Fine-Grained Clustering of Social Media: How Moral Triggers Drive Preferences and Consensus

Emanuele Brugnoli[1,2][a], Pietro Gravino[3][b], Donald Ruggiero Lo Sardo[1,2][c],
Vittorio Loreto[1,2,4][d] and Giulio Prevedello[3][e]

[1] *Sony Computer Science Laboratories Rome, Joint Initiative CREF-SONY, Centro Ricerche Enrico Fermi, Rome, Italy*
[2] *Centro Ricerche Enrico Fermi, Rome, Italy*
[3] *Sony Computer Science Laboratories Paris, Paris, France*
[4] *Physics Department, Sapienza University of Rome, Rome, Italy*

Keywords: Moral Foundation Theory, LLMs, Social Media Analysis, Political Polarization, Clustering Comparison.

Abstract: The increased access to online information provided by social media platforms allows individuals to form and convey their beliefs regarding events in their daily lives. The wide range of interactions carried out in these virtual environments has the power to impact the decisions and behaviours of others, also creating conflict, polarisation, misinformation, and toxic content. When individuals engage in public debates about topics tied to significant societal concerns, these discussions often regard or imply moral values. By analyzing five years of Italian Twitter/X debate on immigration, we show how a language model aware of moral values detects community structures more accurately, better depicting the actual political scenario in Italy.

1 INTRODUCTION

Social media platforms are pivotal in the modern information landscape, acting as robust channels for the swift dissemination of a vast array of information to a global audience. This rapid and extensive online distribution of content is crucial in shaping public opinion on various issues. On these platforms, users play dual roles: they are not only recipients of information but also active contributors, constantly shaping and reshaping the online narrative. This process apparently democratizes information sharing, enabling diverse voices to participate in the collective dialogue, profoundly influencing the formation and propagation of opinions. However, this diversity of interactions on social media can sometimes result in the creation of echo chambers (Cinelli et al., 2021). In such environments, individuals primarily encounter opinions and information aligning with their beliefs (Brugnoli et al., 2021). This phenomenon can deepen ideological divides as users become increasingly entrenched in their viewpoints, potentially intensifying disagreements and societal divisions. While the role of Information Technologies in echo chamber dynamics is still not fully understood (Gravino et al., 2019; De Marzo et al., 2023), this aspect of social media interaction is particularly evident in discussions surrounding social and moral issues.

Building on the principles of the Moral Foundation Theory (Graham et al., 2013), it is evident that moral beliefs, which are not uniform but are based on a diverse range of "irreducible basic elements" described by five moral dyads: care/harm, fairness/cheating, loyalty/betrayal, authority/subversion, and purity/degradation, play a significant role in these discussions. Social media platforms thus emerge as critical arenas for the debate, challenge, and reshaping of societal norms and values. The relative anonymity and lack of face-to-face interaction on these platforms can sometimes lead to more polarized and extreme expressions of moral values, escalating discussions into conflicts.

Despite the recognition of these dynamics, there remains a gap in understanding how the expression of moral values on social media influences group formation and affects user reactions. Our current research addresses this gap by examining the complex inter-

[a] https://orcid.org/0000-0002-5342-3184
[b] https://orcid.org/0000-0002-0937-8830
[c] https://orcid.org/0000-0003-3102-6505
[d] https://orcid.org/0000-0002-2506-2289
[e] https://orcid.org/0000-0002-9857-2351

Brugnoli, E., Gravino, P., Lo Sardo, D., Loreto, V. and Prevedello, G.
Fine-Grained Clustering of Social Media: How Moral Triggers Drive Preferences and Consensus.
DOI: 10.5220/0012595000003636

In *Proceedings of the 16th International Conference on Agents and Artificial Intelligence (ICAART 2024) - Volume 3*, pages 1405-1412
ISBN: 978-989-758-680-4; ISSN: 2184-433X

play of moral beliefs in the Italian Twitter/X discourse on immigration. In this study, we integrate moral domain analysis into social network analysis, revealing nuanced patterns of user interactions and group dynamics related to moral beliefs.

Our methodological approach employs advanced computational techniques, including a BERT-based model specifically fine-tuned to classify Italian tweets according to the expressed moral dyads (Brugnoli et al., 2023b). This allows us to effectively map tweet content onto a framework of moral beliefs, providing insights into how these beliefs influence user preferences and lead to the formation of consensus within online communities. Through this innovative approach, our work contributes significantly to understanding social media dynamics, particularly in the context of morally charged debates.

In our comprehensive analysis of the Italian immigration debate on social media, we employed a novel approach to delineate the intricate relationships between leaders and followers within the network. By "leaders" we mean the main influential actors of the social debate: news outlets or political entities (Brugnoli et al., 2023a), while the "followers" represent the rest of the community of users. By focusing on retweets (RTs) as a primary measure of endorsement, we identified followers' engagement with leaders' content, thus revealing the patterns of influence and information dissemination in the network.

We constructed a bipartite network, segregated into leaders and followers, connected by retweets. Analyzing this configuration through a monopartite projection on the leader layer, we utilized an optimized Louvain algorithm (Blondel et al., 2008) to discern communities based on political leanings. This analysis was further extended by labelling these communities and mapping their influence on the followers, thus inferring followers' political orientations based on their retweet behaviour.

Our study deepened by examining two distinct datasets: one encompassing all tweets from leaders and the other exclusively containing tweets that expressed moral values. This dual approach unveiled critical insights, particularly in distinguishing political groups and ideologies when moral content was considered. Notably, a more nuanced political spectrum emerged from the moral tweets dataset, highlighting the significance of moral values in discerning political affiliations.

Additionally, we observed a consistent clustering of questionable sources among the news outlets within specific communities, indicating potential biases or alignments with certain political ideologies. We conducted a rigorous comparison against a random benchmark to validate our findings, assessing the significance of the overlap between community structures derived from both datasets.

Crucially, we quantified the moral orientations of individuals within these communities using a novel method. By mapping each node onto a 4-dimensional probability simplex, we assigned a "moral vector" to every individual, indicating their likelihood of engaging with content related to specific moral dyads. This mapping, normalized using z-scores and softmax functions, revealed distinct moral configurations within communities, showcasing varied ethical underpinnings across the political spectrum.

2 RESULTS AND DISCUSSION

2.1 Unveiling Real and Nuanced Ideological Divisions with a Moral Focus

Distinguishing accounts between leaders and followers in social networks, particularly in the context of discussions surrounding immigration, offers a valuable framework for understanding how public opinions and attitudes are shaped and how consensus is reached within digital communities. Leaders in social networks play a crucial role in influencing the beliefs and attitudes of others, thus driving the collective consensus on various topics (Dong et al., 2017).

In our study, we focus on identifying these leaders within the context of the immigration debate in Italy over five years from 2018 to 2022. To accomplish this, we analyze a comprehensive set of entities comprising news media outlets and political figures actively participating in this discourse. These news media outlets are further classified based on assessments by independent third-party organizations, which categorize them as either questionable or reliable sources. This categorization is pivotal as it helps in understanding the quality and potential bias of the information disseminated by these outlets, particularly in terms of their reputation for spreading misinformation.

To capture the discourse around immigration, we collected tweets containing specific immigration-related keywords (see Section 3). This keyword-based approach ensures that our dataset is focused and relevant to the topic of interest.

Followers in the immigration debate are then identified through their engagement with the content produced by these leaders. Specifically, we consider retweets as a key metric for this identification, being retweeting unanimously regarded as an endorsement of the content created by others (Becatti et al., 2019). By analyzing retweet patterns, we can discern the followers in the network, and understand how they interact with and propagate the leaders' messages. This interaction is crucial, as it reflects the extent of influence leaders have in shaping the discourse and how followers contribute to the dissemination and reinforcement of these narratives.

With this node configuration, we built a bipartite network whose layers are leaders and followers and whose links represent retweets. If a particular group of followers retweets two different leaders, it suggests that these leaders are likely conveying similar messages or viewpoints. To further analyze these relationships, we employed a monopartite projection on the leader layer (for further information on the construction of the monopartite network refer to Section 3.3). This projection simplifies the network by focusing only on the leaders and the connections inferred through their shared followers. For the analysis of this projected network, we utilized an optimized version of the Louvain algorithm (Blondel et al., 2008). Namely, the nodes' order undergone a random shuffle 1,000 times, and the configuration with the highest modularity value was selected. Once these communities were identified, we assigned labels to each community based on the political leanings of the leaders within them (for a detailed list of leaders and their identified community membership see https://github.com/Sony-CSL-Rome/Italian_information_leaders). The next step in our analysis involved propagating these community labels to the followers in the retweet network (Raghavan et al., 2007). This process allows us to infer the political leanings of the followers based on the leaders they retweet. We assume that followers who predominantly retweet leaders from a particular community are likely to share similar political leanings.

To deepen our analysis, we repeated this labelling process twice with two different datasets. The first dataset included all tweets collected from the leaders, providing a general overview of the network dynamics. The second dataset was more focused, including only those tweets that expressed values aligned with at least one of the moral dyads, as identified by our BERT fine-tuned model. This allowed us to understand not only the political leanings but also the moral underpinnings of the interactions within the network.

Panels (A) and (B) of Figure 1 show a representation of the monopartite RT network of leaders aggregated in terms of communities identified. In these visualizations, the node positions remain constant, providing a consistent framework for comparison. The community labels, assigned *a posteriori*, reflect the political leanings of the leaders within each community. This approach allows us to visually and analytically discern the political landscape within the debate.

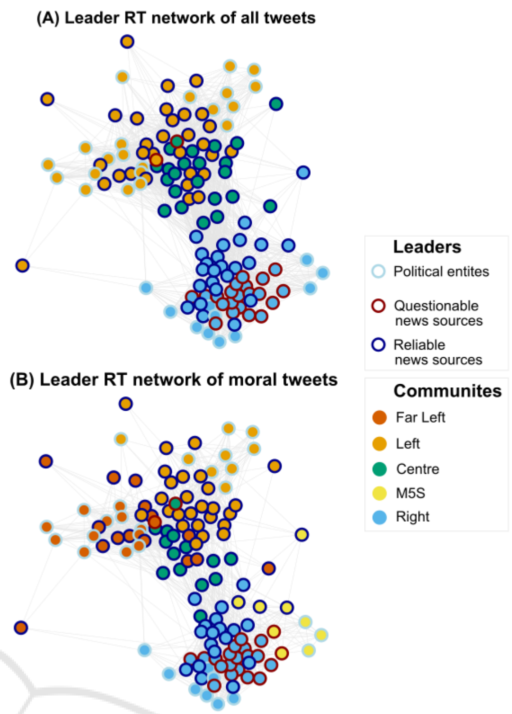

Figure 1: **(a)** Monopartite RT network of leaders obtained considering all tweets they produced. **(B)** Monopartite RT network of leaders obtained considering only moral tweets. Node positions are preserved. Node colours refer to communities. Nodes frame colours refer to the different types of leaders: political entities (azure), questionable news sources (dark red), and reliable news sources (dark blue).

The comparison between the two representations – one leveraging all retweets in the followers-leaders interaction network and the other limiting retweets to those expressing moral values – reveals critical insights into the structure of the debate. When all retweets are considered, the network appears to reflect a two-party system. This system is accompanied by a minority of news sources that do not align with any political faction, indicating a certain level of unbiased or neutral reporting within the media landscape.

However, when the analysis is refined to include only retweets that express moral values, a more nuanced and realistic separation of political entities emerges. This approach successfully identify the disunity of the Italian left, distinguishing the Far Left from the Left, as well as the Five Stars Movement (M5S) from the right-wing bloc (simply called Right), offering a clearer understanding of the complex political spectrum within the debate. It underscores the significance of moral values in discerning finer distinctions between political groups and ideologies.

An interesting observation from both network configurations is the placement of questionable sources. In both cases, these sources tend to cluster within the same community. This consistency suggests a potential alignment or affinity of questionable sources with specific political leanings or ideologies within the debate.

Then, we exploit the contingency table (Brier, 1980) associated with the the two network representations to compute the community overlap. This calculation is designed to demonstrate how the community structure emerging from the moral tweet analysis represents a more detailed, fine-grained version of the broader community structure obtained by analyzing all tweets.

In Figure 2, we present a visualization that illustrates this relationship. This panel specifically shows the maximum percentage of nodes from each community in representation (B) of Figure 1 - which is based on moral tweets - that correspond to a single community in representation (A) of the same figure, derived from all tweets. This comparison is crucial as it reveals the degree of alignment or divergence between the two community structures, providing insights into how the inclusion of moral content refines our understanding of the network's dynamics.

To ensure the robustness of our findings, we benchmark these results against a random model. This is achieved by shuffling the order of the nodes in representation (A) $10,000$ times, thereby generating a wide range of random community structures. For each shuffled version, we calculate the community overlap with the structure of representation (B). By comparing our observed overlaps to this confidence interval derived from the random benchmark, we can assess the significance of our findings.

This comparison against a randomized benchmark is essential for two main reasons. First, it allows us to determine whether the observed overlap between the two community structures is statistically significant or merely a product of random chance. Second, it provides a baseline to understand the extent to which focusing on moral tweets enhances the resolution of community detection in social media networks. If the observed overlap significantly exceeds the random benchmark, it underscores the value of including moral dimensions in network analyses, particularly in understanding complex social and political discussions.

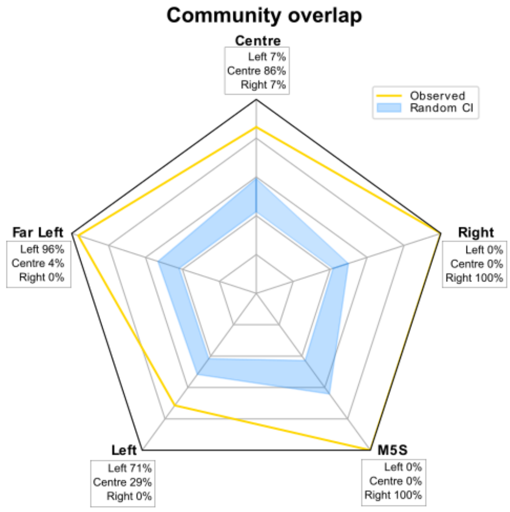

Figure 2: Community overlap between the Leader RT networks of all tweets (A) and moral tweets (B), respectively (see Figure 1). Radar shows the maximum percentage of nodes of a community of (B) that fall in exactly one community of (A). Results are also compared with the confidence interval of a random benchmark ($10,000$ reshuffling).

2.2 Diverse Ideologies Align with Varied Moral Configurations

To characterize the moral configuration of the emerged communities, we developed a method to quantify the moral orientation of both leaders and followers. This method involves mapping each individual in the network onto a 4-dimensional probability simplex, a mathematical space used to represent probabilities of different outcomes. In our context, the vertices of the simplex (the 5 standard unit vectors in $\mathbb{R}^5$) correspond to the five moral dyads as outlined in the Moral Foundations Theory. These dyads are care/harm, fairness/cheating, loyalty/betrayal, authority/subversion, and purity/degradation. Namely, we assigned to each individual in the retweet network its "moral vector", whose i-th component represents the probability to tweet (if the node is a leader) or retweet (if the node is a follower) content related to the i-th moral dyad. For example, if a leader frequently tweets about issues related to fairness/cheating, this moral dyad will have a higher probability in their moral vector compared to the other dyads. To ensure that our analysis of these moral vectors is consistent and comparable across all components, we further process the data using a statistical normalization technique. Specifically, we convert the raw probabilities into z-scores. This transformation allows us to compare how strongly an individual aligns with a particular moral dyad relative to the average alignment within the network. Finally we normalize using softmax in order to

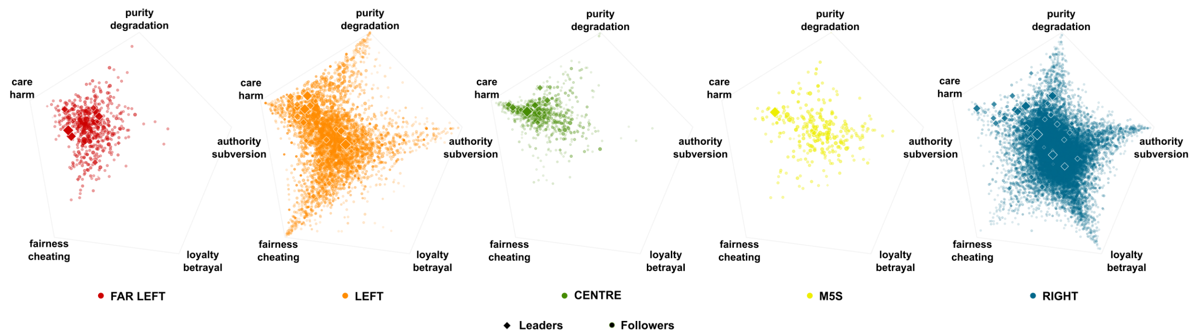

Figure 3: Moral configuration of both leaders (diamonds) and followers (circles) of the retweet network, divided by community membership. The position of a node represents how the propensity to share moral content is distributed between the five moral dyads. The node size indicates the corresponding activity regarding tweets if the node is a leader or retweets if the node is a follower.

better appreciate the behaviours in which moral values are over-represented.

Figure 3 clearly shows how nodes within the same community share similar moral beliefs and how different communities exhibit very different moral configurations. Namely, the Right community is almost exclusively represented in the region to the right of the diagonal identified by the vertices fairness/cheating and purity/degradation. Conversely, the Far Left, Left, and Centre communities, each following different moral patterns, mainly occupy the region to the left of this diagonal. In other words, care/harm seems to be not a pivotal dyad in the moral configuration of Right, as well as loyalty/betrayal and authority/subversion for Far Left, Left, and Centre. Instead, nodes belonging to the M5S community are mainly distributed along the diagonal authority/subversion - care/harm. Beyond visual inspection we assess how significantly the distributions differ through a MANOVA test (Stevens, 2012). The results show the differences are extremely significant, p-value $\ll 0.05$ (see Appendix for further information on the results of the MANOVA).

3 METHODS

3.1 Data Collection

Our study focuses on the social discussion around immigration in Italy, particularly its representation on Twitter/X. To capture a comprehensive view of the online immigration debate, we combined lists from external organizations, encompassing a wide range of news media and political groups active in Italy from 2018 to 2022. We then utilized the Twitter/X API to search for tweets from these sources, filtering for content containing specific keywords (Poletto et al., 2017), i.e. *immigrat* (immigrant*), immigrazion* (immigration*), migrant*, stranier* (foreigner*), profug* (refugee*), ong (ngo)*. These terms were chosen for their neutrality, avoiding bias towards any particular stance on immigration. We also gathered data on the retweeters of these tweets, focusing on tweets with significant interaction. Additionally, we used a binary classification from prior studies on disinformation spread (Gravino et al., 2022; Pennycook and Rand, 2019) to determine the credibility of Twitter/X accounts, categorizing them as either questionable or reliable sources. Unlike news outlets, whose reliability can be more commonly agreed upon in academic literature, political entities do not have a widely accepted measure of reliability. Therefore, we did not apply the reliability label to political entities in our research. Table 1 shows a breakdown of the dataset, which is the same used in (Brugnoli et al., 2023b).

Table 1: Breakdown of the Twitter dataset.

Category	Accounts	Tweets	Retweets
Questionable news outlets	76 (14.7%)	23,033 (13.1%)	345,624 (20.5%)
Reliable news outlets	403 (78.1%)	130,398 (74.1%)	362,595 (21.5%)
Political entities	37 (7.2%)	22,507 (12.8%)	976,033 (58.0%)
Total	516 (100.0%)	175,938 (100.0%)	1,684,252 (100.0%)

3.2 Modelling Morality

To explore the influence of moral beliefs on content creation and user engagement on Twitter, we fine-tuned a BERT-based model (Devlin et al., 2018) to classify, limited to the topic immigration, tweets in Italian according to the moral dyad expressed, as defined by the Moral Foundation Theory (Graham et al., 2013). Of note, moral values are considered as dyads of opposing poles (e.g., care/harm), instead of considering the poles as two separate labels (e.g., care and harm). Then, every post was annotated by which one of the five dyads was mostly expressed, or if no dyad was present. The fine-tuned model is the same used in (Brugnoli et al., 2023b) to which the readers can refer for further details about the training set (Stranisci et al., 2021) and the performance scores.

3.3 Leader Networks

By distinguishing between the selected Twitter/X accounts (leaders) and the general audience who retweeted their content (followers), we naturally define a biadjacency matrix A whose entries a_{fl} indicate the number of times the follower f retweeted the leader l (Holme et al., 2003). Then, to make the connections between different leaders explicit, the bipartite network of followers and leaders, which is completely defined by its biadjacency matrix A, is projected on the corresponding layer (Saracco et al., 2017). In other words, by exploiting the relations established between followers and leaders when the former retweet the latter, we can obtain an adjacency matrix for each layer. The adjacency matrix for the leaders tells us how similar are the population of followers between each couple of leaders. This operation is straightforwardly implemented through the matrix product $A^L = {}^tA \cdot A$. If we indicate with L the total number of users and with F the total number of accounts, the dimensions of A and its transpose are $F \times L$ and $L \times F$, respectively. This implies that A^L results in a symmetric $L \times L$ matrix whose generic element $a^L_{ll'}$, with $l \neq l'$, represents the strength of the link between the leaders l and l'. In addition, to provide a fair representation of all the leaders, thus reducing popularity bias and size effects, we consider the following normalization procedure. Let T be the vector of total retweets per page, namely the column sums of A^L. We set $T^L = (T^t \cdot T)^{\circ -1}$, where $X^{\circ -1}$ denotes the Hadamard inverse of X (Horn and Johnson, 2012). Then, we consider $\mathcal{A}^L = T^L \odot A^L$ as the normalized adjacency matrix of the corresponding leader network, where $\odot$ denotes the Hadamard product. The diagonal of $\mathcal{A}^L$ is set to zero in order to discard loops.

3.4 Clustering Comparison

A clustering $\mathcal{C}$ refers to a way of dividing a set of data points D into pairwise disjoint non-empty subsets $C_1, C_2, \ldots, C_I$ called clusters. Namely,

$$\mathcal{C} = \{C_1, C_2, \ldots, C_I\} \text{ such that } \begin{cases} C_i \neq \emptyset, \\ C_i \cap C_j = \emptyset, \\ \bigcup_{i=1}^{I} C_i = D. \end{cases}$$

Let the cardinality of D and C_i be n and n_i, respectively. Let a second clustering of the same set D be $\mathcal{C}' = \{C'_1, C'_2, \ldots, C'_{I'}\}$, with $|C'_{i'}| = n'_{i'}$. It is straightforward to observe that

$$n = \sum_{i=1}^{I} n_i = \sum_{i'=1}^{I'} n'_{i'}.$$

Most methods for comparing clusterings can be explained through the use of a contingency table associated with the pair of clusterings $\mathcal{C}$ and $\mathcal{C}'$ (Meilă, 2007). This table is essentially a matrix N with dimensions $I \times I'$, where each element $n_{ii'}$ represents the count of data points belonging to the intersection of clusters C_i from $\mathcal{C}$ and $C'_{i'}$ from $\mathcal{C}'$. The cluster sizes in respective clusterings are the row and column totals of N, that is,

$$|C_i| = n_i = \sum_{i'=1}^{I'} n_{ii'} \quad \text{and} \quad |C_{i'}| = n'_{i'} = \sum_{i=1}^{I} n_{ii'}.$$

In simpler terms, the contingency table provides a systematic way to compare how data points are distributed across the clusters of the two different clusterings.

4 CONCLUSIONS

Our comprehensive analysis of the Italian immigration debate on social media has underscored the critical role of moral values in shaping the dynamics of digital discourse. Through a novel approach that combined Network Analysis with the Moral Foundations Theory and Natural Language Processing, we were able to dissect the intricate relationships and influence patterns between leaders (news outlets and political entities) and followers in this debate.

Our key findings reveal that the structure of the debate on social media is not only influenced by political affiliations but is also deeply rooted in moral values. By examining retweet patterns, we were able to map out a complex landscape where leaders and followers formed distinct communities based on shared political and moral orientations. The use of a bipartite network, coupled with a monopartite projection and the application of the optimized Louvain algorithm, allowed us to identify and label these communities effectively.

The distinction between the datasets - one encompassing all tweets and the other focused on tweets with moral content - brought to light the nuanced nature of the discourse. The more granular analysis of moral tweets led to a clearer differentiation of political entities and ideologies, highlighting the pivotal role of moral values in distinguishing between seemingly similar political groups.

Furthermore, our observation of the consistent clustering of questionable sources within specific communities sheds light on the potential biases in information dissemination and the echo chambers that can arise as a result. This finding emphasizes the need for critical examination of source credibility in social media discourse.

By mapping individuals onto a 4-dimensional probability simplex and assigning them moral vectors, we were able to characterize the moral landscape of the debate quantitatively. This approach illuminated the varied moral underpinnings within each community and revealed how different communities prioritize different moral dyads.

Our study's visualization, particularly in Figure 3, effectively illustrates these moral configurations, clearly representing how different communities align with specific moral values. This visual evidence reinforces the importance of considering the moral dimension in understanding social media interactions and community formation.

In conclusion, taking the moral domain into account seems to be crucial not only for inferring the community structure of social networks at a finer resolution, but also for understanding where preferences and the resulting consensus are rooted and differentiated. Our findings highlight the need for communication strategies that recognize and leverage the moral dimensions, particularly in politically and socially charged discussions. This approach holds the potential to mitigate polarization and the formation of echo chambers.

As we navigate the complexities of digital discourse in an increasingly interconnected world, our findings offer valuable insights for researchers, policymakers, and communicators alike. They emphasize the importance of a holistic approach to understanding social media dynamics, one that goes beyond political leanings and takes into account the underlying moral values that drive human interactions and consensus-building in the digital age.

ACKNOWLEDGEMENTS

This work has been supported by the Horizon Europe VALAWAI project (grant agreement number 101070930).

REFERENCES

Becatti, C., Caldarelli, G., Lambiotte, R., and Saracco, F. (2019). Extracting significant signal of news consumption from social networks: the case of Twitter in Italian political elections. *Palgrave Communications*, 5(91).

Blondel, V. D., Guillame, J.-L., Lambiotte, R., and Lefebvre, E. (2008). Fast unfolding of communities in large networks. *Journal of Statistical Mechanics: Theory and Experiment*, 10008(6).

Brier, S. S. (1980). Analysis of contingency tables under cluster sampling. *Biometrika*, 67(3):591–596.

Brugnoli, E., Cinelli, M., Quattrociocchi, W., and Scala, A. (2021). Recursive patterns in online echo chambers. *Scientific Reports*, 11(1):5909.

Brugnoli, E., Galletti, M., Lo Sardo, R., Prevedello, G., Di Canio, M., and Gravino, P. (2023a). Decoding political social media posts. *Nature Italy*.

Brugnoli, E., Gravino, P., and Prevedello, G. (2023b). Moral values in social media for disinformation and hate speech analysis. In *Proceedings of the VALE workshop at ECAI, (to appear)*.

Cinelli, M., De Francisci Morales, G., Galeazzi, A., Quattrociocchi, W., and Starnini, M. (2021). The echo chamber effect on social media. *Proceedings of the National Academy of Sciences*, 118(9):e2023301118.

De Marzo, G., Gravino, P., and Loreto, V. (2023). Recommender systems may enhance the discovery of novelties. *arXiv preprint arXiv:2312.08824*.

Devlin, J., Chang, M.-W., Lee, K., and Toutanova, K. (2018). Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*.

Dong, Y., Ding, Z., Martínez, L., and Herrera, F. (2017). Managing consensus based on leadership in opinion dynamics. *Information Sciences*, 397-398:187–205.

Graham, J., Haidt, J., Koleva, S., Motyl, M., Iyer, R., Wojcik, S. P., and Ditto, P. H. (2013). Moral foundations theory: The pragmatic validity of moral pluralism. In *Advances in experimental social psychology*, volume 47, pages 55–130. Elsevier.

Gravino, P., Monechi, B., and Loreto, V. (2019). Towards novelty-driven recommender systems. *Comptes Rendus Physique*, 20(4):371–379.

Gravino, P., Prevedello, G., Galletti, M., and Loreto, V. (2022). The supply and demand of news during covid-19 and assessment of questionable sources production. *Nature Human Behaviour*, 6(8):1069–1078.

Holme, P., Liljeros, F., Edling, C. R., and Kim, B. J. (2003). Network bipartivity. *Phys. Rev. E*, 68:056107.

Horn, R. A. and Johnson, C. R. (2012). *Matrix analysis*. Cambridge University Press.

Meilă, M. (2007). Comparing clusterings—an information based distance. *Journal of Multivariate Analysis*, 98(5):873–895.

Pennycook, G. and Rand, D. G. (2019). Fighting misinformation on social media using crowdsourced judgments of news source quality. *Proceedings of the National Academy of Sciences*, 116(7):2521–2526.

Poletto, F., Stranisci, M., Sanguinetti, M., Patti, V., and Bosco, C. (2017). Hate speech annotation: Analysis of an italian twitter corpus. In *Proceedings of the Fourth Italian Conference on Computational Linguistics CLiC-it 2017*, pages 263–268.

Raghavan, U. N., Albert, R., and Kumara, S. (2007). Near linear time algorithm to detect community structures in large-scale networks. *Phys. Rev. E*, 76:036106.

Saracco, F., Straka, M. J., Di Clemente, R., Gabrielli, A., Caldarelli, G., and Squartini, T. (2017). Inferring monopartite projections of bipartite networks: an entropy-based approach. *New Journal of Physics*, 19(053022).

Stevens, J. P. (2012). *Applied multivariate statistics for the social sciences*. Routledge.

Stranisci, M., De Leonardis, M., Bosco, C., and Patti, V. (2021). The expression of moral values in the twitter debate: a corpus of conversations. *IJCoL. Italian Journal of Computational Linguistics*, 7(7-1, 2):113–132.

APPENDIX

Table 2: Significance of moral value distribution difference in identified clusters. Summary statistics of the MANOVA regression.

Intercept	Value	Num DF	Den DF	F Value	Pr > F
Wilks' lambda	-0.00	5.00	19369.00	-1.76×10^{17}	1.00
Pillai's trace	1.00	5.00	19369.00	-1.76×10^{17}	1.00
Hotelling-Lawley trace	-4.55×10^{13}	5.00	19369.00	-1.76×10^{17}	1.00
Roy's greatest root	-4.55×10^{13}	5.00	19369.00	-1.76×10^{17}	1.00
Group	Value	Num DF	Den DF	F Value	Pr > F
Wilks' lambda	0.51	20.00	64240.66	732.29	0.00
Pillai's trace	0.54	20.00	77488.00	602.08	0.00
Hotelling-Lawley trace	0.89	20.00	42604.36	861.67	0.00
Roy's greatest root	0.78	5.00	19372.00	3026.45	0.00