I Want to Break Free! Recommending Friends from Outside the Echo Chamber

ANTONELA TOMMASEL, antonela.tommasel@isistan.unicen.edu.ar ISISTAN (CONICET/UNCPBA), Argentina
JUAN MANUEL RODRIGUEZ, juanmanuel.rodriguez@isistan.unicen.edu.ar ISISTAN (CONICET/UNCPBA), Argentina
DANIELA GODOY, daniela.godoy@isistan.unicen.edu.ar ISISTAN (CONICET/UNCPBA), Argentina

Recommender systems serve as mediators of information consumption and propagation. In this role, these systems have been recently criticized for introducing biases and promoting the creation of echo chambers and filter bubbles, thus lowering the diversity of both content and potential new social relations users are exposed to. Some of these issues are a consequence of the fundamental concepts on which recommender systems are based on. Assumptions like the homophily principle might lead users to content that they already like or friends they already know, which can be naïve in the era of ideological uniformity and fake news. A significant challenge in this context is how to effectively learn the dynamic representations of users based on the content they share and their echo chamber or community interactions to recommend potentially relevant and diverse friends from outside the network of influence of the users' echo chamber. To address this, we devise `FRediECH` (a Friend RecommenDer for breakIng Echo CHambers), an echo chamber-aware friend recommendation approach that learns users and echo chamber representations from the shared content and past users' and communities' interactions. Comprehensive evaluations over Twitter data showed that our approach achieved better performance (in terms of relevance and novelty) than state-of-the-art alternatives, validating its effectiveness.

CCS Concepts: • **Information systems** → **Social recommendation**; • **Computing methodologies** → *Neural networks.*

Additional Key Words and Phrases: link prediction, echo chambers, social media, diversity, filter bubbles, social network analysis

ACM Reference Format:
Antonela Tommasel, Juan Manuel Rodriguez, and Daniela Godoy. 2021. I Want to Break Free! Recommending Friends from Outside the Echo Chamber. In *Fifteenth ACM Conference on Recommender Systems (RecSys '21), September 27-October 1, 2021, Amsterdam, Netherlands.* ACM, New York, NY, USA, 16 pages. https://doi.org/10.1145/3460231.3474270

1 INTRODUCTION

Social media have become a ubiquitous part of their millions of users' daily lives and activities by providing new forms of communication and possibilities of interactions. One of the essential characteristics of social platforms is their potential for rapidly disseminating information on a large scale. In recent years, while enabling access to information, social media have fostered the propagation of different forms of misinformation, and even abusive language (collectively known as *online harms*). In turn, these phenomena have been associated with an increased user (and society) polarization regarding politics, science and healthcare, among others [5], leading to the emergence of the so-called "*echo chambers*".

Echo chambers are related to situations in which individuals only consume content or interact with other users expressing their same points of view [8], resulting in selective exposure, biased assimilation, and group polarization. This phenomenon has arisen as a concern not only in political discourses [11] but also in the context of conspiracy theories, in which they could lead to a stronger radicalization, seclusion from society and destructive actions [23].

Recommender systems play an important role as mediators of information propagation. In turn, they are affected by the different forms of online harms, hindering their ability to achieve accurate predictions, thus becoming unintended means for spreading and amplifying harms [10]. This situation stems from the fundamental concepts and assumptions on which recommenders are based on. The premise of *homophily* reduces the diversity (and novelty) of the information and the social relationships to which users are exposed, which could amplify biases and reinforce echo chambers.

Harnessing recommender systems with misinformation- and harm-aware mechanisms becomes essential to mitigate the negative effects of the propagation of online harms and increase the user-perceived quality of recommender systems. In this direction, to reduce (and eventually break) the echo chamber effect, social media platforms have started to apply different strategies [10]. Even though some strategies have been criticized for intensifying inflammatory political rhetoric and misinformation, others have been found to be effective [3].

In this work, we tackle the friend recommendation problem by fostering recommendation diversification in an echo chamber awareness setting. To this end, we devise `FRediECH` (A Friend RecommenDer for breakIng Echo CHambers), an echo chamber-aware friend recommendation approach that relies on implicitly modeling the echo chamber or community membership of users to present them with relevant friend recommendations from outside the influence of their community. Given the dynamic nature of social media interactions, implicitly modeling echo chambers allows providing recommendations to users belonging either to existing communities or to potentially new forming ones. In summary, the contributions of this work are:

- We formulate the echo chamber-aware friend recommendation problem based on implicitly modeling the echo chamber structure.
- We develop `FRediECH` inspired by a graph convolutional network and a Deep & Wide architecture, coupling echo chamber awareness and user representations to balance the relevance, diversity and novelty of friend recommendations.

To support these contributions, we conducted an experimental evaluation over a polarized Twitter data collection, and compared our approach with state-of-the-art recommendation techniques, both in terms of relevance and diversity of recommendations. `FRediECH` produced similar recommendations in terms of relevance to those of the selected baselines while increasing their diversity and novelty. As a result, `FRediECH` allows recommending users who are different among them and from the already known ones, thus effectively helping to reduce the echo chamber effect.

2 RELATED WORKS

Link prediction is one of the classic tasks of recommender systems in social media consisting on inferring which new interactions (i.e., links) among users are likely to be observed in the future [18]. Multiple and diverse techniques have been proposed. The simplest methods consider pairwise similarities based on network structure, neighbourhood features, user-generated content, community structures, or even random walks. If casting the adjacency matrix of a network as a rating matrix, any recommendation technique can be obliviously applied to the link prediction task [13].

Social networking sites have been shown to face two simultaneous effects over users' points of view [12]. First, *echo chambers* by which users tend to consume information aligned to their views. Second, *filter bubbles* where personalization traps users only presenting them with similar content. Both effects induce recommenders to narrow the suggested items' diversity, potentially segregating users and biasing their opinions [20]. Diverse models have been proposed to burst both chambers and bubbles. However, these proposed solutions might not be easy to deploy as they rely on knowing users' opinions and their willingness to accept other viewpoints. Whereas recommendation tasks have traditionally aimed at only maximizing the correctly predicted links [25], there are other relevant qualities to consider, as two correctly predicted links might lead, for example, to different patterns of information propagation, content structure or diversity.

Closely related to this study are the works of Sanz-Cruzado and Castells [25] and Grossetti et al. [12], who proposed re-ranking strategies to enhance the structural diversity of recommendations and mitigate filter bubbles. Masrour et al. [19] proposed an adversarial technique with a Modularity based re-ranking for fostering recommendation fairness regarding a protected user feature. Experimental evaluation using gender as a protected feature showed that the approach could improve fairness without significantly degrading accuracy. Despite fostering fairness, increasing gender diversity might not reduce the filter bubble effect as it does not necessarily promote content nor structural diversity.

Sanz-Cruzado and Castells [25] (referred as SCC) focused on the Twitter friend recommendation problem based on weak links by applying a greedy re-ranking to a well-performing topology-based baseline. The underlying network was based on the content-based interactions between users. Optimization aimed at balancing the initial recommendations' accuracy and a targeted diversity metric (Modularity, Gini Index, or Clustering Coefficient complements). Such metrics depended on explicitly discovering the network's community structure, which might be computationally complex for large graphs. The authors showed that their approach resulted in a diversity enhancement in the flow of information through the network at the expense of reducing recommendations' accuracy.

On the other hand, Grossetti et al. [12] proposed a tweet re-ranking strategy based on defining users' community profiles and computing their similarity with the tweets' community score (referred as CAM, Community Aware Model). Similarity was measured based on the topology structure of the social network, an embedding representation of the shared tweets, and an indicator of how tweets flowed between communities. Experimental evaluation showed that the re-ranking strategy reduced the filter bubble effect while increasing the accuracy of recommendations.

Our research aims to take a step further from re-ranking strategies by proposing a recommendation approach that learns both users and the implicit echo chamber or community representations to jointly optimize the relevance, novelty, and diversity of recommendations. Unlike the described works, our approach does not explicitly rely on discovering communities or knowing protected user features.

3 FREDIECH: A FRIEND RECOMMENDER FOR BREAKING ECHO CHAMBERS

A fundamental challenge for broadening the recommendations outside users' echo chamber is learning the dynamic user and echo chamber representations based on user-user interactions and shared content for both existing and new communities. In general, the diversity of recommendations is increased by re-ranking the results obtained through other techniques according to the network's explicit community structure. However, knowing such precise structure might constrain the approaches as requires defining in advance the criteria to find such communities. In this sense, `FRediECH` aims to implicitly induce the community structure to seamlessly adapt to changes in user interactions and content patterns, striking a balance between relevance and diversity.

We denote $G =< U, E >$ the graph structure of the social network where U is its set of users, and E is a subset of U^2 representing the existing interactions between user pairs in the network. For each user $u_i \in U$, we refer to $\Gamma(u_i)$ as his/her relations in the network. As relations have a direction, we differentiate between $\Gamma_{in}(u_i)$ and $\Gamma_{out}(u_i)$, the incoming and outgoing relations, respectively. We will denote T_{u_i} the set of tweets shared by user u_i.

Problem. The **echo chamber-aware recommendation** problem for a given user u_i and his/her past interactions $(\Gamma(u_i))$ and shared content (T_{u_i}) aims at learning a function $f\left(\Gamma(u_i), T_{u_i}\right) \rightarrow s_{ij}$ for each candidate user $u_j \in U$ where s_{ij} denotes the estimated strength of the potential interaction between u_i and u_j. This estimated strength is a proxy for how likely users are to interact in the future. Thus, the higher the strength, the better the balance between the relevance of u_i and his/her distance to the echo chamber of u_j, and the likelihood of users interacting in the future. Then, these estimations are used for making a ranking of recommendations for each target user.

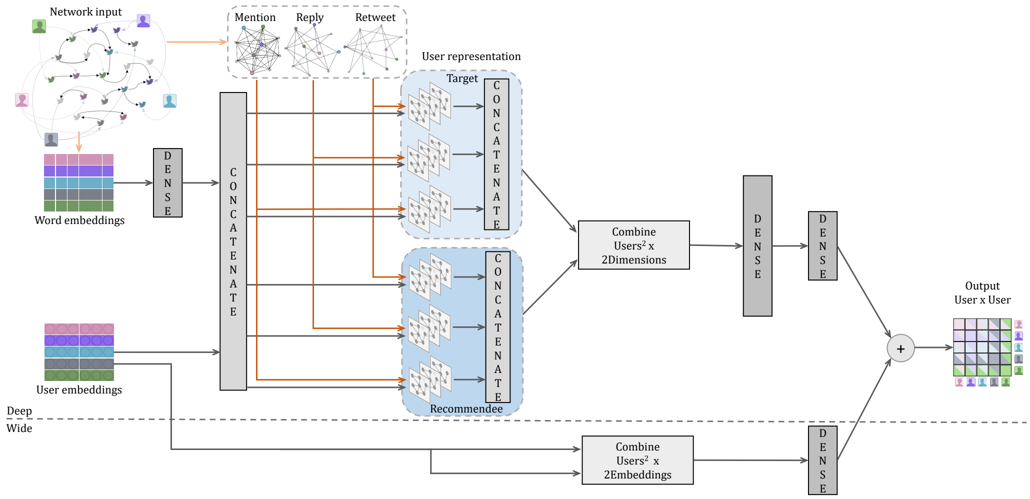

Fig. 1. Schematic diagram of FRediECH

The overall architecture of FRediECH is schematized in Figure 1. It takes as input a social network, and for each user u_i outputs a ranking of users according to the estimated relationship strength.

User representation construction.

Given the dynamic content-based nature of echo chambers, instead of focusing on the topological follower/followee graph, FRediECH builds on the dynamic conversational interaction network of Twitter. Thus, edges between users u_i and u_j represent that user u_i replied, mentioned, or retweeted a tweet shared by user u_j. Edges are directed and weighted based on the number of interactions (e.g., the number of mentions between two users). Additional weight considerations could be introduced based on the number of reactions that conversations received or their length.

Graph Convolutional Networks (GCN) allow neural networks to represent nodes in a graph based on their characteristics and those of the adjacent ones (i.e., the users with whom they interact). Considering a recommendation scenario, from the target's point of view, this includes the characteristics of their interactions. Conversely, for the potential recommendees, it includes characteristics related to how they present themselves in terms of the content they share and others interact with. This is represented in Figure 1 by the bluish boxes. The model is based on the GCN as defined by Kipf and Welling [16], and its corresponding implementation in Spektral[1] that defined the graph convolution as:

$$GCN = \sigma(\hat{D}^{-\frac{1}{2}}\hat{A}\hat{D}^{-\frac{1}{2}}XW + b) \tag{1}$$

Here, σ represents the chosen linear activation function. A represents the corresponding adjacency matrix from which $\hat{A} = A + I$ and $\hat{D_{ii}} = \sum_j \hat{A_{ij}}$ are derived. $\hat{D}$ acts as a symmetric normalizer of the $\hat{A}$ matrix to avoid weights linearly scaling as the number of user friends (or adjacencies) increases. X is a matrix representing users' latent features, W is a matrix of trainable weights and b represents the trainable bias vector.

Users are represented by a trainable embedding of their intrinsic characteristics and BERT embeddings, which are concatenated in X (the matrices in the left part of Figure 1). For each user, the BERT[2][7] embeddings without fine tuning for the last 15 shared tweets are averaged and passed through a dense layer to match their size to that of the other embeddings. Each tweet was represented by its average pooled embeddings. For FRediECH to learn users' latent

[1] https://graphneural.network/
[2] https://huggingface.co/bert-base-uncased

characteristics based on their interactions, we defined three parallels GCNs. Each GCN allows learning the specific contributions (weights) of a particular interaction type to the complete set of users' relations. Then, the outputs of the three GCNs are concatenated to generate an intermediate user representation based on the combined interaction types.

Model prediction.

As mentioned, the end goal of `FRediECH` is to recommend novel and diverse users based on estimated interaction strengths. These estimated strengths are asymmetric to accommodate for the asymmetric nature of Twitter relations. Thus, as the bluish blocks in Figure 1 show, `FRediECH` uses two GCN blocks to extract users' latent features based on trainable embeddings of a user and his/her interactions. The blocks extract the representation of the target user and the potential recommendees, respectively. These two representations are then concatenated and passed through two dense layers to estimate the strength of the interaction. Despite the networks are architecturally the same, they are not siamese as weights are not shared between them.

Finally, the resulting embeddings of the target user and the potential recommendees are concatenated and passed through a dense layer with one unit. The output of this layer is added to the model's output based on the GCNs and the two dense layers. This is inspired by the Deep & Wide architecture [4]. While the Deep part (the GCNs) can better generalize over the unseen interactions, the Wide part can better learn the implicit characteristics of users by linear operations, facilitating the training process. In the final step, for each user, `FRediECH` ranks the recommendations in descending order based on the estimated strength. Even though `FRediECH` can consider different interaction types, it combines the contribution of the different interactions into a unique strength value.

Due to hardware limitations, it is unfeasible to estimate all interaction strengths simultaneously, as the corresponding adjacency matrices and their combinations require a large amount of memory in the order of U^2. For this reason, estimates are computed for a user pair at the time, which allows generating matrix representations based only on the target user and his/her neighbours, as all non-adjacent user representations in the GCN would be multiplied by 0, rendering them irrelevant.

Model training.

We define a loss function (Eq.2) based on the distance between users ($d(u_i, u_j)$) and the number of interactions (Y_{ij}). The logarithm aims to reduce the influence of users with many interactions. Then, the scalar multiplier prevents an interaction with a weight of 1 to become zero.

$$L\left(Y, \hat{Y}\right) = \frac{\sum_{i,j} d\left(u_i, u_j\right)^{\beta} \left(\hat{Y_{ij}} - log_2\left(2Y_{ij}\right)\right)^2}{|E|} \tag{2}$$

While it is expected that existing relations between users in the same echo chamber will have a high intensity, the distance definition aims at weighting the actual loss of the network in a way that interactions between users belonging to different echo chambers carry a higher weight than interactions between users in the same echo chamber. As a result, the function favours the diversity of recommendations by learning the structure of echo chambers or communities without explicitly finding them. In turn, this allows for more freedom in the echo chamber definition and more sensitivity to changes in the network. In this sense, β allows tuning the preference of whether recommendations belong to the same group, i.e., whether to receive closer and more accurate recommendations, or farther and more diverse recommendations. For the purpose of the performed evaluation, β was set to 1.

The $d(u_i, u_j)$ is based on the cosine similarity between users. To avoid creating a dependency between the model and the distance function, distance is computed over a new 10-dimensional embedding (e_i) introduced to represent users. These embeddings were defined to capture the implicit community structure and were trained before the main model. The intuition is that users with similar interaction patterns will be represented by similar embeddings. The loss function (Eq. 3) was based on the combined network of the three interaction types and the interactions Y_{ij} between users. $f(x)$ (Eq. 4) is a parameterized weighting function inspired by GloVe's loss function [21] setting α to 3/4. While GloVe uses a word co-occurrence matrix, we replace words with users and define co-occurrences as the number of user interactions.

$$L_{dis} = \sum_{i,j=1}^{E} f\left(Y_{ij}\right)\left(e_i^T e_j - ln\left(Y_{ij}+1\right)\right) \tag{3}$$

$$f\left(x\right) = \begin{cases} \left(\frac{x}{x_{max}}\right)^{\alpha} & if x < x_{max} \\ 1 & otherwise \end{cases} \tag{4}$$

Then, based on the cosine similarity of all edges in the training set, we computed μ_{cos} and σ_{cos} to define $d(u_i, u_j)$ as Eq. 5 shows. This definition standardizes the scores to the range $[-2 * \sigma_{cos}, 2 * \sigma_{cos}]$.

$$d\left(u_i, u_j\right) = \begin{cases} 0.1 & if\, d_{raw}\left(u_i, u_j\right) < 0.1 \\ 0.9 & if\, d_{raw}\left(u_i, u_j\right) > 0.9 \\ d_{raw}\left(u_i, u_j\right) & otherwise \end{cases} \tag{5}$$

$$d_{raw}\left(u_i, u_j\right) = 1 - \frac{cos\left(e_i, e_j\right) - \left(\mu_{cos} - 2\sigma_{cos}\right)}{4\sigma_{cos}} \tag{6}$$

$d_{raw}(u1, u2)$ (Eq. 6) aims at avoiding the vanishing gradient problem. For the great majority of user pairs, it returns a value between zero and 1 when $\mu_{cos} - 2\sigma_{cos} \leq cos(u_i, u_j) \leq \mu_{cos} + 2\sigma_{cos}$. $d(u_i, u_j)$. Clipping values were set to ensure that user interactions contribute to the loss function without each having too much influence. The inferior value (0.1) was chosen to avoid examples with zero or small contributions. On the other hand, the superior was chosen for symmetry around the 0.5, which represents the average cosine similarity between user pairs.

Finally, following Word2Vec, we applied a negative sampling with a 1:1 ratio, where Y_{ij} is randomly set to 0.5 (and the logarithm to zero). In this sense, unlike the BPR [22] framework or siamese architectures, negative and positive samples are independently given to the network, which has to learn that positive instances are strongly related. In contrast, negative instances decrease the strength of the interaction. As `FRediECH` aims at recommending users from outside their area of influence (i.e., maximize their novelty and diversity), the distance function was set to $d_{neg}(u_i, u_j) = 1 - d(u_i, u_j)$. This modification prevents the negative sampling from largely penalizing users who do not belong to the same echo chamber and are unlikely to have interacted in the training set but could be relevant recommendations.

4 EXPERIMENTAL SETTINGS

To evaluate the performance of `FRediECH` we conducted extensive experiments aiming at answering the following research questions:

- **RQ1**. How does `FRediECH` perform when compared with other state-of-the-art recommender techniques?
- **RQ2**. How do the key components of `FRediECH` contribute to the recommendation performance?

To this end, we describe the data collection used, the baselines selected from the literature and the evaluation metrics.

	avg (± std)	25% - 50% - 75% quantiles
#Users	6,442	
#Tweets	7,016,552	
Tweets per User	1089.188 (± 1413.743)	11 - 81 - 2865
Relations per User	680.884 (± 1071.555)	6 - 42 - 1168
Mentions per User	460.298 (± 733.093)	5 - 34 - 792
Replies per User	87.693 (± 190.687)	2 - 12 - 83
Retweets per User	399.440 (± 353.426)	124 - 316 - 585

Table 1. Data collection details

4.1 Data collection

Evaluation was based on the `obamacare`[3][11] data collection, which includes tweets related to the `#obamacare` and `#aca` hashtags in Twitter, shared between May 2008 and October 2017. The estimated polarity of users is also included, based on the model by Barberá et al. [1]. Originally, positive polarities were associated with republicans, while negative polarities indicated a democrat leaning. To align with the `obamacare` topic, in this study, positive scores will be associated with democrats (as they were pro Obamacare), while negative ones with republicans (as they were against Obamacare).

Tweets were retrieved using the *Faking it!* tool[4]. For each tweet, we retrieved its content, replies and its retweets. From the original set of tweets, we were able to retrieve approximately 8 million public tweets belonging to 8,164 users, and 585,524 adjacent users (users that were mentioned or replied to but that did not write any tweet on the original collection)[5]. From such set, we kept users with at least one relation and that belonged to the largest connected component (LCC) of the retrieved interaction graph. This selection ensures that each user can be both source and destination of information content. The final data collection used in the performed evaluation comprised 6,442 users (representing 79% of the originally retrieved users). Table 1 summarizes the basic statistics of the collected tweets.

Echo chambers are characterized by users mostly interacting with others with similar views [6], in this case, their political leaning. Following Garimella et al. [11] and Cota et al. [6] we quantified the polarization of users in the LCC, relying on the relation between user leaning (i.e., production leaning) and that of the users with whom they interacted (i.e., consumption leaning) to assess the existence of echo chambers.

Figure 2 shows the conversational interaction graphs of users in the LCC. Green nodes represent democrats, while grey nodes represent republicans. Despite each interaction lead to diverse topological structures, in all cases, users are grouped based on their leaning, with a few small mixed groups with users having leanings close to zero. Users seemed to be more likely to reply to users with the same leaning. Figure 3 shows the relation between the political leaning of users and the average information consumption leaning per interaction type. The colour represents the number of users (density) in the space, the lighter the area, the higher the density of users in such area. Marginals show the distribution of user leaning. As observed, despite user behaviour differs according to user leaning and the interaction type, in all cases, positive correlations were found between users' production and consumption leaning.

In average, 89% of the interactions of republican users were with other republicans Conversely, democrats interacted with users on a wider range of democrat and neutral leanings. In the case of replies and mentions, democrats also engaged in conversations with republicans. In summary, the graph and leaning analyses showed the existence of groups of users interacting with other users with similar leanings, which allows inferring the existence of echo chambers.

[3]The data collection can be obtained upon request from the authors of [11].
[4]Available at: https://github.com/knife982000/FakingIt
[5]The final retrieved set of tweet IDs, their metadata and the resulting graphs are available at: https://github.com/tommantonela/frediech_recsys2021. As per Twitter TOS, the shared graphs only include user and tweet IDs involved in the interactions. No user information was included in the analysis.

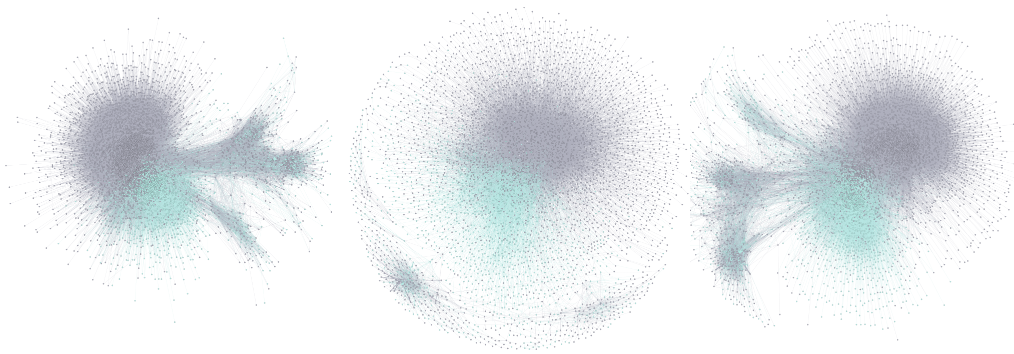

(a) Mentions network (b) Replies network (c) Retweets network

Fig. 2. Conversational graphs and user leaning
Green nodes represent users with positive leaning, while grey nodes represent users with negative leaning.

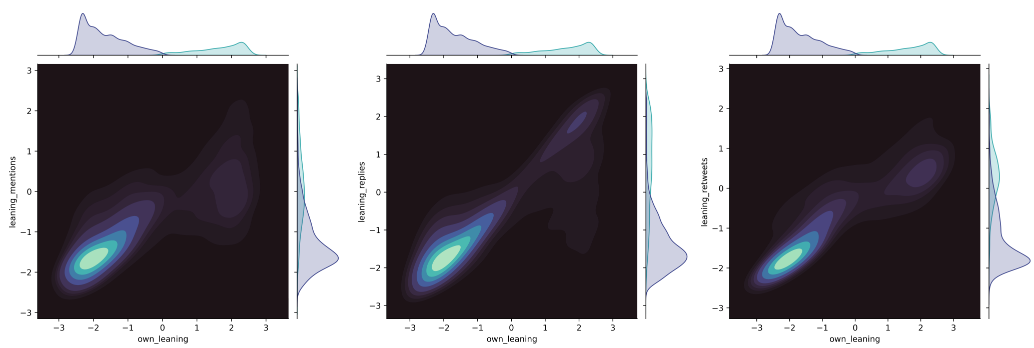

(a) Mentions network (b) Replies network (c) Retweets network

Fig. 3. Distribution of own and consumption polarity for the different types of interactions

4.2 Baselines

The performance of FRediECH was compared to 11 different recommendation techniques. First, a trivial and non-personalized reference baseline: **popularity** (i.e., users with the highest degree were recommended), and a **random** recommendation, as a lower bound reference. To this end, 10 rounds of recommendations were performed with different random seeds, averaging their results. Second, we included two traditional user recommendation techniques:

• **Topological-based techniques**. We considered four commonly used metrics based on neighbourhood overlap [18], namely Jaccard similarity, Adamic-Adar, Common Neighbours and Resource Allocation. In all cases, the NetworkX[6] implementations were used. Three alternative neighbourhoods definitions were considered: a graph only containing incoming edges, a graph only containing outgoing edges, and an undirected transformation of the graph.

• **Content-based techniques**. Users were recommended based on the cosine similarity of their representations given by either the centroid of their TD-IDF or Word2Vec vectors. User representations were obtained from i) the tweets they shared with others, ii) the tweets shared by the users with whom they interacted (i.e., the tweets shared by the users at the end of their outgoing links), or iii) the tweets shared by the users that interacted with them (i.e., the tweets

[6]https://networkx.org/

shared by the users at the beginning of their incoming links). The number of tweets considered in the representation was also varied, including either all tweets or only the last 15.

Third, we considered the closely related works of (see Section 2 for more details):

• **SCC** [25]. We experimented with the defined re-ranking metrics (Modularity, and Gini Index complements).

• **CAM** [12]. We adapted the technique to the friend recommendation problem. In this case, the target users' community profile was compared to the item/user community score vectors. Unlike in the original study, the graph was not based on the follower/followee relationship but on the content-derived interactions.

In both cases, communities were found using the Louvain algorithm [2]. The base recommender was set to the simple Adamic Adar ranking metric or Implicit MF [15], which were the best performing techniques in the original studies.

Finally, we adapted several traditional and state-of-the-art user-item recommendation techniques for friend recommendation:

• **Implicit MF** [15]. A top-performing matrix factorization technique based on a factor model tailored for implicit feedback settings.

• **NeuralCF** [14]. A deep-learning collaborative filtering model for recommendation based on fusing matrix factorization with multi-layer perceptron to learn the user–item interaction function.

• **GraphRec** [9]. A graph neural network jointly representing social interactions, item features and ratings.

• **Diffnet** [27]. A deep-learning influence propagation model that simulates how users are influenced by social propagation processes. Users and items are represented by embeddings comprising both the collaborative and the content information.

• **Mult-VAE** [17]. A neural generative model based on variational autoencoders with multinomial conditional likelihood.

When available, the original implementations were used. When corresponding, parameters were either optimized according to the procedures described in the original studies or set to the default best-performing ones in the original studies. Then, for each baseline, we selected the configuration achieving the highest results. Tweets were slightly pre-processed by replacing URLs, and removing symbols and numbers[7].

4.3 Evaluation Metrics

Evaluating recommendations only based on the number of correct predictions provides a partial perspective of their performance [25] as there might be other qualities that would enhance their value, such as their structural and global impact in the network and its evolution. Recommendations play a substantial role in shaping network growth, allowing to enhance desirable properties during its evolution [25] beyond the short-term recommendations [26]. In this sense, diversity has been studied based on its implications on network shaping and how it can promote the common good [26], for example, by mitigating the existence of echo chambers or filter bubbles. Then, diversity and novelty can measure how recommendations help users in outing from their known community, thus broadening their social experiences [25].

In terms of *relevance*, the quality of recommendations was evaluated based on three ranking-oriented metrics: Precision@k, Recall@k and nDCG@k. On the other hand, variations of intra-list dissimilarities [24] were used to assess the *diversity* (i.e., differences within the recommended list) and *novelty* (i.e., differences between the known users and the recommended ones). To better capture the effect of recommendations in social groups (in this case, echo chambers),

[7] More details and implementations can be found at: https://github.com/tommantonela/frediech_recsys2021.

both diversity and novelty were assessed considering both the individual and community point of view. Also, the community point of view could allow assessing the fairness of recommendations across the multiple groups.

Eq 7 and 8 define individual diversity and novelty, where R_u represents the recommendations made for user u and F_u the users with whom u has already interacted. Community diversity and novelty are not computed per user but per community. In this regard, R_u is replaced by $R_{c_i} = \bigcup_{u \in c_i} R_u$, which represents the set of recommendations made for all users in community c_i, and F_u was replaced by $F_{c_i} = \bigcup_{u \in c_i} F_u$, which represents the set of users with whom users in c_i have already interacted. Communities were discovered based on the the Louvain algorithm [2].

$$IndDiversity\,(u) = \frac{1}{|R_u|\,(|R_u|-1)} \sum_{i \in R_u} \sum_{j \in R_u} d_m\,(i,j) \tag{7}$$

$$IndNovelty\,(u) = \frac{1}{|R_u|\,|F_u|} \sum_{i \in R_u} \sum_{j \in F_u} d_m\,(i,j) \tag{8}$$

Dissimilarities ($d_m\,(i,j)$) were measured using the euclidean distance of structural and content-based vectors. Based on the work by Grossetti et al. [12], structural distance was computed over the community profile of users, defined as a normalized vector $CP_u = (r_{(c_1)}, \ldots, r_{(c_n)})$ where $r_{(c_i)}$ represents the interactions rate with users in community c_i. On the other hand, content-based distance was computed based on the Word2Vec representation of users' own tweets.

To evaluate the effect of recommendations on the diversity and novelty of the resulting network, we considered the extended network $G' =< U, E' >$ that would result if users accepted and interacted with all recommended friends (i.e., perfect precision). Then, $E' = E \cup \hat{E}$, where $\hat{E}$ represents the full set of recommendations as incorrectly predicted links were not excluded from the diversity and novelty analysis. Despite assuming that all recommendations will be accepted might not be realistic [25], it is useful for assessing how recommendations may shape the future network.

All evaluations were performed over the same data partitions and evaluated using the same set of metrics. A cut-off threshold was defined to select the top-K recommended users, where K was set to 10, as approximately half of the users had 10 or more interactions in the test set. Recommendations were considered correct if they appeared in the test set.

4.4 Implementation Details

The model was implemented on TensorFlow. The optimizer was set to Adam with a learning rate of $1e-3$, $\beta_1 = 0.9$ and $\beta_2 = 0.999$. The dimension of the user and BERT embeddings was set to 64. The GNC and the deep leaning had 32 units. The only pre-trained component was BERT, while `FRediECH` was trained end-to-end from random states. Hyper-parameter optimization was focused on the dimension of the intermediate layers and embeddings (with a maximum size of 64 to avoid overfitting). To prevent the explosion of the parameter combinations to try, output size remained constant across layers. Batch size was set to 20 to reduce memory consumption (in each batch for each user the embeddings of adjacent users are required). The learning process was stopped once no loss changes were observed, reaching convergence after 4 epochs.

The model was trained on a Dell Inspiron7559 with 16Gb RAM, a i7-6700HQ and a NVidia GeForce 960 GTX 4Gb. In this hardware, training and recommendations took approximately 6 and 8 hs, respectively. Scalability might be limited by the size of the embedding tables. In this regard, making recommendations for users that already have many interactions (e.g., celebrities or politicians) could hinder the model training as it would require operating with high-dimensional matrices. To avoid this, for example, temporal profiles could be introduced to limit the analysis to the last N interactions.

5 EVALUATION

Evaluation was performed in an offline setting, based on a temporal graph split into training and test data. The interactions (replies, retweets and mentions) made before August 30 2017 were used as the training set (approximately the 80% of all interactions), while the remaining interactions were used as the test set. For each metric, we report the average score across all users (or communities) and the corresponding standard deviation. For all but the group diversity/novelty metrics, we also report the results of a paired statistical analysis (considering an alpha value of 0.01).

5.1 RQ1. Comparison with state-of-the-art techniques

Table 2 presents the evaluation results for `FRediECH` and the selected baselines. For each metric, the best results are shown in bold, and the second-best are underlined. Diffnet results are not included as recommendation relevance was close to zero in several of the performed runs. For NeuralCF, no difference was observed between the optimized parameters and the default ones. The reported GraphRec results were obtained with the default parameters, while for Mult-VAE they were obtained with the optimized parameters. In both cases significant differences were observed favouring the selected configuration. The Table also includes the diversity and novelty of the original graph. In this case, metrics were computed considering the test set as the set of recommended users.

The accuracy-based results for the random recommender were lower than those of the other baselines (with differences up to a 150%, 433% and 50% for precision, recall and nDCG). Conversely, it achieved high diversity and novelty results. Random recommendations showed statistically significant differences regarding diversity and novelty when compared to all techniques but GraphRec and `FRediECH`, which achieved significantly better results.

As observed, there is a trade-off between the relevance of recommendations, and their diversity and novelty. In general, techniques achieving high relevance also achieved low diversity and novelty scores, as it is the case of popularity. These observations imply that users tend to follow popular users from their own echo chambers or communities. The topological baselines also achieved high precision and low diversity, which is expected as recommendations are based on user neighbourhood. For the four metrics, the highest results reported in the Table were observed when considering the incoming links (as Sanz-Cruzado and Castells [25] also reported). Adamic-Adar and Common Neighbours achieved similar results, with non statistically significant differences. As the Table shows, most of the the differences observed for diversity and novelty were statistically significant, favouring `FRediECH`.

ImplicitMF was one of the best performing techniques in terms of recall and nDCG, when representing users by their outgoing relations. In addition, ImplicitMF was statistically superior to NeuralCF in terms of relevance and novelty and to GraphRec in terms of relevance. Mult-VAE achieved similar results to those of NeuralCF, showing good nDCG, but lower diversity and novelty when compared to GraphRec and `FRediECH`. In all cases, significant differences favouring `FRediECH` were observed for at least one diversity and novelty metric.

The results reported for the content-based baselines correspond to a user representation based on the TF-IDF centroid of the content shared by the users that interacted with them, which achieved the highest diversity and novelty results (as Sanz-Cruzado and Castells [25] also reported). The highest relevant results were observed when representing users based on their own content. While considering the full tweet set increased the diversity of recommendations, using only the last 15 increased their relevance. These observations could relate to the broad period covered by the data collection, in which conversation topics (and user interests) could have shifted. Finally, the embedding-based representations allowed increasing the relevance scores while decreasing their diversity and novelty.

Regarding the diversity-based baselines, while SCC achieved higher relevance and structural novelty, CAM achieved higher content diversity. For SCC, while Gini Index obtained the most diverse and novel results (as reported in the

	Precision	Recall	nDCG	Structural dissimilarities				Content-based dissimilarities			
				Ind. Diversity	Ind. Novelty	Comm. Diversity	Comm. Novelty	Ind. Diversity	Ind. Novelty	Comm. Diversity	Comm. Novelty
FRediECH	0.152 ± 0.08	0.183 ± 0.28	**0.685** ± 0.24	0.888 ± 0.07	**0.992** ± 0.11	**0.927** ± 0.05	**0.938** ± 0.09	0.618 ± 0.08	**0.842** ± 0.13	**0.65** ± 0.04	**0.723** ± 0.08
Random	0.113** ± 0.04	0.053** ± 0.13	0.459** ± 0.2	0.732 ± 0.13	0.699** ± 0.17	0.726 ± 0.05	0.797 ± 0.16	0.462 ± 0.13	0.46** ± 0.2	0.457 ± 0.05	0.578 ± 0.22
Popularity	**0.281** ± 0.26	0.22 ± 0.25	**0.686** ± 0.26	0.369** ± 0.05	0.559** ± 0.2	0.391 ± 0.04	0.673 ± 0.19	0.223** ± 0.03	0.365** ± 0.24	0.231 ± 0.02	0.519 ± 0.3
Topology-based Adamic-Adar	0.27 ± 0.27	**0.285** ± 0.29	0.632 ± 0.26	0.359** ± 0.11	0.431** ± 0.12	0.517 ± 0.17	0.653 ± 0.18	0.247** ± 0.06	0.281** ± 0.07	0.304 ± 0.13	0.413 ± 0.21
Topology-based Jaccard	0.191 ± 0.17	0.249 ± 0.3	0.567 ± 0.24	0.364** ± 0.16	0.453** ± 0.14	0.592 ± 0.16	0.667 ± 0.14	0.272 ± 0.09	0.296** ± 0.08	0.371 ± 0.18	0.399 ± 0.12
Topology-based RA	0.272 ± 0.27	0.27 ± 0.29	0.642 ± 0.26	0.367** ± 0.11	0.436 ± 0.12	0.573 ± 0.18	0.643 ± 0.17	0.248** ± 0.06	0.282** ± 0.07	0.363 ± 0.21	0.387 ± 0.14
Topology-based Common Neighbours	0.259 ± 0.26	**0.302** ± 0.3	0.619 ± 0.25	0.356** ± 0.1	0.424** ± 0.11	0.564 ± 0.18	0.633 ± 0.16	0.25** ± 0.06	0.282** ± 0.07	0.363 ± 0.21	0.389 ± 0.14
Content-based Full Tweets	0.115** ± 0.04	0.053** ± 0.13	0.439** ± 0.18	0.726 ± 0.12	0.698** ± 0.17	0.727 ± 0.06	0.797 ± 0.17	0.449** ± 0.12	0.452** ± 0.2	0.434 ± 0.07	0.581 ± 0.24
Content-based 15 Tweets	0.246 ± 0.22	0.22 ± 0.27	0.584 ± 0.24	0.428** ± 0.16	0.491** ± 0.16	0.629** ± 0.13	0.69 ± 0.17	0.313** ± 0.12	0.331** ± 0.14	0.394 ± 0.13	0.42 ± 0.14
SCC	0.259 ± 0.25	0.252 ± 0.3	0.597 ± 0.24	0.35** ± 0.13	0.496** ± 0.15	0.469 ± 0.2	0.621 ± 0.2	0.24** ± 0.1	0.289** ± 0.1	0.317 ± 0.13	0.382 ± 0.13
CAM	0.228 ± 0.18	0.158 ± 0.22	0.513 ± 0.21	0.345** ± 0.12	0.424** ± 0.12	0.53 ± 0.17	0.647 ± 0.16	0.256** ± 0.06	0.275** ± 0.07	0.336 ± 0.12	0.39 ± 0.13
Implicit MF	0.271 ± 0.23	0.252 ± 0.28	0.654 ± 0.25	0.401** ± 0.12	0.435** ± 0.12	0.559 ± 0.12	0.643 ± 0.15	0.249** ± 0.07	0.271** ± 0.08	0.358 ± 0.15	0.418 ± 0.18
NeuralCF	0.251 ± 0.25	0.262 ± 0.31	0.579 ± 0.25	0.351** ± 0.11	0.419** ± 0.12	0.566 ± 0.22	0.647 ± 0.18	0.238** ± 0.07	0.273** ± 0.08	0.38 ± 0.16	0.407 ± 0.15
GraphRec	0.103** ± 0.02	0.183** ± 0.28	0.389** ± 0.13	**0.935** ± 0.11	0.842** ± 0.13	0.739 ± 0.07	0.895 ± 0.11	**0.662** ± 0.07	0.612** ± 0.12	0.634** ± 0.07	0.65 ± 0.07
Mult-VAE	0.26 ± 0.22	0.254 ± 0.28	0.627 ± 0.24	0.413** ± 0.12	0.433** ± 0.11	0.607 ± 0.13	0.637 ± 0.12	0.266** ± 0.1	0.277** ± 0.09	0.394 ± 0.15	0.41 ± 0.14
Original graph	- -	- -	- -	0.325 ± 0.2	0.418 ± 0.14	0.581 ± 0.16	0.603 ± 0.12	0.214 ± 0.13	0.268 ± 0.09	0.345 ± 0.16	0.376 ± 0.13

Table 2. Relevance, diversity and novelty recommendation results comparison for $k = 10$.
** indicates statistically significant differences favouring FRediECH

Table), Modularity achieved the most relevant ones. In the case of CAM, the reported results correspond to those based on ImplicitMF, which achieved similar diversity and novelty than those with Adamic-Adar, but higher relevance. These observations agree with those in [25]. Nonetheless, the achieved diversity and novelty of both techniques were close to those of the original network, thus failing to significantly improve the diversity or novelty of recommendations.

FRediECH achieved the highest diversity and novelty results, followed by GraphRec. The only exception was individual diversity, in which GraphRec outperformed FRediECH. In terms of relevance, FRediECH also significantly outperformed GraphRec. As the Table shows, most of the differences favouring FRediECH were statistically significant. FRediECH had lower precision and recall than several of the baselines, showing small significant differences in their favour. nDCG results showed that even when recommending non relevant users, the relevant ones were ranked high.

In average, the diversity/novelty improvements of FRediECH regarding the simpler and state-of-the-art baselines were of 47% and 44%, respectively. In both cases, the maximum improvements were observed for individual novelty (60%). A similar tendency was observed for the improvements regarding the original structure of the graph, with a maximum improvement of 67% for individual novelty. In general, the novelty of recommendations was higher than their diversity. This implies that even though recommendations could include similar users, such users were different from the already known ones, thus extending users' neighbourhood. In particular, novelty was higher for the structural distance, which implies that recommended users belong to other communities, but still shared similar content. This is

	Precision	Recall	nDCG	Structural dissimilarities				Content-based dissimilarities			
				Ind. Diversity	Ind. Novelty	Comm. Diversity	Comm. Novelty	Ind. Diversity	Ind. Novelty	Comm. Diversity	Comm. Novelty
FRediECH	0.152 ± 0.08	0.183 ± 0.28	0.685 ± 0.24	**0.888** ± 0.07	**0.992** ± 0.11	**0.927** ± 0.05	0.938 ± 0.09	**0.618** ± 0.08	0.842 ± 0.13	**0.65** ± 0.04	0.723 ± 0.08
FRediECH$_{\text{NO-NS}}$	0.149** ± 0.07	0.172 ± 0.28	0.553** ± 0.27	0.726 ± 0.15	0.82** ± 0.13	0.845 ± 0.03	0.852 ± 0.12	0.538 ± 0.16	0.58** ± 0.14	0.6 ± 0.04	0.633 ± 0.15
FRediECH$_{\text{NO-WIDE}}$	0.152 ± 0.08	0.189 ± 0.29	0.685 ± 0.24	**0.888** ± 0.07	**0.993** ± 0.11	0.845 ± 0.03	0.966 ± 0.09	**0.618** ± 0.08	0.853 ± 0.15	0.606 ± 0.06	**0.76** ± 0.1
FRediECH$_{\text{NO-WIDE-NO-NS}}$	0.134 ± 0.06	0.172 ± 0.28	0.609 ± 0.28	0.597** ± 0.09	0.82** ± 0.13	0.728 ± 0.07	0.852 ± 0.12	0.484 ± 0.15	0.58 ± 0.14	0.6 ± 0.04	0.633 ± 0.15
FRediECH$_{\text{DUAL}}$	**0.169** ± 0.1	0.192 ± 0.28	0.561 ± 0.25	0.73 ± 0.16	0.937** ± 0.12	0.762 ± 0.07	0.912 ± 0.09	0.589 ± 0.15	0.777** ± 0.14	0.646 ± 0.06	0.718 ± 0.07
FRediECH$_{\text{NO-BERT}}$	**0.16** ± 0.09	0.193 ± 0.28	0.56 ± 0.26	0.596** ± 0.12	0.97** ± 0.13	0.708 ± 0.06	0.936 ± 0.09	0.529 ± 0.16	0.804** ± 0.15	0.627 ± 0.06	0.735 ± 0.07
FRediECH$_{\text{MENTION}}$	0.14 ± 0.07	0.182 ± 0.28	0.544 ± 0.28	0.541** ± 0.1	**0.993** ± 0.11	0.698 ± 0.07	0.93 ± 0.11	0.51** ± 0.12	0.838 ± 0.13	0.627 ± 0.06	0.71 ± 0.09
FRediECH$_{\text{REPLY}}$	0.103** ± 0.02	**0.203** ± 0.31	**0.732** ± 0.34	0.509** ± 0.12	**0.99** ± 0.09	0.643 ± 0.08	**0.99** ± 0.11	0.396** ± 0.12	**0.916** ± 0.16	0.543 ± 0.08	**0.761** ± 0.15
FRediECH$_{\text{RETWEET}}$	0.146 ± 0.07	0.193 ± 0.29	0.567 ± 0.28	0.646** ± 0.15	**0.99** ± 0.11	0.724 ± 0.07	0.941 ± 0.09	0.525** ± 0.15	0.838 ± 0.13	0.625 ± 0.06	0.728 ± 0.08
FRediECH$_{\text{MENTION-REPLY}}$	0.136 ± 0.06	0.176 ± 0.27	0.547 ± 0.29	0.651 ± 0.15	**0.99** ± 0.11	0.741 ± 0.06	0.932 ± 0.1	0.531** ± 0.15	0.845 ± 0.13	0.633 ± 0.06	0.714 ± 0.08
FRediECH$_{\text{MENTION-RETWEET}}$	0.159 ± 0.09	0.184 ± 0.28	0.542 ± 0.26	0.627** ± 0.08	0.96 ± 0.12	0.732 ± 0.07	0.916 ± 0.1	0.554** ± 0.14	0.805** ± 0.14	0.645 ± 0.06	0.725 ± 0.07
FRediECH$_{\text{REPLY-RETWEET}}$	**0.162** ± 0.09	0.183 ± 0.27	0.55 ± 0.26	0.69 ± 0.12	0.947** ± 0.12	0.762 ± 0.04	0.909 ± 0.08	0.579 ± 0.13	0.783** ± 0.12	0.648 ± 0.04	0.704 ± 0.05

Table 3. Ablation study of FRediECH
** indicates statistically significant differences favouring FRediECH

expected given the topic focused nature of the data collection. Nonetheless, similar novelty results were observed when considering a community division based on user leaning. Finally, unlike group recommendation systems that aim to suggest users satisfying the individual preferences of all members in a group, echo chamber-aware recommendations aim to recommend users that disrupt group homogeneity and thus would not be of interest for all members. In this sense, the group novelty scores indicate that both FRediECH and GraphRec were able to produce recommendations from outside the community of influence, thus fostering the openness of such community.

In summary, results showed that FRediECH (despite the trade-off with precision) satisfactory increased the diversity and novelty of recommendations, when measured in terms of individual users and the communities they belong to.

5.2 RQ2. Ablation Study

To verify the effectiveness and contribution of each component of FRediECH, we performed an ablation study considering the following variants:

- FRediECH$_{\text{NO-NS}}$. Remove the negative sampling from the described model.
- FRediECH$_{\text{NO-WIDE}}$. Remove the wide component of the architecture.
- FRediECH$_{\text{NO-WIDE-NO-NS}}$. Remove the wide component of the architecture and the negative sampling.
- FRediECH$_{\text{DUAL}}$. Different embeddings are used for representing the target and recommended users, which are processed by different GCNs.
- FRediECH$_{\text{NO-BERT}}$. Remove the textual embeddings from the described model.
- FRediECH$_{\text{MENTION}}$, FRediECH$_{\text{REPLY}}$ and FRediECH$_{\text{RETWEET}}$. Only one interaction type is considered.
- FRediECH$_{\text{MENTION-REPLY}}$, FRediECH$_{\text{MENTION-RETWEET}}$, and FRediECH$_{\text{REPLY-RETWEET}}$. The described model includes pairs of interactions.

For the purpose of these analyses, relations were removed from both the training and test sets, and a new model was trained from scratch for each evaluation. Table 3 presents the results for the different `FRediECH` variations. When varying aspects related to `FRediECH` architecture (i.e., the first 4 variants), it can be observed that relevance was not greatly affected by the modifications. For example, even when $\texttt{FRediECH}_{\text{DUAL}}$ achieved higher recall and nDCG than `FRediECH`, those improvements were not significant. Instead, diversity and novelty showed more variability. First, negative sampling had a positive and significant effect over the diversity and novelty of recommendations, as the results of $\texttt{FRediECH}_{\text{NO-NS}}$ show. Second, removing the Wide component ($\texttt{FRediECH}_{\text{NO-WIDE-NO-NS}}$) achieved similar individual diversity and novelty, while slightly decreasing group diversity and increasing group novelty. Third, removing the Wide component and the negative sampling ($\texttt{FRediECH}_{\text{NO-WIDE-NO-NS}}$) significantly decreased both the diversity and novelty of recommendations. According to the results, the impact over diversity of removing the negative sampling was higher for the non-Wide model. Fourth, even though $\texttt{FRediECH}_{\text{DUAL}}$ increased the relevance of recommendations, it significantly decreased their diversity and novelty. The effects were higher for structural diversity and novelty. This shows that considering extra embeddings for user representation could introduce noise in the recommendation process.

Regarding the variations of the data fed to the model (i.e., the last 7 variations), even though the dataset was focused on only one topic, including content allowed to significantly increase the novelty and diversity of recommendations. In particular, differences were higher for the structural analysis, showing that content helped `FRediECH` in searching for recommendations farther away from users.

The distribution of interactions differed according to their type, with replies accounting for the lowest quantity, and mentions and retweets presenting a more balanced distribution. In general, only considering one interaction significantly decreased diversity and novelty. The only exception was $\texttt{FRediECH}_{\text{REPLY}}$ that improved relevance and novelty, while decreasing diversity. In this sense, even though recommendations were novel regarding the known user interactions, they seemed to be already known in their community, which in the long run might lead to its strengthen. In addition, differences between $\texttt{FRediECH}_{\text{REPLY}}$, $\texttt{FRediECH}_{\text{MENTIONS}}$ and $\texttt{FRediECH}_{\text{RETWEETS}}$, favouring $\texttt{FRediECH}_{\text{REPLY}}$ might imply that despite involving fewer interactions, replies might carry more relevant information (or indicators of future interactions) than mentions and retweets.

When considering pairs of interactions, precision and recall slightly increased for those combinations including retweets, while diversity and novelty decreased. Even though $\texttt{FRediECH}_{\text{REPLY}}$ improved community content novelty, it was not the case for the combinations including it. Thus, it can be stated that interactions carry different weights in fostering user interaction, implying the need for different mechanisms for adequately leveraging them.

In summary, results showed that each component of `FRediECH` significantly contributed to its performance. Nonetheless, more studies of the interaction types and their interplay in the quality of recommendations are needed.

6 CONCLUSIONS

This work presented `FRediECH`, an echo chamber-aware recommender system aiming at recommending novel and diverse users in social networks. Unlike other approaches in the literature, `FRediECH` does not rely on explicitly discovering communities nor on knowing user protected features. Instead, it learns users and their implicit echo chamber or community representations to jointly optimize the relevance and diversity of recommendations. Results showed that `FRediECH`, despite the small variations in relevance, increased the diversity and novelty of recommendations, when compared to other works in the literature. These improvements were observed in relation to both individual users and the communities they belong to. In addition, the contribution of each component was verified in an ablation study.

Several aspects could be tackled as future works. First, the study was based on a politics-focused dataset on a specific social network. In this sense, additional evaluations should be performed over different datasets varying the domain and time period to truly assess the usefulness and generalizability of `FRediECH`. Second, as the ablation study showed, more analyses should be made regarding the relevance of each type of interaction, and their contribution to the final recommendations. Third, `FRediECH` could be enriched to provide explanations to recommendations, and thus better guide users in broadening their interactions. Finally, a qualitative study should be perform to assess the perceived relevance, usefulness and diversity of recommendations.

REFERENCES

[1] Pablo Barberá, John T Jost, Jonathan Nagler, Joshua A Tucker, and Richard Bonneau. 2015. Tweeting from left to right: Is online political communication more than an echo chamber? *Psychological science* 26, 10 (2015), 1531–1542.

[2] Vincent D Blondel, Jean-Loup Guillaume, Renaud Lambiotte, and Etienne Lefebvre. 2008. Fast unfolding of communities in large networks. *Journal of statistical mechanics: theory and experiment* 2008, 10 (2008), P10008.

[3] Leticia Bode and Emily K Vraga. 2015. In related news, that was wrong: The correction of misinformation through related stories functionality in social media. *Journal of Communication* 65, 4 (2015), 619–638.

[4] Heng-Tze Cheng, Levent Koc, Jeremiah Harmsen, Tal Shaked, Tushar Chandra, Hrishi Aradhye, Glen Anderson, Greg Corrado, Wei Chai, Mustafa Ispir, et al. 2016. Wide & deep learning for recommender systems. In *Proceedings of the 1st workshop on deep learning for recommender systems*. 7–10.

[5] Uthsav Chitra and Christopher Musco. 2020. Analyzing the Impact of Filter Bubbles on Social Network Polarization. In *Proceedings of the 13th International Conference on Web Search and Data Mining* (Houston, TX, USA) *(WSDM '20)*. ACM, New York, NY, USA, 115–123. https://doi.org/10.1145/3336191.3371825

[6] Wesley Cota, Silvio C Ferreira, Romualdo Pastor-Satorras, and Michele Starnini. 2019. Quantifying echo chamber effects in information spreading over political communication networks. *EPJ Data Science* 8, 1 (2019), 1–13.

[7] J. Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. In *NAACL-HLT*.

[8] Arthur Edwards. 2013. (How) do participants in online discussion forums create 'echo chambers'? The inclusion and exclusion of dissenting voices in an online discussion forum about climate change. *Journal of Argumentation in Context* 2 (2013), 127–150.

[9] Wenqi Fan, Yao Ma, Qing Li, Yuan He, Eric Zhao, Jiliang Tang, and Dawei Yin. 2019. Graph Neural Networks for Social Recommendation. In *The World Wide Web Conference* (San Francisco, CA, USA) *(WWW '19)*. ACM, New York, NY, USA, 417–426. https://doi.org/10.1145/3308558.3313488

[10] Miriam Fernández and Alejandro Bellogín. 2020. Recommender Systems and Misinformation: The Problem or the Solution?. In *Proceedings of the Workshop on Online Misinformation- and Harm-Aware Recommender Systems co-located with 14th ACM Conference on Recommender Systems (RecSys 2020), Rio de Janeiro, Brazil, September 25, 2020 (CEUR Workshop Proceedings, Vol. 2758)*, Antonela Tommasel, Daniela Godoy, and Arkaitz Zubiaga (Eds.). CEUR-WS.org, 40–50.

[11] Kiran Garimella, Gianmarco De Francisci Morales, Aristides Gionis, and Michael Mathioudakis. 2018. Political Discourse on Social Media: Echo Chambers, Gatekeepers, and the Price of Bipartisanship. In *Proceedings of the 2018 World Wide Web Conference* (Lyon, France) *(WWW '18)*. International World Wide Web Conferences, Republic and Canton of Geneva, CHE, 913–922. https://doi.org/10.1145/3178876.3186139

[12] Quentin Grossetti, Cédric du Mouza, and Nicolas Travers. 2019. Community-Based Recommendations on Twitter: Avoiding the Filter Bubble. In *Web Information Systems Engineering – WISE 2019*, Reynold Cheng, Nikos Mamoulis, Yizhou Sun, and Xin Huang (Eds.). Springer International Publishing, Cham, 212–227.

[13] John Hannon, Mike Bennett, and Barry Smyth. 2010. Recommending Twitter Users to Follow Using Content and Collaborative Filtering Approaches. In *Proceedings of the Fourth ACM Conference on Recommender Systems* (Barcelona, Spain) *(RecSys '10)*. ACM, New York, NY, USA, 199–206. https://doi.org/10.1145/1864708.1864746

[14] Xiangnan He, Lizi Liao, Hanwang Zhang, Liqiang Nie, Xia Hu, and Tat-Seng Chua. 2017. Neural Collaborative Filtering. In *Proceedings of the 26th International Conference on World Wide Web* (Perth, Australia) *(WWW '17)*. International World Wide Web Conferences, Republic and Canton of Geneva, CHE, 173–182. https://doi.org/10.1145/3038912.3052569

[15] Yifan Hu, Yehuda Koren, and Chris Volinsky. 2008. Collaborative filtering for implicit feedback datasets. In *2008 Eighth IEEE International Conference on Data Mining*. IEEE, 263–272.

[16] Thomas N. Kipf and Max Welling. 2017. Semi-Supervised Classification with Graph Convolutional Networks. In *5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24-26, 2017, Conference Track Proceedings*. OpenReview.net.

[17] Dawen Liang, Rahul G Krishnan, Matthew D Hoffman, and Tony Jebara. 2018. Variational autoencoders for collaborative filtering. In *Proceedings of the 2018 world wide web conference*. 689–698.

[18] David Liben-Nowell and Jon Kleinberg. 2003. The Link Prediction Problem for Social Networks. In *Proceedings of the Twelfth International Conference on Information and Knowledge Management* (New Orleans, LA, USA) *(CIKM '03)*. ACM, New York, NY, USA, 556–559. https://doi.org/10.1145/956863.956972

[19] Farzan Masrour, Tyler Wilson, Heng Yan, Pang-Ning Tan, and Abdol-Hossein Esfahanian. 2020. Bursting the Filter Bubble: Fairness-Aware Network Link Prediction. In *The Thirty-Fourth AAAI Conference on Artificial Intelligence, AAAI 2020, The Thirty-Second Innovative Applications of Artificial Intelligence Conference, IAAI 2020, The Tenth AAAI Symposium on Educational Advances in Artificial Intelligence, EAAI 2020, New York, NY, USA, February 7-12, 2020*. AAAI Press, 841–848.

[20] Tien T. Nguyen, Pik-Mai Hui, F. Maxwell Harper, Loren Terveen, and Joseph A. Konstan. 2014. Exploring the Filter Bubble: The Effect of Using Recommender Systems on Content Diversity. In *Proceedings of the 23rd International Conference on World Wide Web* (Seoul, Korea) *(WWW '14)*. ACM, New York, NY, USA, 677–686. https://doi.org/10.1145/2566486.2568012

[21] Jeffrey Pennington, Richard Socher, and Christopher D Manning. 2014. Glove: Global vectors for word representation. In *Proceedings of the 2014 conference on empirical methods in natural language processing (EMNLP)*. 1532–1543.

[22] Steffen Rendle, Christoph Freudenthaler, Zeno Gantner, and Lars Schmidt-Thieme. 2009. BPR: Bayesian Personalized Ranking from Implicit Feedback. In *Proceedings of the Twenty-Fifth Conference on Uncertainty in Artificial Intelligence* (Montreal, Quebec, Canada) *(UAI '09)*. AUAI Press, Arlington, Virginia, USA, 452–461.

[23] Marten Risius, Okan Aydingül, and Maximilian Haug. 2019. Towards an Understanding of Conspiracy echo Chambers on Facebook. In *27th European Conference on Information Systems - Information Systems for a Sharing Society, ECIS 2019, Stockholm and Uppsala, Sweden, June 8-14, 2019*, Jan vom Brocke, Shirley Gregor, and Oliver Müller (Eds.).

[24] Dimitris Sacharidis. 2019. Diversity and Novelty in Social-Based Collaborative Filtering. In *Proceedings of the 27th ACM Conference on User Modeling, Adaptation and Personalization* (Larnaca, Cyprus) *(UMAP '19)*. ACM, New York, NY, USA, 139–143. https://doi.org/10.1145/3320435.3320479

[25] Javier Sanz-Cruzado and Pablo Castells. 2018. Enhancing Structural Diversity in Social Networks by Recommending Weak Ties. In *Proceedings of the 12th ACM Conference on Recommender Systems* (Vancouver, British Columbia, Canada) *(RecSys '18)*. ACM, New York, NY, USA, 233–241. https://doi.org/10.1145/3240323.3240371

[26] Javier Sanz-Cruzado, Sofía M. Pepa, and Pablo Castells. 2018. Structural Novelty and Diversity in Link Prediction. In *Companion Proceedings of the The Web Conference 2018* (Lyon, France) *(WWW '18)*. International World Wide Web Conferences, Republic and Canton of Geneva, CHE, 1347–1351. https://doi.org/10.1145/3184558.3191576

[27] Le Wu, Peijie Sun, Yanjie Fu, Richang Hong, Xiting Wang, and Meng Wang. 2019. A Neural Influence Diffusion Model for Social Recommendation. In *Proceedings of the 42nd International ACM SIGIR Conference on Research and Development in Information Retrieval* (Paris, France) *(SIGIR'19)*. ACM, New York, NY, USA, 235–244. https://doi.org/10.1145/3331184.3331214