets, and quote tweets. Generally, though, they represent an explicit, functional reference by one author to another, often in order to attract the attention of, respond to, or on-share the content posted by the target of the interaction. The networks that result from this mapping of collective interaction patterns can then be utilized to identify densely connected clusters of actors that are distinguished by their higher degrees of homophily, and that may in some cases—depending on the specific interactional affordances incorporated into the network map—be interpreted as communities that represent shared affiliations, interests, or ideologies. In the context of polarization studies, these clusters (and their relative distance from one another) might then in turn indicate more or less pronounced division and opposition; if such opposition is regularly and consistently expressed in ways that match the symptoms outlined above, then this would point strongly to the presence of destructive polarization within the discourse.

Network analysis was necessarily drawn to interactions since they represent an explicit and deliberate interactive act, and—especially in contemporary social media platforms—are also clearly and formally encapsulated and observable in the data such platforms provide(d) to researchers and other analysts. Maps of such interactions are therefore comparatively simple to produce, and do offer valuable insights into the datasets they draw on—for example, they often indicate clusters of accounts that may be interpreted as representing specific political or other interests because of their distinct interactional patterns. Not only was Twitter generally overrepresented in scholarly research on social networks because of the exceptionally powerful and accessible Application Programming Interface (API) it provided (especially through its dedicated Academic API, in the final years before the platform's acquisition and enshittification by Elon Musk; Doctorow, 2022), but a large proportion of Twitter studies therefore also included network maps of various levels of sophistication, showing the interactions amongst the accounts captured in their datasets.

However, this prominence of network analysis in social media (and especially in Twitter) research always obscured several substantial limitations. The first is not inherently a problem *with* interaction network analysis, but a potentially insurmountable challenge *for* such analysis: unlike the considerable accessibility and interpretability of interactions network data from Twitter, many other major social media platforms never provided comparable interactions data, in part also because their communicative affordances structured engagement between accounts in very

https://doi.org/10.1017/nws.2026.10040 Published online by Cambridge University Press

different and much more hierarchical ways. On platforms like Facebook or Reddit, for instance, direct communicative interaction between accounts takes place exclusively or predominantly through more or less sophisticated comment threads: a straightforward network visualization of such threads would produce a series of branching root systems originating from each original post as their common point of origin, for instance, but would not lend itself straightforwardly to the detection of account clusters with similar networking practices.

Indeed, for platforms like Facebook even such root-and-branch structures of comment threads were difficult to produce because its most prominent data access platform, CrowdTangle, offered access only to posts on public pages and groups, but not to the comments attached to these posts. This has changed only very recently, with the transition from the now decommissioned CrowdTangle to its purported replacement, the Meta Content Library. Further, on Facebook and other platforms (and notably unlike Twitter), most other direct interactions between accounts are private or at least semi-private, and therefore not included in the data made available to researchers and other analysts; while for Twitter it was possible to map @mentions, retweets, and even follower-followee relations between accounts participating in public debates (Bruns et al., 2017; Williams et al., 2015), comparable analyses were never feasible for most other platforms.

A second limitation of such conventional interaction-centric network mapping relates to the difficulties in interpreting its results. Returning to Twitter as the most prominent example of such analysis, each major interactional affordance had one or more possible interpretations: while retweets were often (but not always) meant to endorse and amplify another account's post, quote tweets might be used for the same purpose but were also often utilized to critique or reference a post without responding directly to its author, and @mentions could be used for a wide variety of purposes from unqualified support to fundamental opposition and attack. Studies that combined all such types of interaction into a single network map could still surface some relevant information (such as the most targeted accounts overall, for instance), but all too often produced mere "furballs" that show high levels of interconnection between all meaningfully active participants; alternatively, some scholars chose to produce and analyze the networks constructed by these different platform affordances separately, and as a consequence detected diametrically opposed patterns for each affordance (e.g., Williams et al., 2015).

The analysis (and visualization) of such networks is further complicated by the lack of reciprocity in many such interactions: on Twitter, an @mention from one account to another does not guarantee an equal and opposite response. Interaction networks are therefore usually *directed* networks (where a connection—a network edge—from account A to account B does not imply an equivalent edge from B back to A); when visualizing such a network, then, should A and B be placed close to each other or not? When identifying tightly interconnected clusters of accounts in the network, should A and B be assigned to the same cluster or not? Popular network analysis tools such as the open-source tool Gephi (Bastian et al., 2009) provide some answers to these questions (by default, Gephi shows edges in directed networks as curved in a clockwise direction, for instance); popular cluster detection algorithms like Louvain (Blondel et al., 2008), including its implementation in Gephi, however, treat all networks as undirected and thereby necessarily overlook what could be important network features. Cluster detection approaches that preserve directionality are available, but considerably less widely used—especially by researchers who draw on network science methods but do not have extensive network science training themselves. Our point here is not to criticize these researchers and their tool choices, of course—rather, we highlight these challenges as issues that much published research, and especially cross-disciplinary studies that use network visualization as a means to an end rather than as an end in itself, tend to ignore. The methodology we outline in the present article is therefore designed in particular to help such projects at the nexus of communication studies and network science avoid the pitfalls that arise from an inexpert application of standard network analysis tools.

The third and final limitation to note here is that direct, explicit *interactions* between accounts—in the form of @mentions, retweets, quote tweets, comments, on-shares, etc.—do of

https://doi.org/10.1017/nws.2026.10040 Published online by Cambridge University Press

course only constitute one distinct and not necessarily representative subset of all the public communicative *actions* that participants in a given debate may engage in. Prominent actors, for instance, may simply make their contributions in the form of public statements, without explicitly interacting with others, if they are confident in the reach that their posts will achieve on their own. Similarly, some—especially minority—actors may deliberately refrain from interaction with others in order not to attract unwanted attention, for instance from potentially aggressive and abusive supporters of dominant groups. Yet others may deliberately seek to create the impression that they are acting independently, even though privately they are coordinating their messaging with other like-minded actors. Conceptually, any such actions, too, are communicative practices, as they are performed with the awareness of others (Warde, 2016), even if not interacting with them directly; these actions are equal to *inter*actions in their contribution to and impact on the overall social network, and indeed many of these initial communicative actions without an explicit interlocutory target (i.e., directed instead towards an implicit, imagined audience within the social network) are likely to spark more explicitly targeted interactions as members of that audience choose to respond to the initial speaker, and subsequently also to each other.

Previous research (e.g., Baele & Brace, 2025; Hellsten & Leydesdorff, 2020; Radicioni et al., 2021) has already begun to address the limitations of a focus solely on interactions by constructing networks not based on interactions between accounts, but on similarities of language or platform affordance use. However, where such studies are constructed on a single type of action (e.g., common hashtags only), they inevitably overlook more complex combinations of diverse action and interaction types. Whatever each actor's reasons for acting in a specific way, for the purpose of analyzing public discourse and identifying the various agonistic and antagonistic groups that engage in it, and that may be more or less destructively polarized against one another, it would be desirable to incorporate all of their public communicative actions into the analysis, including but not limited to their explicit interactions with others.

This is the intention at the core of the practice mapping approach (Bruns et al., 2025), as an extension of and complement to conventional interaction network mapping. It considers both actions and interactions to be part of an account's repertoire of *practices*, and in doing so builds on a theorization of practices as developed by the interdisciplinary body of scholarship known as practice theory. Practice theory's main aim is to answer the question "Why people do what they do?" (Nicolini, 2012, 163). By studying diverse "offline" communities ranging from medical professionals (Gherardi, 2019) to fans of Nordic walking (Pantzar & Shove, 2010), these scholars demonstrated that practices, or sustained patterns of action, come into being as actors share goals, values, knowledge, and material environments with relevant others. Some practice theorists refer to loose groups of people undertaking practices together as "communities of practice" (Lave & Wenger, 1991).

As we argue elsewhere (Kasianenko, 2026), the same elements are relevant to sustained doings by users of digitally networked environments. The affordances of digital platforms allow their users to perform various actions on their own (but with awareness of relevant others who may see the traces of these actions), or engage in interactions with other accounts that may or may not be members of the same community of practice. Depending on the communicative affordances enabled by a given platform, then, and on the platform user community's adoption and adaptation of these practices, what we consider to be an account's practices and incorporate into our analysis can encompass a broad range of types and targets of actions and interactions. They may include:

Actions:

- Topics addressed by the account's posts and/or comments, e.g., determined by topic modeling.
- Language choices of the account's posts and/or comments, e.g., extracted through word embeddings.

https://doi.org/10.1017/nws.2026.10040 Published online by Cambridge University Press

- Specific individuals and organizations addressed, e.g., identified through named entity recognition.
- Sentiment and emotive language employed, e.g., assessed through sentiment analysis.
- External content referenced, e.g., at the level of distinct URLs or overall domains shared.
- Audiovisual materials included in posts, e.g., as embedded images or videos.
- Other manually or computationally extracted/coded content features, dependent on the case study.

Interactions:

- Direct mentions of other accounts, using the specific affordances available on the platform.
- On-sharing of content from other accounts, with or without further contextualization or commentary.
- Responses to content from other accounts, in the form of replies or comments to their posts.
- Likes, reactions, up- and downvotes and other formalized nonverbal reactions to other posts.

This list is indicative only, of course; its full extent depends both on the practices likely to be encountered in any given dataset, and on the specific interests explored by a concrete research project. By further extension, even more static account attributes—such as profile descriptions and self-categorizations—might be included here, especially if they point to or are likely to affect the range of communicative practices an account is most likely to engage in. Account descriptions could be coded for the presence of explicit statements indicating political orientation, for instance, or for more indirect markers of personal identity—such as emoji. Whatever the set of distinct actions and interactions that is operationalized for a given study, the network map that results from this process would then serve as the basis for the identification of any clusters—groups of accounts—that are similar in these practices, and distinct from other groups. This has a range of applications, but to the extent that these practices are relevant to the study of polarization, in particular, it serves as a key step towards examining how each such group expresses its own viewpoints and engages with those of others—through the full set of actions and interactions considered, rather than only through its direct interactions via the platform's available communicative affordances.

Practice mapping operationalizes any and all such practice attributes by first converting them into individual vectors describing their use by each account, and then by executing a pairwise comparison of these vectors for each pair of accounts. This produces a normalized rating of the similarity between each pair of accounts for each practice; these similarities can then be analyzed separately or combined into an aggregate similarity rating for each pair of accounts. These similarity ratings, finally, serve as input into the practice mapping network analysis and visualization: here, they represent the strength of the (undirected) edges connecting each pair of actors, and thereby enable the detection of clusters of especially similar actors, and the positioning of these clusters in relation to each other. This diverges from studies that leverage bipartite graphs (e.g., of accounts to domains, or accounts to hashtags) by "folding in" trace data on multiple types of actions into an undirected monopartite visualization of similarities between accounts. It is equivalent to combining multiple affinity networks, derived from both actions and interactions, into a unified vector embedding space. We demonstrate this process in detail through the case study of Australia's Voice to Parliament referendum campaign on Facebook, but step through the process here in a more abstract sense to outline the reasoning behind this approach. (Appendix A outlines an implementation of this approach in a series of Google BigQuery SQL queries.)

https://doi.org/10.1017/nws.2026.10040 Published online by Cambridge University Press

The process begins by extracting all desired practice attributes and converting them to individual vectors for each account. For instance, we may apply topic modeling to the entire corpus of posts in the dataset to identify a set number of topics—say, 100. We could then calculate the affinity of each post by a given account with these 100 topics, and combine these per-post affinities into an aggregate pattern that represents the account's overall topical preferences. This would produce a 100-dimensional vector of topic affinities for each account. (In practice, it will be desirable to exclude largely inactive accounts from this analysis, and to ignore posts that are too short for reliable topic modeling. Such exclusion of inactive accounts is motivated by practice mapping's attention to sustained, rather than intermittent, patterns of action.)

Similarly, for a given account we might extract all its mentions of other accounts, count how often the account engaged with each of these others, and convert this information into a new vector whose dimensionality is determined by the total number of accounts that receive mentions in the dataset. We may also wish to normalize the length of this vector (i.e., rescale the length of each vector to 1.0 while keeping the relative contributions of its components the same) in order to emphasize relative rather than absolute similarity between practices: in this scenario, an account which mentions target A 100 times and target B 200 times would be regarded as highly similar in its communicative practices to an account which mentions target A 200 times and target B 400 times—both show very similar preferences for whom to engage with, even if at different levels of volume. (Again, though, highly inactive accounts should be excluded—an account mentioning A one time and B two times may only accidentally appear similar to the other two.)

Each of these vectors describes one specific attribute of an account's practices, then: its choice of topics, its targeting of other accounts, its sharing of information from external sources, etc. Each of these attribute vectors can then be compared with the equivalent vector describing each other account's practices, using various available vector similarity measures (cosine similarity, Euclidean similarity, dot product, etc.). By incorporating multiple vectors representing a wider array of "sayings and doings"—practices—on a platform, our similarity measures enrich the resultant network with nuanced attributes from the trace data. As a measure which is comparatively unaffected by differences in vector lengths—and therefore generally ignores small differences in the *strength* of practices as long as they are similar in *style*—we will utilize cosine similarity here.

Cosine similarity comparisons between any two vectors result in a normalized value between 0 (entirely unrelated: vector angle 90°) to 1 (entirely identical: vector angle 0°); the crossover point from predominantly unrelated to predominantly related is located at $cos(45°)$, or approximately 0.707. We note that negative cosine similarity values (vector angles above 90°) are not possible in our operationalization, as they would imply vectors pointing in opposite directions—however, as defined here practice attributes can only be present or absent: we do not position one action (e.g., mentioning a specific political leader) as the precise opposite of another action (e.g., mentioning their chief opponent). Importantly, too, these similarity comparisons are symmetrical: the comparison of A with B produces the same similarity rating as the comparison of B with A. This substantially simplifies our subsequent analysis, removing the problems with managing directionality that conventional interaction network mapping continues to face.

Nonetheless, this pairwise comparison between all accounts for each practice attribute to be analyzed constitutes the most computationally intensive element of the practice mapping process, and for each attribute scales quadratically with the number of accounts to be analyzed, at a rate of $n^2/2$. For datasets comprising large numbers of individual accounts, therefore, it is advisable to filter for those accounts that are especially active (and therefore most likely to model common practices across the entire dataset), or to otherwise extract a representative or random sample from the data. As we discuss in the concluding section of this article, once clusters of accounts with similar practices have been identified from a practice mapping of this smaller sample, further accounts from the full dataset can be added to the practice map without the need for full pairwise comparisons with all other accounts.

https://doi.org/10.1017/nws.2026.10040 Published online by Cambridge University Press

In any further use of these similarity ratings, individually or in combination, it is also crucial to remember that any similarity ratings below the mid-point of *cos(45°)* or ∼0.707 predominantly represent *dissimilarity* rather than *similarity*. Especially when used in subsequent network analysis and visualization, where the similarity ratings will be operationalized as edge weights, the inclusion of such ratings may substantially distort the analysis, as even such lower weight values would still serve to attract two network nodes to each other, rather than repel them from one another as they actually should. It is therefore highly advisable to exclude any similarity ratings below this threshold value; the remaining similarity ratings (now between ∼0.707 and 1) could also be normalized to use the full value scale between 0 and 1. Alternatively, where network analysis tools support negative edge weights that indicate repulsion between two nodes, the ratings could be normalized to a scale from −1 to + 1, with original values around 0.707 revised to a new midpoint of 0. (To demonstrate that the application of such a threshold does not falsify the subsequent analysis, Appendix B compares practice mapping results for a broad range of thresholds, as applied to our case study.)

The pairwise comparisons between accounts thus result in several parallel similarity ratings for each pair of accounts, describing for instance similarities in their choices of topics, their mentions of other accounts, and their utilization of external sources. Across the entire dataset, each of these already describes a practice mapping network, for this specific attribute: in other words, they represent a network of affinities between all accounts, which could be used to mathematically and visually analyze this network, and in particular to detect clusters of accounts with especially strong mutual affinities. However, at this stage it is also possible to combine these *individual* practice similarity ratings into an *overall* similarity assessment that incorporates all of the practices selected for analysis. This also enables a relative weighting of practices, to distinguish between more generic and more distinct practice patterns. For instance, for a dataset addressing a specific public debate, differences in language choices may be comparatively minor if all actors use the same issue-specific jargon; differences in what information sources they share or whom they interact with may be more pronounced, however. In combining the various individual similarities, then, researchers may wish to add weighting factors, justified by the research context, to the calculation. For a pair of accounts A and B, this might take the following form:

$$combined_similarity_{A,B} = \alpha \times language_similarity_{A,B} + \beta \times mention_similarity_{A,B} + \gamma \times source_similarity_{A,B} + \ldots$$

Weighting the various similarity components as we combine them into a single similarity measure necessarily privileges some activity components over others; that is the point of the weighting. This is justifiable if it is known *a priori* that one component is more important than another in the context of the present analysis; if that is not the case, it would be advisable to proceed with equal weighting first, evaluate which practice components are common across the entire dataset or occur in distinct combinations only, and on the basis of that evaluation iteratively explore possible weighting choices. Much as with choices of filtering, processing, and visualization parameters in conventional interaction network analysis of communication data, there is no one "correct" choice here, and scholars will instead need to clearly justify the particular choices they made in their specific context. (Appendix C explores the contributions of the individual practice attributes to the combined similarity score, and the consequences of different approaches to weighting these components, for our case study.)

Whether utilizing the similarity ratings for individual practice attributes, or this combined similarity across a bundle of attributes, these ratings represent a network of affinities between accounts, as we have noted. Based in pairwise comparisons, these affinities are symmetrical: A's affinity with B is identical to B's affinity with A. This substantially simplifies subsequent analysis: it enables the use of cluster detection algorithms like Louvain (Blondel et al., 2008) that do not

https://doi.org/10.1017/nws.2026.10040 Published online by Cambridge University Press

account for edge directions or reciprocity; and it similarly allows for simplified network visualizations that have no need to distinguish between distinct directed edges connecting each pair of accounts.

When we apply network analysis and visualization tools to these practice mapping networks, then, what they reveal are patterns in, and clusters of, accounts that act in substantially similar ways. Whether these constitute genuine *communities of practice* remains yet to be seen, however: contrary to the overly loose use of the term "community" in computer science terms like "community detection," to qualify as a genuine community of like-minded actors in the more precise sense used by media and communication scholars the accounts accumulated into a cluster would also need to show mutual awareness of each other, commitment to a common purpose, and adherence to a shared system of beliefs. The presence of such collective attributes can be determined only through further investigation—practice mapping therefore is only a means to an end, and not an end in itself; it must be followed by further qualitative and/or quantitative analysis of the practice clusters it has detected. Previous studies have identified practices of online communities through ethnographic methods, surveys, discourse-analytical or mixed-methods techniques (Hillman et al., 2023; Mendes et al., 2023; Trillò et al., 2022). However, to our knowledge, practice mapping is the first scalable and interpretable approach for studying practices through social network analysis—a methodological and conceptual goal signposted by scholars interrogating the usability of the methodology for making sense of digital trace data (Howison et al., 2011). In the following section, we demonstrate its practical deployment and analytical utility in the context of a case study of political debate on Facebook during the Australian Voice to Parliament referendum in 2023.

4. Facebook debates about Australia's Indigenous Voice to Parliament referendum

A referendum on the establishment of an Indigenous Voice to Parliament, representing Australia's First Nations people in legislative decision-making processes, was a key promise made in the election-night victory speech of incoming Prime Minister Anthony Albanese, from the centrist Australian Labor Party, on 21 May 2022. The referendum required all Australians to vote on whether they agreed to a proposed change to the Australian Constitution to recognize Aboriginal and Torres Strait Islander people as the First Nations people of Australia, and to establish a Voice to Parliament—essentially a committee of Indigenous Australians who would advise government on laws that impact First Nations people. The proposal stemmed from the 2017 Uluru Statement from the Heart, which was created by 250 Aboriginal and Torres Strait Islander delegates who met at the foot of Uluru in Central Australia on the lands of the Anangu people to pen a 439-word affirmation about having a Voice in the Australian Parliament (Fredericks & Bradfield, 2021).

This was a historic moment for Indigenous people, but also the beginning of a wave of confusion, misinformation, mistrust, and racism, experienced by everyday Australians and some politicians at all levels of government, who have seen little or no change from the Commonwealth when it comes to Indigenous people's issues and rights. The most significant prior vote on Indigenous recognition occurred in May 1967, when Australians voted in a referendum to count Aboriginal people on the national census and give the Australian Parliament power to make laws with respect to Indigenous people. The 1967 referendum asked the Australian people: "Do you approve the proposed law for the alteration of the Constitution entitled 'An Act to alter the Constitution so as to omit certain words relating to the people of the Aboriginal race in any state and so that Aboriginals are to be counted in reckoning the population'?" It resulted in a resounding 90.77% Yes vote; notably, there was no No campaign at the time, as both sides of Parliament supported a Yes vote.

The Voice referendum itself was held on 14 October 2023; as with all Australian elections and referenda, participation was compulsory for all voters. Although opinion polls in early 2023 still

https://doi.org/10.1017/nws.2026.10040 Published online by Cambridge University Press

suggested that some two thirds of Australian voters were in favor of the constitutional recognition of the country's Indigenous peoples, the eventual referendum was defeated by a 60% No vote, with none of the country's six states and only one of its two territories returning a majority Yes vote (McAllister & Biddle, 2024).

In the documentary *Exhausted Voices: Past, Present, Future* (Johannessen, 2025), Aboriginal Embassy co-founder Professor Gary Foley predicted the failure of the Voice referendum, stating that Australians today were more distant than they had ever been when it comes to Aboriginal people. He suggests that after the 1967 referendum and in the early 1970s, a *white backlash* began: media propaganda, especially also from Rupert Murdoch's News Corporation, which dominates the Australian media landscape (and which had supported the Aboriginal Land Rights movement in 1972), substantially affected government, bureaucratic, and public attitudes. This also enabled the rise of far-right political actors such as Pauline Hanson and her One Nation party, creating a space for racist attitudes in Australian politics and thereby facilitating the defeat of the Voice to Parliament referendum.

As Pattel-Gray (1998) argues, racism is an endemic and chronic problem and must be addressed and solved. The Jumbunna Institute's annual *Call It Out* report for 2023–24 highlights abuse on social media as a major problem for Aboriginal and Torres Strait Islanders. The report (Coombes et al., 2025: pp. 73–74) shows that 42% of abuse reports were related to the use of social media, including intense personal abuse about the Voice referendum. Online comments supporting Indigenous people were targeted by racist responses; this was the case especially for comments on videos by *Sky News Australia* on Facebook. Overall, Facebook featured heavily as a social media site spreading racist messages tied to the Voice and the referendum.

Such abuse was not limited to online spaces: other reports document how Indigenous and non-Indigenous people who supported the Voice with "Yes" signs in their front yards were targeted by "No" voters; one case reported graffiti on a pathway that read "Earn Your Voice" alongside a Nazi symbol. Similarly, offline and online abuse and problematic activities targeted not only the Voice referendum, but extended to Indigenous rights and recognition in general: Aboriginal journalist and reporter Natalie Ahmat (in Johannessen, 2025) and her colleagues at the National Indigenous Television (NITV) network have noted that the mis-, dis-, and malinformation circulating in mainstream media was more directed at Aboriginal and Torres Strait Islanders, and not about the proposed change to the Constitution.

We chose the Voice referendum as a case study for this article as referenda necessarily condense complex political choices to a binary Yes/No vote; potentially, and in line with the symptoms outlined above, this renders them especially susceptible to destructive polarization dynamics. Space available for this article does not permit us to provide a detailed analysis of all developments over the nearly year-long referendum campaign (we will address this in other publications); here, we present a brief analysis of communicative patterns amongst referendum campaigners on Facebook as an application of the practice mapping approach that illustrates its utility in identifying discourse alliances and antagonisms, and thereby serving as a key step towards the subsequent diagnosis of destructive polarization.

A referendum is inherently polarizing through its reduction of a complex question to a Yes/No choice; however, this does not mean that its analysis through the methodological lens of practice mapping and the conceptual framework of destructive polarization is superfluous, or the results of such an analysis predictable. The overall referendum framework does not mean, for instance, that we will already know *a priori* which actors and groups are engaged in the referendum debate, how they are positioning themselves in relation to the referendum question and to each other, and to what extent they are showing symptoms of destructive polarization in their actions. Any one actor or group may choose to push back against the referendum process's inherent erasure of complexities by providing a detailed argument for or against the proposal being put forward, or lean into this erasure by employing simplistic rhetoric that further encourages voters to abstain from any deep engagement with the matter at hand. Additionally, as we have pointed out, the

https://doi.org/10.1017/nws.2026.10040 Published online by Cambridge University Press

symptoms of destructive polarization are most problematic when they occur together: while the erasure of complexities inherent in the referendum question is necessarily present to some extent, it becomes more destructive when it is coupled with rhetoric that uses emotions to exclude other valid actors and their perspectives, or dismisses and discredits the information they provide, for example. Our analysis—and of the use of practice mapping as a means to that analysis—serves to assess the polarized or polarizing practices and stances of the distinct actors and groups within the debate, and to illuminate whether such polarization was symmetrical, or driven by specific political interest groups.

For this case study, we focus on public debates on Facebook for several reasons. First, in spite of the outsized attention scholarly research has often paid to Twitter (now X), Facebook remains the far more widely used platform in Australia, especially also reaching a more general and politically disinterested audience (Park et al., 2024: 95). Combined with Australia's compulsory voting system, this means that it is likely to serve as key source of information for ordinary voters needing to make up their minds about the referendum question they are required to vote on; indeed, we find substantial campaigning activity on Facebook. Second, the structure of data available from Facebook (at the time of the referendum from CrowdTangle; now from the Meta Content Library, which offers comparable data) positions it as a valuable application case for the practice mapping approach: Facebook data contain very limited information about direct communicative interactions between public actors on the platform, and therefore do not lend themselves well to conventional network analysis. Practice mapping offers a powerful new framework for applying network analysis tools to the networks of similarity between the broader range of communicative actions that it supports—including, but not limited to, interaction through content on-sharing between spaces.

Our data for this case study were gathered from CrowdTangle, which provided posts and metadata from public pages, public groups, and public verified profiles on the platform (collectively referred to as public spaces from here on). We queried CrowdTangle for posts matching any of "Voice to Parliament", "Uluru Statement", (Voice) AND (Referendum), VoteNo, VoteYes, voicereferendum, ulurustatement, yes23 for the period from 1 January 2023 to 13 October 2023 (the day before the referendum), resulting in 89,094 posts from 9,483 Facebook spaces. In the context of the Voice referendum, we focussed our analysis on four distinct attributes to describe the communicative practices of each space.

First, in light of the political and news-related nature of the referendum, we drew on content sharing patterns: for each space, we constructed a vector to describe how often it had shared content from any distinct domain (not including Facebook's own domains), and normalized the length of the vector against the total number of times it had shared external content. We did so at the level of the domain (e.g., *abc.net.au* for the Website of the Australian Broadcasting Corporation (ABC), the main Australian public service media organization) rather than the distinct URL (e.g., *abc.net.au/news/123456789*, for a specific article on that site) as we assumed that most domains—and especially those of mainstream and alternative news outlets, political parties, referendum campaign organizations, etc.—would already reflect a general supportive, neutral, or oppositional stance towards the Voice proposal. In addition to internal Facebook domains (such as *facebook.com*, *fb.me*), we excluded from this analysis any generic domains belonging to other social media platforms (e.g., Twitter, YouTube), short URL services (*bit.ly*, *ow.ly*, *t.co*), or content aggregators (Apple News), whose presence across the data does not indicate strong similarities in communicative styles.

Second, for one of these excluded domains we examined similarities at the level of individual content items shared in posts: we extracted the unique video IDs from any YouTube links shared in posts, and constructed vectors for each space that described how often it had shared which distinct YouTube videos; these were again normalized against the total number of videos shared by a given space. We focussed on YouTube here as it is a particularly prominent source of external content in the dataset, and as prominent YouTube videos shared on Facebook originated both

https://doi.org/10.1017/nws.2026.10040 Published online by Cambridge University Press

from major Australian news outlets and from fringe political influencers. (This approach could also be extended to other key content sources, of course—including bit.ly or TikTok, for instance.)

Third, as the only indicator of direct interactions between different Facebook spaces available in our data, we also filtered the shared links to internal Facebook URLs to retrieve any evidence of content on-sharing—i.e., of instances where one space reposts the content posted by another space, with or without added commentary. This is the closest equivalent to retweeting available on Facebook, and while it is not especially common amongst public spaces on this platform, it does indicate direct attention to the content published by other spaces, and may therefore point to what a space considers to be its fellow travelers or persistent antagonists in the referendum campaign. CrowdTangle does not provide this information in any immediately accessible way; rather, the numerical ID or alphanumerical handle of the source Facebook space can be extracted with some difficulty from the internal Facebook URL pointing to the shared content that is associated with the sharer's post. Having extracted this information, we against constructed a vector for each space that represented such sharing activity, normalized by the total number of on-shares posted by the space.

Finally, for each space we used Google's Vertex AI to create text embeddings for any of its posts that contained more than 50 characters (i.e., excluding very short posts that contain insufficient text to reliably identify their language choices); we then averaged these text embeddings across all of the posts by a given space, in order to arrive at an aggregate text embedding result per space. Text embeddings can be understood as dense, high-dimensional vector representations of text, learned by language models such that texts with similar semantic content occupy nearby positions in the vector space. Embeddings are returned by Vertex AI in the form of high-dimensional, normalized embedding vectors, enabling us to treat them in an equivalent way to the other vectors produced for the practice mapping process.

For each of these four attribute vector types—domain sharing, YouTube video sharing, Facebook post on-sharing, and content word embeddings—we then conducted a pairwise comparison of vectors across the Facebook spaces in the dataset, using cosine similarity as our similarity measure. To exclude spaces that engaged only peripherally with the Voice referendum, for our subsequent analysis we selected only those spaces that made at least three posts related to the Voice referendum during our target timeframe; this reduced the total population of spaces to 3,545 (37% of the initial dataset), which collectively posted 81,727 (92%) of the posts in our original dataset.

We note here that the populations of spaces addressed by each vector comparison varies slightly: for instance, spaces that did not share any YouTube videos during the timeframe are necessarily not represented in the YouTube vector comparison. This is acceptable (these spaces are inherently dissimilar to any spaces that did share YouTube videos), and indeed enables pairs of spaces that are highly similar across all four vector categories to stand out especially strongly.

As noted above, the per-attribute similarities between spaces can be combined into an overall similarity rating in a variety of ways, for instance by weighting similarities in particular attributes (e.g., the rarer practice of content on-sharing) more highly than similarities in more generic attributes (e.g., in language choices as represented by word embeddings). For our purposes in this illustrative case study, we simply calculated the average across all four similarity calculations. In this, we treated a space's similarity with others as zero for that practice if it did not engage in a particular practice such as YouTube video sharing (this applies also for any pair of spaces that *both* did not engage in that specific practice, since we cannot determine whether, for instance, both spaces deliberately *chose* not to share YouTube videos, or merely did not encounter any such videos that they deemed shareworthy). We then filtered this network of average similarities between spaces for a minimum value of 0.707, in order to remove network ties that predominantly indicate *dis*similarity, and removed outliers by excluding spaces that did not meet a minimum weighted degree threshold (i.e., spaces whose communicative practices showed few similarities to others in the dataset). (As noted, Appendix B shows how alternative minimum similarity settings

https://doi.org/10.1017/nws.2026.10040 Published online by Cambridge University Press

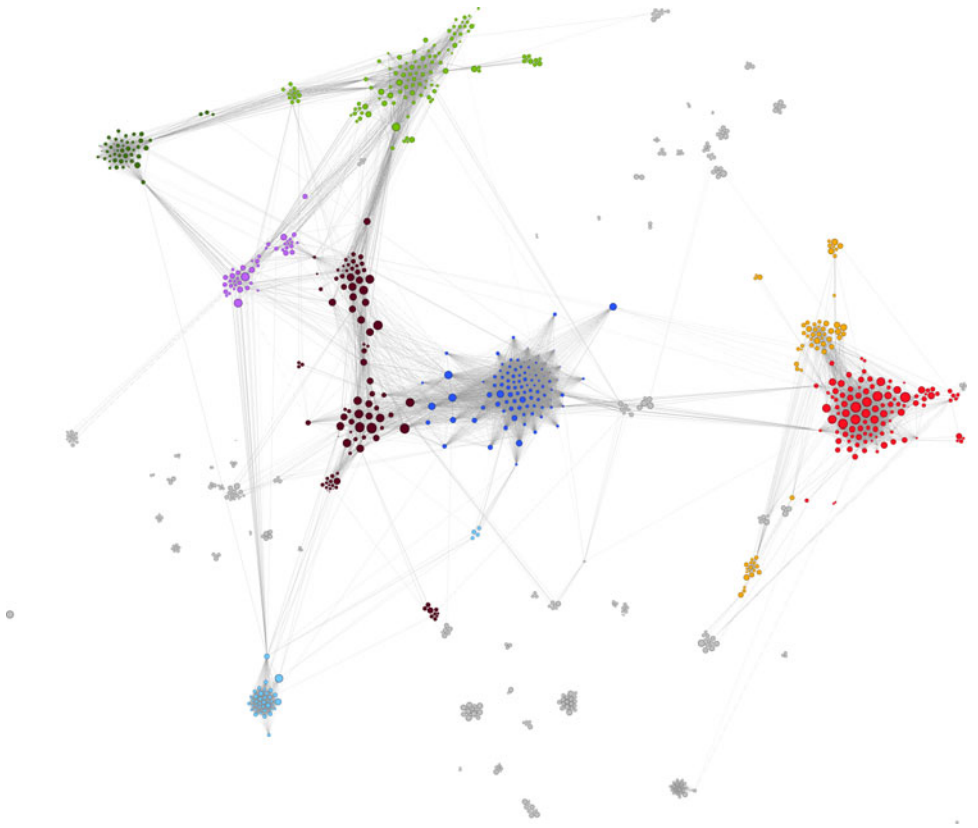


Figure 1. Practice mapping network for the Voice to Parliament referendum debate on Facebook. Node sizes indicate volume of posts per space; node colors represent Louvain cluster detection results.

affect the outcome of this filtering process; it demonstrates clearly that values other than 0.707 change the sparseness of the resulting network, but not its overall structure. Appendix C further explores how the individual practice attributes contribute to the combined similarity score, and demonstrates two alternative approaches to weighting these components.)

These filtering steps are designed to retain only those spaces that represent common patterns of practice within the dataset; how such filters are applied will vary for each case study, of course. In our case, the final population included 982 spaces meeting these criteria (10% of the initial dataset), accounting for 54,195 (61%) of the original volume of posts. These percentages indicate that our approach selected the most active contributors to the Voice debate on Facebook. We visualized the network of similarities for this population of spaces using the Force Atlas 2 layout algorithm (Jacomy et al., 2014) and the Louvain cluster detection algorithm (Blondel et al., 2008) in Gephi (Bastian et al., 2009). The result is show in Figure 1; node sizes indicate the number of posts contributed by each space.

The resulting network shows a number of distinct structures: most importantly, a notable separation between a large and highly active cluster of spaces on the right, accompanied by a closely connected smaller cluster, and a collection of smaller and generally less active clusters on the left, which are clearly distinct but relatively strongly interconnected. There are only very few nodes that serve as connectors between these two sides of the network, and, because of their sparse connections to either side, none of them show significant levels of betweenness centrality. Moderate betweenness centrality exists only for those nodes that connect the various clusters on the left or the right, respectively. To avoid any misunderstanding, we stress again that this network does *not*

https://doi.org/10.1017/nws.2026.10040 Published online by Cambridge University Press

represent a network of direct interactions between these spaces; rather, it maps the similarities and differences between these spaces' communicative practices over the course of the referendum campaign, for practices combining domain sharing, YouTube video sharing, Facebook post on-sharing, and language choice.

As noted above, practice mapping is not an end in itself, however, nor are the network visualization and community detection algorithms applied here. As others have cautioned (see e.g., Blondel et al., 2024; Jacomy et al., 2024), the statistical and visual outputs of such algorithms should not be taken at face value, in part due to the overfitting tendencies of modularity optimization models or the possibility of misinterpreting visual representations of a network and distances between clusters in such visuals. Therefore, now that we have identified several clusters of accounts with distinct communicative practices within our dataset, the critical interpretive work must begin. While the network structure in Figure 1 with its relative affinities and distances suggests the possibility of various discursive alliances and antagonisms between the clusters of spaces it contains, we must first verify that this structure represents a meaningful and explainable result. To do so, we engaged in a process of qualitative close reading: we identified the spaces in each cluster that were most active (by number of posts) and/or most central (by weighted degree within the cluster, indicating high similarity with other cluster members), and reviewed their titles, descriptions, and a sample of their posts. This usually revealed clear patterns relating to the communicative identity of each cluster, and its stance towards the Voice referendum, and showed the network of clusters on the left to be broadly supportive of the Yes campaign, with the large cluster and its smaller companion on the right representing the No campaign. It also highlighted a notable absence—that of distinct clusters representing Indigenous actors, whose voices were arguably the most important to be reflected in the public debate on the Voice Referendum (see also Pereira et al., 2025, who report a similar absence of Indigenous voices in a debate where they are the most affected party, in the Brazilian context).

Overall, clusters at the center or on the periphery were generally more neutral (though somewhat more similar in communicative practices to the Yes than the No side), with two such clusters forming around the reporting of Australia's two public service media networks, the ABC and the Special Broadcasting Service (SBS)—the latter of which also has a strong focus on addressing Indigenous communities, and operates the National Indigenous TV (NITV) channel. Meanwhile, central to the prominent No cluster were several Facebook spaces associated with *Sky News Australia*, a conservative and controversial cable TV and online news channel operated by News Corporation (cf. Copland et al., 2025); the smaller cluster accompanying it was strongly affiliated with the far-right One Nation party. On the Yes side, clusters formed around a number of distinct leading Facebook spaces, including the official Yes23 campaign page and a page dedicated to the Uluru Statement from the Heart, and also represented various party-political and other interests. Figure 2 presents an annotated map of the similarity network.

Available space does not permit us to present a full discussion of these observations, nor a further exploration of the specific communicative practices that contributed to these patterns. However, as indicated in Figure 2 our initial reading is of the network of clusters on the left of the map as a loose discursive alliance between groups that broadly support the Voice to Parliament initiative, if for a variety of reasons (from strong support for Indigenous recognition and reconciliation to strong opposition to the conservative forces from the Liberal and National Parties (LNP) that were prominent in the No camp). Conversely, the No side appears considerably more singular and unified both in its identity and in its communicative practices, and is in both aspects clearly antagonistic towards the network of Yes campaigners.

If these initial observations are supported by further in-depth analysis, they point to considerable—and expected—polarization between the Yes and No camps, as well as to milder differences between the various groups comprising the Yes side. They suggest that partisan sorting and the move from issue-based to identity-based and affective polarization is more advanced

https://doi.org/10.1017/nws.2026.10040 Published online by Cambridge University Press

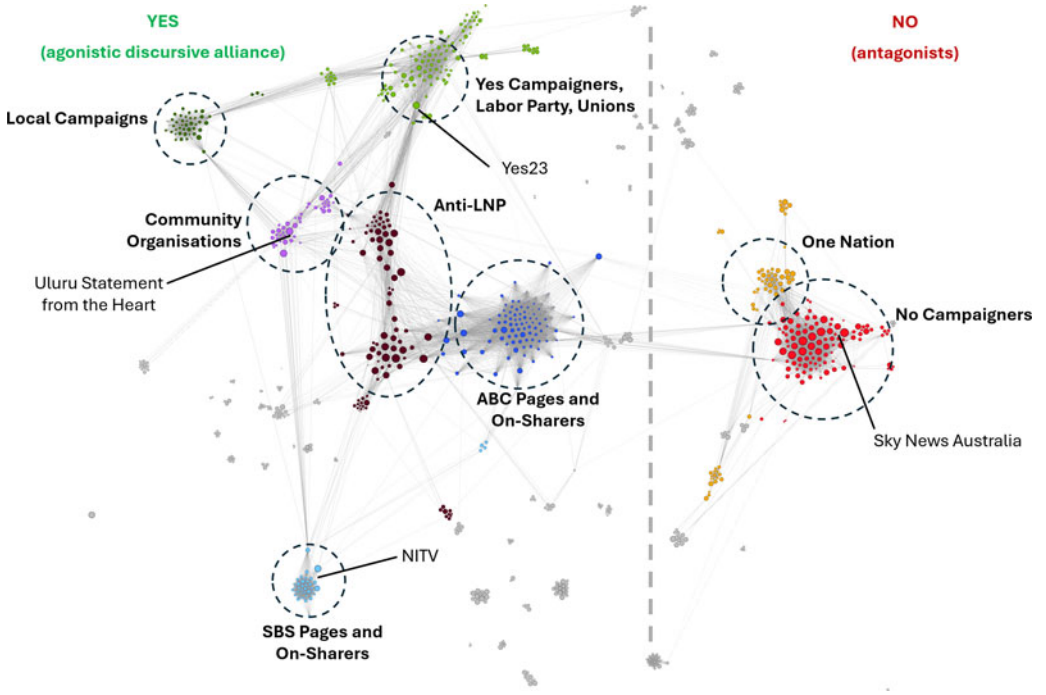


Figure 2. Annotated practice map, showing selected cluster descriptions (circles) and key Facebook pages (lines).

on the No side than amongst the various distinct groups supporting a Yes vote. This supposition would be further bolstered by similar analyses of public debates in Australia on different issues: if these produced networks where at least some of the same actors appeared in similar positions of alliance with and antagonism to each other, this would strongly point to partisan sorting beyond the level of individual issues, and thus to ideological and potentially affective rather than merely issue-based polarization. Further analysis of the communicative practices embraced by these different groups might then be employed to determine whether any such practices show symptoms of destructive polarization: in this sense, practice mapping itself serves as a necessary step towards the discovery of distinct groups of communicative actors, and enables the isolation of these groups' practices for further in-depth analysis, using qualitative, quantitative, or—most likely—mixed methods. In turn, any distinct communicative practice patterns identified through such analysis also serve to validate the results of the practice mapping itself.

To illustrate at least in brief how this might proceed for the present case, and to thereby validate our preliminary interpretation of the practice map, we conducted a zero-shot classification of the content posted by members of each of the major clusters identified by our practice mapping. This employs a universal classifier built on Natural Language Inference (NLI) principles (Laurer et al., 2023). In NLI, a language model receives two texts: a premise (the user's text) and a hypothesis (a statement formulated as a possible label, e.g., "this text is polarizing"). The model then estimates whether the hypothesis is entailed (true), contradicted (false), or unrelated to the premise. For zero-shot classification, the problem is simplified into binary entailment vs. non-entailment, and any label can be expressed as a hypothesis such as "This text is divisive" vs. "This text is not divisive." The classifier produces a probability for entailment, which can be interpreted as the likelihood that the label applies. The particular model we employ—deberta-v3-large-zeroshot-v2.0 (Laurer et al., 2023)—was trained on five major NLI datasets (MNLI, ANLI, FEVER-NLI, WANLI, LingNLI) and 28 additional classification datasets covering sentiment, stance, toxicity, bias, factuality, emotion, and topics. Together, these account for 389 distinct classes. The training

https://doi.org/10.1017/nws.2026.10040 Published online by Cambridge University Press

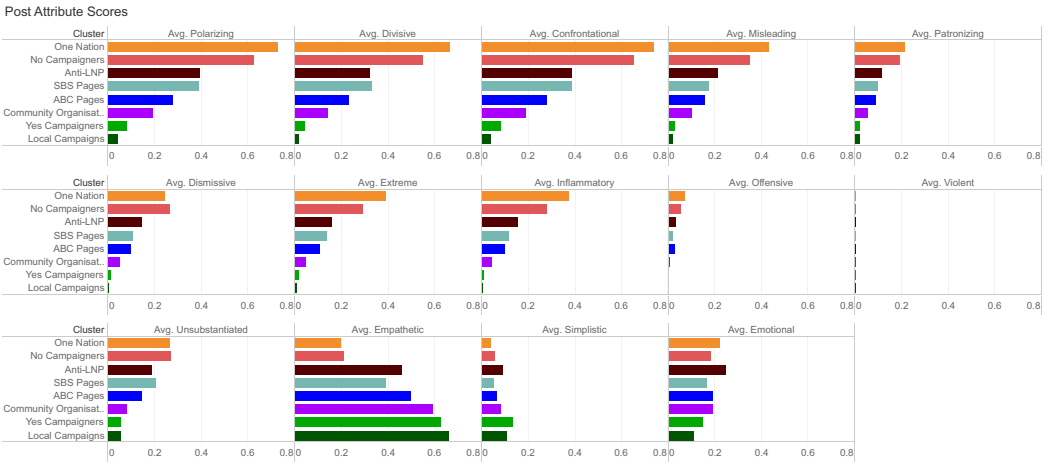


Figure 3. Average post attribute scores per cluster, following zero-shot classification of post texts.

strategy of this classifier makes it a "universal" model: it does not need to be retrained for new tasks but instead evaluates any user-defined label that can be expressed as a hypothesis.

For our case, we selected a set of 14 attributes that serve as linguistic proxies for the forms of polarization and symptoms outlined in Section 2. These include indicators of discursive antagonism (polarizing, divisive, confrontational, inflammatory), symptoms of destructive polarization such as the *discrediting of information* (misleading, unsubstantiated, dismissive, patronizing) and the *erasure of complexities* (simplistic, emotional), as well as markers of affective polarization and its counter-indicators (extreme, offensive, violent, empathetic), some of which indicate an *exclusion through emotions*. We processed each post text with the classifier multiple times, each corresponding to a label set, where the outputs are probabilities for entailment and non-entailment—whether a hypothesis is "true" or "not true" for a given label. Each label was a hypothesis for a content classification in its affirmative and negative form—e.g., "this text is divisive," "this text is not divisive;" "this text is polarizing," "this text is not polarizing." We averaged these per-post results for the full set of posts published by each of the clusters identified through the practice mapping. The result is shown in Figure 3, with cluster colors matching those in the practice maps above. Note that we present this as a proof of concept within our broader practice mapping approach: the aim is not to deliver a definitive measure of polarization, but to illustrate how automated content classification can be combined with cluster mapping to highlight tentative indicators of polarization dynamics.

In spite of the substantial levels of abstraction and the necessarily reductive labels employed by the zero-shot classifier, Figure 3 shows clearly divergent patterns between the communicative practices of the various clusters. The One Nation and No Campaigners clusters, in particular, are marked by comparatively polarizing, divisive, confrontational, misleading, patronizing, dismissive, extreme, inflammatory, and unsubstantiated communication styles, while lacking in empathy; a similar pattern, though less extreme, is exhibited also by the Anti-LNP cluster opposing the leading conservative actors in the No camp. Perhaps more surprisingly, the clusters associated with ABC and SBS Pages and their on-sharers show somewhat similar tendencies, especially for SBS; this, however, may be explained by their role as largely neutral reporters of a referendum debate featuring such polarized and polarizing language from the No camp and its political adversaries: in reporting the statements of such actors, these media outlets would then also appear to the classifier to use this language themselves.

https://doi.org/10.1017/nws.2026.10040 Published online by Cambridge University Press

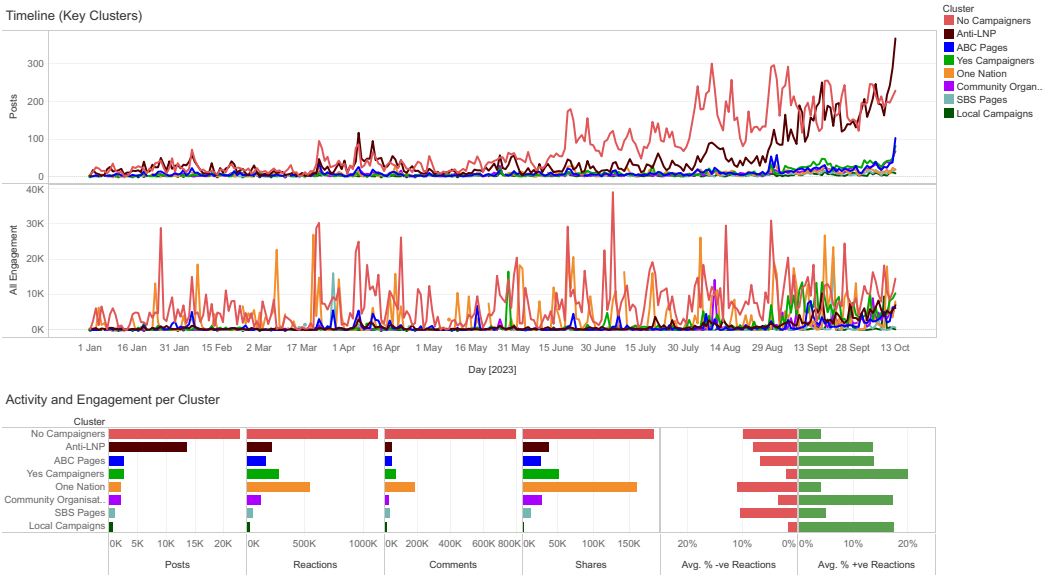


Figure 4. Posts by and engagement with Facebook spaces in the key clusters.

Conversely, the Yes Campaigners and Local Campaigns clusters' communicative patterns showed almost diametrically opposite patterns, with both especially low on the defining attributes of the No cluster, and very high on empathetic language. This suggests a very different campaigning style, presenting a Yes vote as a simple choice for empathy with Australia's Indigenous peoples—but in doing so perhaps failing to defuse the attempts by the No side to confuse and complicate the issue and lead voters to endorse the status quo. In combination, the use of empathetic but comparatively unemotional language by some Yes campaigners, and of inflammatory, extreme, and emotional language by the two No clusters also suggests a considerable asymmetry in affective polarization between these different sides of the debate.

Notably, the No Campaigners cluster also differs vastly from all other clusters in its activity over time, and in user engagement with its posts (Figure 4). The cluster was considerably more active since the start of 2023, and ramped up its posting further from mid-March (when the referendum design was revealed) and again from mid-June onwards (after the Constitution Alteration bill triggering the referendum was passed by federal parliament). While the Anti-LNP cluster largely matched the No Campaigners cluster during the early months of 2023, it did not boost its activities in mid-year, and along with other clusters increased its activity substantially only from late August 2023 onwards, when the official Yes and No campaigns were launched. Overall, engagement with such posts is considerably larger for the No Campaigners and One Nation clusters than any of the other clusters; this is especially notable for One Nation, where a relatively modest amount of posts generated outsized levels of engagement activity (and especially post on-sharing). Finally, posts by the No Campaigners and One Nation, but also from SBS Pages and On-Sharers, received substantially more negative ("angry" and "sad") than positive ("love" and "care") Facebook reactions; the opposite is true for the Yes Campaigners cluster and several others.

These results illustrate and validate the utility of the practice mapping method as a foundation for further in-depth analysis: practice mapping has enabled us to identify key groups of Facebook spaces that clearly act in highly divergent ways, and receive substantially different patterns of engagement for these activity practices. In the context of our conceptual framework,

https://doi.org/10.1017/nws.2026.10040 Published online by Cambridge University Press

subsequent steps in the analysis should include a systematic search for the symptoms of destructive polarization both in the statements made by the leading actors in the various clusters in the Voice to Parliament practice map, and in the identity of these leading actors themselves—patterns pointing to marked differences in the use of divisive and empathetic language in posts by these distinct groups already hint at these symptoms of destructive polarization. These further steps are necessarily outside the scope of the present article, however.

5. Further outlook

This article has introduced the practice mapping approach as a methodological framework for the detection and analysis of similarities in the communicative practices of actors in social media debates. It used the 2023 Australian Voice to Parliament referendum debate on Facebook as its case study, aiming in particular to provide the foundations for an identification of potential symptoms of destructive polarization between the actors involved in this debate. The utility of practice mapping is not limited to this particular application, however: more generally, it presents an extension of and alternative to conventional social media interaction network mapping approaches, especially also in contexts where social media platforms and their datasets are not primarily structured around account-to-account *interaction* networks. While the method is still reliant on digital trace data, and thus, as other such methods, will surface practices of more active participants more prominently, it opens the possibilities for mixed-methods network analyses of public debates on a larger number of digital platforms. Indeed, in other studies employing the practice mapping method we have already applied it successfully to data from Reddit, both to examine similarities and differences in the communicative practices of participants in various subreddits, and on that basis also to assess such similarities at the level of whole subreddit communities.

The practice mapping approach as we have outlined here in brief (and see Bruns et al., 2025, for a more detailed discussion) is new and innovative, and as such also open to further extensions. First, as we have noted above there is no obvious limit to the range of attributes that may be considered and operationalized as practices: in addition to the post and account attributes already present in the data, further such attributes could be generated through manual and computational coding or the connection with additional datasets (e.g., describing an account's political, professional, or personal characteristics). Such attributes may also be highly specific to the particular case being studied, of course.

Second, there are substantial opportunities for further experimentation with the combination of attributes—and similarity ratings for each attribute—ahead of their use in the practice mapping itself. Similarities on various attributes could be weighted differently ahead of analysis and visualization, and indeed different such weighting combinations might produce considerably different outcomes. This presents an opportunity as much as a challenge: an interactive experimentation with various weights might help researchers in understanding which attributes contribute especially strongly to the shape of the practice map and are therefore especially diverse in the underlying data. Our exploration of practice attribute contributions and weighting options in Appendix C points towards these opportunities.

Third, the most computationally intensive aspect of the practice mapping process is the pairwise comparison of practice attribute vectors across all selected actors. This complete comparison between all actors may not be necessary to produce a representative map, however; similar to other network capture, analysis, and visualization processes (Münch et al., 2021), to identify the broad network structure it may be sufficient to work with a random or purposive sample of the entire actor population. Further testing is required to determine the reliability of such a simplified approach; additionally, as we have noted, once the overall structure of the network and the attributes of its key clusters have been identified, it may then be possible to add in a larger subset

https://doi.org/10.1017/nws.2026.10040 Published online by Cambridge University Press

of the actor population by systematically calculating their similarities only with the centroids of those clusters, rather than with all other actors, thereby considerably reducing the computational load required.

Fourth, while the case study presented here, and many other network analyses, were undertaken to produce a static map covering the entire timeframe represented by the dataset, this necessarily ignores any evolution in practices over time. As Figure 4 shows, even such time-insensitive practice mapping can sometimes detect temporal factors (e.g., if clusters of actors with specific practices are active at different points within the overall timeframe), but it would be desirable to better integrate such temporal aspects into the practice mapping itself. This could be done by creating separate practice maps for distinct periods; by recalculating specific actors' similarities with the clusters in a static practice map within rolling windows of time and thereby tracing their affinities with these clusters over the course of the entire timeframe; or by incorporating activity levels over time into a separate, temporal practice attribute vector that is used to calculate an additional measure of similarity for each pair of actors (assessing whether such actors were consistently active during similar periods).

Finally, while our approach as outlined here has championed a particular similarity measure—cosine similarity, as applied to vectors of normalized length—there are several alternative similarity metrics that could also be explored (and systematically compared with each other in order to determine their use cases and limitations). Beyond the similarity or distance metrics to compare vectors that we have used, several more sophisticated tools to reduce the dimensionality of the embedded vectors are also available. Various off-the-shelf dimensionality reduction techniques, such as Principal Component Analysis (PCA), have already been used to study social networks (Kurucz et al., 2008; Vilenchik et al., 2019). With particular attention given to how our affinity vectors act as mathematical metaphors for actor practices, these computational techniques can be operationalized for practice mapping, too. Non-linear dimensionality reduction techniques such as Uniform Manifold Approximation and Projection (UMAP; see McInnes et al., 2018) even open up possibilities to augment our methods with even more complex vectors, albeit with different downstream interpretations for the clusters that emerge from such analysis. We therefore offer this article not as a definitive guide to the practice mapping approach, but as an indication and illustration of its potential, and an encouragement for the further development of the method.

Supplementary material. The supplementary material for this article can be found at https://doi.org/10.1017/nws.2026.10040.

Acknowledgement. This research was funded by the Australian Research Council through the Australian Laureate Fellowship project *Determining the Drivers and Dynamics of Partisanship and Polarisation in Online Public Debate*.

Competing interests. None.

Funding statement. This research was supported by the Australian Research Council through the Laureate Fellowship project *Determining the Drivers and Dynamics of Partisanship and Polarisation in Online Public Debate* (grant number FL210100051).

Data availability statement. Facebook data used for this research were provided courtesy of CrowdTangle, which does not permit the public sharing of datasets.

References

Adamic, L. A., & Glance, N. (2005). The political blogosphere and the 2004 U.S. election: Divided They Blog. In Adibi, J., Grobelnik, M., Mladenic, D., & Pantel, P. eds. *Proceedings of the 3rd International Workshop on Link Discovery (LinkKDD '05)* (pp. 36–43). ACM. https://doi.org/10.1145/1134271.1134277.

https://doi.org/10.1017/nws.2026.10040 Published online by Cambridge University Press

Almond, G. A., & Verba, S. (1963 [2015]). *The Civic Culture: Political Attitudes and Democracy in Five Nations*. Princeton: Princeton University Press.

Baele, S. J., & Brace, L. (2025). Socio-semantic network analysis for extremist and terrorist online ecosystems. *Studies in Conflict & Terrorism*. https://doi.org/10.1080/1057610X.2025.2480622.

Barnes, R. (2022). Loving to hate: Fandom fuelling polarized behaviour. In *Fandom and Polarization in Online Political Discussion: From Pop Culture to Politics* (pp. 61–86). Cham: Palgrave Macmillan, https://doi.org/10.1007/978-3-031-14039-6_4.

Bastian, M., Heymann, S., & Jacomy, M. (2009). Gephi: An open source software for exploring and manipulating networks. In William W. C., Nicolas N., Natalie G., & Matthew H. eds. *Proceedings of the Third International ICWSM Conference* (pp. 361–362). Association for the Advancement of Artificial Intelligence. http://www.aaai.org/ocs/index.php/ICWSM/09/paper/viewFile/154/1009/.

Benson, J. (2023). Democracy and the epistemic problems of political polarization. *American Political Science Review*, *118*(4), 1719–1732. https://doi.org/10.1017/S0003055423001089.

Blondel, V. D., Guillaume, J.-L., Lambiotte, R., & Lefebvre, E. (2008). Fast unfolding of communities in large networks. *Journal of Statistical Mechanics: Theory and Experiment*, *2008*(10), P10008. https://doi.org/10.1088/1742-5468/2008/10/P10008.

Blondel, V., Guillaume, J.-L., & Lambiotte, R. (2024). Fast unfolding of communities in large networks: 15 years later. *Journal of Statistical Mechanics: Theory and Experiment*, *2024*(10), 10R001. https://doi.org/10.1088/1742-5468/ad6139.

Bruns, A., Kasianenko, K., Suresh, V. P., Dehghan, E., & Vodden, L. (2025). Untangling the furball: A practice mapping approach to the analysis of multimodal interactions in social networks. *Social Media + Society*, *11*(2). https://doi.org/10.1177/20563051251331748.

Bruns, A., Moon, B., Münch, F., & Sadkowsky, T. (2017). The Australian Twittersphere in 2016: Mapping the follower/Followee network. *Social Media + Society*, *3*(4), 1–15. https://doi.org/10.1177/2056305117748162.

Coombes, L., Allison, F., Cunneen, C., & Selcuk, A. (2025). *Call It Out Racism Register 2023-2024: "If you don't think racism exists, come take a walk with us."* Sydney: Jumbunna Institute for Indigenous Education and Research, University of Technology Sydney. https://callitout.com.au/wp-content/uploads/2025/02/Jumbunna-Call-It-Out-Annual-Report-2023-2024-Final.pdf.

Copland, S., Bruns, A., & Graham, T. (2025). Sky News Australia as network propaganda: How a niche cable channel became an international right-wing propaganda machine. *Media International Australia*. https://doi.org/10.1177/1329878X251406899.

Doctorow, C. (2022). How monopoly enshittified Amazon. *Pluralistic*. https://pluralistic.net/2022/11/28/enshittification/#relentless-payola.

Esau, K., Choucair, T., Vilkins, S., Svegaard, S. F. K., Bruns, A., O'Connor-Farfan, K. S., & Lubicz-Zaorski, C. (2025a). Destructive polarization in digital communication contexts: A critical review and conceptual framework. *Information, Communication & Society*, *28*(8), 1499–1520. https://doi.org/10.1080/1369118X.2024.2413127.

Esau, K., Meyer, H., Farjam, M., Bruns, A., Rauxloh, H., & Brüggemann, M. (2025b). Polarised media framing of climate movements: A comparative mixed-methods analysis of Australia and Germany. *Media International Australia*. https://doi.org/10.1177/1329878X251377504.

Esau, K. (2025). The quality of connections: Deliberative reciprocity and inclusive listening as antidote to destructive polarization online. *Social Media + Society*, *11*(2). https://doi.org/10.1177/20563051251332421.

Fiorina, M. P. (2017). *Unstable Majorities: Polarization, Party Sorting, and Political Stalemate*. Stanford: Hoover Press.

Fiorina, M. P., & Abrams, S. J. (2008). Political polarization in the American public. *Annual Review of Political Science*, *11*(1), 563–588. https://doi.org/10.1146/annurev.polisci.11.053106.153836.

Fredericks, B., & Bradfield, A. (2021). More than a thought bubble. . .: The Uluru Statement from the Heart and Indigenous Voice to Parliament. *M/C Journal*, *24*(1). https://doi.org/10.5204/mcj.2738.

Gherardi, S. (2019). *How to Conduct a Practice-Based Study: Problems and Methods* (2nd ed.). Cheltenham: Edward Elgar Publishing.

Hellsten, I., & Leydesdorff, L. (2020). Automated analysis of actor–topic networks on Twitter: New approaches to the analysis of socio-semantic networks. *Journal of the Association for Information Science and Technology*, *71*(1), 3–15. https://doi.org/10.1002/asi.24207.

Hetherington, M. J. (2009). Putting polarization in perspective. *British Journal of Political Science*, *39*(2), 413–448. https://doi.org/10.1017/S0007123408000501.

Hillman, T., Seredko, A., Osborne, T. K., & Nivala, M. (2023). From answering for points to commenting for others. *Proceedings of the Tenth ACM Conference on Learning @ Scale*, 92–98. https://doi.org/10.1145/3573051.3593385.

Howison, J., Wiggins, A., & Crowston, K. (2011). Validity issues in the use of social network analysis with digital trace data. *Journal of the Association for Information Systems*, *12*(12), 767–797. https://doi.org/10.17705/1jais.00282

Iyengar, S., Sood, G., & Lelkes, Y. (2012). Affect, not ideology: A social identity perspective on polarization. *Public Opinion Quarterly*, *76*(3), 405–431. https://doi.org/10.1093/poq/nfs038.

Jacomy, M., Venturini, T., Heymann, S., & Bastian, M. (2014). ForceAtlas2, a continuous graph layout algorithm for handy network visualization designed for the Gephi software. *PLoS One*, *9*(6), e98679. https://doi.org/10.1371/journal.pone.0098679.

https://doi.org/10.1017/nws.2026.10040 Published online by Cambridge University Press

Jacomy, M., Ficozzi, M., & Munk, A. K. (2024). Sub-optimal recall in visual cluster retrieval: When clusters look like bridges. In *CEUR Workshop Proceedings*, *3834*, 1075–1085.

Johannessen, D. (Director). (2025). Exhausted voices: Past, present, future [Documentary].

Kasianenko, K. (2026). Elements of practices in digital publics: A model for identifying collective doings on social media platforms. *Communication Theory*, *36*(1). https://doi.org/10.1093/ct/qtaf026.

Kekkonen, A., & Ylä-Anttila, T. (2021). Affective blocs: Understanding affective polarization in multiparty systems. *Electoral Studies*, *72*, 102367. https://doi.org/10.1016/j.electstud.2021.102367.

Kligler-Vilenchik, N., Baden, C., & Yarchi, M. (2020). Interpretative polarization across platforms: How political disagreement develops over time on Facebook, Twitter, and WhatsApp. *Social Media + Society*, *6*(3). https://doi.org/10.1177/2056305120944393.

Kreiss, D., & McGregor, S. C. (2023). A review and provocation: On polarization and platforms. *New Media & Society*, *26*(1), 556–579. https://doi.org/10.1177/14614448231161880.

Kubin, E., & von Sikorski, C. (2021). The role of (social) media in political polarization: A systematic review. *Annals of the International Communication Association*, *45*(3), 188–206. https://doi.org/10.1080/23808985.2021.1976070v.

Kurucz, M., Benczúr, A., & Pereszlényi, A. (2008). Large-scale Principal Component Analysis on LiveJournal friends network. In *Proceedings of SNAKDD, 2008* (pp. 74–77). Association for Computing Machinery.

Lave, J., & Wenger, E. (1991). *Situated Learning: Legitimate Peripheral Participation*. Cambridge: Cambridge University Press.

Laurer, M., van Atteveldt, W., Casas, A., & Welbers, K. (2023). Building efficient universal classifiers with natural language inference. arXiv: 2312.17543 [cs.CL]. https://doi.org/10.48550/arXiv.2312.17543.

Leifeld, P., & Brandenberger, L. (2019). Endogenous coalition formation in policy debates. arXiv: 1904.05327 [CS.Stat]. https://arxiv.org/abs/1904.05327.

Lelkes, Y. (2016). Mass polarization: Manifestations and measurements. *Public Opinion Quarterly*, *80*(S1), 392–410. https://doi.org/10.1093/poq/nfw005.

Lerman, K., Feldman, D., He, Z., & Rao, A. (2024). Affective polarization and dynamics of information spread in online networks. *npj Complexity*, *1*(8). https://doi.org/10.1038/s44260-024-00008-w.

Levendusky, M. (2009). *The Partisan Sort: How Liberals Became Democrats and Conservatives Became Republicans*. Chicago: University of Chicago Press.

Mason, L. (2015). "I disrespectfully agree": The differential effects of partisan sorting on social and issue polarization. *American Journal of Political Science*, *59*(1), 128–145. https://doi.org/10.1111/ajps.12089.

McAllister, I., & Biddle, N. (2024). Safety or change? The 2023 Australian Voice referendum. *Australian Journal of Political Science*, *59*(2), 141–160. https://doi.org/10.1080/10361146.2024.2351018.

McCoy, J., & Somer, M. (2019). Toward a theory of pernicious polarization and how it harms democracies: Comparative evidence and possible remedies. *The Annals of the American Academy of Political and Social Science*, *681*(1), 234–271. https://doi.org/10.1177/0002716218818782.

McCoy, J., Rahman, T., & Somer, M. (2018). Polarization and the global crisis of democracy: Common patterns, dynamics, and pernicious consequences for democratic polities. *American Behavioral Scientist*, *62*(1), 16–42. https://doi.org/10.1177/0002764218759576.

McInnes, L., Healy, J., Sail, N., & Großberger, L. (2018). UMAP: Uniform manifold approximation and projection. *Journal of Open Source Software*, *3*(29), 861. https://doi.org/10.21105/joss.00861.

Mendes, K., Hollingshead, W., Nau, C., Zhang, J., & Quan-Haase, A. (2023). The evolution of #MeToo: A comparative analysis of vernacular practices over time and across languages. *Social Media + Society*, *9*(3). https://doi.org/10.1177/20563051231196692.

Münch, F. V., Thies, B., Puschmann, C., & Bruns, A. (2021). Walking through Twitter: Sampling a language-based follow network of influential Twitter accounts. *Social Media + Society*, *7*(1). https://doi.org/10.1177/2056305120984475.

Nicolini, D. (2012). *Practice Theory, Work, and Organization: An Introduction*. Oxford: Oxford University Press.

Pantzar, M., & Shove, E. (2010). Understanding innovation in practice: A discussion of the production and reproduction of nordic walking. *Technology Analysis & Strategic Management*, *22*(4), 447–461. https://doi.org/10.1080/09537321003714402.

Park, S., Fisher, C., McGuinness, K., Lee, J. Y., McCallum, K., & Cai, X. (2024). *Digital News Report: Australia*. Canberra: News and Media Research Centre. https://apo.org.au/node/326816.

Pattel-Gray, A. (1998). *The Great White Flood: Racism in Australia; Critically Appraised from an Aboriginal Historico-Theological Viewpoint*. Atlanta: Scholars Press.

Pereira, M. A., Rocha, B. A., & Lourenço, D. R. (2025). ¿Dónde está el protagonismo político indígena? Análisis de disputas discursivas sobre causas indígenas brasileñas en twitter. *América Latina Hoy*, *95*, e31875–e31875. https://doi.org/10.14201/alh.31875.

Radicioni, T., Saracco, F., Pavan, E., & Squartini, T. (2021). Analysing Twitter semantic networks: The case of 2018 Italian elections. *Scientific Reports*, *11*(1). https://doi.org/10.1038/s41598-021-92337-2.

Trillò, T., Hallinan, B., & Shifman, L. (2022). A typology of social media rituals. *Journal of Computer-Mediated Communication*, *27*(4), zmac011. https://doi.org/10.1093/jcmc/zmac011.

https://doi.org/10.1017/nws.2026.10040 Published online by Cambridge University Press

Vilenchik, D., Yichye, B., & Abutbul, M. (2019). To interpret or not to interpret PCA? This is our question. *Proceedings of the International AAAI Conference on Web and Social Media*, *13*(1), 655–658. https://doi.org/10.1609/icwsm.v13i01.3265.

Warde, A. (2016). *The Practice of Eating*. Polity Press.

Williams, H. T. P., McMurray, J. R., Kurz, T., & Lambert, F. H. (2015). Network analysis reveals open forums and echo chambers in social media discussions of climate change. *Global Environmental Change*, *32*, 126–138. https://doi.org/10.1016/j.gloenvcha.2015.03.006.

Yarchi, M., Baden, C., & Kligler-Vilenchik, N. (2021). Political polarization on the digital sphere: A cross-platform, over-time analysis of interactional, positional, and affective polarization on social media. *Political Communication*, *38*(1–2), 98–139. https://doi.org/10.1080/10584609.2020.1785067.

Cite this article: Bruns A., Esau K., Kasianenko K., Dehghan E., Dos Santos Choucair T., Padinjaredath Suresh V. and Johannessen D. (2026). Beyond interaction networks: Practice mapping as a step towards the detection of destructive polarization in social networks. *Network Science* **14**, 1–25. https://doi.org/10.1017/nws.2026.10040

https://doi.org/10.1017/nws.2026.10040 Published online by Cambridge University Press