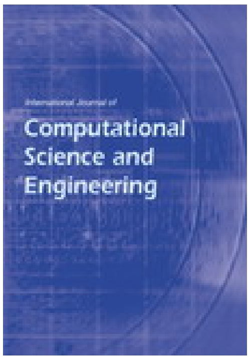

International Journal of Computational Science and Engineering

ISSN online: 1742-7193 - ISSN print: 1742-7185
https://www.inderscience.com/ijcse

Harmony amidst division: leveraging genetic algorithms to counteract polarisation in online platforms

Zaka Ul Mustafa, Muhammad Amir, Manal Mustafa, Muhammad Adnan Anwar

DOI: 10.1504/IJCSE.2025.10068868

Article History:

Received:	21 March 2024
Last revised:	30 July 2024
Accepted:	31 October 2024
Published online:	15 January 2025

Harmony amidst division: leveraging genetic algorithms to counteract polarisation in online platforms

Zaka Ul Mustafa* and Muhammad Amir

International Islamic University Islamabad,
Islamabad, Pakistan
Email: zaka.phdee105@iiu.edu.pk
Email: m.amir@iiu.edu.pk
*Corresponding author

Manal Mustafa

Zaman Technologies Pvt Limited, Pakistan
Email: manalmustafazia@gmail.com

Muhammad Adnan Anwar

Department of Mathematics (CEMAT),
IST – Instituto Superior Técnico,
Ulisboa, Portugal
Email: adnan.anwar@tecnico.ulisboa.pt

Abstract: While facilitating global communication, social media platforms often exacerbate ideological polarisation, contributing to echo chambers and filter bubbles that hinder constructive discourse. This paper presents a novel approach using genetic algorithms (GAs) to mitigate polarisation in online social networks. Unlike traditional methods that offer static, symptom-based solutions, the GA-driven model dynamically identifies and modifies influential nodes to reduce polarisation while preserving network integrity. Optimising interactions among key nodes enhances diversity and fosters more inclusive dialogue. The model is validated using real-world datasets, including Polbooks and Polblogs networks, demonstrating significant polarisation reduction with minimal disruption to the network structure. The approach advances the theoretical understanding of polarisation mitigation and provides practical applications for designing recommendation systems that promote healthier, more diverse online communities. The study emphasises the potential of interdisciplinary methodologies, integrating computational algorithms with social science insights to create a more balanced and engaging digital environment.

Keywords: echo chamber; filter bubble; digital social networks; social network analysis; polarisation; diversity; community discovery.

Reference to this paper should be made as follows: Ul Mustafa, Z., Amir, M., Mustafa, M. and Anwar, M.A. (2025) ‘Harmony amidst division: leveraging genetic algorithms to counteract polarisation in online platforms’, *Int. J. Computational Science and Engineering*, Vol. 28, No. 7, pp.1–17.

Biographical notes: Zaka Ul Mustafa obtained his BSc in Electrical Engineering from the University of Engineering and Technology (UET) Taxila, Pakistan in 2000. He also completed his Master of Engineering from the National University of Sciences and Technology (NUST) in 2004. He had been into academia since 2002. He taught in departments of Electrical Engineering at UET Taxila and NUST Pakistan. He is also an author of numerous research papers. His interest area wireless networks and social networks. Currently, he is pursuing his PhD in Social Networks from Islamic International University Islamabad, Pakistan. Besides academic activities he is also an entrepreneur.

Muhammad Amir received his BS in Electrical Engineering from the NWFP University of Engineering and Technology, Peshawar, Pakistan in 1997. He earned his PhD in Electronic Engineering from the International Islamic University, Islamabad, in 2013. Since 2014, he has been a Professor at the Department of Electrical Engineering, International Islamic University, Islamabad, Pakistan. He has served in various key academic and administrative roles. He has authored numerous publications and supervised multiple research projects, with his interests spanning electrical energy systems, advanced signal processing, and computational techniques. He is a reviewer for several reputable journals and a recipient of multiple research grants.

Manal Mustafa received her Bachelor's in Electrical Engineering from the Lahore University of Management Sciences, Lahore, Pakistan in 2020. She worked for an e-commerce company in Pakistan. Her research areas are social networks, data engineering and machine learning.

Muhammad Adnan Anwar received his PhD in Computational Mathematics from Lahore University of Management Sciences (LUMS). He is currently a post-doctoral researcher at the Center for Computational and Stochastic Mathematics (CEMAT), Instituto Superior Técnico, Portugal. He is actively involved in numerous industrial and academic research projects, driving innovation and advancing knowledge in his field.

1 Introduction

Social networking services (SNS) platforms are now a big part of our digital world. More people are using them than ever before. These platforms, which started as places to chat and share, now collect a lot of information about how we behave online. This information can be clear as well as hidden. For example, it can be clear similar to the 'likes' we give on social media. Whereas, it can be hidden like the patterns of our online chats (Ko et al., 2022).

This information is very important for today's online recommendation systems. These systems look at our online actions to suggest things we might like. Many of these systems use the hidden patterns to guess our likes and dislikes (Beheshti et al., 2020). When they use both clear and hidden information, they can understand us better (Abbasi-Moud et al., 2021; Liu and Lee, 2010; Yang et al., 2017; Amato et al., 2019; Capdevila, 2016; Tarus, 2017; Choi et al., 2012). But there's a problem. These systems often show us things that we already know or like. This can create 'echo chambers' where we only hear similar views. This can lead to more division on social media. The Cambridge Analytica issue (Kaufman, 2018) showed how this information can be misused. This research looks at ways to make online spaces more open and welcoming for everyone.

1.1 Research gap

Current approaches to mitigating polarisation often focus on detecting and measuring it rather than providing robust solutions. Existing methods, such as edge addition for exposure or opinion flipping, address symptoms rather than root causes, and their effectiveness varies across different network structures. Moreover, many methods fail to account for the dynamic and complex nature of real-world social networks which this research has done.

1.2 Background and impact

Polarisation in social media manifests as ideological divides, leading to decreased common ground and increased social distance between groups. This study aims to address this issue by leveraging genetic algorithms (GAs) to reduce polarisation efficiently and effectively. GAs are well-suited for this task due to their adaptability and robustness in navigating complex problem spaces. By identifying and influencing key nodes within the network, GAs can significantly mitigate polarisation while maintaining the network's structural integrity.

1.3 Novelty of the research

This research introduces a novel approach using GAs to reduce polarisation in social networks, addressing the root causes rather than just the symptoms. Unlike traditional methods like edge addition or opinion flipping, which are static and sometimes ethically contentious, our GA-based approach dynamically adapts to changes in network structure and user behaviour. It focuses on identifying and optimising influential nodes and connections, crucial for managing the complex dynamics of social media environments. GAs efficiently explore large networks to find optimal solutions, avoiding local optima and ensuring minimal intervention. The method is highly scalable, requiring only minor modifications to the network structure, thereby preserving its integrity while promoting diverse interactions. Its versatility allows application across various social media platforms and datasets, making it a robust tool for tackling polarisation in different contexts. This approach is more targeted and ethically sound, enhancing network diversity without disrupting natural discourse. Validation with real-world datasets demonstrates significant reductions in polarisation metrics, confirming the method's effectiveness and potential to foster more inclusive and diverse online communities.

The application of GAs to reduce polarisation in social networks has significant implications across various industries. In social media, it enhances user engagement and satisfaction by promoting diverse and inclusive environments, while also improving content moderation. E-commerce platforms can use this approach to offer more varied product recommendations and optimise ad targeting, broadening customer choices and campaign effectiveness. News platforms benefit by balancing diverse viewpoints, fostering a more informed public discourse. In corporate settings, this method encourages open dialogue and innovation by minimising group think. The key contributions include the introduction of a GA-based approach for dynamic polarisation reduction, a comprehensive fitness function for precise interventions, and minimal, ethical intervention. This approach is practical and scalable, offering benefits across social media, e-commerce, news, and corporate sectors. Future research should explore additional factors and larger datasets to validate and expand its effectiveness.

1.4 Key terms and definitions

In addressing the complex phenomenon of social media polarisation, this research focuses on following key concepts that are critical for understanding the dynamics at play:

1.4.1 Polarisation

In our study, we focus on 'ideological polarisation', which refers to the extent to which individuals or groups within a social network adopt increasingly divergent views, leading to a decrease in common ground and an increase in social distance between the groups. Ideological polarisation encompasses both attitude polarisation, which relates to the positions individuals or groups hold on specific issues, and affective polarisation, which relates to how these groups feel about each other. Our primary concern is with how these forms of polarisation manifest within social networks and can be mitigated through strategic interventions.

1.4.2 Diversity

We define diversity in the context of social networks as the range of different opinions, backgrounds, and perspectives present within a network. A high level of diversity indicates a broad spectrum of views and a healthy dialogue across different segments of the network. Our goal is to enhance this diversity, thereby enriching the network's resilience to polarisation and encouraging a more inclusive exchange of ideas.

1.4.3 Echo chambers

An echo chamber in the context of our research refers to a situation within a social network where information, ideas, or beliefs are amplified or reinforced by communication and repetition inside a defined system. In such systems, members are exposed primarily to viewpoints that mirror their own, leading to a decrease in exposure to and acceptance of differing viewpoints.

1.4.4 Filter bubbles

We use the term 'filter bubble' to describe the state of intellectual isolation that can occur when algorithms selectively guess what information a user would like to see based on the user's past behaviour, location, or other factors. This can limit a user's exposure to a wider range of perspectives, thereby reinforcing pre-existing biases and contributing to polarisation.

1.5 Problem of polarisation

Social media has changed how we talk to each other. It helps people from different places and cultures connect (Gillani et al., 2018). If we think of social media like electronic signals, we can see some patterns. Just like signals have different parts, social media has things like the number of links, levels of division, and how fast information spreads. By thinking this way, we can use ideas from electronics to solve social media problems. We want to find and fix problems in how information moves, especially problems caused by division. We don't just think about solutions; we also test them to make sure they work well.

1.6 Effects of polarisation

Social media can be good and bad. It can help share new ideas but can also trap people in places where they only hear things they agree with Donkers and Ziegler (2021). Many people believe that having different views is good. Groups with different views often do better than groups where everyone thinks the same (Sobkowicz, 2023). A study Garcia (2023) explores the impact of Facebook's news feed algorithms on political polarisation. It was found that these algorithms can filter partisan political news to users, thereby reinforcing existing beliefs and contributing to ideological divides. The research highlights the need for algorithmic transparency and modifications to reduce polarisation effects on social media platforms. Another recent study conducted by researchers from Princeton University (Tokita et al., 2021) investigates how users' self-curation of their social media feeds leads to the formation of polarised networks. By analysing Twitter data, the study found that users tend to follow and unfollow accounts based on perceived trustworthiness and alignment with their views, inadvertently creating echo chambers that amplify polarisation. This research underscores the natural emergence of polarised discourse even without explicit partisan identities being known. So, we want to make sure the online world is balanced. To do this, we need to understand and reduce the effects of division.

1.7 Solution space

Fighting division is important both online and in the real world. In real life, things like university exchange programs and cultural events help bring people together (Akçay, 2018). Online, recommendation systems can help, especially when they find patterns in what we do or show us new things (Sun, 2023; Kotkov et al., 2016b; Maksai et al., 2015; Zhang et al., 2012; Kotkov et al., 2016a, 2017). To find good solutions, we first need to understand why division happens.

1.8 Proposed solution

The challenge of polarisation in social networks is not just about recognising its presence but actively working towards its reduction. Building upon the established metric from previous research to quantify polarisation, this study introduces a novel approach to address the issue. Harnessing the power of GAs, a method has been devised to curtail polarisation effectively. GAs, inspired by the process of natural selection, are adept at finding solutions to optimisation and search problems. In the context of this research, they are employed to identify and modify key nodes and connections within the network, aiming to reduce polarisation.

Figure 1 A polarised network (see online version for colours)

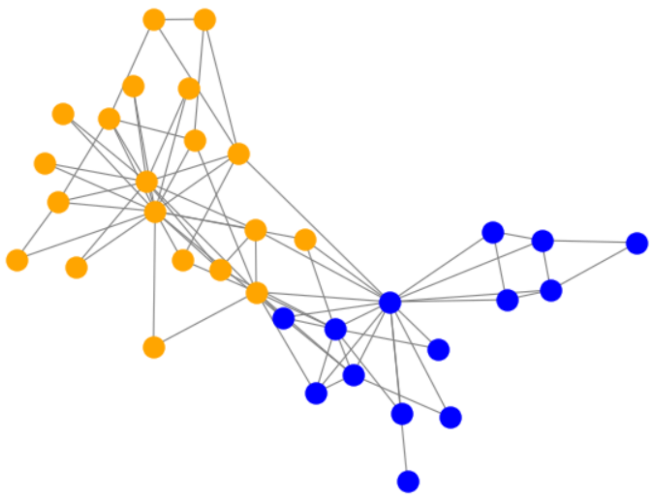

Figure 2 An unpolarised network (see online version for colours)

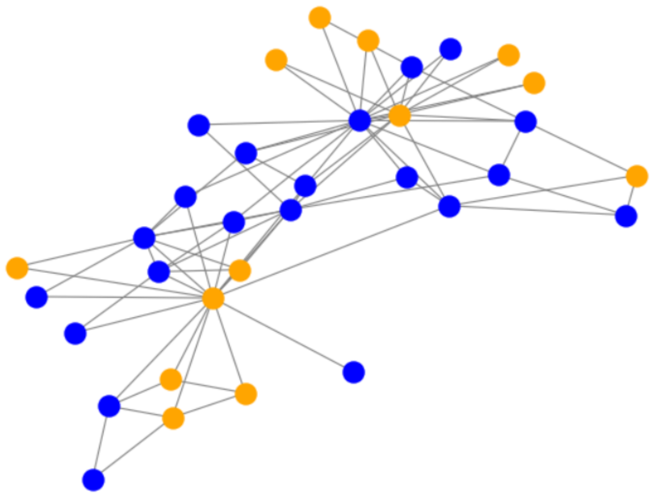

The primary advantage of this method is its efficiency. Instead of a brute-force approach that might involve extensive computations and modifications, the GA seeks to achieve the desired reduction in polarisation in the shortest possible time. This is crucial for real-world applications where timely interventions can prevent the further spread of polarisation.

The algorithm operates by identifying influential nodes in the network that contribute significantly to polarisation. By moderating the influence of these nodes or altering their connections, the overall polarisation of the network can be reduced. The iterative nature of the GA ensures that with each generation, the solution is refined, leading to an optimal strategy for polarisation reduction. In essence, the proposed solution offers a dynamic, efficient, and targeted approach to combat polarisation in social networks. By focusing on key influencers within the network and leveraging the optimisation capabilities of GAs, this method promises a significant reduction in polarisation, fostering a more inclusive and diverse online community. Figures 1 and 2 show what a polarised and non-polarised network looks like.

In Figure 1, a clear division is evident between the blue and green nodes, indicating a high degree of polarisation as they form distinct communities. However, upon observing Figure 2, strategic interventions have been applied to key nodes from both communities. This strategic approach aims to bridge the gap between the two polarised groups. The results stemming from this intervention strategy are notably impactful and underscore its potential in addressing polarisation.

1.9 Datasets

We use many sources to procure and generate specific datasets which will be used for this research. We explored the dataset collection method on various levels and attempted to use real-world dataset for validation of our claims. This extensive testing helped us evaluate results in different network situations and configurations to get rid of any potential biases. We have stated the limitation and scope of the datasets where required. For example, the number of nodes in the selected dataset versus the practical size of the social media platforms. After working on the detection and quantification of the polarisation problem, we finalised the following sources.

1 *Polbooks:* A book network related to US politics, available on 'amazon.com' (Krebs, 2004), is represented where individual books are depicted as nodes. Links between nodes indicate books that are often bought together. The categorisation of these books is as follows: liberal (43), conservative (49), and neutral (13). Books labelled as neutral are arbitrarily allocated to either the liberal or conservative group. A plot of the network is shown in Figure 3(a).

2 *Polblogs:* A network showcasing links between weblogs focused on US politics from 2005 is presented Adamic and Glance (2005). These blogs are identified as either Liberal or Conservative. We overlook the direction of the links and focus on the most significant connected segment. This analysis yields two groups with 636 and 586 blogs respectively. Details about these datasets can be found

in Table 1. All connections in the network are considered bidirectional, and every link is assigned a weight of 1. A plot of the network is shown in Figure 3(b).

Figure 3 Graphical representation of, (a) Polbooks dataset (b) Polblogs dataset (see online version for colours)

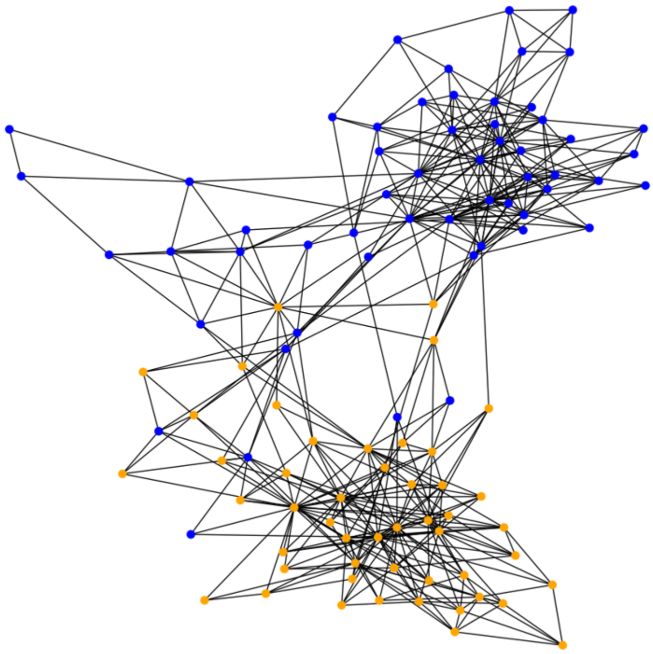

(a)

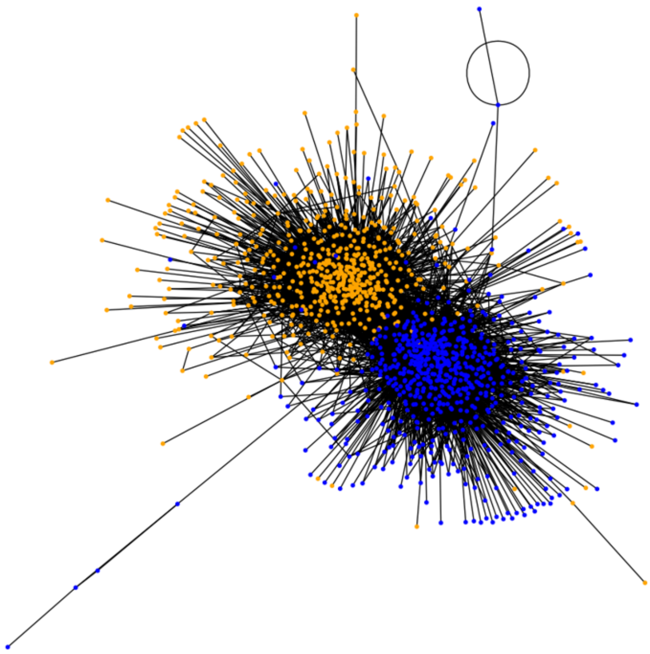

(b)

We want to point out here that we used data from different online platforms as our focus was on getting network data. The structural properties of online networks, such as degree distribution and community formation, are consistent across different types of platforms, supporting the relevance of our approach to social media. Since social media has the similar data structure where there is a network structure comprising of nodes and edges, this method can be easily applied to any social media. Here in our datasets, nodes represent books and blogs and edges determine association whereas in social media nodes represent users and edges determine friendships.

Table 1 Dataset summary statistics

Dataset name	*Number of nodes*	*Number of edges*
Polbooks	105	441
Polblogs	1,222	16,717

2 Literature review

In the evolving landscape of social media, understanding the dynamics of filter bubbles and resultant polarisation is crucial. While this research does not propose a new method to quantify polarisation, reviewing existing methods provides a foundation to understand the current state of the art and contextualises our approach to address the issue using GAs.

2.1 Diversity measures

In the work of Matakos and Gionis (2018) the analysis of social networks is approached by modelling the networks as graphs, where individuals are represented as nodes, and the connections between them are denoted as edges. Additionally, each individual is attributed a binary opinion value to signify their stance on a matter. The assumption here is that these opinions are binary, falling into one of two categories (i.e., $s_i \in -1, 1$), and that both the weight of the edges and the cost attributed to the nodes are uniform, set to 1. Their methodology for calculating diversity essentially involves tallying the edges that connect nodes of opposing opinions, representing two distinct communities. However, this model's reliance on binary opinions may not fully encapsulate the complexity of real-world views, which often exist on a spectrum ranging from mild to extreme, including neutral positions

The research by Vendeville et al. (2023) investigates echo chambers on Twitter during the 2017 French Elections. They focused on users' feeds, which are made up of tweets from accounts the users follow, known as 'leaders'. These tweets were classified into various political categories. The study calculated the proportion of tweets supporting a specific political party ('s') in a user's ('n') feed. This was done by assessing how frequently tweets about 's' appeared in 'n's feed. They also developed a score for each user that reflects the political diversity of their feed. This diversity metric is based on the mix of political content in a user's feed. A score of 0 indicates a complete echo chamber, where the feed is dominated by one political viewpoint, suggesting a lack of diversity. Conversely, a score of 1 signifies that the feed has an equal representation of all political parties, indicating a highly diverse and balanced range of political content.

2.2 Polarisation measures

Several methods have been proposed to detect and quantify polarisation.

Akoglu's (2014) research approaches the issue by treating it as a task of classifying nodes within edge-signed bipartite opinion networks. This method is tailored to networks that fit the bipartite graph model, which may not cover all scenarios.

Next work Guerra et al. (2013) we discuss, frames the issue of measuring a social network's polarisation level as a boundary challenge, dividing a graph into two communities, G1 and G2, each with its own boundary. It differentiates between 'boundary nodes', which connect to the opposite community, and 'internal nodes', which have no direct connections outside their community. Polarisation is assessed by comparing the number of edges boundary nodes have within their community versus with the opposite community. This approach suggests polarisation is higher when boundary nodes connect more within their community and lower when they connect more with the opposite community, offering a polarisation scale from -0.5 to 0.5. However, this method may not accurately reflect network structure nuances, limiting its applicability.

Another study Morales et al. (2015) introduces the 'centre of gravity' concept to gauge network polarisation, dividing nodes into 'elites', with fixed opinions, and 'listeners', whose opinions evolve to reflect the average viewpoint of their neighbours over time. Polarisation is measured by the distance between the average opinions of these two groups, with A^+ representing the fraction with positive opinions and A^- the fraction with negative opinions. This approach quantitatively assesses how opinions within a network diverge.

Hohmann et al. (2023) propose a nuanced approach to measuring network polarisation, employing a generalised Euclidean (GE) distance metric that accounts for network structure, reflecting the effort needed to bridge differing opinions within the network. This method not only measures the ideological divergence among individuals but also incorporates the network's structural aspects, such as community formations and their connectivity. By evaluating the 'resistance' between nodes, it offers a comprehensive view of polarisation, considering the interplay between opinion diversity and network organisation, thereby providing a more holistic understanding of polarisation dynamics.

Next, the Friedkin and Johnsen model Matakos et al. (2017) assesses polarisation by assigning a fixed internal opinion and a variable expressed opinion to each user, influenced by their own views and those of their network. The model calculates expressed opinions based on the likelihood of encountering similar viewpoints during a network exploration, with polarisation determined by the magnitude of expressed opinions. This approach provides insights into the degree of echo chambers within networks, using a mathematical formula to quantify polarisation based on the aggregation of individual opinions and their alignment within the social structure.

After this, Cinus et al. (2022) analyse how social media recommendation algorithms, like 'People You May Know', influence echo chambers and polarisation by fostering similar connections. They assess three algorithms against two opinion dynamics models to see how user opinions evolve with interactions. The study measures echo chamber effects and polarisation using metrics that consider opinion correlation among connected users and the likelihood of opinion diversity crossing through random walks. However, it notes limitations in fully capturing polarisation due to overlooking community size and node degree. Another research Vicario et al. (2016) investigates how users in echo chambers, with uniform opinions, interact and express emotions over time, utilising growth models to analyse community evolution. It tracks user engagement through comments, categorising users by interaction levels and sentiments, and identifies trends in emotional expression linked to community activity. Findings indicate both science and conspiracy communities exhibit similar growth, with sentiment analysis revealing a trend towards negativity as engagement increases, highlighting differences in sentiment polarisation between science-oriented and conspiracy-oriented users as their activity levels change.

2.3 Controversy measures

The study Garimella et al. (2015) introduces a method to measure network polarisation by evaluating edge significance through edge betweenness centrality, which reflects an edge's role in connecting node pairs via the shortest paths. It assesses polarisation by comparing betweenness centrality of edges within and between two ideologically opposed communities, using the Kullback-Leibler divergence. The resulting betweenness centrality controversy score, ranging from 0 to 1, indicates the network's polarisation degree, with higher scores signifying greater polarisation

In summary, while various methods exist to quantify polarisation, this research leverages these foundational understandings to focus on devising a solution using GAs. The goal is to effectively reduce polarisation, building upon the insights gained from these established methods.

2.4 Recommendation systems

Recommendation systems, by tailoring content based on user preferences, have revolutionised our online interactions. However, (Zhang and Hurley, 2008) highlight the inadvertent creation of 'filter bubbles' that can limit our exposure to diverse content. Recognising this potential pitfall, efforts have been made to diversify recommendation algorithms for a richer online experience (Gao et al., 2022). Yet, as another research points out (Castells et al., 2011), striking a balance between personalisation and diversity remains challenging. In the quest for solutions, an innovative approach is the 'serendipity-based' recommendations by Murakami et al. (2009). This method focuses on showing users the content that is a bit different from what they usually watch and like.

2.5 Genetic algorithms

Setting the stage for our exploration, recommendation systems, by tailoring content based on user preferences, have revolutionised our online interactions. However, there's a growing concern about their potential to inadvertently create 'filter bubbles' and 'echo chambers'. Addressing this concern, researchers have looked into various methods to diversify recommendations and reduce polarisation. One such promising approach is the use of GAs. Diving into the mechanics of GAs, inspired by the process of natural selection, are optimisation techniques that have found applications in diverse domains, including social networks. The essence of GAs lies in evolving a population of candidate solutions over generations through processes like selection, crossover, mutation, and reproduction (Holland, 1992). When applied to social networks, the primary objective is often influence maximisation: selecting a subset of influential nodes to either maximise the spread of information or minimise polarisation.

Building on this foundation, Bucur and Iacca (2016) delved into the challenging problem of influence maximisation using GAs. Their findings suggest that GAs, with simple genetic operators, can identify influential nodes comparable to known heuristics without making assumptions about the network's structure. Transitioning from this foundational work, Bucur et al. (2018) further explored surrogate-assisted multi-objective evolutionary algorithms for the same problem. Their approach, which utilised an approximate model of influence propagation, underscored the importance of precision in model selection.

Exploring cloud-based solutions, Chen et al. (2020) introduced a cloud computing-based solution for targeted influence maximisation. Their tag-aware IC model, which considers user characteristics and similarities, achieved significant improvements in speed and storage efficiency. Qin et al. (2022) took a different approach by integrating community and topic features into an IC model, resulting in enhanced stability and efficiency.

Shifting our focus to prioritised influence, Pham et al. (2020) introduced the influence maximisation with priority (IMP) problem, focusing on influencing potential users with priority during influence diffusion campaigns. Their algorithms, integrated greedy (IG) and integrated greedy sampling (IGS), provided efficient solutions with theoretical guarantees.

Building on the concept of influential nodes, Wang et al. (2021) proposed a method based on discrete moth-flame optimisation. Their approach, which considers the valuation of neighbour nodes, was found effective in real-world social networks. Sivaganesan's (2021) algorithm, on the other hand, leverages semantic metrics like user interests to identify influential nodes. Talukder et al. (2019) tackled the reverse influence maximisation problem, proposing a Knapsack-based solution for optimised seeding costs. Aghaee and Kianian's (2020) GIN algorithm reduced the search space for influential nodes, selecting seeds with the highest expected diffusion value. The challenge of scalability in identifying influential nodes becomes evident as social networks expand. Bucur and Iacca's (2016) work stands out as they demonstrated the efficiency of GAs in this context. Pal et al. (2014) introduced new centrality measures, diffusion degree and maximum influence degree, to pinpoint top influential individuals. Liu et al. (2017) and Song et al. (2017) proposed incremental approaches using GAs to track influential nodes in evolving networks.

In conclusion, GAs have shown significant potential in addressing the challenges posed by influence maximisation in social networks. As these networks continue to grow and evolve, the role of GAs in identifying and leveraging influential nodes becomes even more crucial. This research aims to further explore this potential, focusing on the application of GAs in influence maximisation and its implications in various real-world scenarios.

2.6 Methods of polarisation reduction

Reducing polarisation is essential because it directly affects how well our society works together and supports a healthy democracy (Keefer and Knack, 2002). When people are deeply divided, it is harder for us to find common ground and make decisions that benefit everyone. This is not just about politics; it is about making sure our community remains strong and united. Our study looks into this issue not just to learn more about it, but to find ways to bring people closer together, ensuring our society stays connected and democratic processes work as they should. Several strategies have been proposed to mitigate polarisation:

1. *Edge addition for exposure:* Garimella et al. (2017) suggests adding new edges in a network graph to expose users to contrasting views. Building on this idea of exposure.
2. *Balanced information diffusion:* An approach Garimella et al. (2017) selects seed nodes to ensure balanced information exposure after propagation. This method of balanced exposure leads us to another perspective on network structure.
3. *Optimal graph structure:* Musco et al. (2018) seeks a graph structure that minimises polarity and potential disagreement from connecting differing opinions. While restructuring the network is one approach, another method focuses on influencing the opinions directly.
4. *Opinion flipping:* In Matakos and Gionis (2018), the method involves flipping the opinions of nodes, especially those with high influence, to break filter bubbles. However, its real-world applicability is debatable. Diving deeper into the realm of recommendations.

5 *Diversity maximisation:* The study Vendeville et al. (2023) optimises content recommendations to maximise newsfeed diversity. This approach to diversification brings us to the concept of moderating opinions.

6 *Opinion moderation:* Another approach (Matakos et al., 2017) identifies key individuals whose moderated opinions can significantly reduce network polarisation. They propose two strategies: moderating internal opinions through education and external opinions via incentives.

2.7 Novel contributions

This study introduces a novel application of GAs to reduce polarisation in social networks by dynamically identifying and modifying key nodes and connections. Unlike traditional methods such as edge addition or opinion flipping, which often address only the symptoms of polarisation, our GA-based approach adapts to changes in network structure and user behaviour, making it robust and scalable. The method uses a comprehensive fitness function that integrates various centrality metrics, including degree, closeness, eigenvector, PageRank, and Katz centralities, along with opinion values, providing a nuanced understanding of the network. This approach requires minimal intervention, focusing on the most influential nodes and edges, thereby preserving the network's structural integrity while respecting ethical considerations.

In conclusion, our GA-based method demonstrates superior adaptability, scalability, and ethical soundness compared to traditional methods. It has shown significant reductions in polarisation in real-world datasets like Polblogs and Polbooks, highlighting its potential for diverse and dynamic social media platforms. This research advances the field of social network analysis by offering a more effective and ethical tool for promoting diversity and inclusivity in online communities. Future work could enhance this approach by integrating additional factors such as sentiment analysis or user engagement metrics to further understand and mitigate polarisation.

3 Proposed polarisation metric

Having delved into the extensive literature surrounding the measurement of polarisation, the application of GAs, and the various methods proposed to mitigate polarisation, it becomes evident that the field is ripe for innovative approaches that can effectively address the challenges of polarisation in social networks. The intricate nature of polarisation, coupled with the dynamic and vast landscape of social networks, necessitates a method that is both robust and adaptable.

Among the myriad of polarisation measures discussed, the measure proposed by Mustafa et al. (2023) stands out due to its comprehensive approach and its ability to capture the nuances of polarisation in diverse network structures. So in this research, we leverage their polarisation measure β as a benchmark to test our proposed method. Our method aims to not only reduce polarisation but also ensure that the network remains resilient to future polarisation surges. By aligning our approach with the insights gained from the literature review, we aspire to contribute a solution that is both novel and effective in the realm of social network polarisation.

The decision to employ GAs in addressing the issue of social media polarisation is grounded in several key attributes that make GAs particularly suited for this complex and dynamic problem. GAs are renowned for their adaptability and robustness in navigating complex problem spaces. The inherent flexibility of GAs, derived from their evolutionary search mechanisms, allows for the exploration of vast and intricate solution landscapes. This is particularly relevant for social media polarisation, where the problem space is characterised by a high degree of complexity and dynamism. The ability of GAs to adaptively search for optimal or near-optimal solutions without requiring a precise mathematical formulation of the problem makes them an ideal choice for tackling polarisation, where the underlying dynamics and interactions within social networks are often nonlinear and unpredictable.

A critical aspect of mitigating social media polarisation involves identifying and influencing key nodes within the network that contribute significantly to the spread of polarised content. GAs excel in this domain by efficiently pinpointing these influential nodes through evolutionary strategies that mimic natural selection processes. By evaluating a population of potential solutions over successive generations, GAs effectively identify network nodes whose modification (e.g., through the addition or alteration of connections) can lead to a significant reduction in polarisation. This process leverages various network centrality measures, integrating them into the GA's fitness function to assess the potential impact of each solution on reducing polarisation levels.

The dynamic and evolving nature of social networks demands solutions that can adapt to changing conditions and structures. GAs offer this flexibility, making them highly suitable for real-world applications. Unlike static or rule-based algorithms, GAs can evolve in response to changes in the network's topology, user behaviour, and content flow. This adaptability ensures that the strategies developed for polarisation reduction remain effective over time, even as the network grows and evolves.

The application of GAs to social media networks is particularly compelling due to the alignment of GAs' strengths with the characteristics of these networks. Social media platforms are vast, dynamic, and complex, with intricate patterns of user interactions and information dissemination. GAs' ability to handle large datasets, coupled with their proficiency in optimising multidimensional objectives, makes them uniquely capable of devising effective strategies for enhancing content diversity and fostering healthier, more inclusive online dialogues.

Our study also contributes to the field by bridging the gap between computational intelligence and social science. By applying a computational algorithm to a pressing social issue, we demonstrate the potential of interdisciplinary research to offer novel solutions to complex problems. This cross-disciplinary approach not only enhances our understanding of social media polarisation but also opens up new avenues for applying advanced computational techniques to social science research, offering a model for future studies.

4 Mathematical method

GAs have been effectively utilised to solve a variety of problems in different domains. Our aim in this research is to mitigate the pervasive problem of polarisation that characterises present-day social networks. Utilising GAs, this section identifies nodes of high influence within a given network. Following this identification, strategic interventions are proposed to foster greater diversity and equilibrium within the network. This approach commences by establishing the social network as an undirected graph $G = (V, E)$, where V signifies the set of individual nodes and E denotes the set of edges characterising relationships. Each node i within V is attributed with an opinion value o_i, denoting the individual's stance on a given subject. Positive opinion values symbolise affiliation with one group, while negative values signify alignment with another.

4.1 Research subjects

The research subjects for this study are diverse datasets representing social networks with inherent polarisation. These datasets are chosen to evaluate the effectiveness of the proposed GAs approach in reducing polarisation. The datasets used are:

1. *Polbooks dataset:* This dataset represents a network of books related to US politics, categorised into liberal, conservative, and neutral. Each node in this network represents a book, and edges between nodes represent books frequently bought together. The categorisation of these books helps simulate ideological polarisation in a controlled environment.
2. *Polblogs dataset:* This dataset represents a network of political blogs from the 2004 US election period, categorised as either liberal or conservative. Nodes represent individual blogs, and edges represent hyperlinks between them. This dataset helps simulate real-world ideological divides present in online platforms.

These datasets are selected for their relevance and representativeness of real-world social networks. By using these well-documented and widely studied datasets, the study ensures the applicability and generalisability of the findings.

4.2 Data collection and preparation

The datasets were collected from publicly available sources:

- *Polbooks dataset:* Sourced from Amazon.com, this dataset includes 105 nodes (books) and 441 edges (frequently bought together pairs). The books are categorised into 43 liberal, 49 conservative, and 13 neutral, with neutral books being arbitrarily assigned to either the liberal or conservative group.
- *Polblogs dataset:* Sourced from networks.skewed.de, this dataset includes 1,222 nodes (blogs) and 16,717 edges (hyperlinks). The blogs are categorised into 636 liberal and 586 conservative.

4.3 Graph initialisation and opinion assignment

We commence by establishing the social network as an undirected graph $G = (V, E)$, where V signifies the set of individual nodes and E denotes the set of edges characterising relationships. Each node i within V is attributed with an opinion value o_i, denoting the individual's stance on a given subject. Positive opinion values symbolise affiliation with one group, while negative values signify alignment with another.

4.4 Node selection using GAs

To identify influential nodes that could mitigate polarisation, we employ a GA for optimisation.The parameters for the GAs were carefully selected based on their proven efficacy in similar optimisation problems. The population size was set to 100 to balance computational efficiency with solution diversity. A higher population size was avoided to reduce computational overhead, while a lower size might limit the algorithm's ability to explore the solution space effectively. The number of generations was set to 200, which is sufficient for convergence in our preliminary experiments. Mutation rate was set at 0.01 to introduce necessary variations without disrupting the convergence process significantly. These parameters were optimised through a series of pilot tests, ensuring a robust and efficient search process.

1. *GA parameters:* The GA is configured with essential parameters, including the size of the population (population_size), the number of generations (num_generations), and the probability of mutation (mutation_rate).
2. *Initial population generation:* The initial population of candidate solutions is generated by sorting nodes based on their degree centrality and selecting the top nodes. These nodes possess a higher likelihood of exerting influence within the network.
3. *Fitness calculation:* The fitness of each solution within the population is computed. This involves evaluating the summation of various centrality

metrics, including degree centrality, closeness centrality, eigenvector centrality, PageRank, and Katz centrality, weighted by their corresponding opinion values.

The fitness function incorporates multiple centrality metrics, each chosen for its unique ability to capture different aspects of node influence within the network. Degree centrality measures the immediate connectivity, closeness centrality reflects the average shortest path to all other nodes, eigenvector centrality accounts for the influence of connected nodes, PageRank scores the probability of visiting a node, and Katz centrality measures the influence considering both direct and indirect connections. By combining these metrics, the fitness function provides a comprehensive evaluation of node influence, essential for identifying key nodes that can effectively reduce polarisation. The fitness function, which quantifies the fitness of a solution (set of nodes) in the GA, can be represented mathematically as follows:

Let S be a solution containing a subset of nodes from the network, and let n be the number of nodes in S. The fitness function $F(S)$ can be defined as:

$$F(S) = \sum_{i=1}^{n} o(i) \cdot \left(\sum_{i=1}^{n} (\mathrm{DC}(i) + \mathrm{CC}(i) + \mathrm{EC}(i) + \mathrm{PR}(i) + \mathrm{KC}(i)) \right)$$

where

$o(i)$	represents the opinion value of node i.
$\mathrm{DC}(i)$	is the degree centrality of node i.
$\mathrm{CC}(i)$	is the closeness centrality of node i.
$\mathrm{EC}(i)$	is the eigenvector centrality of node i.
$\mathrm{PR}(i)$	is the PageRank score of node i.
$\mathrm{KC}(i)$	is the Katz centrality of node i.

We are making the assumption that the terms: $\sum_{i=1}^{n} o(i)$ and $\sum_{i=1}^{n} (\mathrm{DC}(i) + \mathrm{CC}(i) + \mathrm{EC}(i) + \mathrm{PR}(i) + \mathrm{KC}(i))$ won't be zero. Lower fitness values indicate solutions with greater potential for reducing polarisation.

4 *Parent selection:* A probabilistic approach is adopted to select parents from the population. The probability of selection is proportional to the fitness score of each solution.

Parents are selected for crossover based on their fitness scores. The selection probability of a solution S is determined by its fitness value relative to the total fitness of the population. Mathematically, the selection probability $P_{\text{select}}(S)$ can be defined as:

$$P_{\text{select}}(S) = \frac{F(S)}{\sum_{S' \in \text{population}} F(S')}$$

where

- $F(S)$ is the fitness value of solution S.
- $F(S')$ is the fitness value of solution S'.
- The summation is performed over all solutions in the population.

These equations provide a mathematical foundation for understanding the mutation and parent selection processes within the GA framework.

5 *Crossover and mutation:* Crossover is performed between pairs of parents to generate offspring. Subsequently, a mutation operation is applied to offspring with a probability determined by the mutation rate. Mutation involves adding a node to the solution, enhancing the diversity of the population.

The mutation operation introduces diversity into the population by randomly adding a node to a solution with a certain probability. Mathematically, the mutation can be represented as follows:

Let S be a solution (set of nodes), and S_m be the mutated solution obtained from S. The mutation operation is applied with a probability of p_{mutate}. If mutation occurs, a random node i is added to the solution:

$$S_m = \begin{cases} S \cup \{i\}, & \text{with probability } p_{\text{mutate}} \\ S, & \text{otherwise} \end{cases}$$

where

- p_{mutate} is the mutation probability.
- S is the original solution.
- S_m is the mutated solution.

6 *Replacement strategy:* Offspring replace the least fit solutions in the population. This process ensures that the population evolves towards more optimal solutions.

7 *Influential node selection:* Ultimately, the most influential nodes are identified by selecting the solution with the highest fitness score. These nodes are anticipated to have a significant impact on reducing polarisation.

8 *Edge addition:* After selecting the most influential nodes of a given network, we proceed to the next step of solving our problem. To decrease polarisation, we select those influential nodes that belong to opposite opinion groups and add edges between them. This signifies the start of a dialogue between two polarised groups.

4.5 Opinion evolution process

The evolution of opinions is simulated to understand the dynamics of opinion change within the network. Here, we use the method of opinion evolution proposed by Mustafa et al. (2023), which consists of the following steps:

1 *Initialisation:* Each node's initial opinion value is assigned based on a predetermined distribution. Nodes commence with divergent viewpoints.

2 *Opinion evolution:* Over multiple discrete time steps, nodes adjust their opinion values. This adjustment is influenced by the opinions of their neighbouring nodes. The updated opinion of a node i at time step t is determined by a combination of its previous opinion and the average opinions of its neighbours. The parameter α governs the extent to which a node gives weightage to its own opinion and $1 - \alpha$ considers the importance it gives to its neighbours' opinion.

Given that each node i has a fixed innate opinion s_i and an expressed opinion at time (t), $(z_i(t))$:

$$z_i(t+1) = \alpha s_i + (1-\alpha)\frac{\sum_{j \in V} w_{ij} z_j(t)}{\sum_{j \in V} w_{ij}}$$

For opinion evolution, $steps = 3$ and $\alpha = 0.85$.

4.6 Polarisation assessment

To quantify and evaluate the degree of polarisation within the network, we use Polarisation Pointer (Mustafa et al., 2023)

$$\beta = \frac{1}{2}(d \times \max(s + (1-\rho), 0))$$

This pointer integrates opinion values and network structure.

1 *Polarisation metric:* The metric characterises polarisation under varying conditions by iteratively introducing edges between nodes holding opposing opinions. For a range of parameter values k, edges are formed between positive and negative opinion nodes. This encourages interaction among influential nodes.

2 *Group opinion means and polarisation parameter:* The computed opinions are employed to ascertain the group means, denoted as gc_plus and gc_minus, for the positively and negatively opinionated nodes, respectively. The difference between these group means is halved to yield the polarisation parameter d. This parameter serves as an indicator of polarisation intensity.

3 *Edge opinion sum and group interaction:* The interaction between groups is quantified by summing the product of opinions for edges linking nodes across different groups. This calculation is integral to the determination of the polarisation metric.

4 *Polarisation calculation and visualisation:* The polarisation metric, incorporating the polarisation parameter d, the sum of the product of opinions, and an attenuation factor ρ, is computed. The relationship between network structure, opinion dynamics, and polarisation is unveiled through visualisation, with polarisation values plotted against parameter values k.

5 Result and discussion

This part talks about what was found when a plan was looked at to make disagreements in social media less intense. The discoveries are explained in detail, especially how certain important nodes in the network affect how disagreements spread. The investigation was meticulously structured, encompassing distinct phases of analysis and experimentation. The initial phase of this study entailed the selection of influential nodes, a pivotal task executed through the careful application of the GA, the intricacies of which have been thoroughly detailed earlier. The algorithm successfully identified nodes that possessed a pronounced potential for mitigating polarisation. These identified nodes were subsequently harnessed as critical components in the ensuing stages of the analysis.

Building upon the identification of influential nodes, a strategy was formulated to produce meaningful inter-group connections within the network. These inter-group connections were established among the most influential nodes that held divergent opinions, orchestrating a deliberate interplay between contrasting stances. The number of connections to make between different groups of opinions, based on all the possible ways they could be connected, was decided. This procedural configuration was methodically tested across datasets, encompassing both synthetic constructs and real-world instances, thus assuring a comprehensive and rigorous assessment. This meticulous evaluation was designed to establish the method's generalisability and effectiveness across a spectrum of scenarios, reinforcing the credibility of the findings made in this work. The discussion revolves around the observations that were made and then details are explained of bringing together pivotal points and connecting opposite groups. Through this study, insights were gained into the mechanisms by which disagreements become less intense in complex social networks.

5.1 Polblogs dataset

Figure 4 shows the effect of adding boundary edges to a social network on the polarisation of the network. The x-axis of the graph represents the percentage of boundary edges added, and the y-axis represents the value of polarisation. The line graph shows that as the percentage of boundary edges added increases, the value of polarisation decreases.

In this case, the boundary edges represent connections between nodes that are in different groups. When these connections are added, it makes it easier for the nodes in the two groups to communicate with each other. This can lead to a decrease in polarisation, as the two groups become more aware of each other's viewpoints. In this case, the percentage decrease in polarisation is approximately 67%.

Figure 4 also shows that the rate of decrease in polarisation slows down as the percentage of boundary edges increases. This is because, as more and more boundary edges are added, the network becomes more connected, and it is more difficult for the two groups to remain isolated from each other.

Figure 4 Effect of adding boundary edges between most influential nodes in a polblogs social network (see online version for colours)

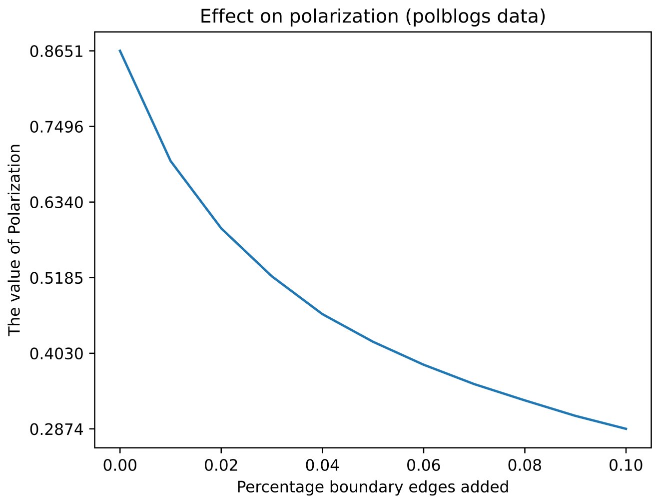

Overall, the image provides evidence that adding boundary edges to a social network can be an effective way to reduce polarisation. However, it is important to note that the effect is not linear, and the rate of decrease in polarisation slows down as more boundary edges are added.

5.2 *Polbooks dataset*

The figure shows the effect of adding boundary edges to a social network on the polarisation of the network. The x-axis of the graph represents the percentage of boundary edges added, and the y-axis represents the value of polarisation. The line graph shows that as the percentage of boundary edges added increases, the value of polarisation decreases.

The percentage difference in polarisation reduction can be calculated by finding the difference between the value of polarisation before the boundary edges were added and the value of polarisation after they were added and then dividing that difference by the initial value of polarisation. In this case, the percentage difference in polarisation reduction is approximately 43%.

The image also shows that the rate of decrease in polarisation slows down as the percentage of boundary edges increases. This is because, as more and more boundary edges are added, the network becomes more connected and it becomes more difficult for the two groups to remain isolated from each other.

Figure 5 Effect of adding boundary edges between most influential nodes in a polbooks social network (see online version for colours)

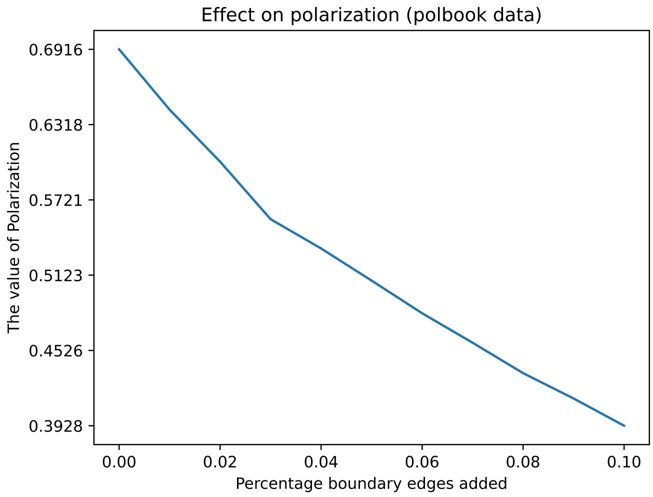

Overall, the image provides evidence that adding boundary edges to a social network can be an effective way to reduce polarisation. However, it is important to note that the effect is not linear, and the rate of decrease in polarisation slows down as more boundary edges are added.

Table 2 Percentage reduction and step change for each dataset

% edges added	*Polblogs*		*Polbooks*	
	Reduction (%)	*Step change (%)*	*Reduction (%)*	*Step change (%)*
0	0.00	–	0.00	–
1	19.41	19.41	6.92	6.92
2	31.34	11.93	12.79	5.87
3	39.26	7.92	19.49	6.70
4	45.18	5.92	22.49	3.00
5	49.23	4.05	26.50	4.01
6	52.39	3.16	30.36	3.86
7	55.88	3.49	34.02	3.66
8	59.86	3.98	37.18	3.16
9	63.05	3.19	40.16	2.98
10	66.45	3.40	43.26	3.10

Results presented in Table 2 validate the efficacy of the proposed method, as evidenced by a consistent trend of reduced polarisation across both datasets under study. To quantify the impact of the interventions, a comparative analysis was conducted on the state of the network before and after the implementation of the proposed changes. This comparison revealed that the method not only reduces polarisation but does so with minimal alterations to the existing network structure. This finding suggests that achieving a more diverse and less polarised network does not necessitate sweeping or disruptive changes. Instead, targeted, minimal interventions can effectively shift the network towards a more balanced state. Moreover, the speed at which the method moves the network towards this balanced state is noteworthy. This rapid transition is particularly beneficial in scenarios where timely decision-making is crucial, such as during political campaigns or public health crises.

It is clear that Polblogs dataset exhibited a remarkable 67% reduction in polarisation, while the Polbooks dataset displayed a 43% reduction by adding just a small fraction of boundary edges. These results are not only significant, but also illustrate how the percentage reduction in polarisation becomes greater when the size of dataset increases. It is important to highlight the challenges associated with obtaining large, diverse datasets that accurately reflect the complexity of real-world social media networks. Despite these constraints, the substantial decreases in polarisation observed in our study offer compelling evidence of the impact of our approach.

Our findings confirm that the application of GAs offers a promising approach to mitigating polarisation in social networks. By identifying and influencing key nodes within the network, our method effectively reduces polarisation while maintaining the network's structural integrity. The results from our experiments on real-world datasets, such as Polblogs and Polbooks, demonstrated significant reductions in polarisation 67% and 43%, respectively – achieved by adding a minimal number of boundary edges. This indicates that our GA-based approach can foster greater diversity and balance without requiring extensive modifications to the network structure.

The success of our approach lies in its adaptability and efficiency. GAs, with their evolutionary mechanisms, can dynamically adjust to changing network conditions, making them particularly suitable for the fluid and complex nature of social media platforms. This adaptability ensures that the strategies developed remain effective over time, even as the network evolves.

Moreover, our study bridges the gap between computational intelligence and social science, highlighting the potential of interdisciplinary research to address pressing social issues. By applying a computational algorithm to the problem of social media polarisation, we provide a novel solution that is both innovative and practical.

In summary, this research successfully meets its goal of developing an effective method for reducing polarisation in social networks. The GA-based approach not only addresses the symptoms of polarisation but also targets its root causes, ensuring a long-term solution. As the landscape of social networks continues to evolve, our findings offer a foundation for future research and interventions aimed at fostering a more inclusive and cohesive digital society.

6 Results and discussion

The primary objective of this study was to develop an efficient, minimally invasive method for reducing polarisation in social networks by leveraging GAs. Through our research, we sought to address the complex issue of social media polarisation, aiming not only to understand its intricacies but also to propose a viable solution that targets the root causes of polarisation.

This section delves into the findings of our research, highlighting the methodological design and steps taken to achieve these results. By systematically applying our proposed GA-based approach to real-world datasets, we aim to demonstrate the effectiveness of our method in reducing polarisation in social networks.

6.1 Methodological design and steps

6.1.1 Graph initialisation and opinion assignment

The social network is represented as an undirected graph $G = (V, E)$, where V denotes the set of nodes (individuals) and E denotes the set of edges (relationships). Each node i is assigned an opinion value o_i, where positive values indicate affiliation with one group and negative values with another.

6.1.2 Node selection using GAs

1. *GA parameters:* Key parameters include population size, number of generations, and mutation rate.
2. *Initial population generation:* The initial population is generated by selecting top nodes based on their degree centrality.
3. *Fitness calculation:* Fitness is evaluated using a weighted sum of centrality metrics (degree, closeness, eigenvector, PageRank, and Katz centrality) and opinion values.
4. *Parent selection and crossover:* Parents are selected based on fitness, and crossover is performed to generate offspring.
5. *Mutation and replacement:* Offspring undergo mutation, and the least fit solutions are replaced to evolve the population.
6. *Influential node selection:* The nodes with the highest fitness scores are selected for intervention.

6.1.3 Edge addition to reduce polarisation

Edges are added between influential nodes with opposing opinions to foster communication between polarised groups.

6.1.4 Opinion evolution process

Opinions evolve over multiple time steps, influenced by neighbouring nodes' opinions. This process is governed by the parameter α.

6.1.5 Polarisation assessment

Polarisation is quantified using the Polarisation Pointer β, which integrates opinion values and network structure.

6.2 Detailed analysis of results

6.2.1 Polblogs dataset

Figure 4 illustrates the effect of adding boundary edges on polarisation. As the percentage of boundary edges increases, polarisation decreases significantly. Specifically, a 67% reduction in polarisation is observed when a small fraction of boundary edges is added. The reduction is maximum when just top 4% of boundary edges are added showing targeted decrease in polarisation. Following observations are made based on Table 2 regarding Polblogs dataset:

- Initial polarisation is high due to strong ideological divides.
- Adding boundary edges facilitates inter-group communication, reducing polarisation.
- The greatest reduction occurs with the initial addition of edges, demonstrating the high impact of targeted interventions.

6.2.2 Polbooks dataset

Figure 5 shows the effect of boundary edges on the Polbooks dataset. A 43% reduction in polarisation is achieved, indicating the effectiveness of our method. Similar to the Polblogs dataset, the reduction rate is maximum when only the top few percent of boundary edges are added. Following observations are made based on Table 2 regarding Polbooks dataset:

- Initial polarisation reflects a divided network of politically themed books.
- The introduction of boundary edges enhances network connectivity, reducing polarisation.
- Results demonstrate that strategic edge addition is a scalable solution across different network sizes.

6.3 Comparison with other studies

While it is common practice to compare results with other state-of-the-art methods, it is important to note that each research work measures polarisation in its unique way, considering different factors and parameters. Consequently, direct comparisons may not accurately reflect the effectiveness of different approaches. Some methods, for example, count the number of boundary edges to determine how polarised a network is, while others, like our approach, perform a detailed analysis of the social network to determine the amount of polarisation.

Our approach uses GAs to identify and modify key nodes within the network to reduce polarisation. This method involves a comprehensive assessment of centrality metrics and opinion values, providing a nuanced understanding of polarisation and its reduction. In contrast, other methods may rely on simpler metrics such as the number of boundary edges, which may not capture the complexity of social networks.

6.4 Scalability and execution performance

6.4.1 Runtime complexity analysis

To ensure our approach is scalable to real-world datasets, we performed a detailed analysis of the execution performance in terms of runtime complexity. Our GA-based method involves several steps, each with its own computational complexity. Here, we outline the key components and their complexities:

1 *Initial population generation:*
   - *Complexity:* $O(|V| \log |V|)$.
   - *Description:* This step involves sorting nodes based on their degree centrality, which has a time complexity of $O(|V| \log |V|)$, where $|V|$ is the number of nodes.

2 *Fitness calculation:*
   - *Complexity:* $O(|V|^2)$.
   - *Description:* The fitness function evaluates the sum of various centrality metrics for each node, which involves traversing the network and computing metrics like degree, closeness, eigenvector, PageRank, and Katz centrality. The overall complexity is $O(|V|^2)$.

3 *Parent selection and crossover:*
   - *Complexity:* $O(|P| \cdot |V|)$.
   - *Description:* Parents are selected based on their fitness scores, and crossover operations are performed to generate offspring. This step's complexity depends on the population size $|P|$ and the number of nodes $|V|$.

4 *Mutation and replacement:*
   - *Complexity:* $O(|P| \cdot |V|)$.
   - *Description:* Offspring undergo mutation, and the least fit solutions are replaced to evolve the population. This step's complexity also depends on the population size $|P|$ and the number of nodes $|V|$.

5 *Edge addition:*
   - *Complexity:* $O(|E|)$.
   - *Description:* Adding edges between influential nodes has a complexity of $O(|E|)$, where $|E|$ is the number of edges in the network.

Overall, the runtime complexity of our GA-based approach is primarily driven by the fitness calculation and evolutionary steps, leading to an approximate complexity of $O(g \cdot |P| \cdot |V|^2)$, where g is the number of generations.

6.4.2 Practical implications and performance

Our approach demonstrates practical efficiency and scalability in real-world scenarios. The key reasons are:

- *Optimisation techniques:* We employ various optimisation techniques within the GA framework, such as efficient parent selection and crossover operations, to reduce unnecessary computations.
- *Parallel processing:* Many operations within the GA, such as fitness evaluations, can be parallelised, significantly reducing runtime on modern multi-core processors.
- *Empirical validation:* In our experiments with the Polbooks and Polblogs datasets, the approach executed within reasonable time frames (e.g., several minutes on standard computing hardware), indicating practical applicability.

6.5 Dataset considerations and generalisability

The datasets used in our study, Polbooks and Polblogs, are standard and widely recognised in the field of social network analysis and polarisation research. These datasets are particularly valuable because they represent real-world networks with binary opinions assigned to nodes, allowing for the examination of polarisation within a network divided into two distinct groups.

6.5.1 Challenges in dataset availability

Acquiring larger real-world datasets that meet these specific criteria – being in the form of a network with nodes and undirected edges and having binary opinions assigned to each node – is exceptionally challenging. During our research, we extensively searched for such datasets but encountered several issues:

Many datasets we found were too small to provide meaningful insights. Other datasets did not meet the necessary criteria for our analysis, such as having binary opinions that could divide the network into two groups. Use of synthetic datasets, we also generated a large synthetic dataset to evaluate our approach. However, the results from this synthetic dataset were "too good to be true", suggesting that the synthetic data did not accurately capture the complexities and nuances of real-world social networks. As a result, we decided to focus on the Polbooks and Polblogs datasets for this study.

6.5.2 Performance on larger datasets

Despite these limitations, our findings indicate a positive trend: our approach appears to perform well on larger datasets. Specifically, we observed greater polarisation reduction in the Polblogs dataset compared to the Polbooks dataset, likely due to the larger size of Polblogs. This suggests that our method will scale effectively and perform even better on larger networks. So, while the datasets used in this study are smaller, they are the most relevant and widely accepted datasets available for this type of research. Our results on these datasets provide a strong foundation, and we look forward to exploring larger datasets in future work to further validate and generalise our findings.

6.6 Implications and future work

Our results indicate that targeted interventions using GAs can significantly reduce polarisation in social networks. The effectiveness of our approach across different datasets suggests its generalisability. Future work could explore:

- Applying this method to larger and more diverse datasets.
- Integrating additional social factors into the model to enhance realism.

In summary, our research presents a novel and effective approach to reducing polarisation in social networks. By leveraging GAs, we can identify and modify key nodes to foster greater diversity and balance, ultimately contributing to healthier and more inclusive online communities.

Data availability statement

Following are the links to the datasets used in this paper:

- PolBooks: (http://www.casos.cs.cmu.edu/computational_tools/datasets/external/polbooks/index11.php).
- PolBlogs: (https://networks.skewed.de/net/polblogs).

References

Abbasi-Moud, Z., Vahdat-Nejad, H. and Sadri, J. (2021) 'Tourism recommendation system based on semantic clustering and sentiment analysis', *Expert Syst. Appl.*, Vol. 167, p.114324.

Adamic, L.A. and Glance, N. (2005) 'The political blogosphere and the 2004 U.S. election: divided they blog', in *LinkKDD '05: Proceedings of the 3rd International Workshop on Link Discovery*, pp.36–43.

Aghaee, Z. and Kianian, S. (2020) 'Influence maximization algorithm based on reducing search space in the social networks', *SN Appl. Sci.*, Vol. 2, No. 12, p.2067.

Akçay, N.C. (2018) *A Comparison of the Cultural Diplomacy Policies of France and the United Kingdom towards African Countries*, Master's thesis, Hacettepe University Graduate School of Social Sciences Turkey.

Akoglu, L. (2014) 'Quantifying political polarity based on bipartite opinion networks', *Proceedings of the 8th International Conference on Weblogs and Social Media, ICWSM 2014*, January, pp.2–11.

Amato, F., Moscato, V., Picariello, A. and Piccialli, F. (2019) 'SOS: a multimedia recommender system for online social networks', *Future Gener. Comput. Syst.*, Vol. 93, pp.914–923.

Beheshti, A., Yakhchi, S., Mousaeirad, S., Ghafari, S.M., Goluguri, S.R. and Edrisi, M.A.(2020) 'Towards cognitive recommender systems', *Algorithms*, Vol. 13, No. 8, p.176.

Bucur, D. and Iacca, G. (2016) 'Influence maximization in social networks with genetic algorithms', in *EvoApplications 2016: Applications of Evolutionary Computation*, pp.379–392.

Bucur, D., Iacca, G., Marcelli, A., Squillero, G. and Tonda, A. (2018) 'Evaluating surrogate models for multi-objective influence maximization in social networks', in *GECCO '18: Proceedings of the Genetic and Evolutionary Computation Conference Companion*, pp.1258–1265.

Capdevila, J., Arias, M. and Arratia, A. (2016) 'GeoSRS: a hybrid social recommender system for geolocated data', *Inf. Syst.*, Vol. 57, pp.111–128.

Castells, P., Vargas, S. and Wang, J. (2011) 'Novelty and diversity metrics for recommender systems: choice, discovery and relevance', in *International Workshop on Diversity in Document Retrieval*.

Chen, S., Yin, X., Cao, Q., Li, Q. and Long, H. (2020) 'Targeted influence maximization based on cloud computing over big data in social networks', *IEEE Access*, Vol. 2, pp.45512–45522.

Choi, S-M., Ko, S-K. and Han, Y-S. (2012) 'A movie recommendation algorithm based on genre correlations', *Expert Syst. Appl.*, Vol. 39, No. 9, pp.8079–8085.

Cinus, F., Minici, M., Monti, C. and Bonchi, F. (2022) 'The effect of people recommenders on echo chambers and polarization', in *Proceedings of the Sixteenth International AAAI Conference on Web and Social Media*, Vol. 16.

Donkers, T. and Ziegler, J. (2021) 'The dual echo chamber: modeling social media polarization for interventional recommending', in *RecSys '21: Proceedings of the 15th ACM Conference on Recommender Systems*, pp.12–22.

Gao, Z., Shen, T., Mai, Z., Bouadjenek, M.R., Waller, I., Anderson, A., Bodkin, R. and Sanner, S. (2022) 'Mitigating the filter bubble while maintaining relevance: targeted diversification with vae-based recommender systems', in *SIGIR '22: Proceedings of the 45th International ACM SIGIR Conference on Research and Development in Information Retrieval*, pp.2524–2531.

Garcia, D. (2023) 'Influence of facebook algorithms on political polarization', *Nature*, Vol. 620, pp.39–41.

Garimella, K., Morales, G.D.F., Gionis, A. and Mathioudakis, M. (2015) 'Quantifying controversy in social media', *CoRR*, abs/1507.05224.

Garimella, K., De Francisci Morales, G., Gionis, A. and Mathioudakis, M. (2017) 'Reducing controversy by connecting opposing views', in *Proceedings of the Tenth ACM International Conference on Web Search and Data Mining, WSDM '17*, ACM, New York, NY, USA, pp.81–90.

Garimella, K., Gionis, A., Parotsidis, N. and Tatti, N. (2017) 'Balancing information exposure in social networks', *CoRR*, abs/1709.01491.

Gillani, N., Yuan, A., Saveski, M., Vosoughi, S. and Roy, D. (2018) 'Me, My echo chamber, and I: introspection on social media polarization', in *WWW '18: Proceedings of the 2018 World Wide Web Conference*, pp.823–831.

Guerra, P.C., Meira Jr., W., Cardie, C. and Kleinberg, R.D. (2013) 'A measure of polarization on social media networks based on community boundaries', in *Proceedings of the International AAAI Conference on Web and Social Media*.

Hohmann, M., Devriendt, K. and Coscia, M. (2023) 'Quantifying ideological polarization on a network using generalized Euclidean distance', *Science Advances*, Vol. 9, No. 9, pp.1–13.

Holland, J.H. (1992) *Adaptation in Natural and Artificial Systems*, The MIT Press, Cambridge, Massachusetts.

Kaufman, E. (2018) 'The social media, the mass media, violence and democracy. The arena for institutional weakening and citizens lack of confidence', *Political Science*, Vol. 25, No. 2028, p.4.

Keefer, P. and Knack, S. (2002) 'Polarization, politics and property rights: links between inequality and growth', *Public Choice*, Vol. 111, No. 1, pp.127–154.

Ko, H., Lee, S., Park, Y. and Choi, A. (2022) 'A survey of recommendation systems: recommendation models, techniques, and application fields', *Electronics*, Vol. 11, No. 1, p.141.

Kotkov, D., Veijalainen, J. and Wang, S. (2016a) 'Challenges of serendipity in recommender systems', in *WEBIST 2016: Proceedings of the 12th International Conference on Web Information Systems and Technologies*, Vol. 2.

Kotkov, D., Wang, S. and Veijalainen, J. (2016b) 'A survey of serendipity in recommender systems', *Knowledge-Based Systems*, Vol. 111, No. C, pp.180–192.

Kotkov, D., Veijalainen, J. and Wang, S. (2017) 'A serendipity-oriented greedy algorithm for recommendations', in *WEBIST 2017: Proceedings of the 13rd International Conference on Web Information Systems and Technologies*, Vol. 1, pp.32–40.

Krebs, V. (2004) *Books About US Politics* [online] http://www.orgnet.com/ (accessed 10 February 2024).

Liu, F. and Lee, H. (2010) 'Use of social network information to enhance collaborative filtering performance', *Expert Syst. Appl.*, Vol. 37, No. 7, pp.4772–4778.

Liu, X., Liao, X., Li, S., Zheng, S., Lin, B., Zhang, J., Shao, L., Huang, C. and Xiao, L. (2017) 'On the shoulders of giants: incremental influence maximization in evolving social networks', *Complexity*, Vol. 1, p.5049836.

Maksai, A., Garcin, F. and Faltings, B. (2015) 'Predicting online performance of news recommender systems through richer evaluation metrics', in *RecSys '15: Proceedings of the 9th ACM Conference on Recommender Systems*, pp.179–186.

Matakos, A. and Gionis, A. (2018) 'Tell me something my friends do not know: diversity maximization in social networks', in *IEEE International Conference on Data Mining (ICDM)*, Singapore, pp.327–336.

Matakos, A., Terzi, E. and Tsaparas, P. (2017) 'Measuring and moderating opinion polarization in social networks', *Data Mining and Knowledge Discovery*, Vol. 31, pp.1480–1505.

Morales, A.J., Borondo, J., Losada, J.C. and Benito, R.M. (2015) 'Measuring political polarization: Twitter shows the two sides of venezuela', *Chaos 25*, Vol. 25, No. 3, pp.1–10.

Murakami, T., Mori, K. and Orihara, R. (2009) 'Metrics for evaluating the serendipity of recommendation lists', in *New Frontiers in Artificial Intelligence*, Vol. 4914, No. 2007, pp.40–46.

Musco, C., Musco, C. and Tsourakakis, C.E. (2018) 'Minimizing polarization and disagreement in social networks', in *Proceedings of the 2018 World Wide Web Conference, WWW '18*, International World Wide Web Conferences Steering Committee, Republic and Canton of Geneva, Switzerland, pp.369–378.

Mustafa, Z.U., Amir, M., Mustafa, M. and Anwar, M.A. (2023) 'From divide to debate: a methodological approach to dismantling polarization in social networks', *Preprints*, Vol. 2023, p.2023091350.

Pal, S.K., Kundu, S. and Murthy, C.A. (2014) 'Centrality measures, upper bound, and influence maximization in large scale directed social networks', *Fundamenta Informaticae*, Vol. 130, pp.317–342.

Pham, C.V., Ha, D.K.T., Vu, Q.C., Su, A.N. and Hoang, H.X. (2020) 'Influence maximization with priority in online social networks', *Algorithms*, Vol. 13, No. 8, p.183.

Qin, X., Zhong, C. and Yang, Q. (2022) 'An influence maximization algorithm based on community-topic features for dynamic social networks', *IEEE Transactions on Network Science and Engineering*, Vol. 9, No. 2, pp.608–621.

Sivaganesan, D. (2021) 'Novel influence maximization algorithm for social network behavior management', *Journal of IoT in Social, Mobile, Analytics, and Cloud*, Vol. 3, No. 1, pp.60–68.

Sobkowicz, P. (2023) 'Social depolarization and diversity of opinions – unified ABM framework', *MDPI*, Vol. 25, No. 4, p.568.

Song, G., Li, Y., Chen, X., He, X. and Tang, J. (2017) 'Influential node tracking on dynamic social network: an interchange greedy approach', *IEEE Transactions on Knowledge and Data Engineering*, Vol. 29, No. 2, pp.359–372.

Sun, N. (2023) 'Overview of definition, evaluation, and algorithms of serendipity in recommender systems', in *Proceedings of the 3rd International Conference on Signal Processing and Machine Learning*, Vol. 6, pp.640–646.

Talukder, A., Alam, M.G.R., Tran, N.H., Niyato, D. and Hong, C.S. (2019) 'Knapsack-based reverse influence maximization for target marketing in social networks', *IEEE Access*, Vol. 7, No. 34, pp.44182–44198.

Tarus, J.K., Niu, Z. and Yousif, A. (2017) 'A hybrid knowledge-based recommender system for e-learning based on ontology and sequential pattern mining', *A Hybrid Knowledge-Based Recommender System for e-Learning Based on Ontology and Sequential Pattern Mining*, Vol. 72, pp.37–48.

Tokita, C.K., Guess, A.M. and Tarnita, C.E. (2021) 'Polarized information ecosystems can reorganize social networks via information cascades', *Proceedings of the National Academy of Sciences*, Vol. 118, No. 50, p.e2102147118.

Vendeville, A., Giovanidis, A., Papanastasiou, E. and Guedj, B. (2023) 'Opening up echo chambers via optimal content recommendation', in *International Conference on Complex Networks and Their Applications XI. Complex Networks 2016 2022*, Vol. 1077, pp.74–85.

Vicario, M.D., Vivaldo, G., Bessi, A., Zollo, F., Scala, A., Caldarelli, G. and Quattrociocchi, W. (2016) 'Echo chambers: emotional contagion and group polarization on Facebook', *Scientific Reports*, Vol. 6.

Wang, L., Ma, L., Wang, C., Xie, N-G., Koh, J.M. and Cheong, K.H. (2021) 'Identifying influential spreaders in social networks through discrete moth-flame optimization', *IEEE Transactions on Evolutionary Computation*, Vol. 25, No. 6, pp.1091–1102.

Yang, B., Lei, Y., Liu, J. and Li, W. (2017) 'Social collaborative filtering by trust', *IEEE Trans. Pattern Anal. Mach. Intell.*, Vol. 39, No. 8, pp.1633–1647.

Zhang, M. and Hurley, N. (2008) 'Avoiding monotony: improving the diversity of recommendation lists', in *RecSys '08: Proceedings of the 2008 ACM Conference on Recommender Systems*, pp.123–130.

Zhang, Y.C., Séaghdha, D., Quercia, D. and Jambor, T. (2012) 'Auralist: introducing serendipity into music recommendation', in *WSDM '12: Proceedings of the Fifth ACM International Conference on Web Search and Data Mining*, pp.13–22.