

Guardian: Evaluating Trust in Online Social Networks with Graph Convolutional Networks

Wanyu Lin, Zhaolin Gao, Baochun Li
University of Toronto






Almost 4.57 billion people were active internet users as of April 2020.

— Statista


Social trust is the basis of
online social networks.

Estimates of **social trust** help indicate to what extent a user could expect someone else to perform given actions, therefore has many applications, such as trust-based recommendations.




Network graph

an example




Can A trust E? And, to what extent?

Large-scale

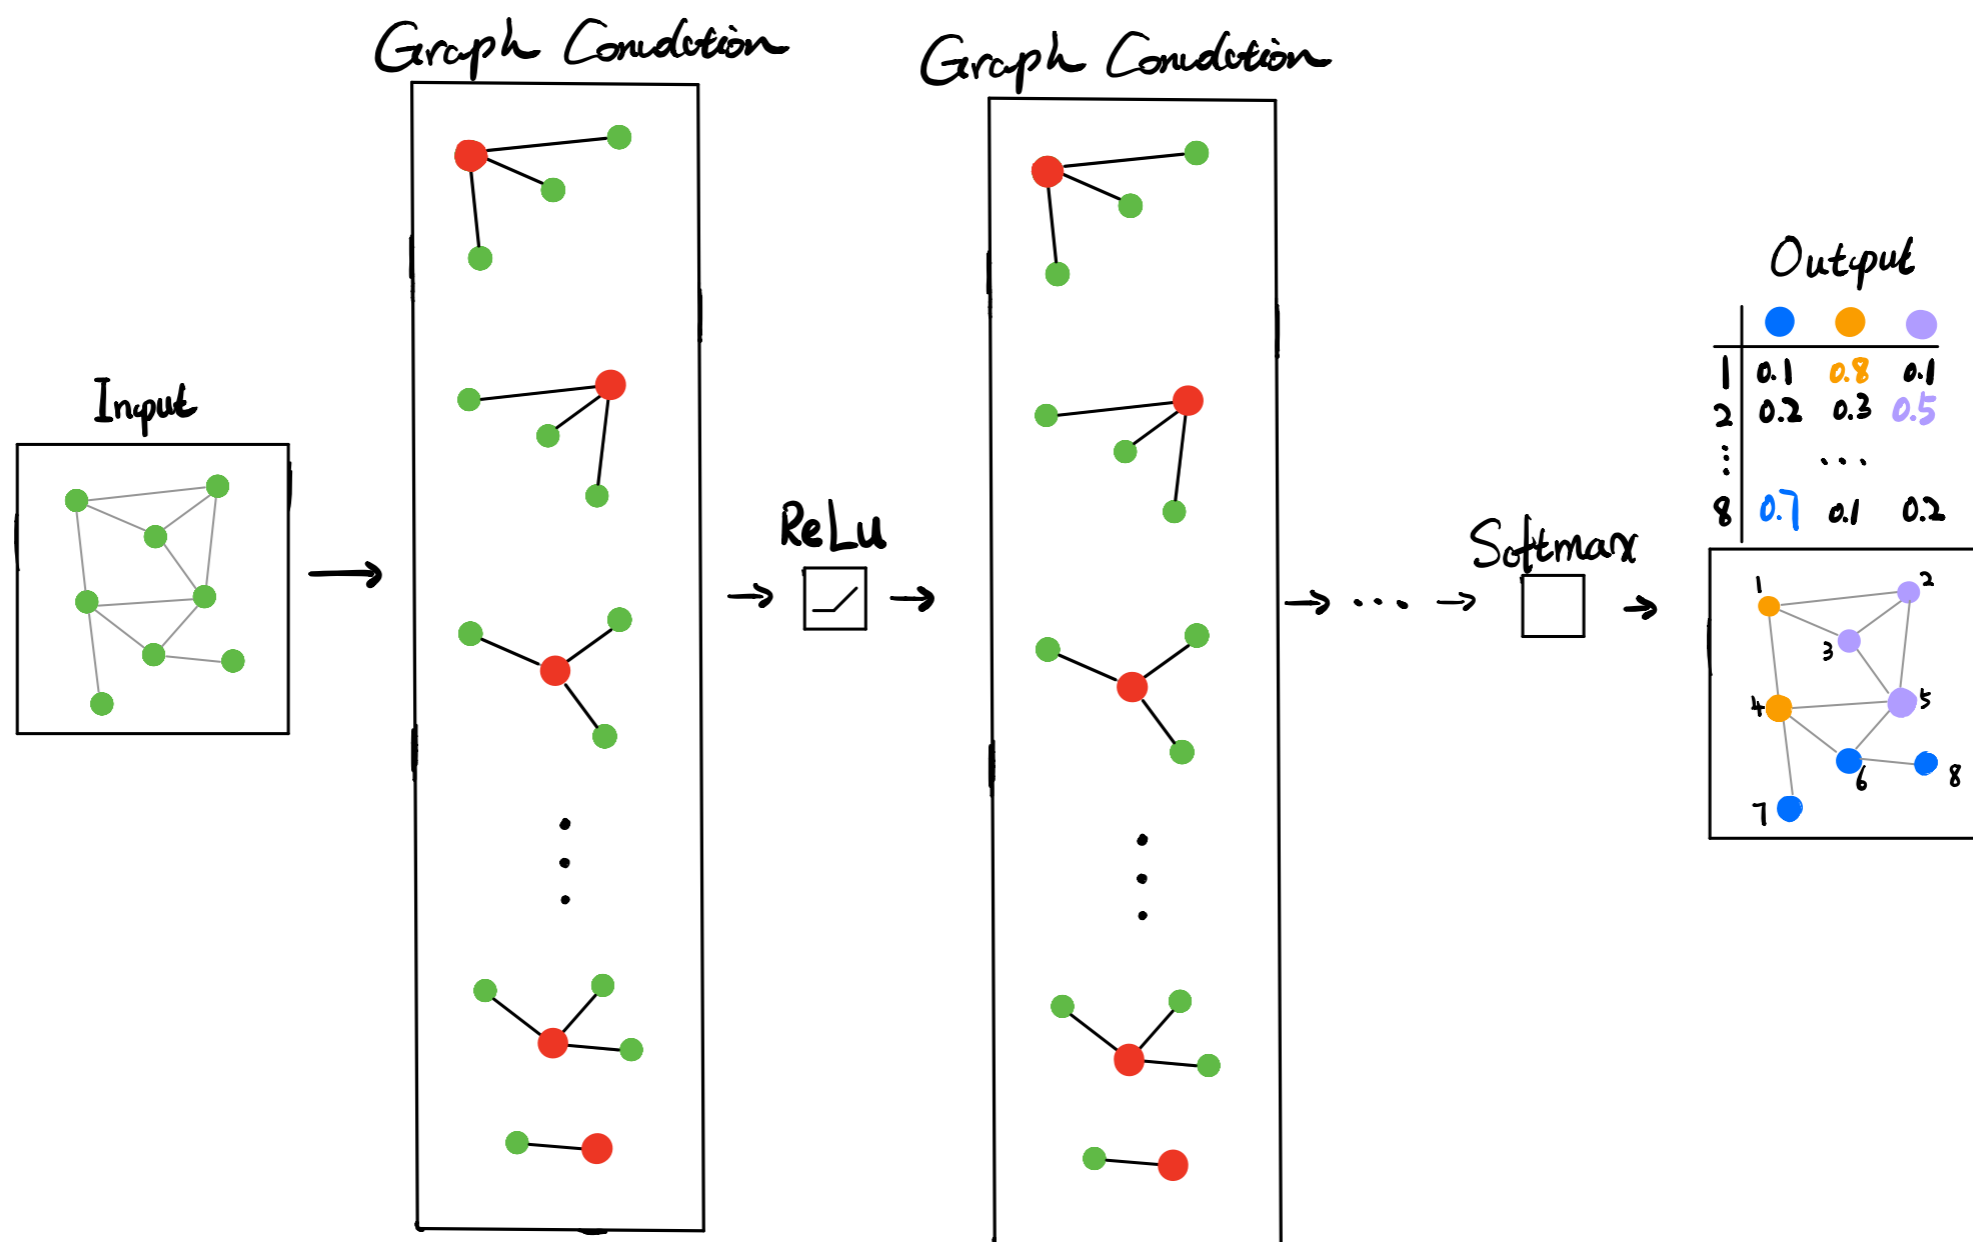


Wait a second ...

Graph convolutional neural networks — an efficient variant of convolutional neural networks on graphs.




Representation learning with graph convolutional networks




The complexity of model parameters
are **independent of the input
graph size.**

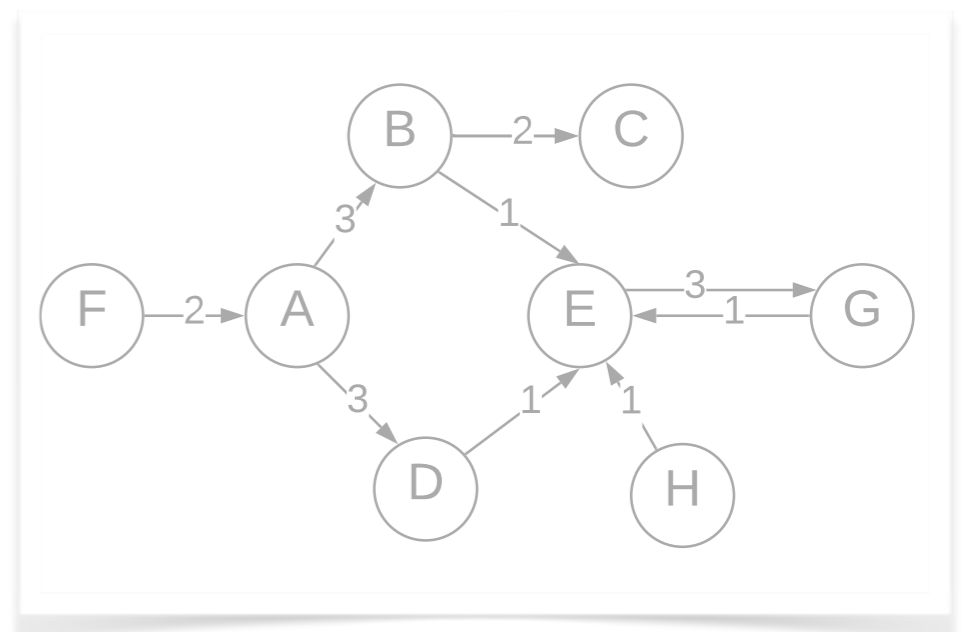
Preliminaries: trust properties




Asymmetry: one user may trust someone else more than she is trusted back.



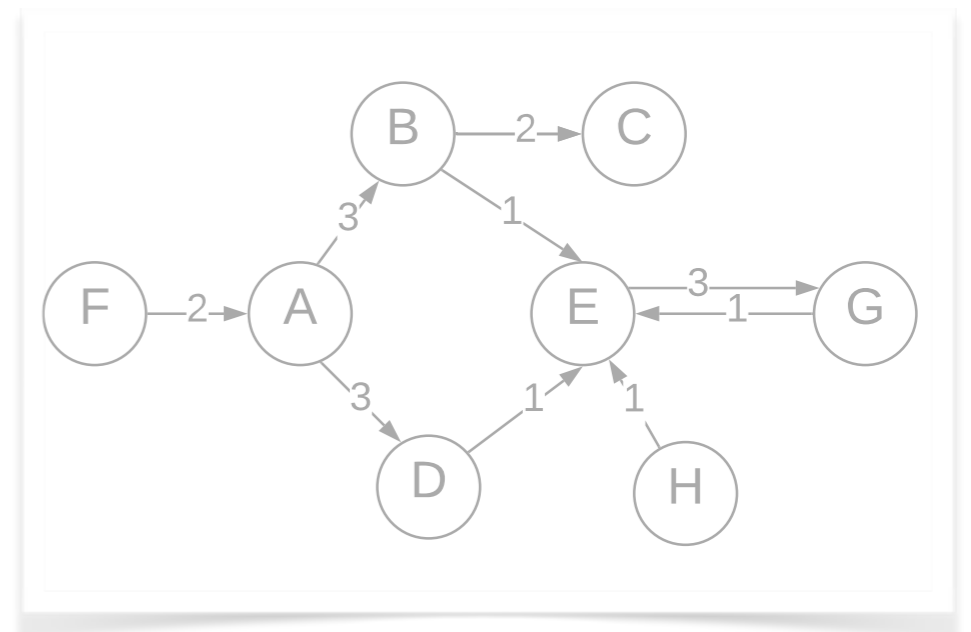
Trust properties





Propagative nature: trust may be passed from one user to another, creating chains of social trust that connects two users who are not connected.



Trust properties



Composable nature: trust needs to be aggregated if several chains of social trust exist.



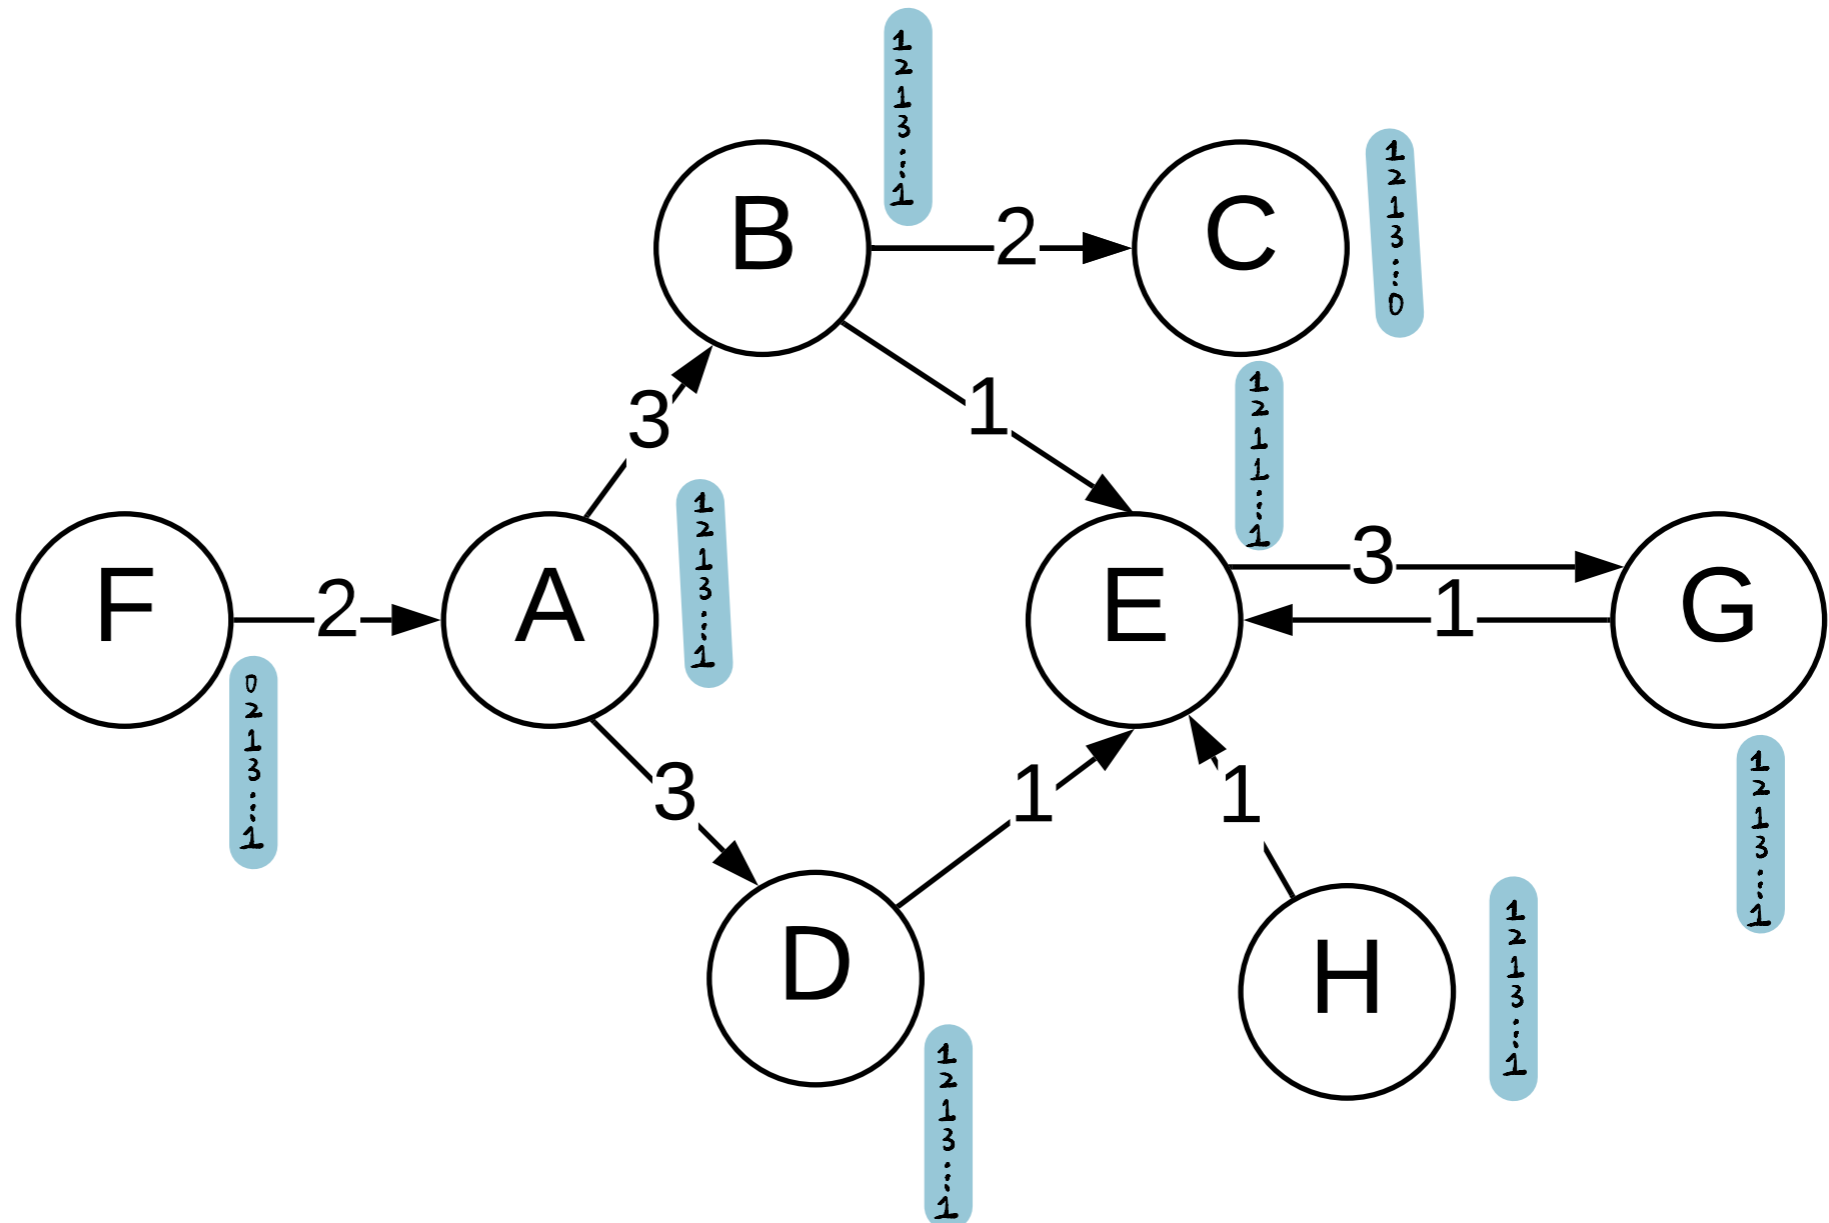
Trust properties

An effective way of evaluating trust should be able to capture these trust properties simultaneously.

Guardian: an end-to-end learning framework for social trust evaluation.



Embedding layer

We use a pre-trained embedding layer that maps each user into a vector.

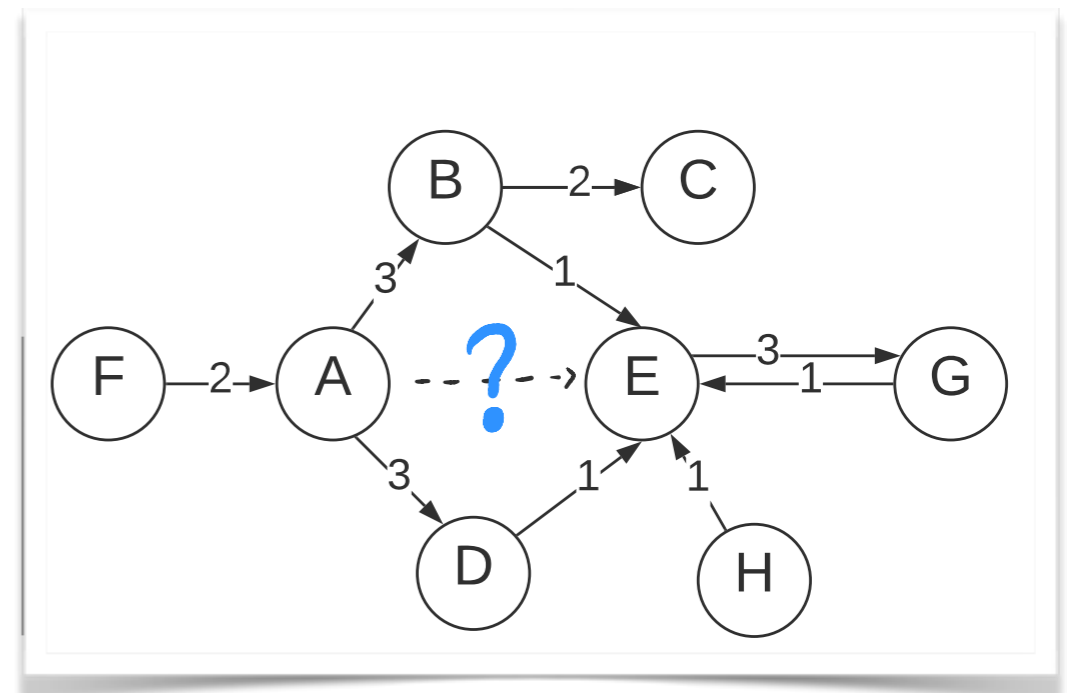




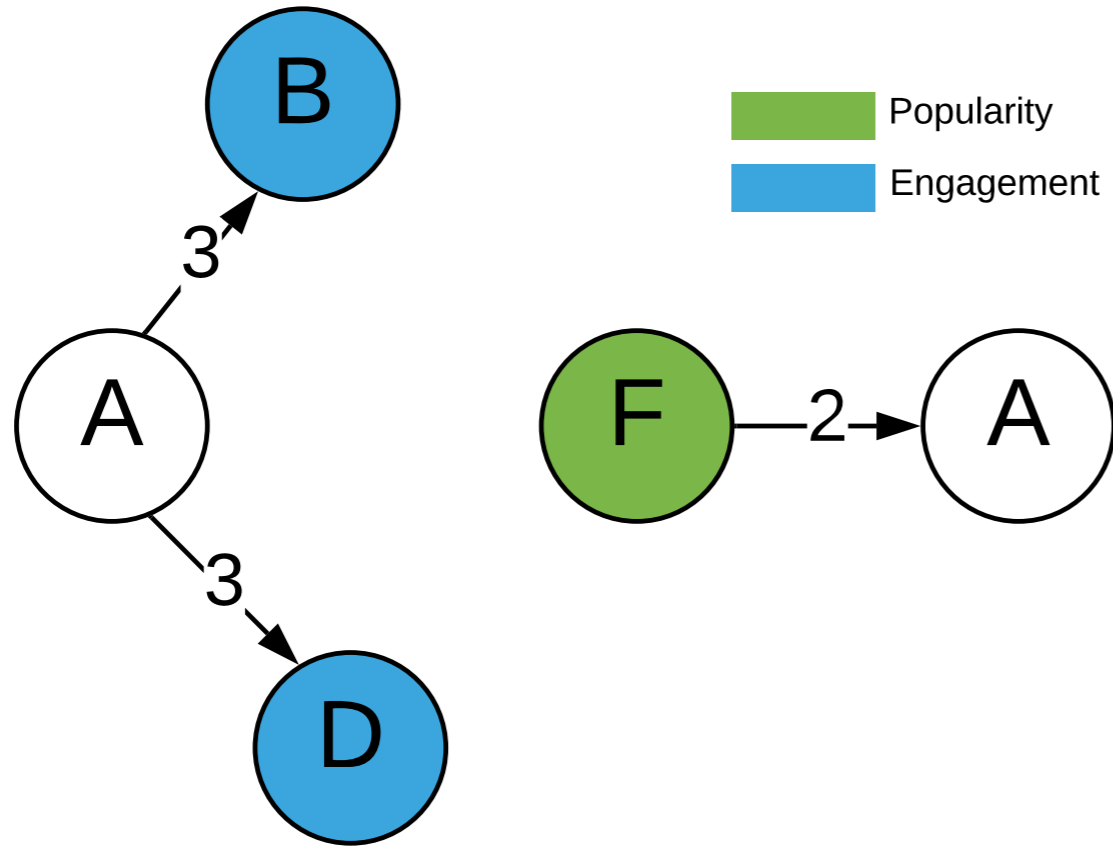
Two types of trust interactions:
popularity trust and engagement trust

Popularity trust: the overall trust of a user endorsed by others (accumulated from the incoming links)

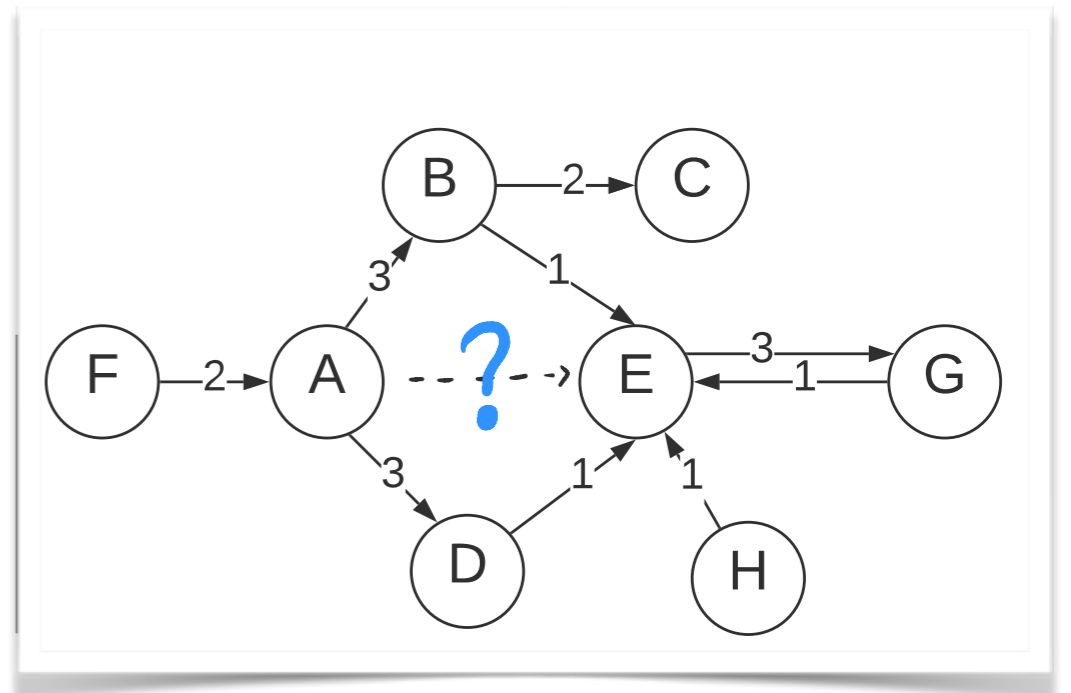


Engagement trust: the willingness of a user to trust others (accumulated from the outgoing links)

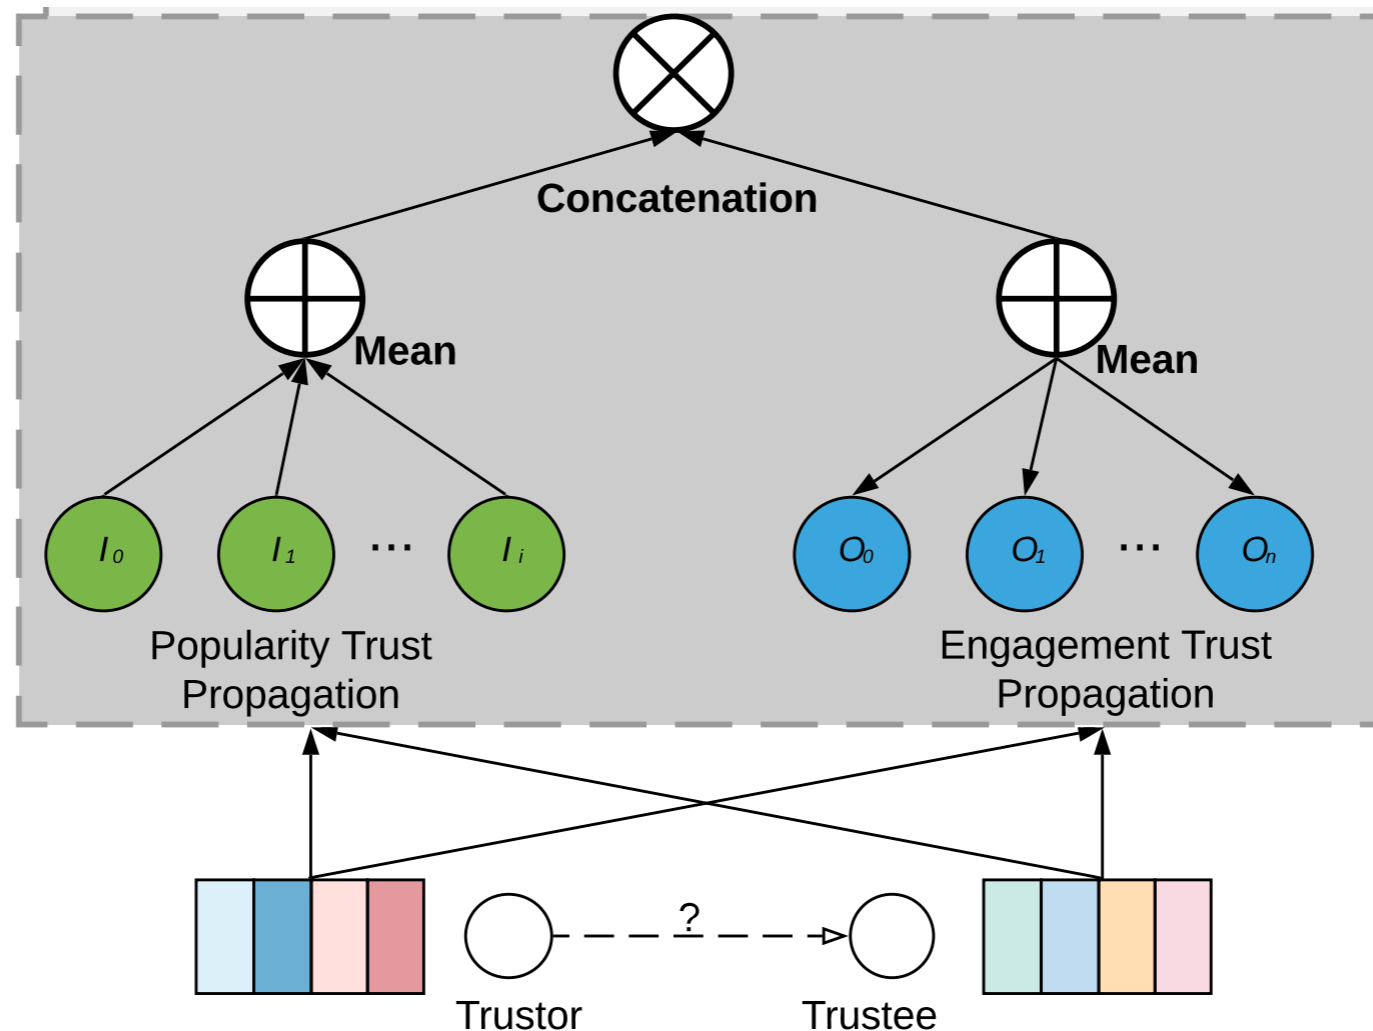




Two types of trust aggregation

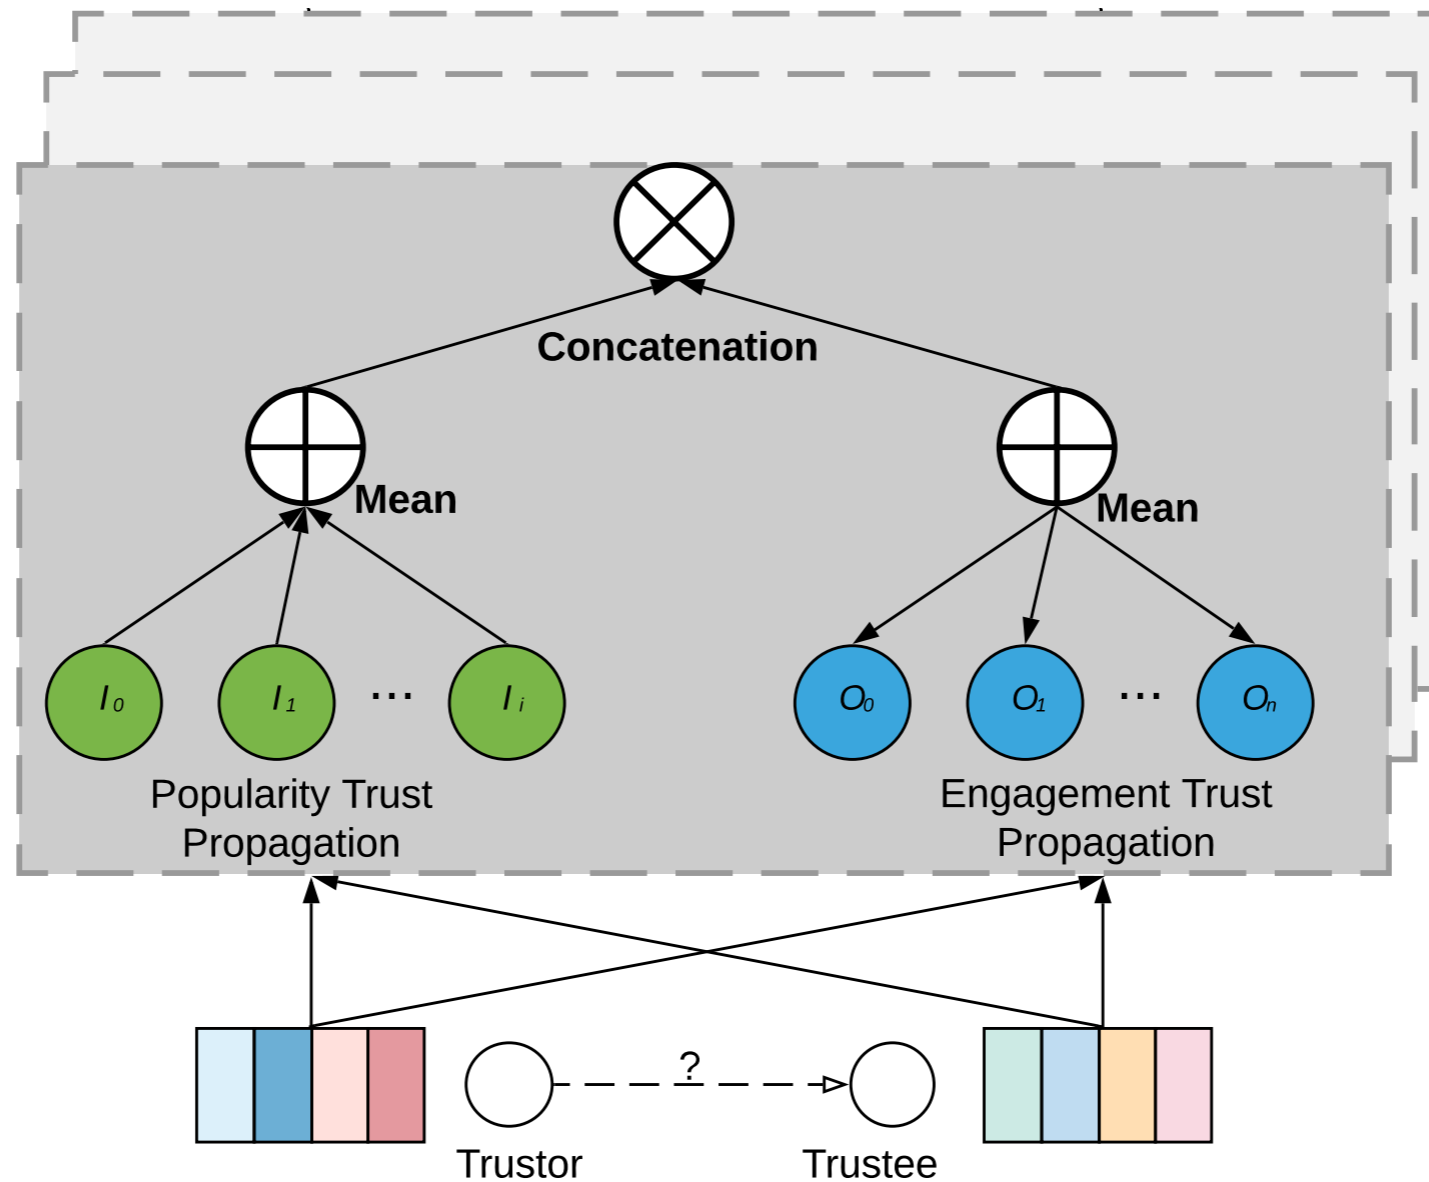


Trust convolutional layer



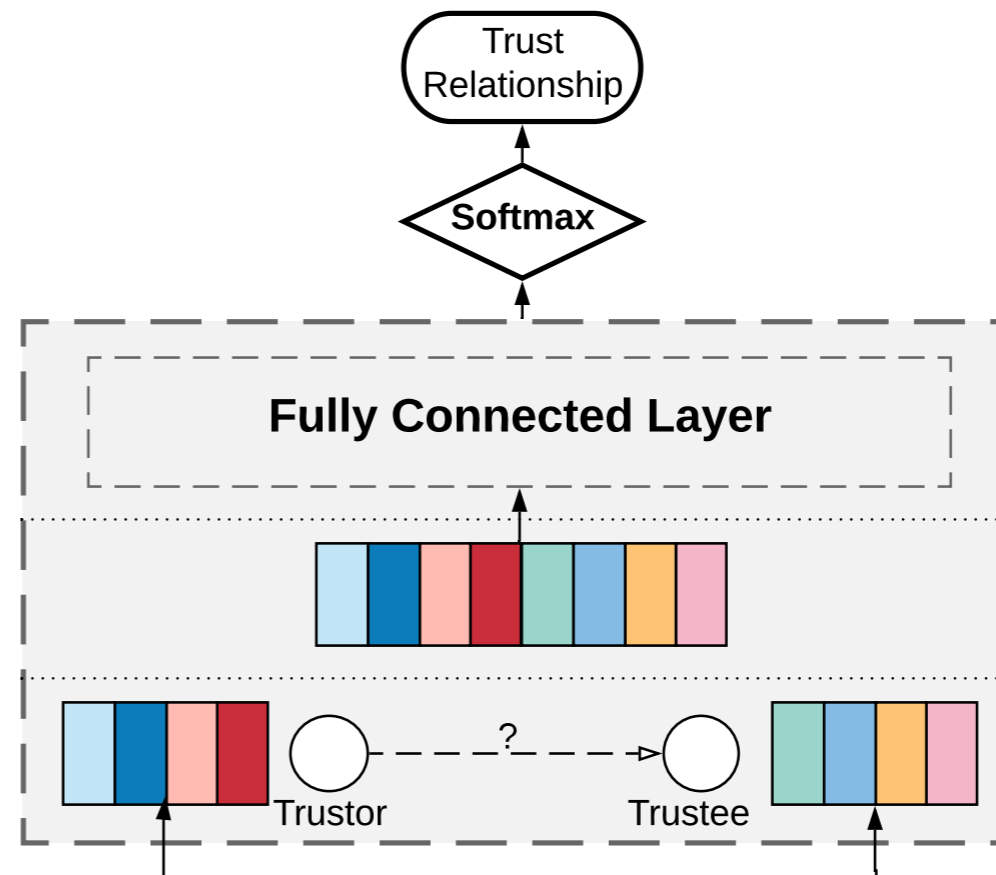
To capture the **composable** and **asymmetric** nature of trust

Stack multiple trust convolutional layers

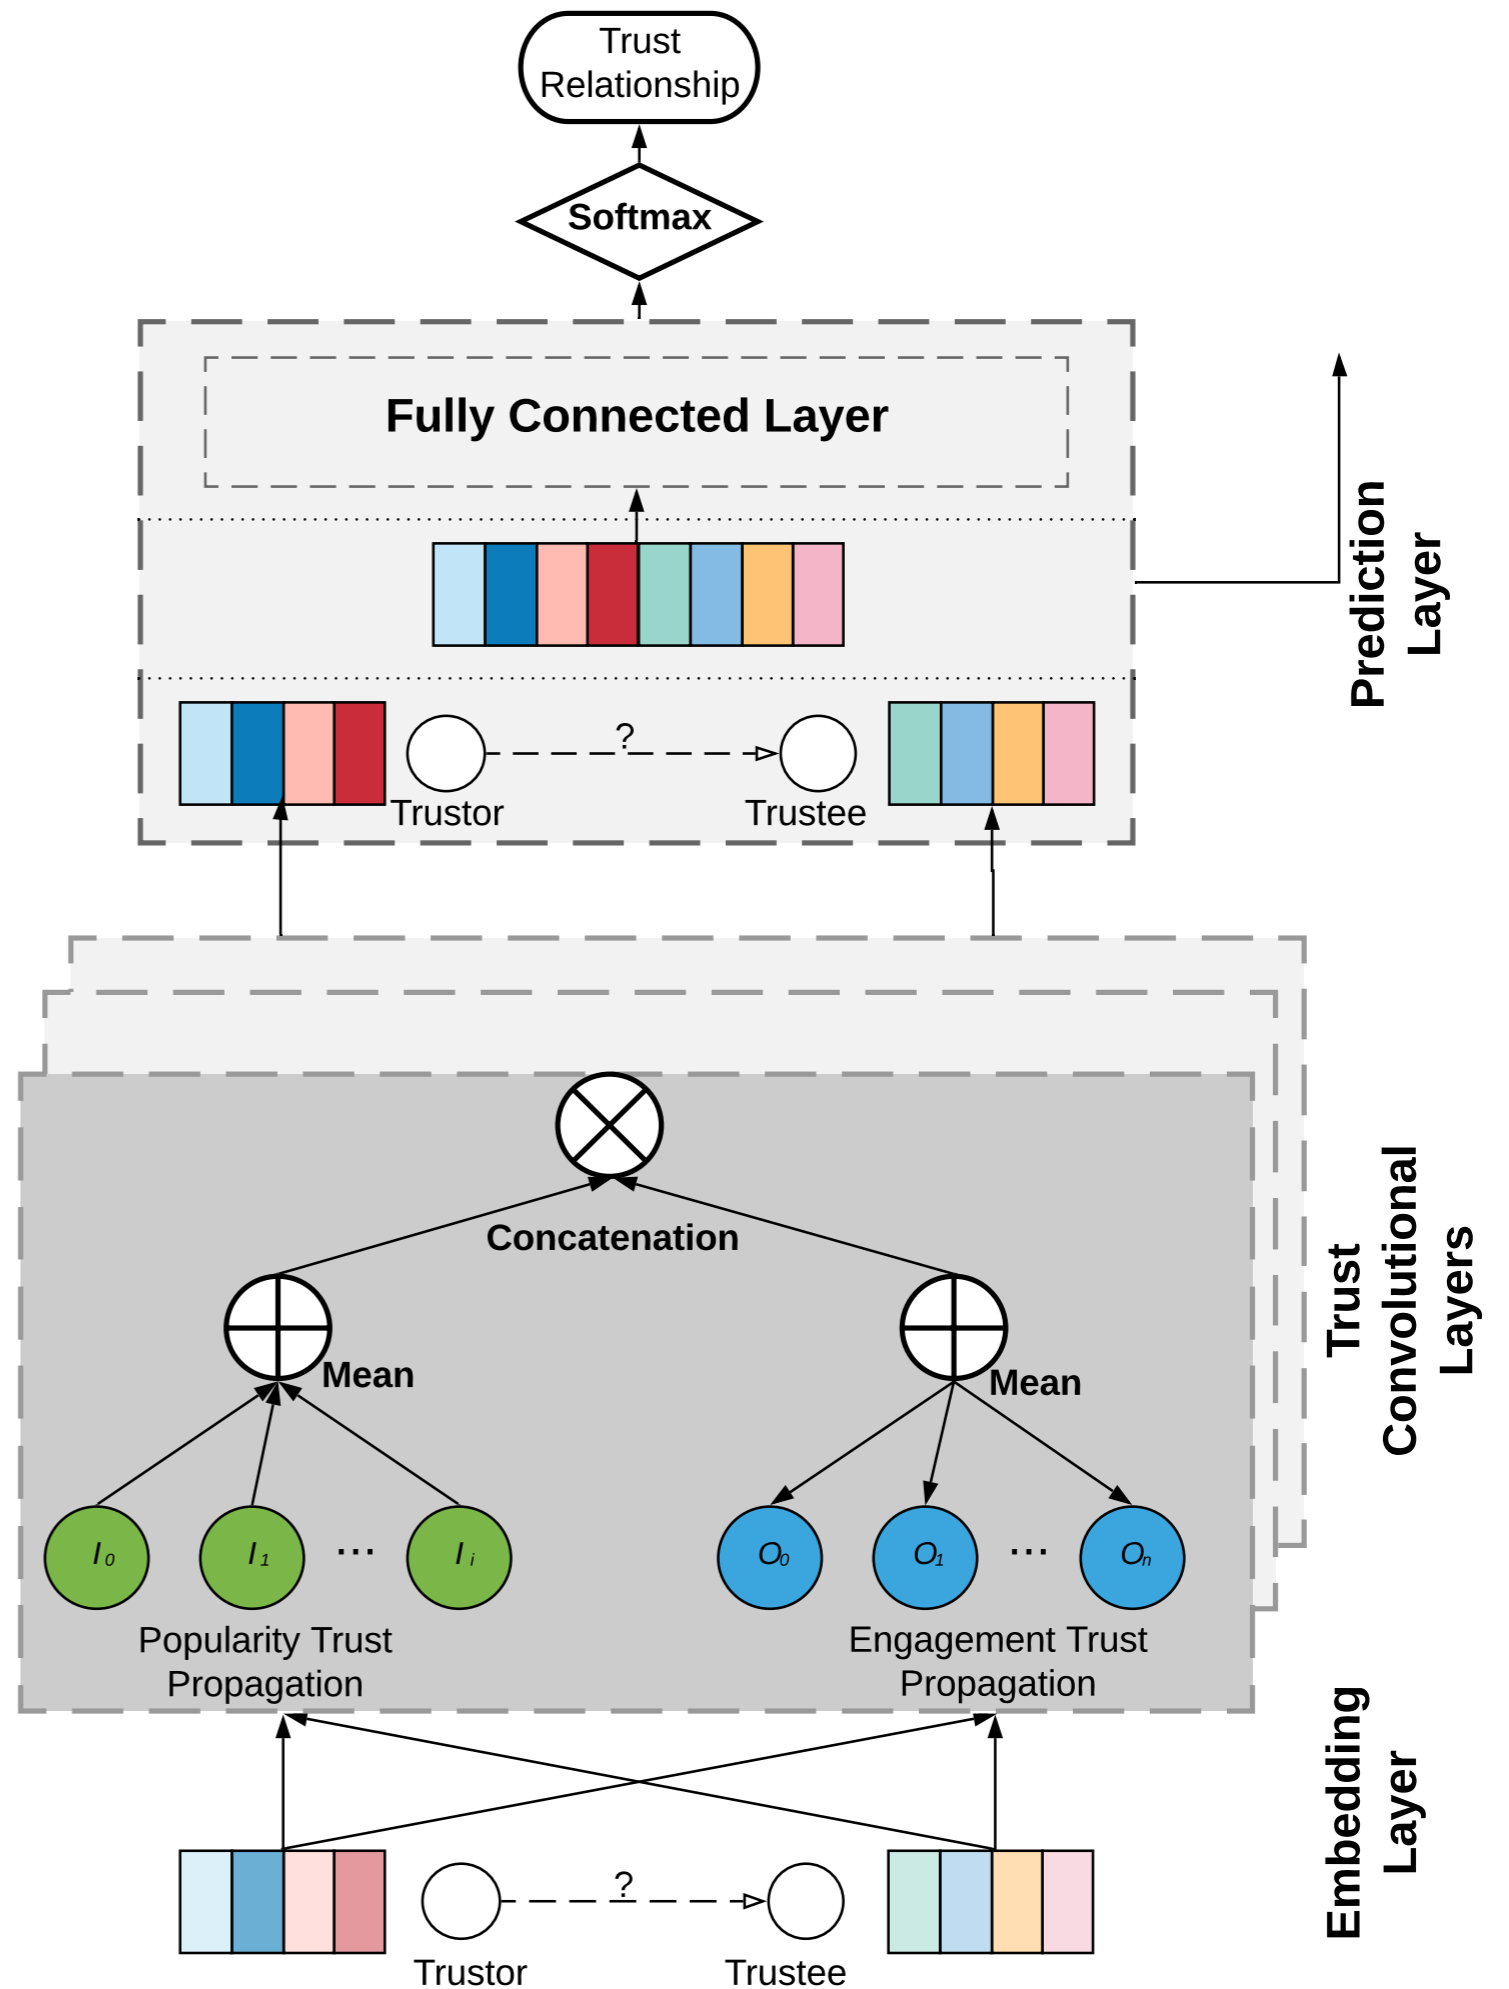


To capture the **propagative nature** of trust

Prediction layer



Guardian



Our experimental results...

Datasets Used

Advogato and Pretty-Good-Privacy (PGP) adopt the concept of the “web of trust”, and both contain four different levels of trust.

DATASET	# OF NODES	# OF EDGES	AVG. DEGREE	DIAMETER
ADVOGATO	6,541	51,127	19.2	4.82
PGP	38,546	317,979	16.5	7.7

Accuracy

Evaluation Accuracy on Advogato

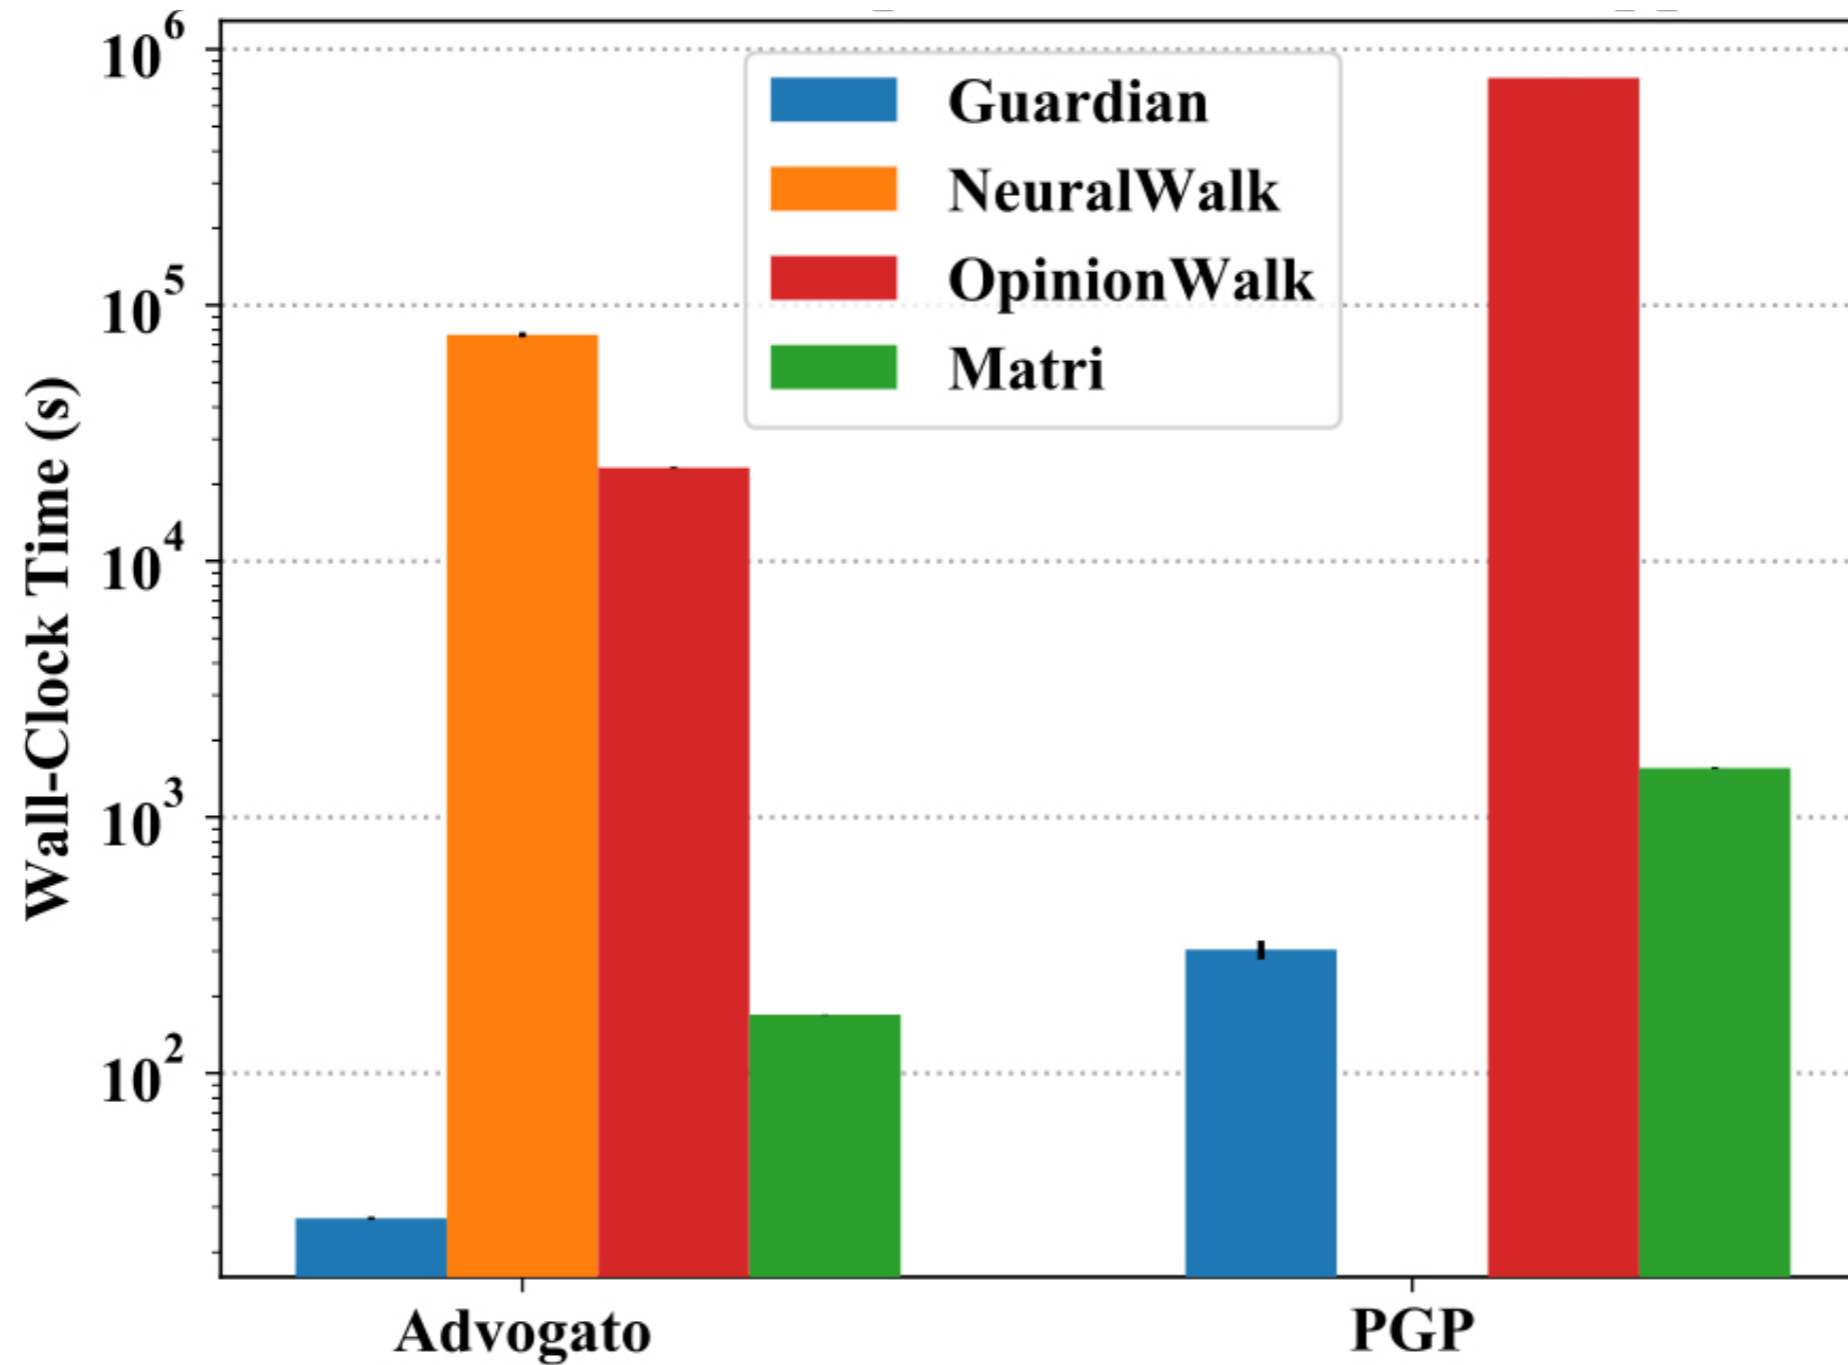
APPROACHES	F1-SCORE	MAE
<i>Guardian</i>	74.3%	0.082
NEURALWALK	74.0%	0.081
OPINIONWALK	64.3%	0.228
MATRI	65.6%	0.127

Accuracy

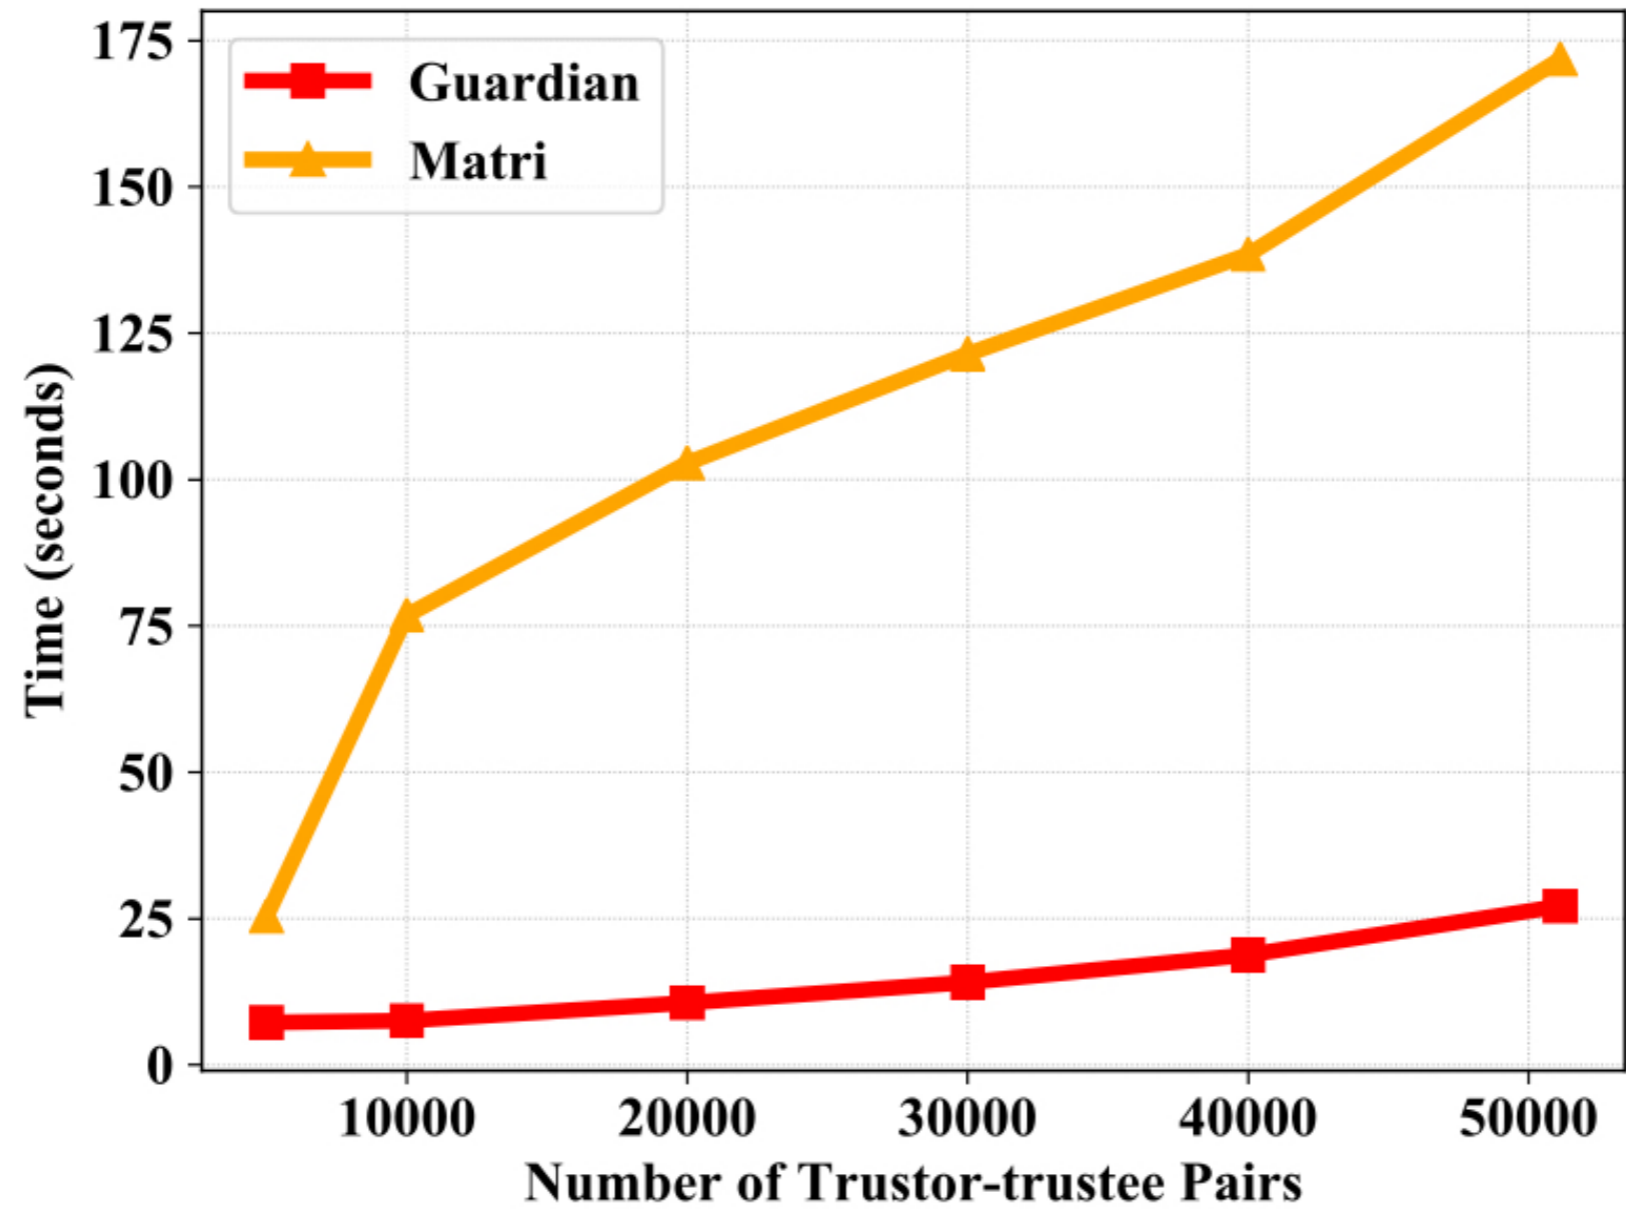
Evaluation Accuracy on PGP

APPROACHES	F1-SCORE	MAE
<i>Guardian</i>	87.1%	0.083
NEURALWALK	—	—
OPINIONWALK	67.3%	0.249
MATRI	68.3%	0.122

Efficiency

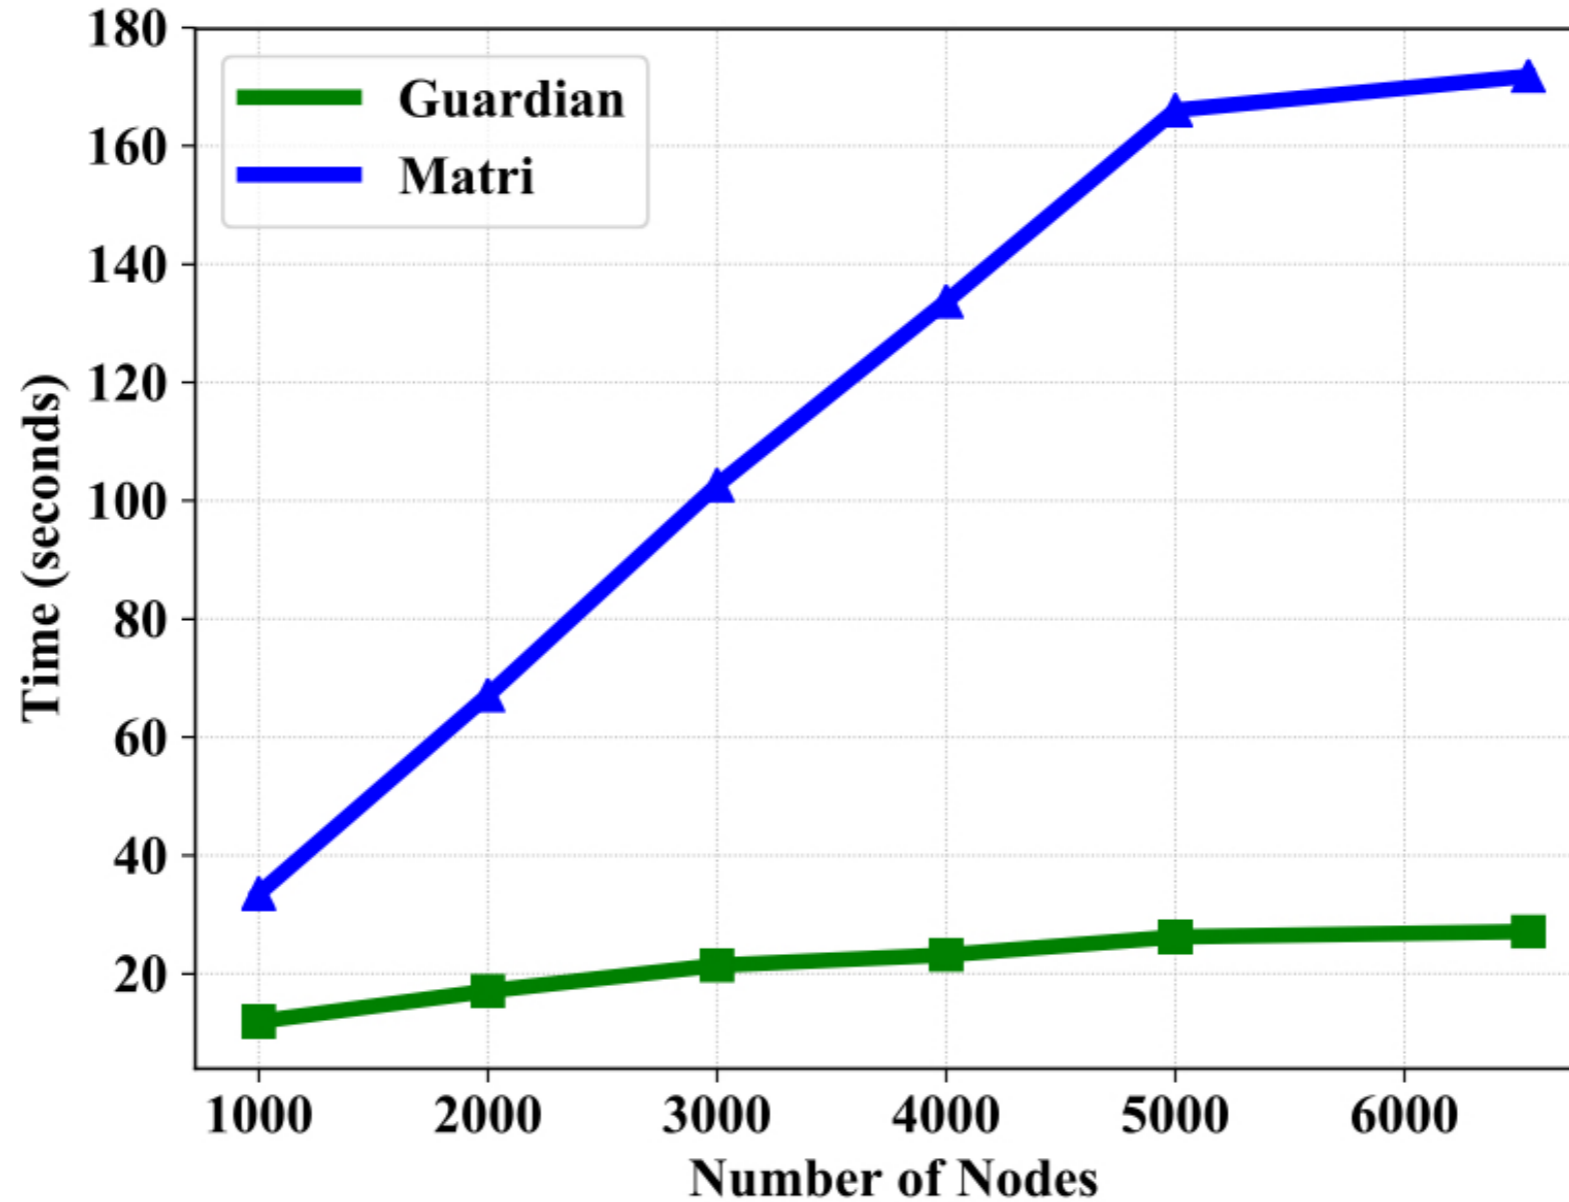


Scalability



Time vs. # of pairs

Scalability



Time vs. # of users

Guardian is an end-to-end learning framework, that can achieve the best possible performance for social trust evaluation in online social networks.



Wanyu Lin, Zhaolin Gao, Baochun Li
University of Toronto