CONICET

ISISTAN

I want to break free! Recommending friends from outside the echo chamber

ANTONELA TOMMASEL

JUAN MANUEL RODRIGUEZ

DANIELA GODOY

Echo chambers & recommendations

- Echo chambers are related to situations in which **individuals only consume content or interact** with other users expressing their **same points of view**.
    - Selective exposure, biased assimilation, and group polarization.

- Echo chambers concern not only political discourses but also conspiracy theories, in which they could lead to a **stronger radicalization, seclusion from society** and **destructive actions**.

- **Recommender systems** play an important role as **mediators of information propagation**.
    - 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 .

- This from the fundamental concepts and assumptions on which recommenders are based on.

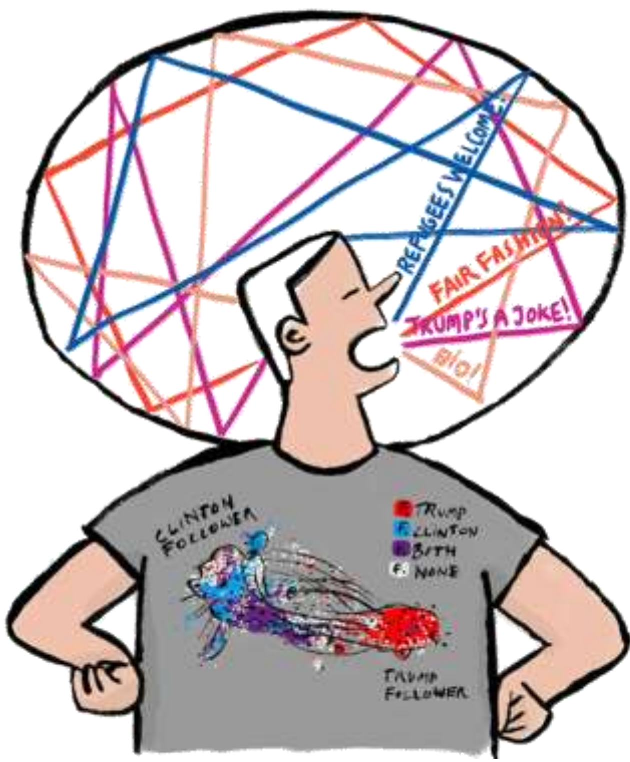

Echo chambers & recommendations

- Echo chambers are related to situations in which **individuals only consume content or interact** with other users expressing their **same points of view**.
  - Selective exposure, biased assimilation, and group polarization.

- Echo chambers concern not only political discourses but also conspiracy theories, in which they could lead to a **stronger radicalization, seclusion from society** and **destructive actions**.

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.

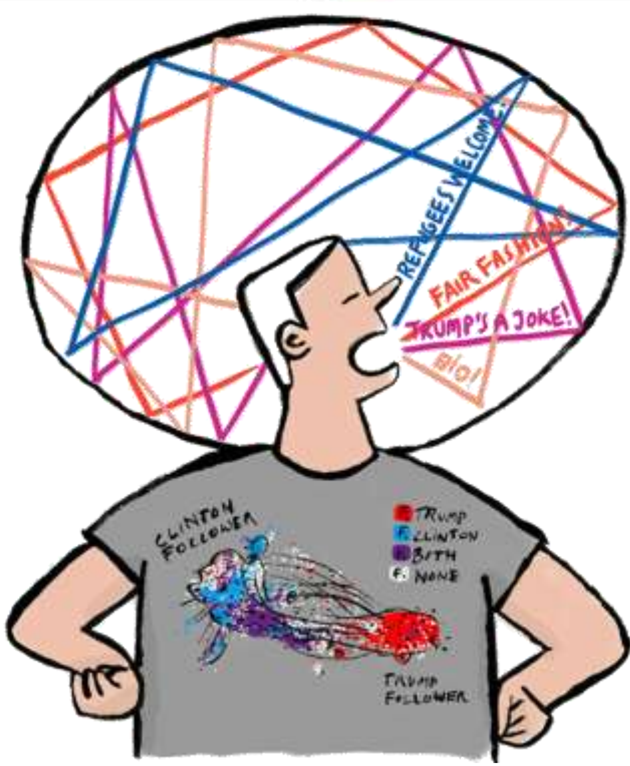

Echo-chamber aware recommendations

The problem

We tackle the **friend recommendation problem** by fostering **recommendation diversification** in an **echo chamber awareness** setting.

We rely on **implicitly modeling the echo chamber** membership of users to present them with **relevant friend recommendations from outside** the influence of their community.

Echo-chamber aware recommendations

The problem

We tackle the **friend recommendation problem** by fostering <u>recommendation diversification</u> in an <u>echo chamber awareness</u> setting.

We rely on <u>**implicitly modeling the echo chamber**</u> membership of users to present them with **relevant friend recommendations from outside** the influence of their community.

`Friend REcommenDer for breakIng Echo CHambers`

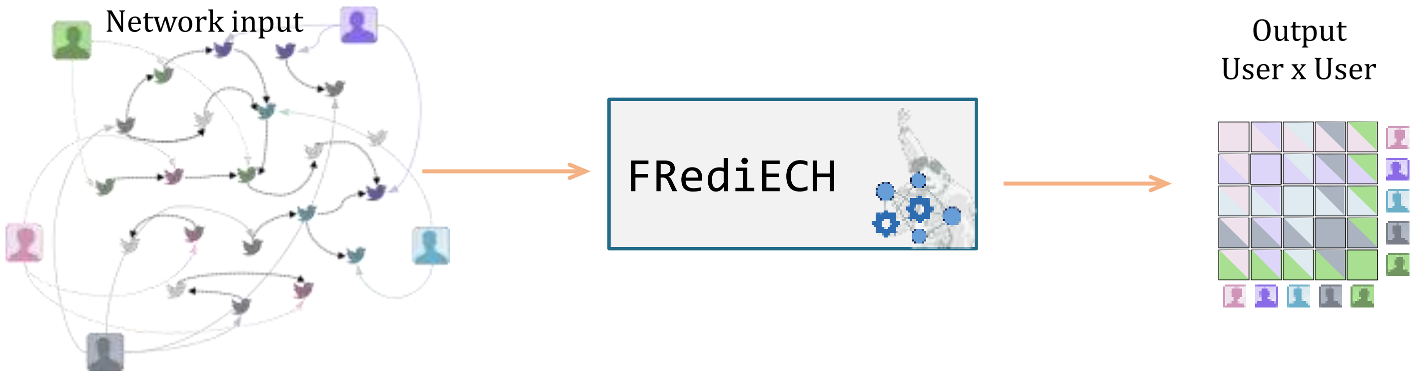

FRediECH

Friend REcommenDer for breakIng Echo CHambers

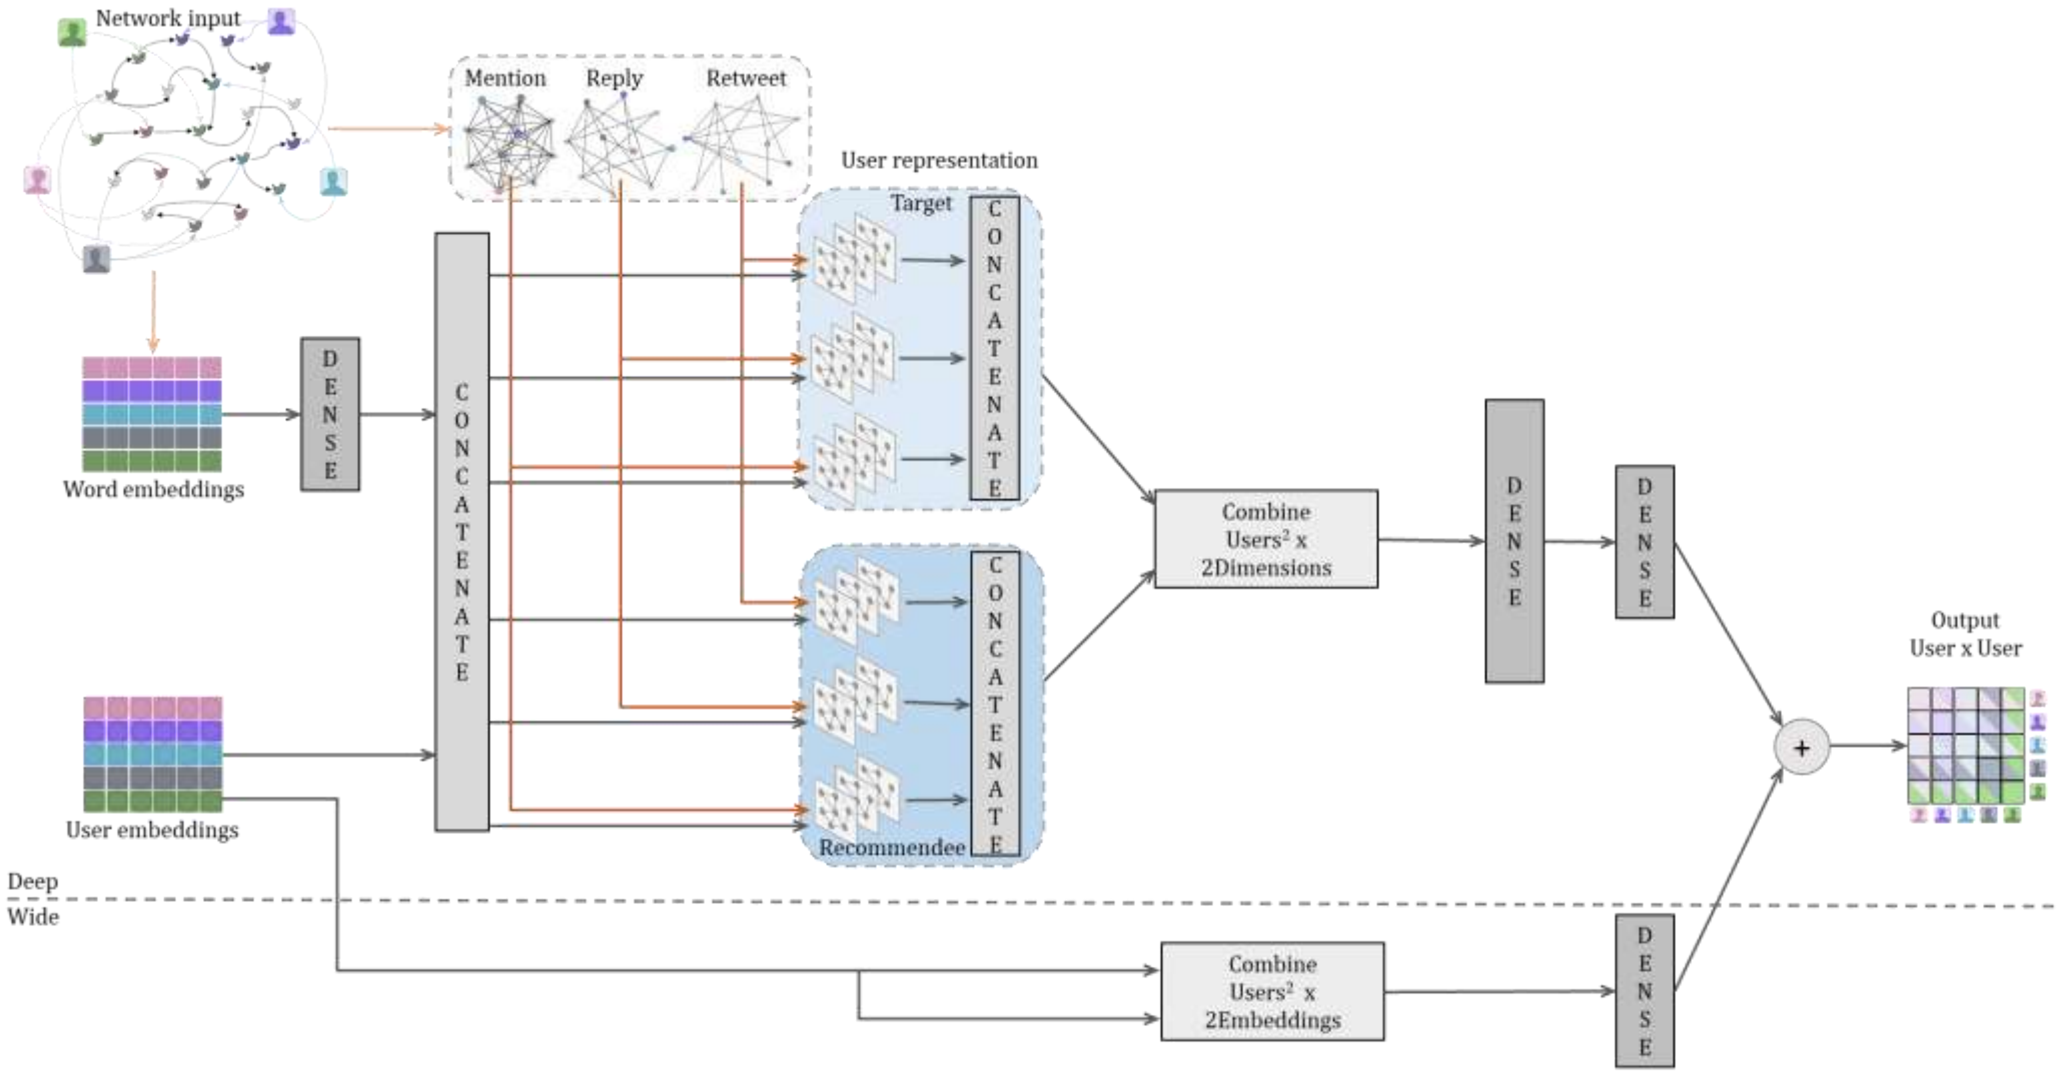

FRediECH

Friend REcommenDer for breakIng Echo CHambers

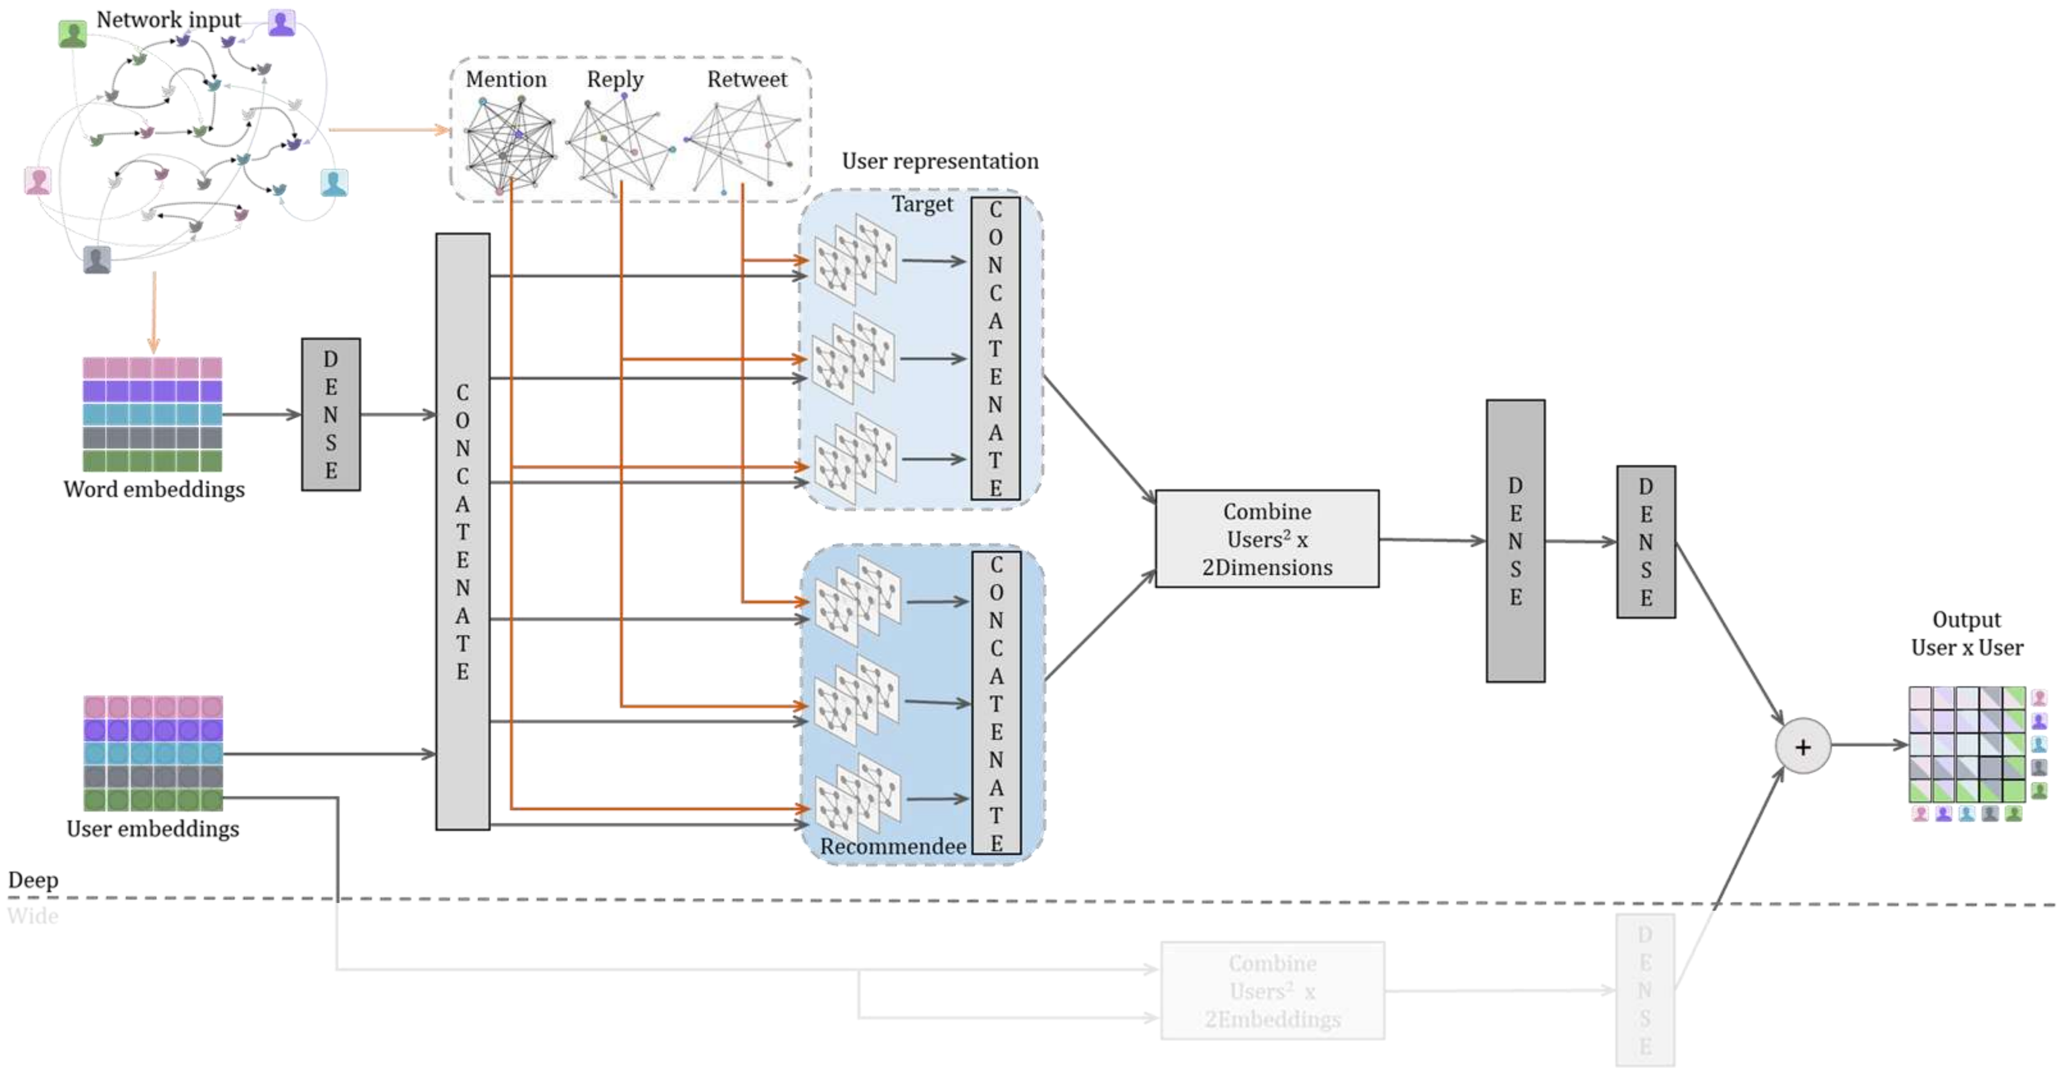

FRediECH

Friend REcommenDer for breakIng Echo CHambers

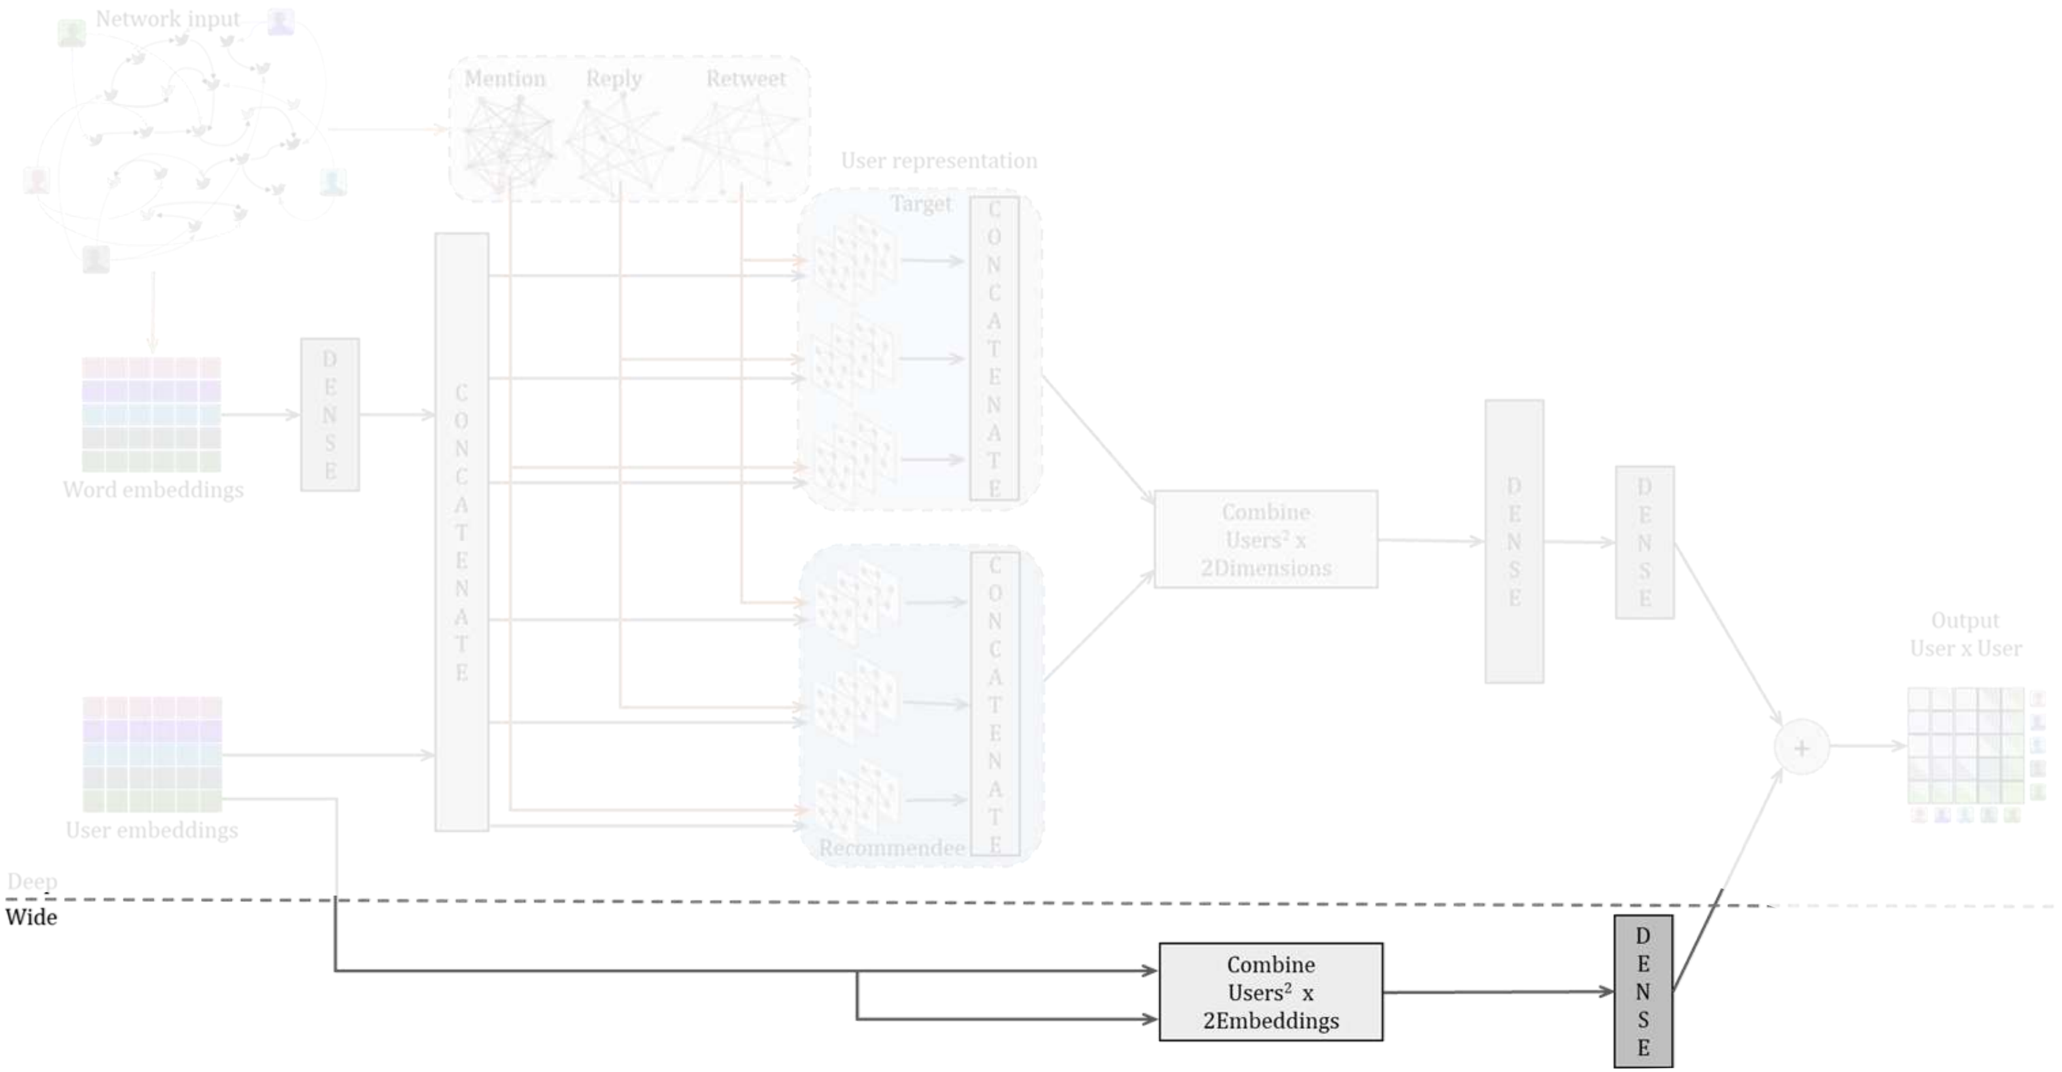

FRediECH

Friend REcommenDer for breakIng Echo CHambers

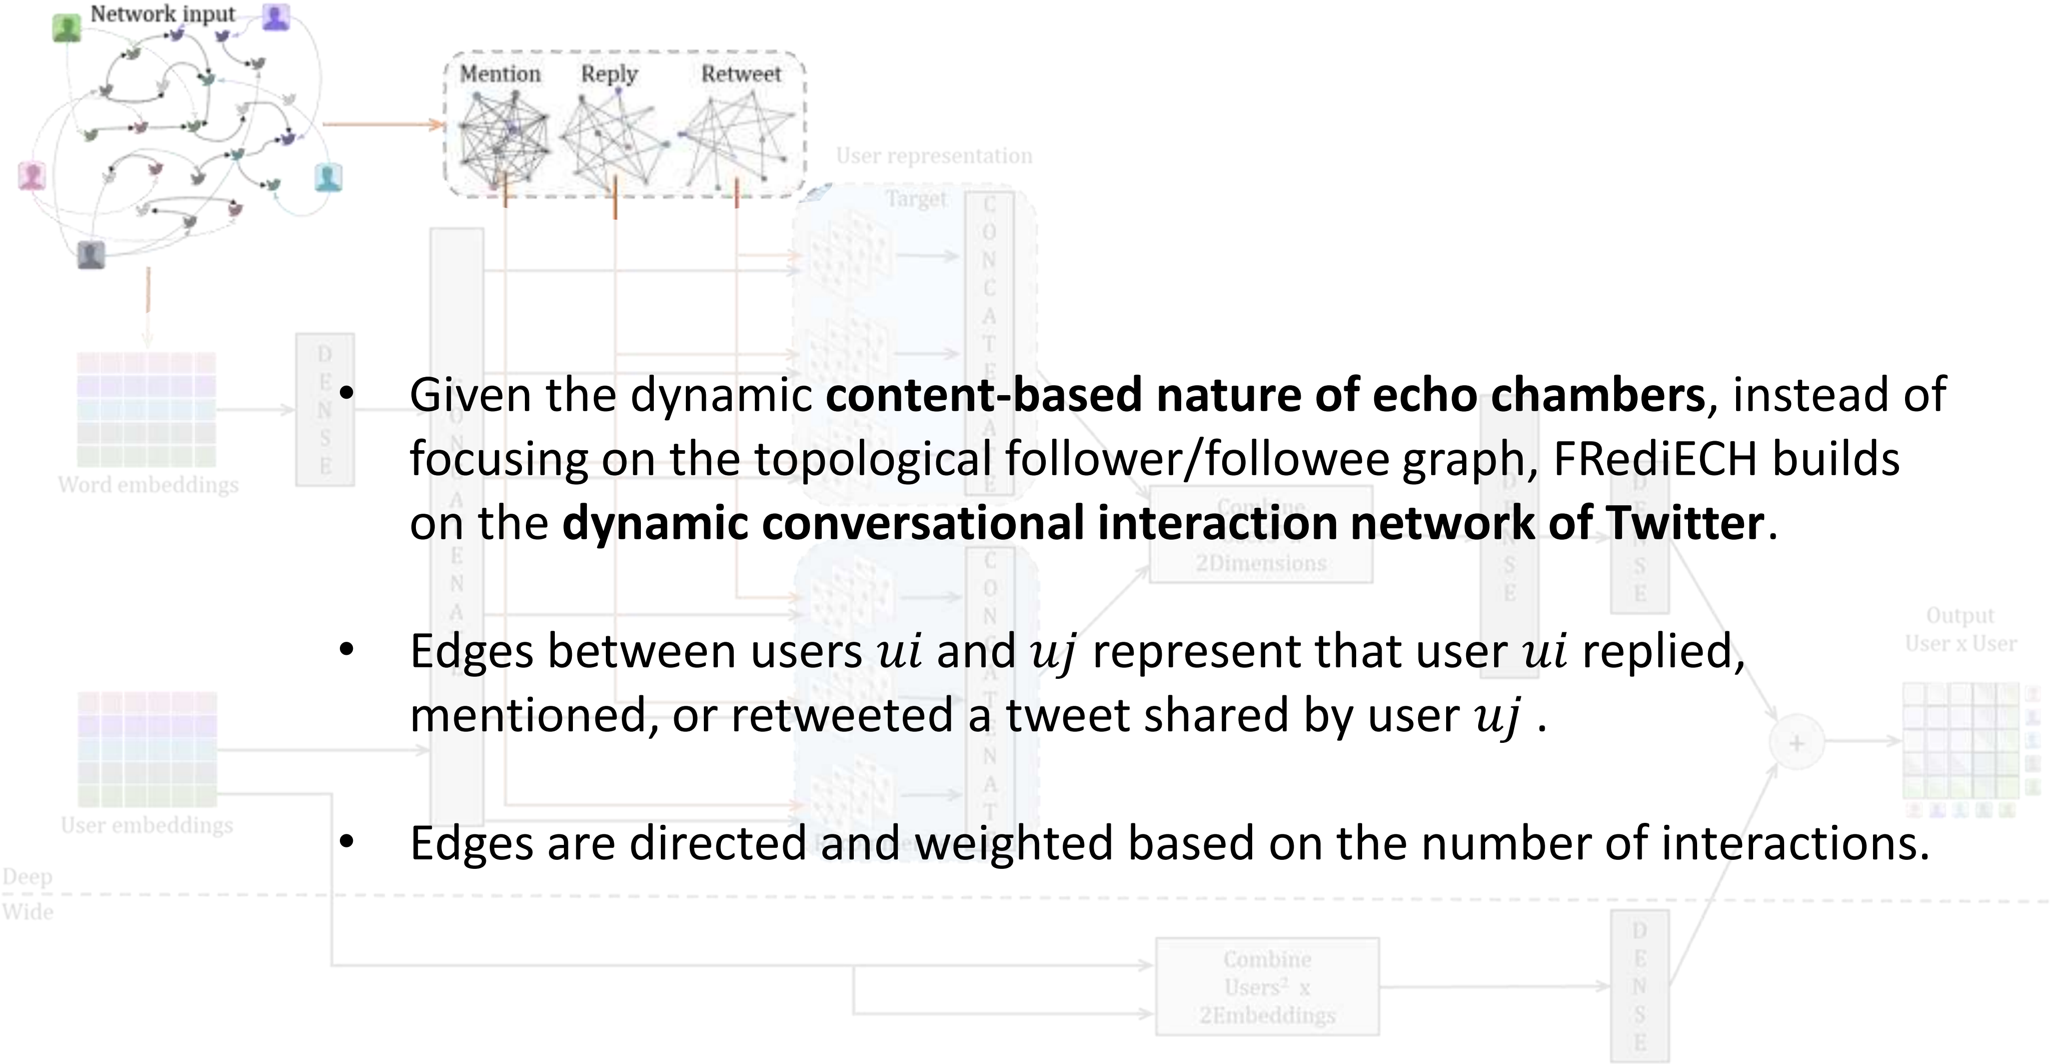

- Given the dynamic **content-based nature of echo chambers**, instead of focusing on the topological follower/followee graph, FRediECH builds on the **dynamic conversational interaction network of Twitter**.

- Edges between users ui and uj represent that user ui replied, mentioned, or retweeted a tweet shared by user uj .

- Edges are directed and weighted based on the number of interactions.

FRediECH

Friend REcommenDer for breakIng Echo CHambers

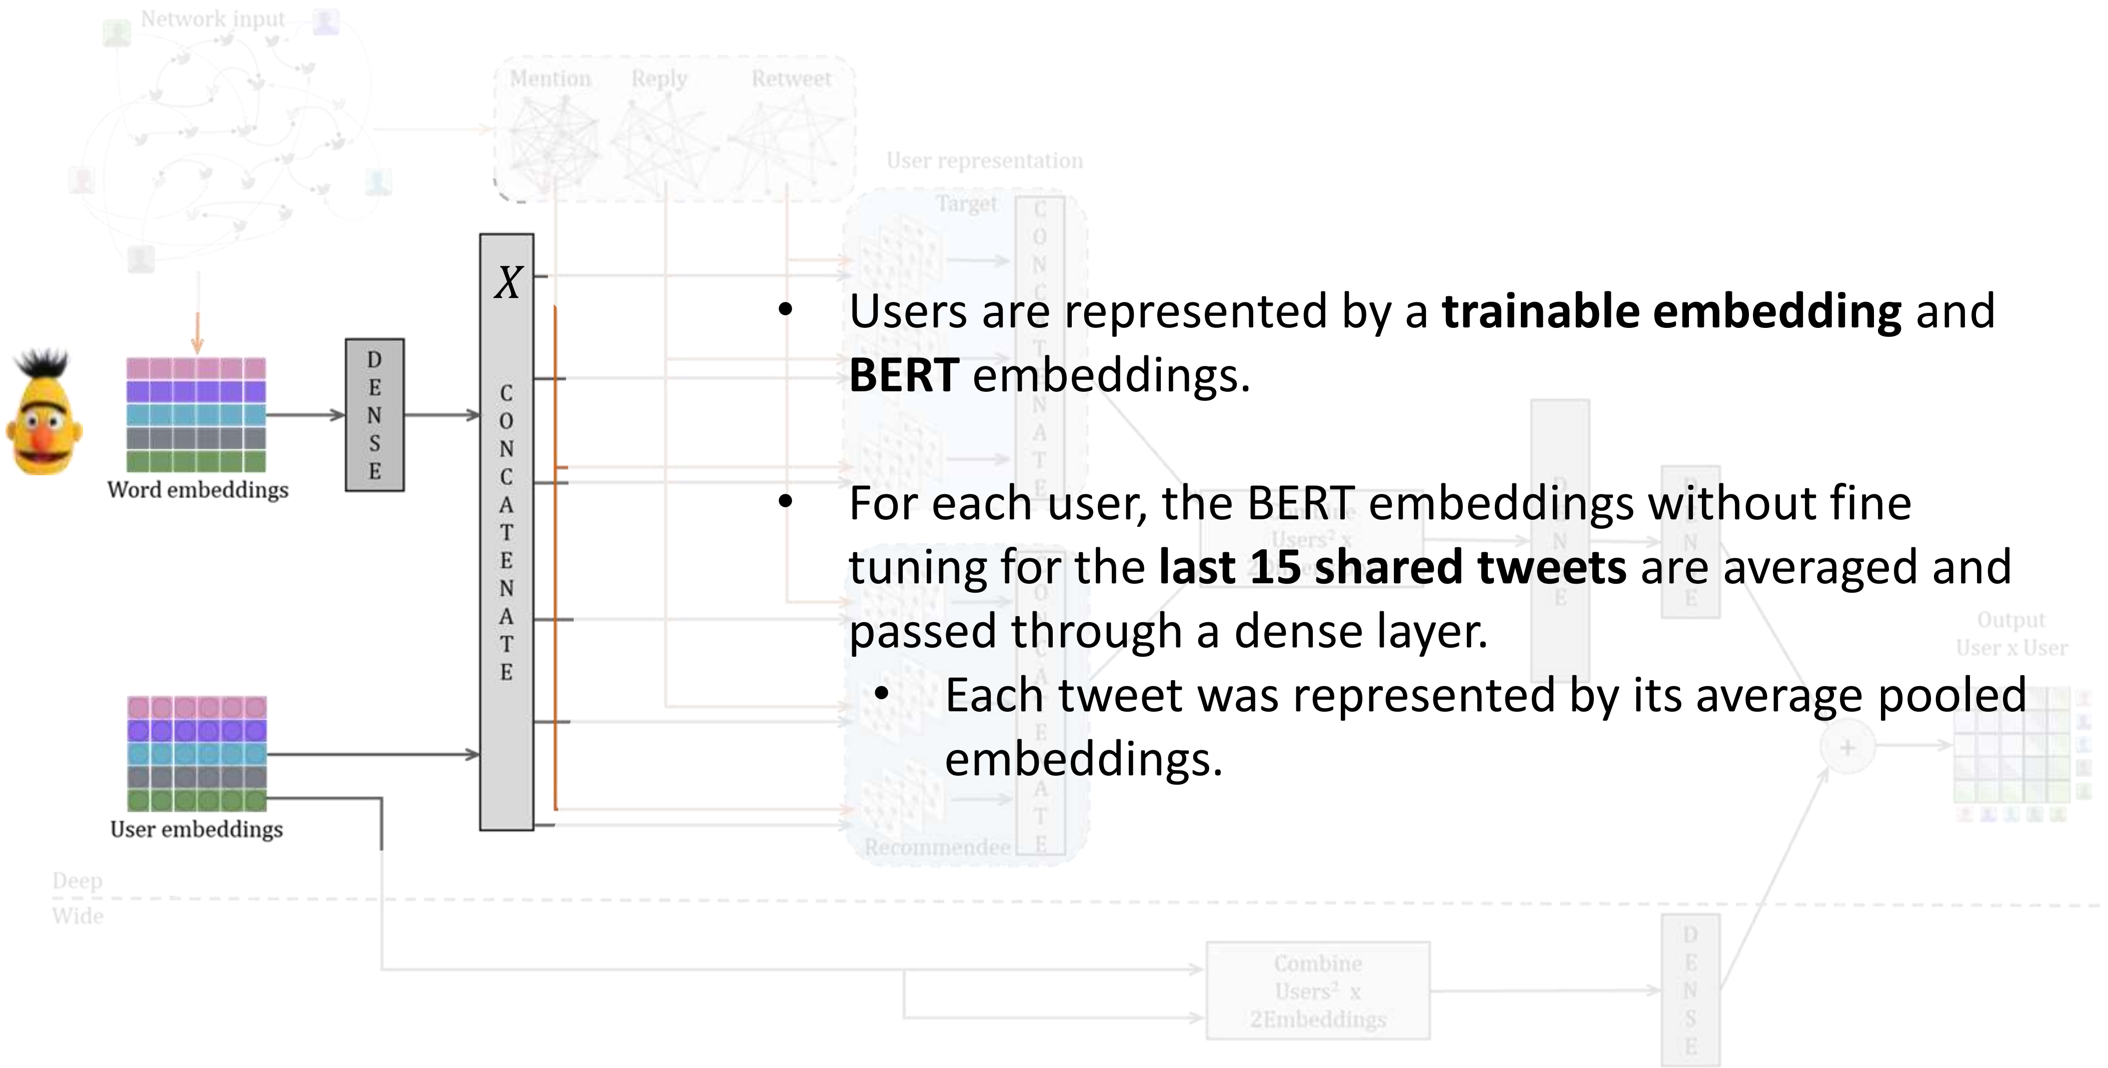

- Users are represented by a **trainable embedding** and **BERT** embeddings.

- For each user, the BERT embeddings without fine tuning for the **last 15 shared tweets** are averaged and passed through a dense layer.
  - Each tweet was represented by its average pooled embeddings.

FRediECH

Friend REcommenDer for breakIng Echo CHambers

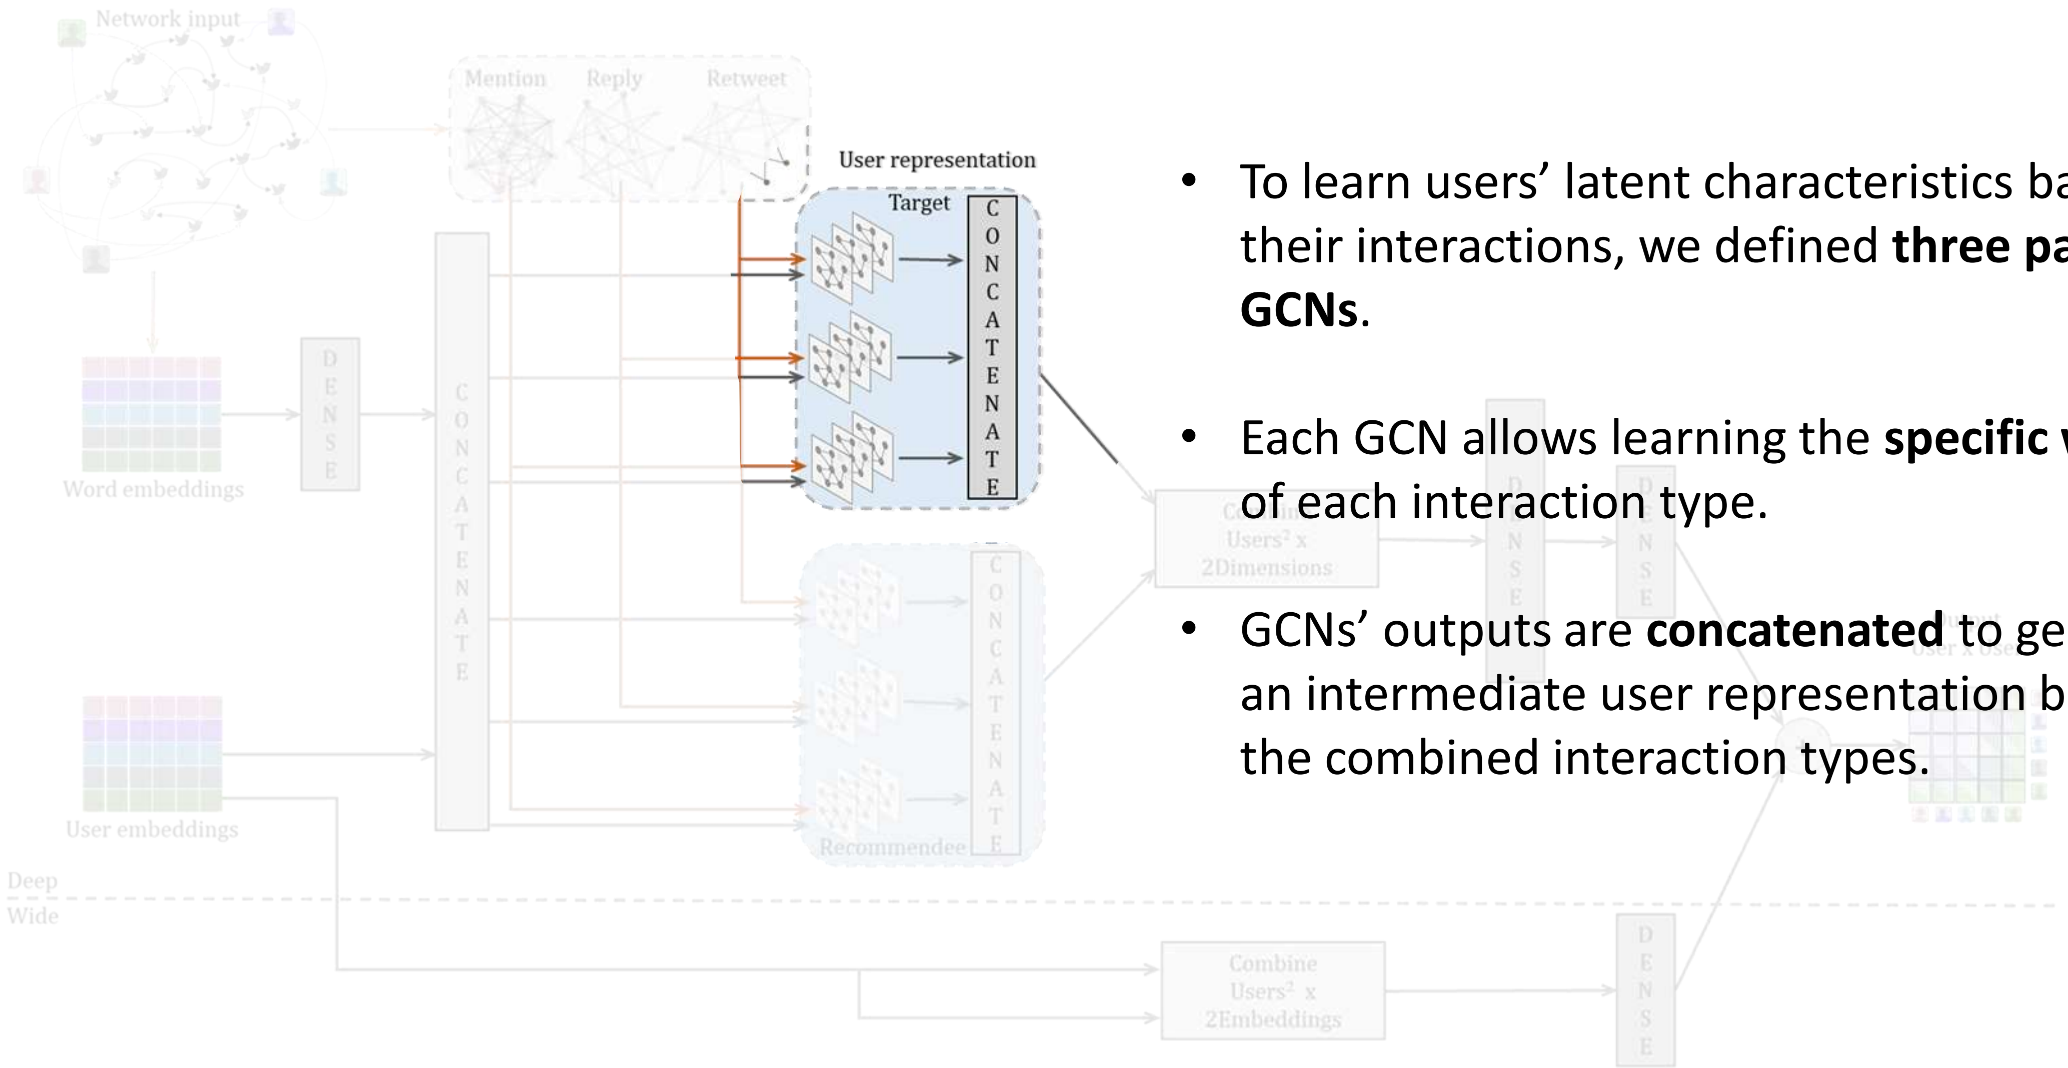

- To learn users' latent characteristics based on their interactions, we defined **three parallels GCNs**.

- Each GCN allows learning the **specific weights** of each interaction type.

- GCNs' outputs are **concatenated** to generate an intermediate user representation based on the combined interaction types.

FRediECH

Friend REcommenDer for breakIng Echo CHambers

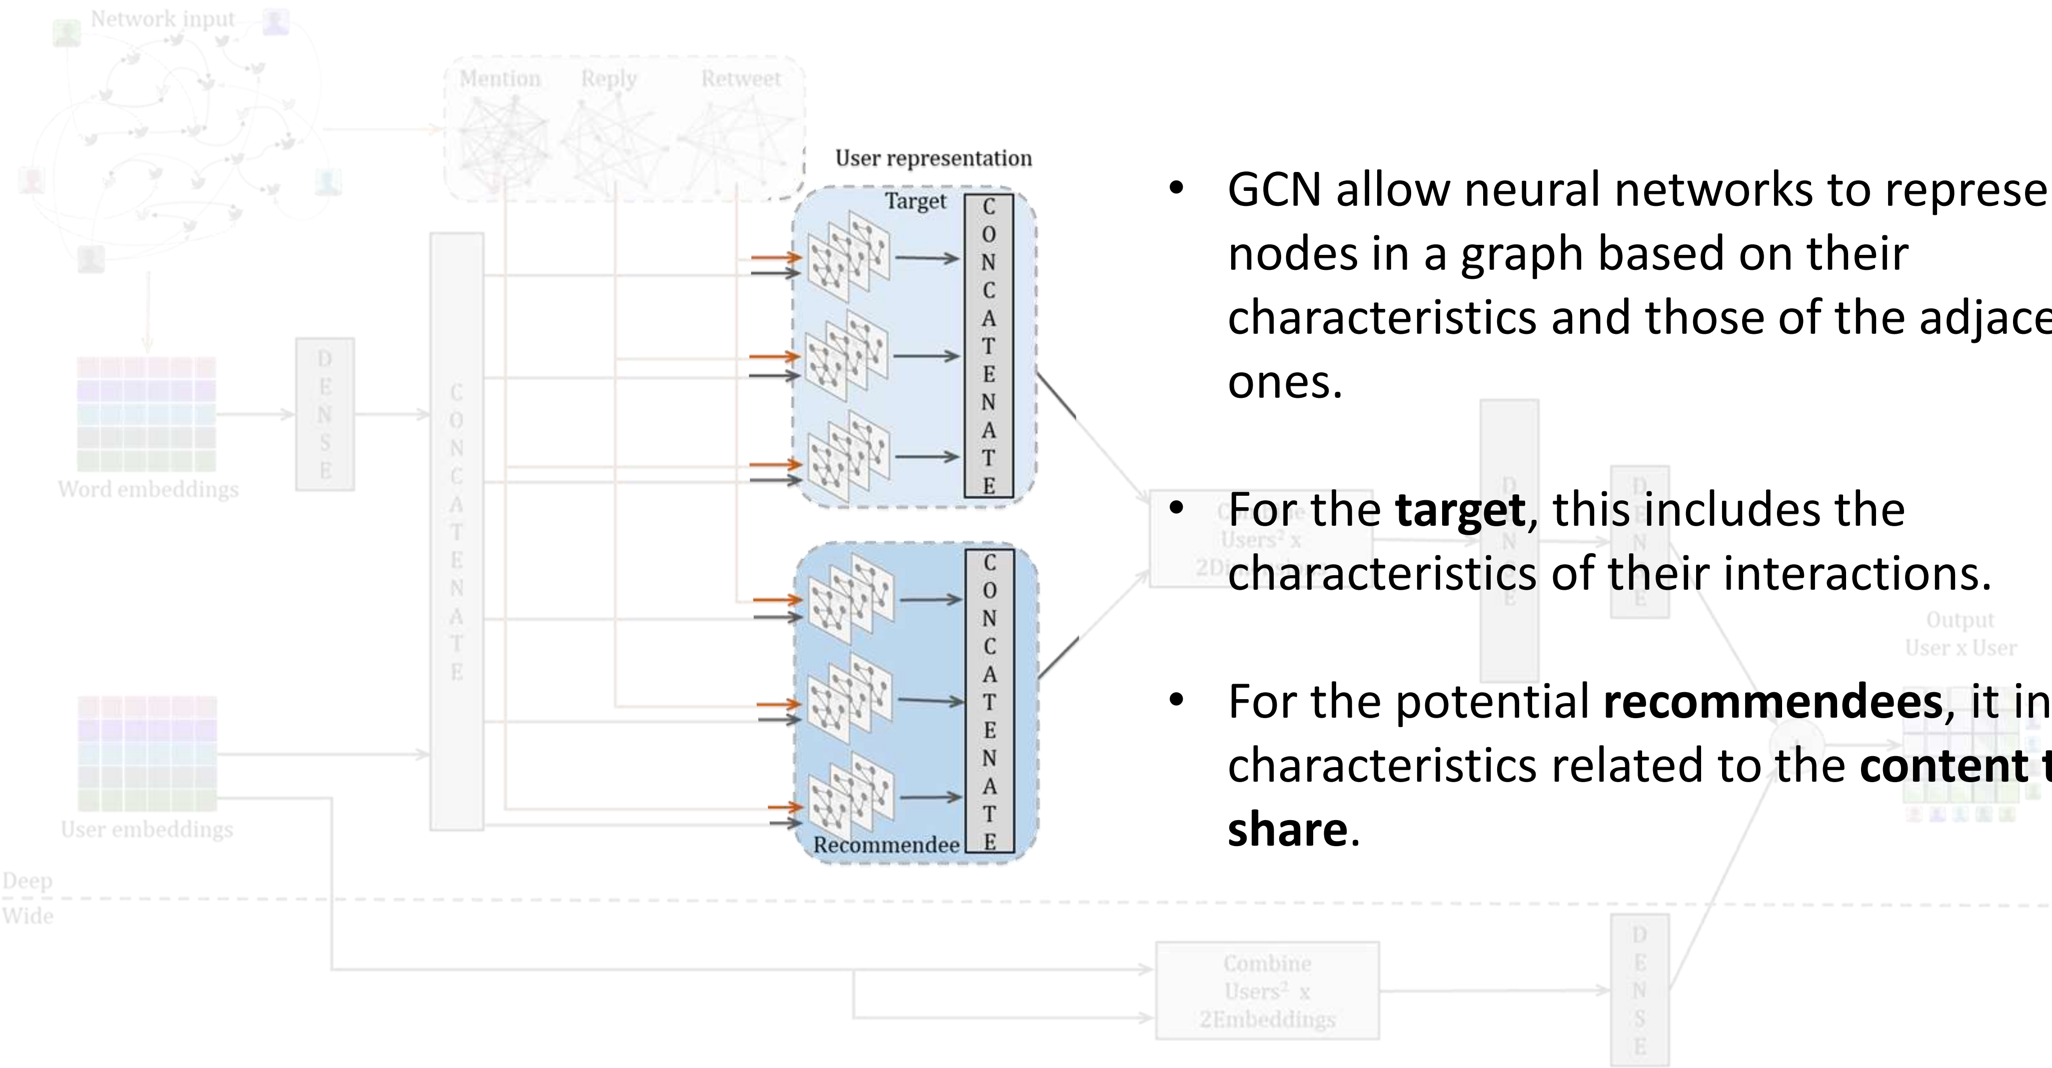

- GCN allow neural networks to represent nodes in a graph based on their characteristics and those of the adjacent ones.
- For the **target**, this includes the characteristics of their interactions.
- For the potential **recommendees**, it includes characteristics related to the **content they share**.

FRediECH

Friend REcommenDer for breakIng Echo CHambers

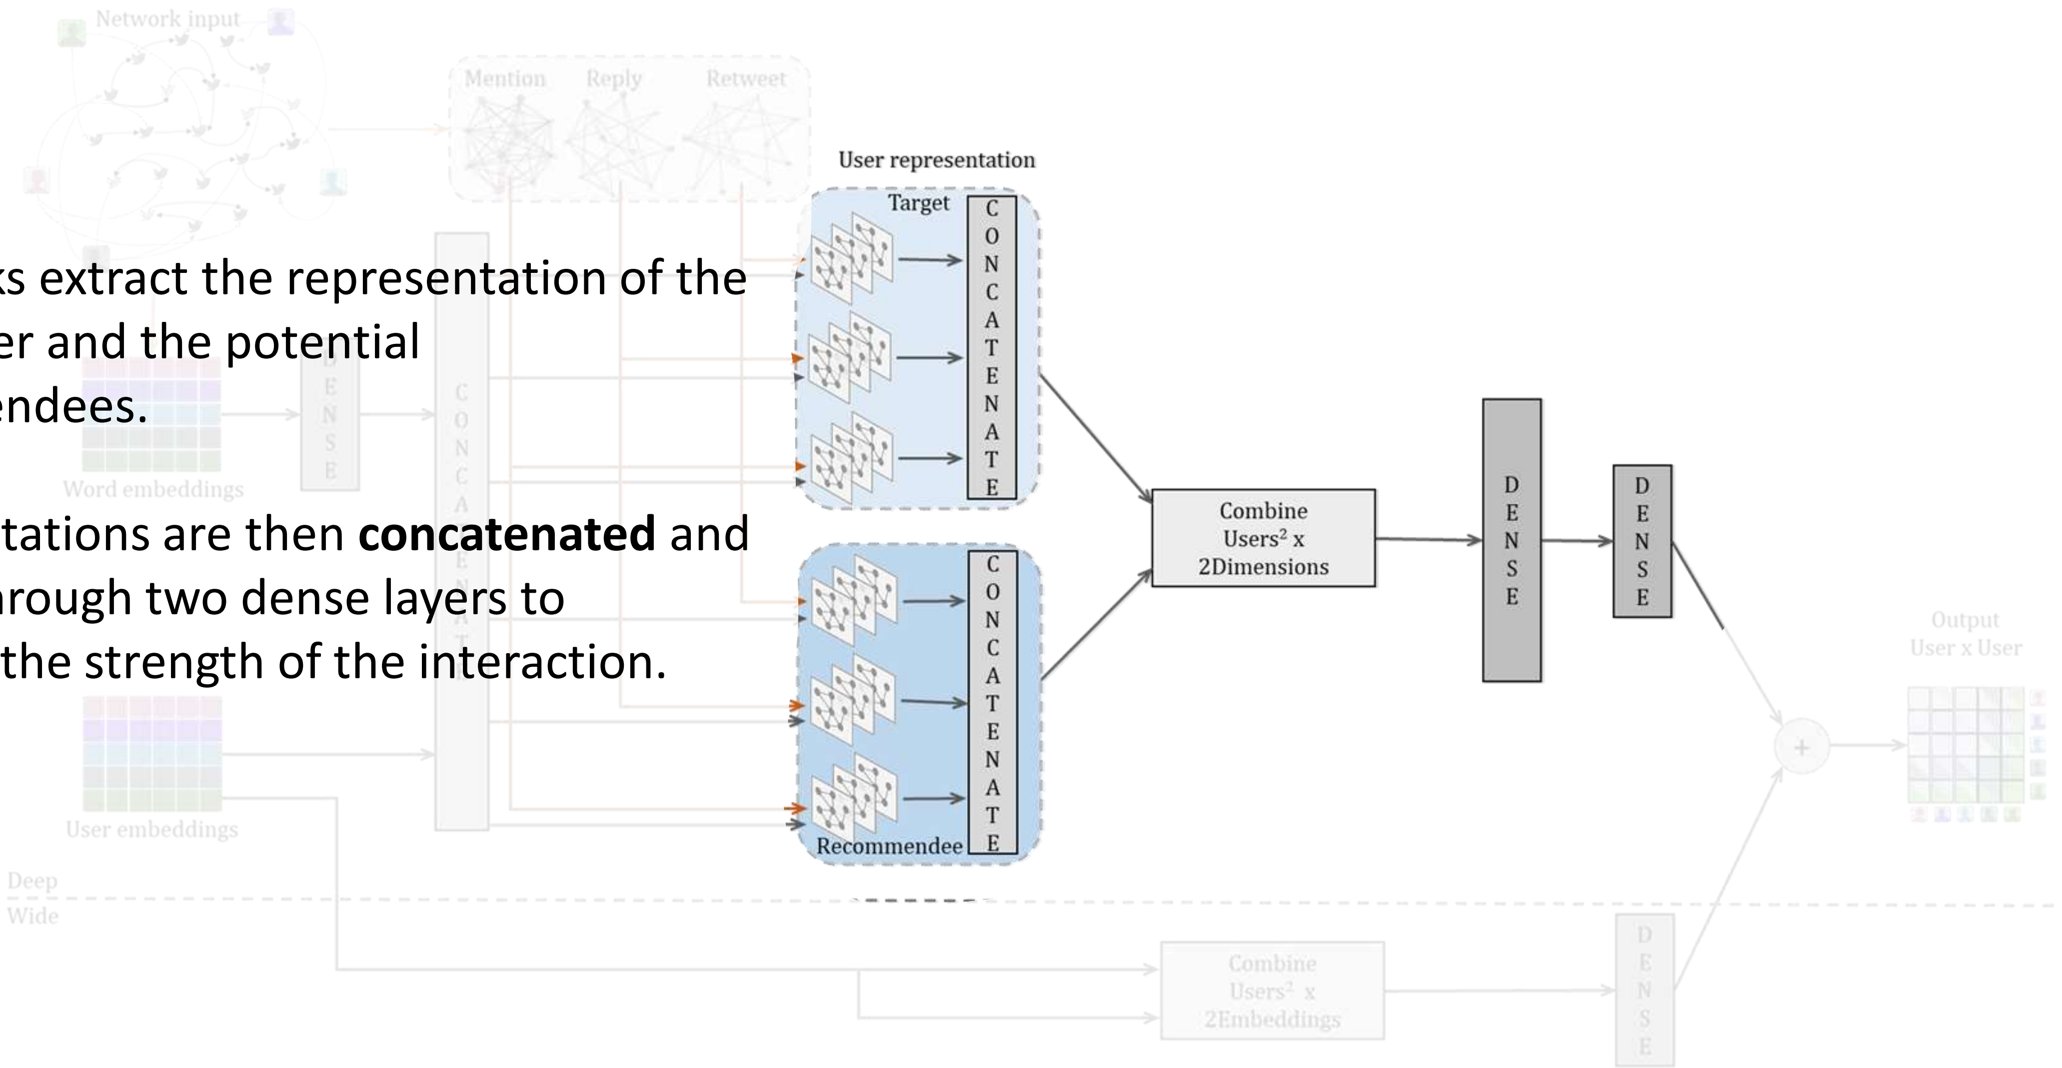

- The blocks extract the representation of the target user and the potential recommendees.

- Representations are then **concatenated** and passed through two dense layers to estimate the strength of the interaction.

FRediECH

Friend REcommenDer for breakIng Echo CHambers

- The resulting embeddings of the target user and the potential recommendees are concatenated and passed through a dense layer with one unit.

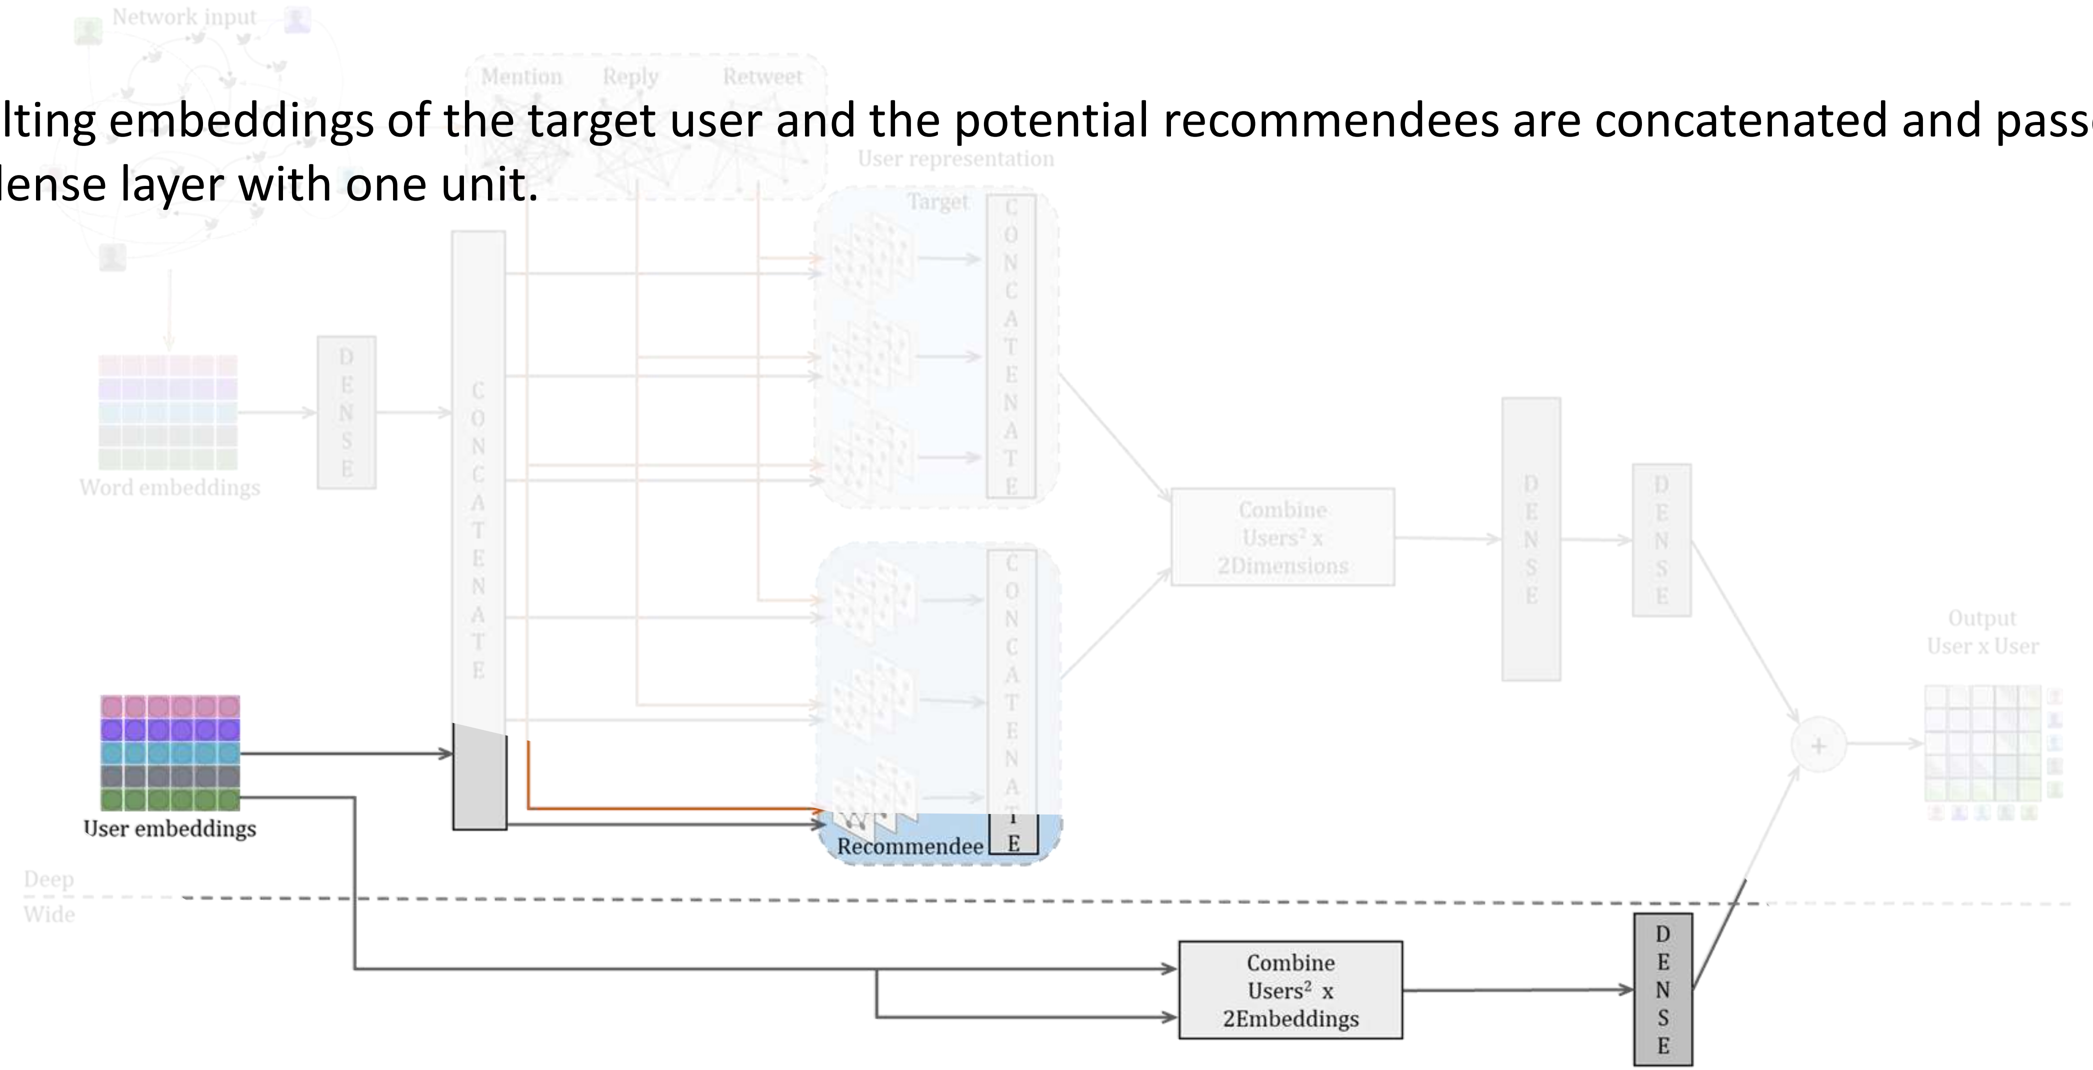

FRediECH

Friend REcommenDer for breakIng Echo CHambers

- 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.

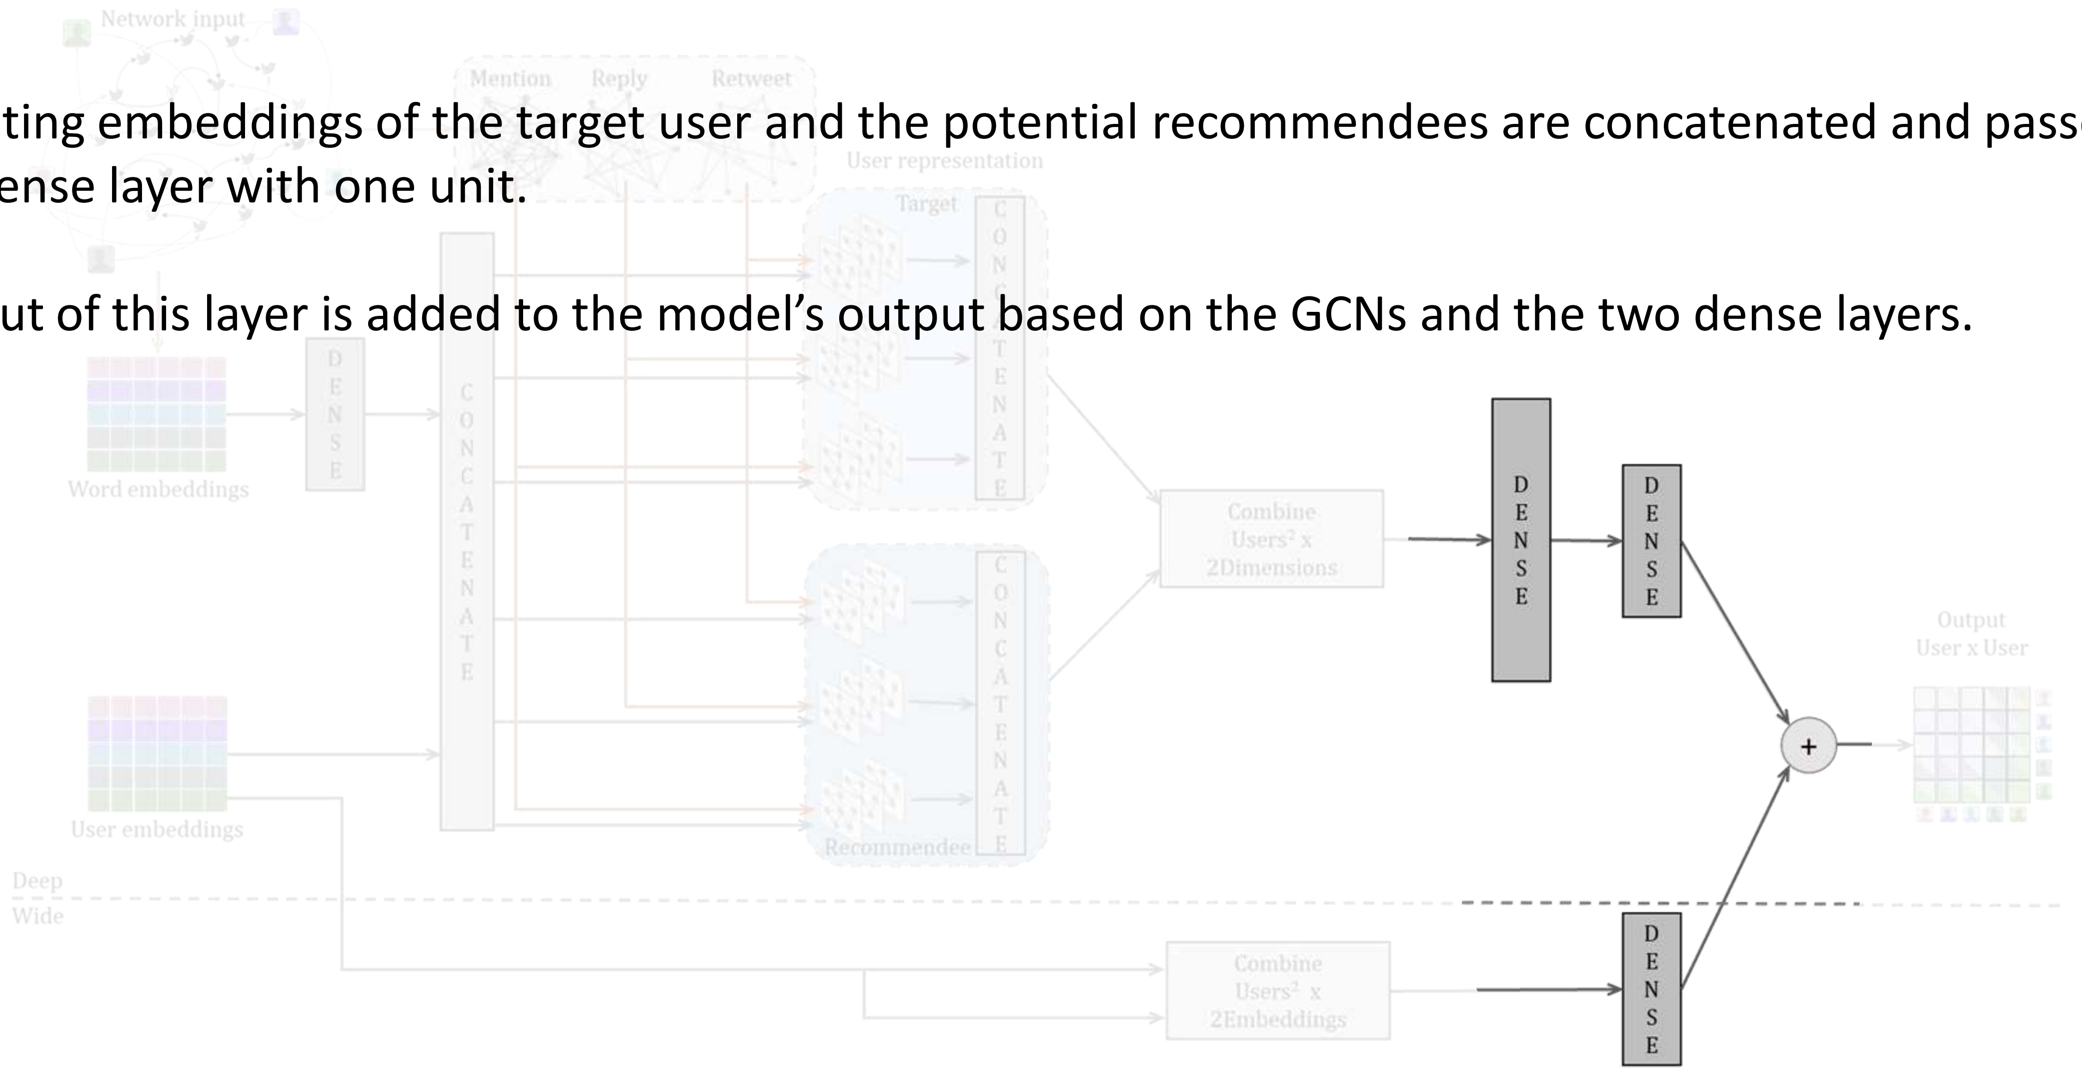

FRediECH

Friend REcommenDer for breakIng Echo CHambers

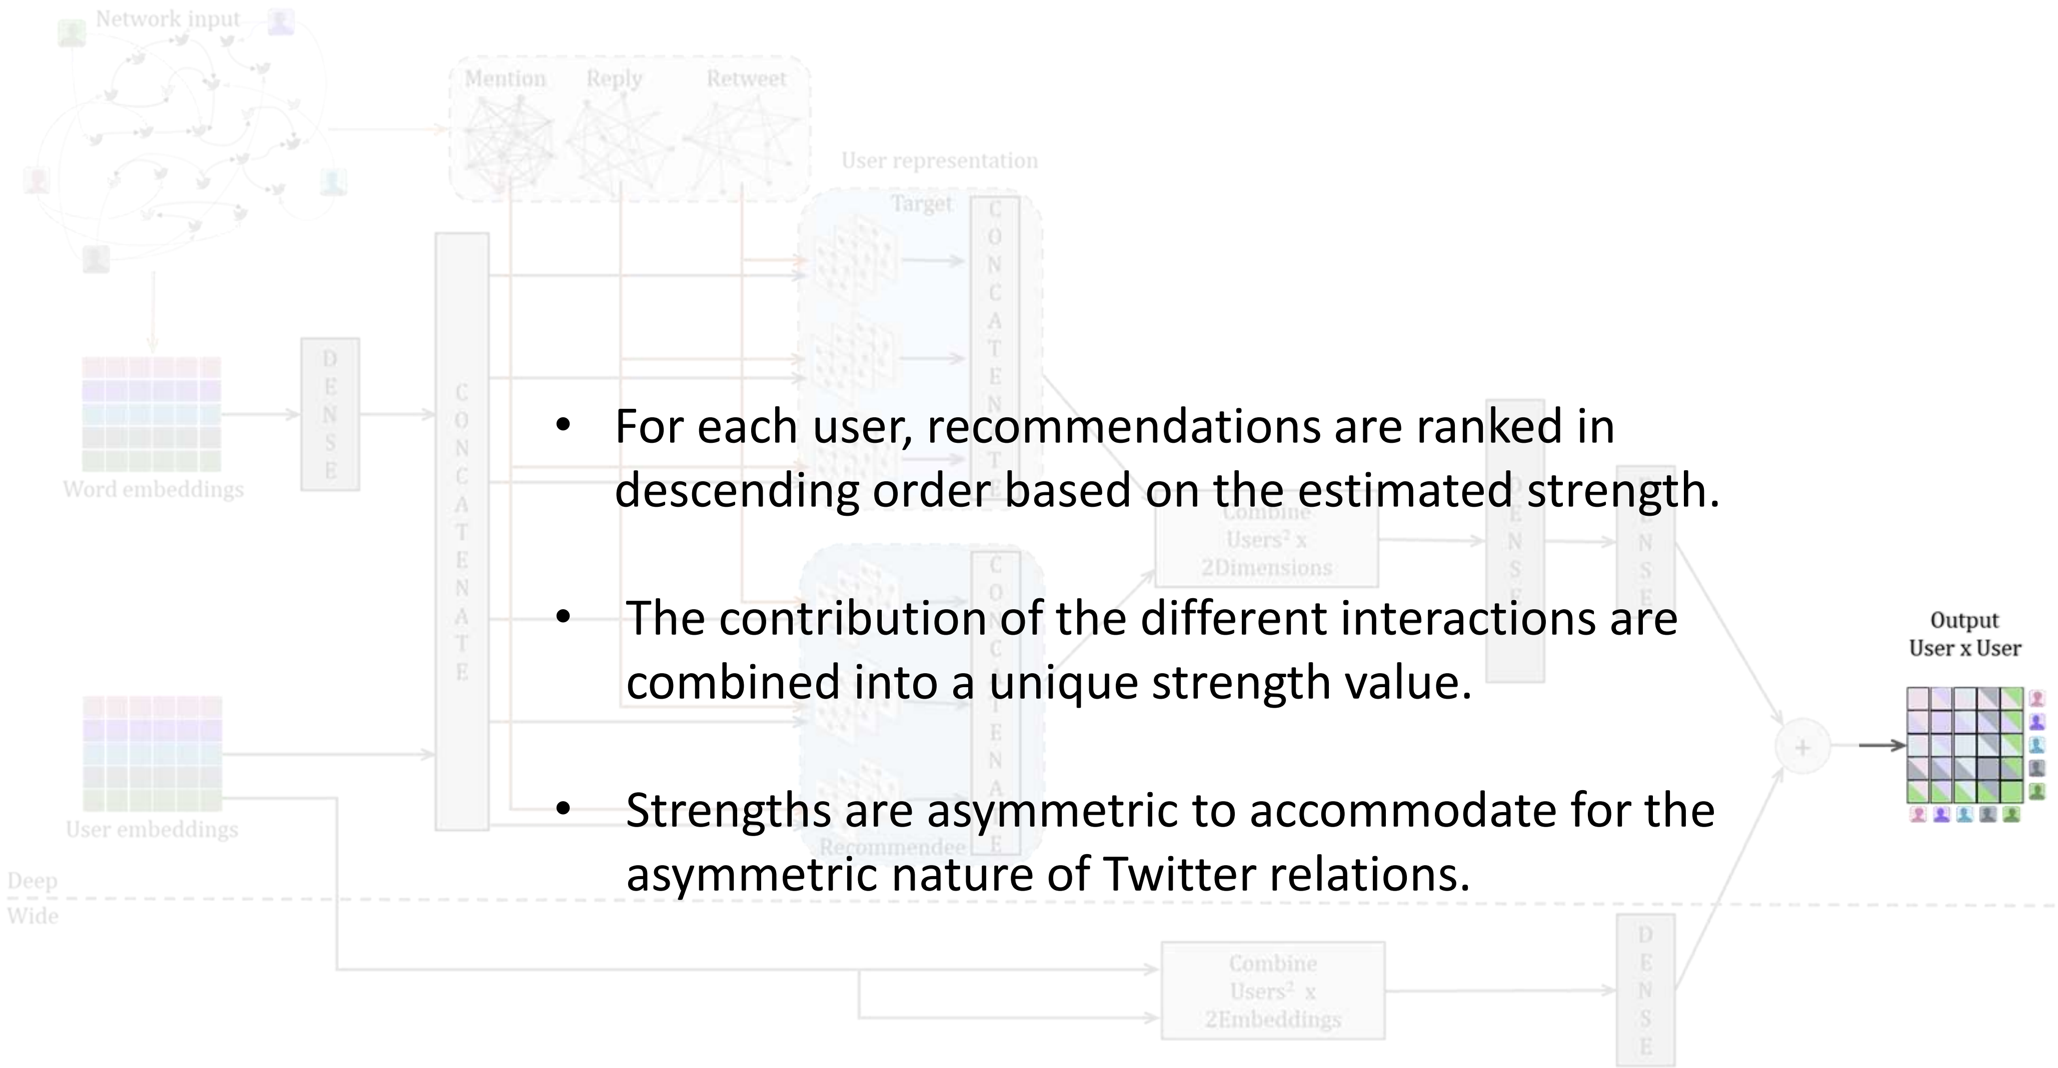

- For each user, recommendations are ranked in descending order based on the estimated strength.
- The contribution of the different interactions are combined into a unique strength value.
- Strengths are asymmetric to accommodate for the asymmetric nature of Twitter relations.

FRediECH

Model training

- Interactions between users belonging to **different echo chambers carry a higher weight** than **interactions** between users in the **same echo chamber**.

- The goal is to favour the diversity of recommendations by **learning the structure of echo chambers** without explicitly finding them.
    - This allows for more freedom in the echo chamber definition and more sensitivity to changes in the network.

FRediECH

Model training

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

- Interactions between users belonging to **different echo chambers carry a higher weight** than **interactions** between users in the **same echo chamber**.

- The goal is to favour the diversity of recommendations by **learning the structure of echo chambers** without explicitly finding them.
  - This allows for more freedom in the echo chamber definition and more sensitivity to changes in the network.

- We define a loss function based on the distance between users ($d\ (ui\ ,uj\ )$) and the number of interactions ($Yi\ j$).
  - The logarithm reduces the influence of users with many interactions.
  - The scalar prevents an interaction with a weight of 1 to become zero.
  - β allows tuning the preference of whether recommendations belong to the same group.

FRediECH

Model training

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

- Interactions between users belonging to **different echo chambers carry a higher weight** than **interactions** between users in the **same echo chamber**.

- The goal is to favour the diversity of recommendations by **learning the structure of echo chambers** without explicitly finding them.
  - This allows for more freedom in the echo chamber definition and more sensitivity to changes in the network.

- We define a loss function based on the distance between users (d $(ui$ $,uj$)) and the number of interactions (Yi j).
  - The logarithm reduces the influence of users with many interactions.
  - The scalar prevents an interaction with a weight of 1 to become zero.
  - β allows tuning the preference of whether recommendations belong to the same group.

- $d(u_i, u_j)$ is based on the cosine similarity over a new 10-dimensional embedding (e_i) representing users.
  - These embeddings were defined to capture the implicit community structure and were trained before the main model.
  - Users with similar interaction patterns will be represented by similar embeddings.
  - This loss function was based on GloVe.

Experimental evaluation

RQ1. How does FRediECH perform when compared with other tecniques?

RQ2. How do the key components of FRediECH contribute to the recommendations' performance?

Experimental evaluation

Data

- We used the obamacare data collection.
  - Tweets related to the `#obamacare` and `#aca` hashtags in Twitter.
  - Spans between May 2008 and October 2017.
  - It includes estimated user polarity.

- Tweets were retrieved using the **Faking it!** tool.

- We retrieved approximately 8 million public tweets belonging to 8,164 users, and 585,524 adjacent users.

- We kept **6,442 users** with at least one relation and that belonged to the largest connected component of the retrieved interaction graph.
  - This selection ensures that each user can be both source and destination of information content.

	Avg (± std)
#users	6,442
#tweets	7,016,552
Tweets per user	1089 (± 1413)
Relations per user	680 (± 1071)
Mentions per user	460 (± 733)
Replies per user	87 (± 190)
Retweets per user	399 (± 353)

Experimental evaluation
Data

- We used the obamacare data collection.
  - Tweets related to the #obamacare and #aca hashtags in Twitter.
  - Spans between May 2008 and October 2017.
  - It includes estimated user polarity.
- Tweets were retrieved using the **Faking it!** tool.
- We retrieved approximately 8 million public tweets belonging to 8,164 users, and 585,524 adjacent users.
- We kept **6,442 users** with at least one relation and that belonged to the largest connected component of the retrieved interaction graph.
  - This selection ensures that each user can be both source and destination of information content.
- We **verified** the existence of echo chambers.

	Avg (± std)
#users	6,442
#tweets	7,016,552
Tweets per user	1089 (± 1413)
Relations per user	680 (± 1071)
Mentions per user	460 (± 733)
Replies per user	87 (± 190)
Retweets per user	399 (± 353)

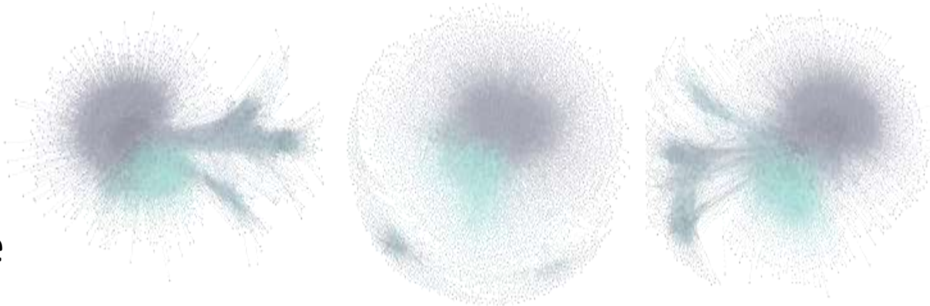

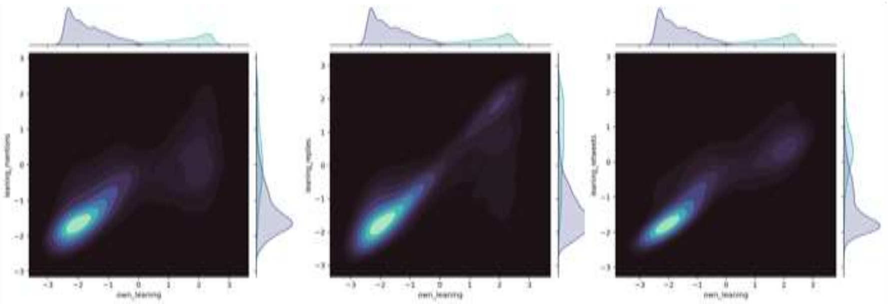

Experimental evaluation

Baselines

Trivial, non-personalized and traditional recommenders.

Random

Popularity

Topology

Content

Adapted traditional and state-of-the-art user-item recommendation techniques.

ImplicitMF

NeuralCF

GraphRec

Diffnet

Mult-VAE

Techniques focused on enhancing the structural diversity of recommendations to mitigate filter bubbles.

SCC

CAM

Experimental evaluation

Evaluation

Relevance

- Precision@k
- Recall@k
- DCG@k

Diversity

- Variations of intra-list dissimilarities were used to assess:
  - **Diversity** (i.e., differences within the recommended list)
  - **Novelty** (i.e., differences between the known users and the recommended ones).
- Individuals and groups.
- Euclidean distance over structural and content-based representations.

- All evaluations were performed over the **same data partitions** and evaluated using the same set of metrics.
- We selected the top-10 recommended users (50% of users have 10 or more interactions).
- Recommendations were considered correct if they appeared in the test set.

- Training set: interactions before August 30 2017 (80% of all interactions)
- Test set: remaining interactions.

Evaluation results

	Precision	Recall	nDCG	Structural dissimilarities			
				Ind Diversity	Ind Novelty	Group Diversity	Group Novelty
FRediECH	0.152	0.183	**0.685**	<u>0.888</u>	**0.992**	**0.927**	**0.938**
Random	0.113**	0.053**	0.459**	0.732	0.699**	0.726	0.797
Popularity	**0.281**	0.22	**0.686**	0.369**	0.559**	0.391	0.673
Topology-based Adamic-Adar	<u>0.27</u>	**0.285**	0.632	0.359**	0.431**	0.517	0.653
Topology-based Jaccard	0.191	0.249	0.567	0.364**	0.453**	0.592	0.667
Topology-based RA	<u>0.272</u>	<u>0.27</u>	0.642	0.367**	0.436	0.573	0.643
Topology-based CN	0.259	0.302	0.619	0.356**	0.424**	0.564	0.633
Content-based Full Tweets	0.115**	0.053**	0.439**	0.726	0.698**	0.727	0.797
Content-based 15 Tweets	0.246	0.22	0.584	0.428**	0.491**	0.629**	0.69
SCC	0.259	0.252	0.597	0.35**	0.496**	0.469	0.621
CAM	0.228	0.158	0.513	0.345**	0.424**	0.53	0.647
Implicit	<u>0.271</u>	0.252	<u>0.654</u>	0.401**	0.435**	0.559	0.643
NeuralCF	0.251	0.262	0.579	0.351**	0.419**	0.566	0.647
GraphRec	0.103**	0.183**	0.389**	**0.935**	<u>0.842</u>**	<u>0.739</u>	<u>0.895</u>
Mult-VAE	0.26	0.254	0.627	0.413**	0.433**	0.607	0.637
Original graph	-	-	-	0.325	0.418	0.581	0.603

** indicates statistically significant differences favouring FRediECH

Evaluation results

	Precision	Recall	nDCG	Structural dissimilarities			
				Ind Diversity	Ind Novelty	Group Diversity	Group Novelty
FRediECH	0.152	0.183	**0.685**	0.888	**0.992**	**0.927**	**0.938**
Random	0.113**	0.053**	0.459**	0.732	0.699**	0.726	0.797
Popularity	**0.281**	0.22	**0.686**	0.369**	0.559**	0.391	0.673
Topology-based Adamic-Adar	0.27	**0.285**	0.632	0.359**	0.431**	0.517	0.653
Topology-based Jaccard	0.191	0.249	0.567	0.364**	0.453**	0.592	0.667
[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	0.643
[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	0.633
[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	0.797
[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	0.69
[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	0.621
[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	0.647
[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	0.643
[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	0.647
[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	0.895
[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	0.637
Original graph	-	-	-	0.325	0.418	0.581	0.603

- Random recommender achieved high diversity and novelty results, but recommendations were less relevant.
- There is a trade-off between the relevance, and diversity and novelty.
- Techniques achieving high relevance also achieved low diversity and novelty scores.
- Statistically significant differences were observed regarding diversity and novelty when compared to all techniques but GraphRec and FRediECH.

** indicates statistically significant differences favouring FRediECH

Evaluation results

	Precision	Recall	nDCG	Structural dissimilarities			
				Ind Diversity	Ind Novelty	Group Diversity	Group Novelty
FRediECH	0.152	0.183	**0.685**	0.888	**0.992**	**0.927**	**0.938**
Random	0.113**	0.053**	0.459**	0.732	0.699**	0.726	0.797
Popularity	**0.281**	0.22	**0.686**	0.369**	0.559**	0.391	0.673
Topology-based Adamic-Adar	0.27	**0.285**	0.632	0.359**	0.431**	0.517	0.653
Topology-based Jaccard	0.191	0.249	0.567	0.364**	0.453**	0.592	0.667
Topology-based RA	0.272	0.27	0.642	0.367**	0.436	0.573	0.643
Topology-based CN	0.259	0.302	0.619	0.356**	0.424**	0.564	0.633
Content-based Full Tweets	0.115**	0.053**	0.439**	0.726	0.698**	0.727	0.797
Content-based 15 Tweets	0.246	0.22	0.584	0.428**	0.491**	0.629**	0.69
SCC	0.259	0.252	0.597	0.35**	0.496**	0.469	0.621

- Topological baselines achieved high precision and low diversity, which is expected as recommendations are based on user neighbourhood.
- Diversity and novelty differences were significant and favoured FRediECH.
- 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.

Evaluation results

	Precision	Recall	nDCG	Structural dissimilarities			
				Ind Diversity	Ind Novelty	Group Diversity	Group Novelty
FRediECH	0.152	0.183	**0.685**	<u>0.888</u>	**0.992**	**0.927**	**0.938**
Random	0.113**	0.053**	0.459**	0.732	0.699**	0.726	0.797
							0.673
							0.653
							0.667
							0.643
							0.633
							0.797
							0.69
							0.621
							0.647
Implicit	<u>0.271</u>	0.252	<u>0.654</u>	0.401**	0.435**	0.559	0.643
NeuralCF	0.251	0.262	0.579	0.351**	0.419**	0.566	0.647
GraphRec	0.103**	0.183**	0.389**	**0.935**	<u>0.842</u>**	<u>0.739</u>	<u>0.895</u>
Mult-VAE	0.26	0.254	0.627	0.413**	0.433**	0.607	0.637
Original graph	-	-	-	0.325	0.418	0.581	0.603

- FRediECH achieved the highest diversity and novelty results, followed by GraphRec.
  - For individual diversity, in which GraphRec outperformed FRediECH.
- In terms of relevance, FRediECH also significantly outperformed GraphRec.
- Most of the differences favouring FRediECH were statistically significant.
- Despite lower precision and recall that other techniques, nDCG results showed that even when recommending non relevant users, the relevant ones were ranked high.

** indicates statistically significant differences favouring FRediECH

Evaluation results

	Precision	Recall	nDCG	Structural dissimilarities			
				Ind Diversity	Ind Novelty	Group Diversity	Group Novelty
FRediECH	0.152	0.183	**0.685**	0.888	**0.992**	**0.927**	**0.938**
Random	0.113**	0.053**	0.459**	0.732	0.699**	0.726	0.797
[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	0.391	0.673
[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	0.517	0.653
[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	0.592	0.667
[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	0.573	0.643
[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	0.564	0.633
Content-based Full Tweets	0.113	0.053	0.459	0.726	0.698	0.727	0.797
Content-based 15 Tweets	0.246	0.22	0.584	0.428**	0.491**	0.629**	0.69
SCC	0.259	0.252	0.597	0.35**	0.496**	0.469	0.621
CAM	0.228	0.158	0.513	0.345**	0.424**	0.53	0.647
Implicit	0.271	0.252	0.654	0.401**	0.435**	0.559	0.643
NeuralCF	0.251	0.262	0.579	0.351**	0.419**	0.566	0.647
GraphRec	0.103**	0.183**	0.389**	**0.935**	0.842**	0.739	0.895
Mult-VAE	0.26	0.254	0.627	0.413**	0.433**	0.607	0.637
Original graph	-	-	-	0.325	0.418	0.581	0.603

- SCC achieved higher relevance and structural novelty.
- CAM achieved higher content diversity.

- Diversity and novelty of both techniques were close to those of the original network, thus failing to significantly improve the quality of the network.

** indicates statistically significant differences favouring FRediECH

Evaluation results

	Precision	Recall	nDCG	Structural dissimilarities			
				Ind Diversity	Ind Novelty	Group Diversity	Group Novelty
FRediECH	0.152	0.183	**0.685**	0.888	**0.992**	**0.927**	**0.938**
Random	0.113**	0.053**	0.459**	0.732	0.699**	0.726	0.797
Popularity	**0.281**	0.22	**0.686**	0.369**	0.559**	0.391	0.673
Topology-based Adamic-Adar	0.27	**0.285**	0.632	0.359**	0.431**	0.517	0.653
Topology[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	0.667
Topology[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	0.643
Topology[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	0.633
Content-[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	0.797
Content-based 15 Tweets	0.246	0.22	0.584	0.428**	0.491**	0.629**	0.69
GraphRec	0.103**	0.183**	0.389**	**0.935**	0.842**	0.739	0.895
Mult-VAE	0.26	0.254	0.627	0.413**	0.433**	0.607	0.637
Original graph	-	-	-	0.325	0.418	0.581	0.603

	Traditional	State-of-the-art	Original structure
Avg. Improvements	47%	44%	
Maximum	60% (individual novelty)		67% (individual novelty)

- In general, the novelty of recommendations was higher than their diversity.
- Novelty was higher for the structural distance, which implies that recommended users belong to other communities, but still shared similar content.

** indicates statistically significant differences favouring FRediECH

Evaluation results

	Precision	Recall	nDCG	Structural dissimilarities			
				Ind Diversity	Ind Novelty	Group Diversity	Group Novelty
FRediECH	0.152	0.183	**0.685**	0.888	**0.992**	**0.927**	**0.938**
Random	0.113**	0.053**	0.459**	0.732	0.699**	0.726	0.797
Popularity	**0.281**	0.22	**0.686**	0.369**	0.559**	0.391	0.673
Topology-based Adamic-Adar	0.27	**0.285**	0.632	0.359**	0.431**	0.517	0.653
Topology-based Jaccard	0.191	0.249	0.567	0.364**	0.453**	0.592	0.667
Topology-based RA	0.272	0.27	0.642	0.367**	0.436	0.573	0.643
Topology-based CN	0.259	0.302	0.619	0.356**	0.424**	0.564	0.633
Content[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]
Content[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]
SCC	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]
CAM	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]
Implicit	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]
NeuralCF	0.251	0.262	0.579	0.351**	0.419**	0.566	0.647
GraphRec	0.103**	0.183**	0.389**	**0.935**	0.842**	0.739	0.895
Mult-VAE	0.26	0.254	0.627	0.413**	0.433**	0.607	0.637
Original graph	-	-	-	0.325	0.418	0.581	0.603

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

** indicates statistically significant differences favouring FRediECH

Summary & conclusions

- We developed FRediECH inspired by a **graph convolutional network** and a **Deep & Wide** architecture, coupling echo chamber awareness and user representations to balance the <u>relevance</u>, <u>diversity</u> and <u>novelty</u> of friend recommendations.

- FRediECH produced **similarly relevant recommendations** to those of the selected baselines while **increasing their diversity and novelty**.

- FRediECH allows recommending **users who are different among them and from the already known ones**, thus effectively **helping to reduce the echo chamber effect**.

- [Data and code] are publicly available.

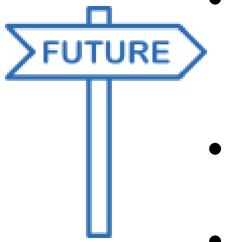

- Evaluations over different datasets varying the domain and time period to truly assess usefulness and generalizability.
- Analyses regarding the relevance of each type of interaction, and their contribution to the final recommendations.
- Explanations to better guide users in broadening their interactions.

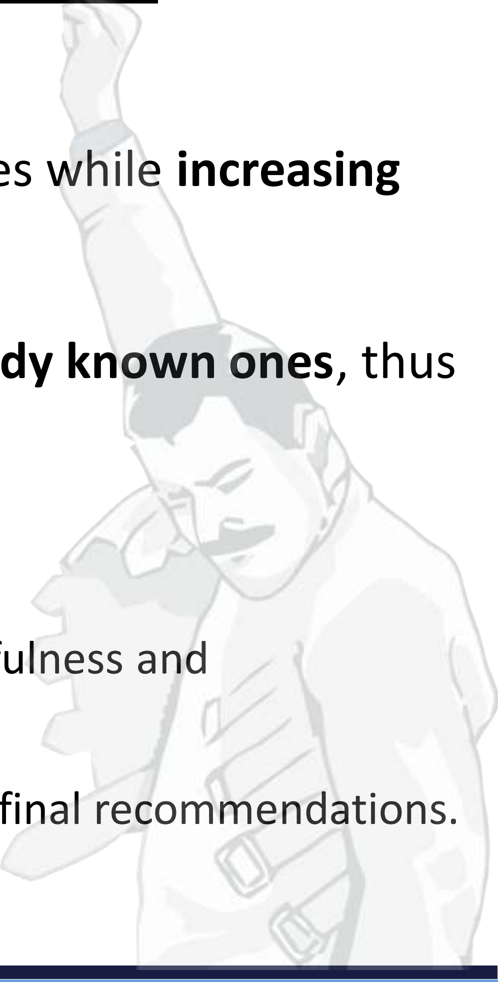

Thanks!

Questions?

antonela.tommasel@isistan.unicen.edu.ar

I want to break free! Recommending friends from outside the echo chamber

ANTONELA TOMMASEL

JUAN MANUEL RODRIGUEZ

DANIELA GODOY

CONICET

ISISTAN

FRediECH

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).
- 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.

Experimental evaluation

Data: Echo chambers

- We quantified the polarization of users in the LCC, relying on the relation between user leaning and consumption leaning to assess the existence of echo chambers.

- Green nodes represent democrats, grey nodes represent republicans.

- The conversational interaction graphs of users in the LCC.
  - 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.

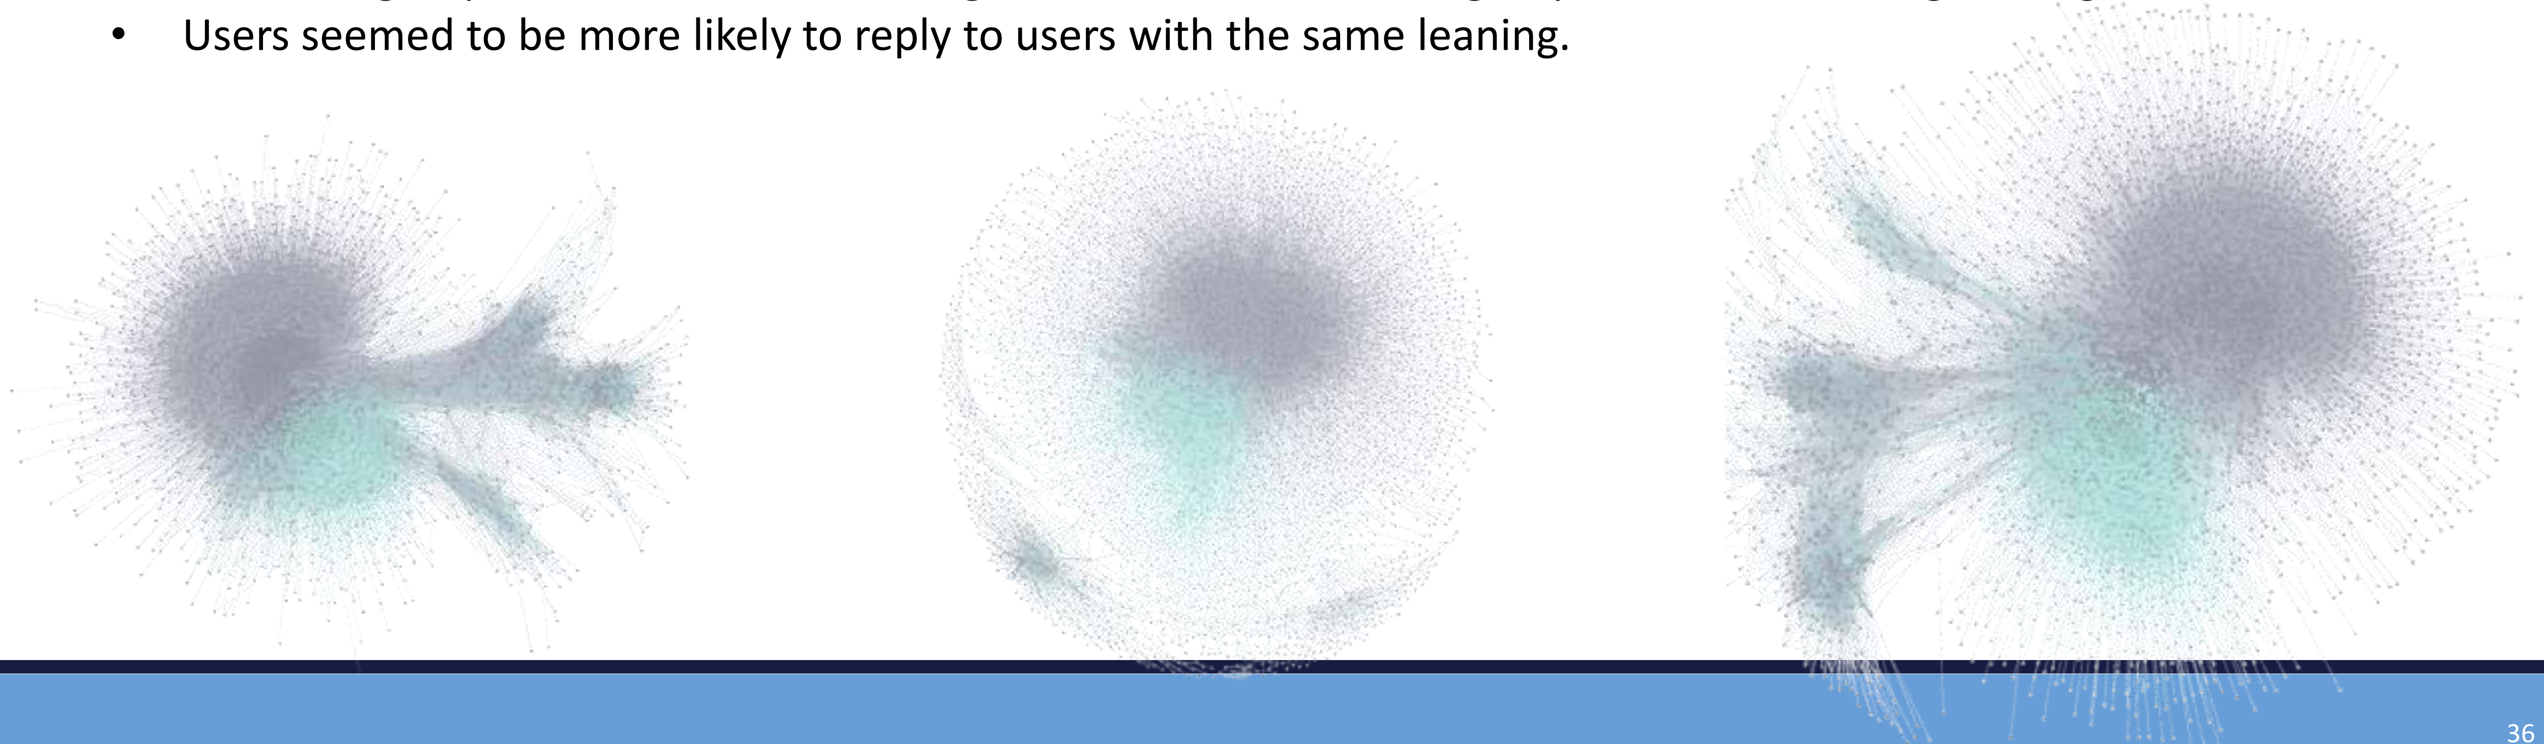

Experimental evaluation

Data: Echo chambers

- We quantified the polarization of users in the LCC, relying on the relation between user leaning and consumption leaning to assess the existence of echo chambers.

- Green nodes represent democrats, grey nodes represent republicans.

- The relation between the political leaning of users and the average information consumption leaning per interaction type.
  - The colour represents user density, the lighter the area, the higher the density of users in such area.
  - Marginals show the distribution of user leaning.

- Positive correlations were found between users' production and consumption leaning.
- In average, 89% of the interactions of republican users were with other republicans
- Democrats interacted with users on a wider range of democrat and neutral leanings.

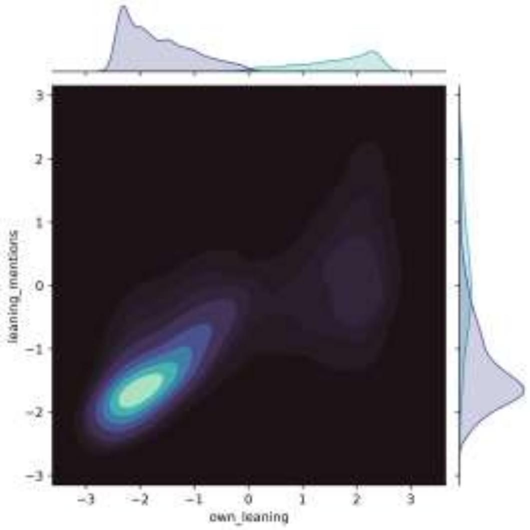

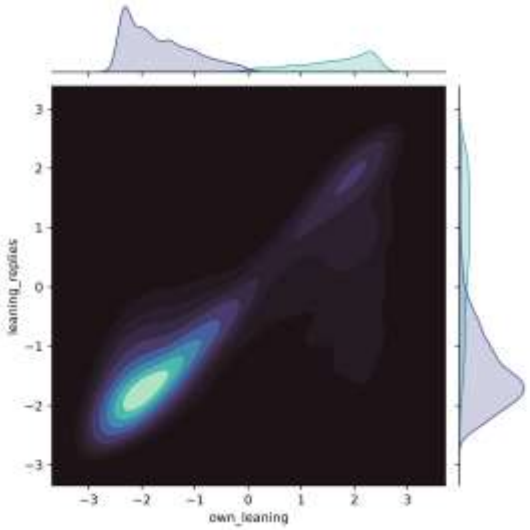

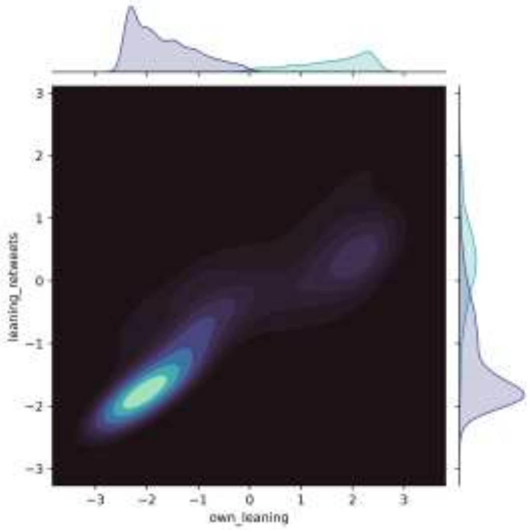

Evaluation results
RQ2. Ablation Study

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

- Relations were removed from both the training and test sets.
- A new model was trained from scratch for each evaluation.

Evaluation results

RQ2. Ablation Study

	Precision	Recall	nDCG	Structural dissimilarities			
				Ind Diversity	Ind Novelty	Group Diversity	Group Novelty
FRediECH	<u>0.152</u>	0.183	<u>0.685</u>	**0.888**	**0.992**	**0.927**	0.938
$\text{FRediECH}_{\text{NO-NS}}$	0.149**	0.172	0.553**	0.726	0.82**	0.845	0.852
$\text{FRediECH}_{\text{NO-WIDE}}$	<u>0.152</u>	0.189	<u>0.685</u>	0.888	**0.993**	<u>0.845</u>	0.966
$\text{FRediECH}_{\text{NO-WIDE-NO-NS}}$	0.134	0.172	0.609	0.597**	0.82**	0.728	0.852
$\text{FRediECH}_{\text{DUAL}}$	**0.169**	<u>0.192</u>	0.561	<u>0.73</u>	0.937**	0.762	0.912
$\text{FRediECH}_{\text{NO-BERT}}$	**0.16**	0.193	0.56	0.596**	<u>0.97</u>**	0.708	0.936
$\text{FRediECH}_{\text{MENTION}}$	0.14	0.182	0.544	0.541**	**0.993**	0.698	0.93
$\text{FRediECH}_{\text{REPLY}}$	0.103**	**0.203**	**0.732**	0.509**	**0.99**	0.643	**0.99**
$\text{FRediECH}_{\text{RETWEET}}$	0.146	<u>0.193</u>	0.567	0.646**	**0.99**	0.724	0.941
$\text{FRediECH}_{\text{MENTION-REPLY}}$	0.136	0.176	0.547	0.651	**0.99**	0.741	0.932
$\text{FRediECH}_{\text{MENTION-RETWEET}}$	<u>0.159</u>	0.184	0.542	0.627**	0.96	0.732	0.916
$\text{FRediECH}_{\text{REPLY-RETWEET}}$	**0.162**	0.183	0.55	0.69	0.947**	0.762	0.909

Evaluation results

RQ2. Ablation Study

	Precision	Recall	nDCG	Structural dissimilarities			
				Ind Diversity	Ind Novelty	Group Diversity	Group Novelty
FRediECH	0.152	0.183	0.685	**0.888**	**0.992**	**0.927**	0.938
$\text{FRediECH}_{\text{NO-NS}}$	0.149**	0.172	0.553**	0.726	0.82**	0.845	0.852
$\text{FRediECH}_{\text{NO-WIDE}}$	0.152	0.189	0.685	0.888	**0.993**	0.845	0.966
$\text{FRediECH}_{\text{NO-WIDE-NO-NS}}$	0.134	0.172	0.609	0.597**	0.82**	0.728	0.852
$\text{FRediECH}_{\text{DUAL}}$	**0.169**	0.192	0.561	0.73	0.937**	0.762	0.912
$\text{FRediECH}_{\text{NO-BERT}}$	**0.16**	0.193	0.56	0.596**	0.97**	0.708	0.936
$\text{FRediECH}_{\text{MENTION}}$	0.14	0.182	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]
$\text{FRediECH}_{\text{REPLY}}$	0.103**	**0.203**	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]

- Relevance was not greatly affected by the modifications.
- Diversity and novelty showed more variability.

	Relevance	Diversity	Novelty
$\text{FRediECH}_{\text{NO-NS}}$	=	↑	↑
$\text{FRediECH}_{\text{NO-WIDE}}$	=	=	=
$\text{FRediECH}_{\text{NO-WIDE-NO-NS}}$	=	↓	↓
$\text{FRediECH}_{\text{DUAL}}$	↑	↓	↓

Evaluation results

RQ2. Ablation Study

- FRediECH$_{NO-BERT}$: including content allowed to significantly increase the novelty and diversity of recommendations.
- In general, only considering one interaction significantly decreased diversity and novelty (except for FRediECH$_{REPLY}$)
- Pairs of interactions. Precision and recall slightly increased while diversity and novelty decreased.
- Interactions might carry different weights, implying the need for different mechanisms for adequately leveraging them.

	Precision	Reca					
FRediECH	0.152	0.18					
FRediECH$_{NO-NS}$	0.149**	0.17					
FRediECH$_{NO-WIDE}$	0.152	0.18					
FRediECH$_{NO-WIDE-NO-NS}$	0.134	0.17					
FRediECH$_{DUAL}$	**0.169**	0.192	0.561	0.73	0.937**	0.762	0.912
FRediECH$_{NO-BERT}$	**0.16**	0.193	0.56	0.596**	0.97**	0.708	0.936
FRediECH$_{MENTION}$	0.14	0.182	0.544	0.541**	**0.993**	0.698	0.93
FRediECH$_{REPLY}$	0.103**	**0.203**	**0.732**	0.509**	**0.99**	0.643	**0.99**
FRediECH$_{RETWEET}$	0.146	0.193	0.567	0.646**	**0.99**	0.724	0.941
FRediECH$_{MENTION-REPLY}$	0.136	0.176	0.547	0.651	**0.99**	0.741	0.932
FRediECH$_{MENTION-RETWEET}$	0.159	0.184	0.542	0.627**	0.96	0.732	0.916
FRediECH$_{REPLY-RETWEET}$	**0.162**	0.183	0.55	0.69	0.947**	0.762	0.909

Evaluation results

RQ2. Ablation Study

	Precision	Recall	nDCG	Structural dissimilarities			
				Ind Diversity	Ind Novelty	Group Diversity	Group Novelty
FRediECH	0.152	0.183	0.685	**0.888**	**0.992**	**0.927**	0.938
FRediECH$_{NO-NS}$	0.149**	0.172	0.553**	0.726	0.82**	0.845	0.852
FRediECH$_{NO-WIDE}$	0.152	0.189	0.685	0.888	**0.993**	0.845	0.966
FRediECH[illegible]	0.134	0.172	0.609	0.597**	0.82**	0.728	0.852
FRe[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]
FRe[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]
FRe[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]
FRe[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]
FRe[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]
FRe[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]
FRediECH$_{MENTION-RETWEET}$	0.159	0.184	0.542	0.627**	0.96	0.732	0.916
FRediECH$_{REPLY-RETWEET}$	**0.162**	0.183	0.55	0.69	0.947**	0.762	0.909

Results showed that each component
significantly contributed to performance.

More studies of the interaction types and their interplay
in the quality of recommendations are needed.

Evaluation results

RQ2. Ablation Study

Architectural

Data available to the model

Evaluation results

RQ2. Ablation Study

Architectural

- Relevance was not greatly affected.
- Diversity/novelty showed more variability.
- Differences favouring the original FRediECH were statistically significant.

Data available to the model

- Including content allowed to significantly increase the novelty and diversity of recommendations.
- In general, only considering one interaction significantly decreased diversity and novelty.
- Interactions might carry different weights, implying the need for different mechanisms for adequately leveraging them.

Evaluation results

RQ2. Ablation Study

Architectural

Data available to the model

Results showed that each component significantly contributed to performance.

More studies of the interaction types and their interplay in the quality of recommendations are needed.