Friendship paradox

How many friends do you have?

You have more friends than I do
(at least on average)

Sadamori Kojaku, Jisung Yoon, Isabel Constantino, and Yong-Yeol Ahn. Residual2Vec: Debiasing graph embedding with random graphs. NeurIPS (2021)

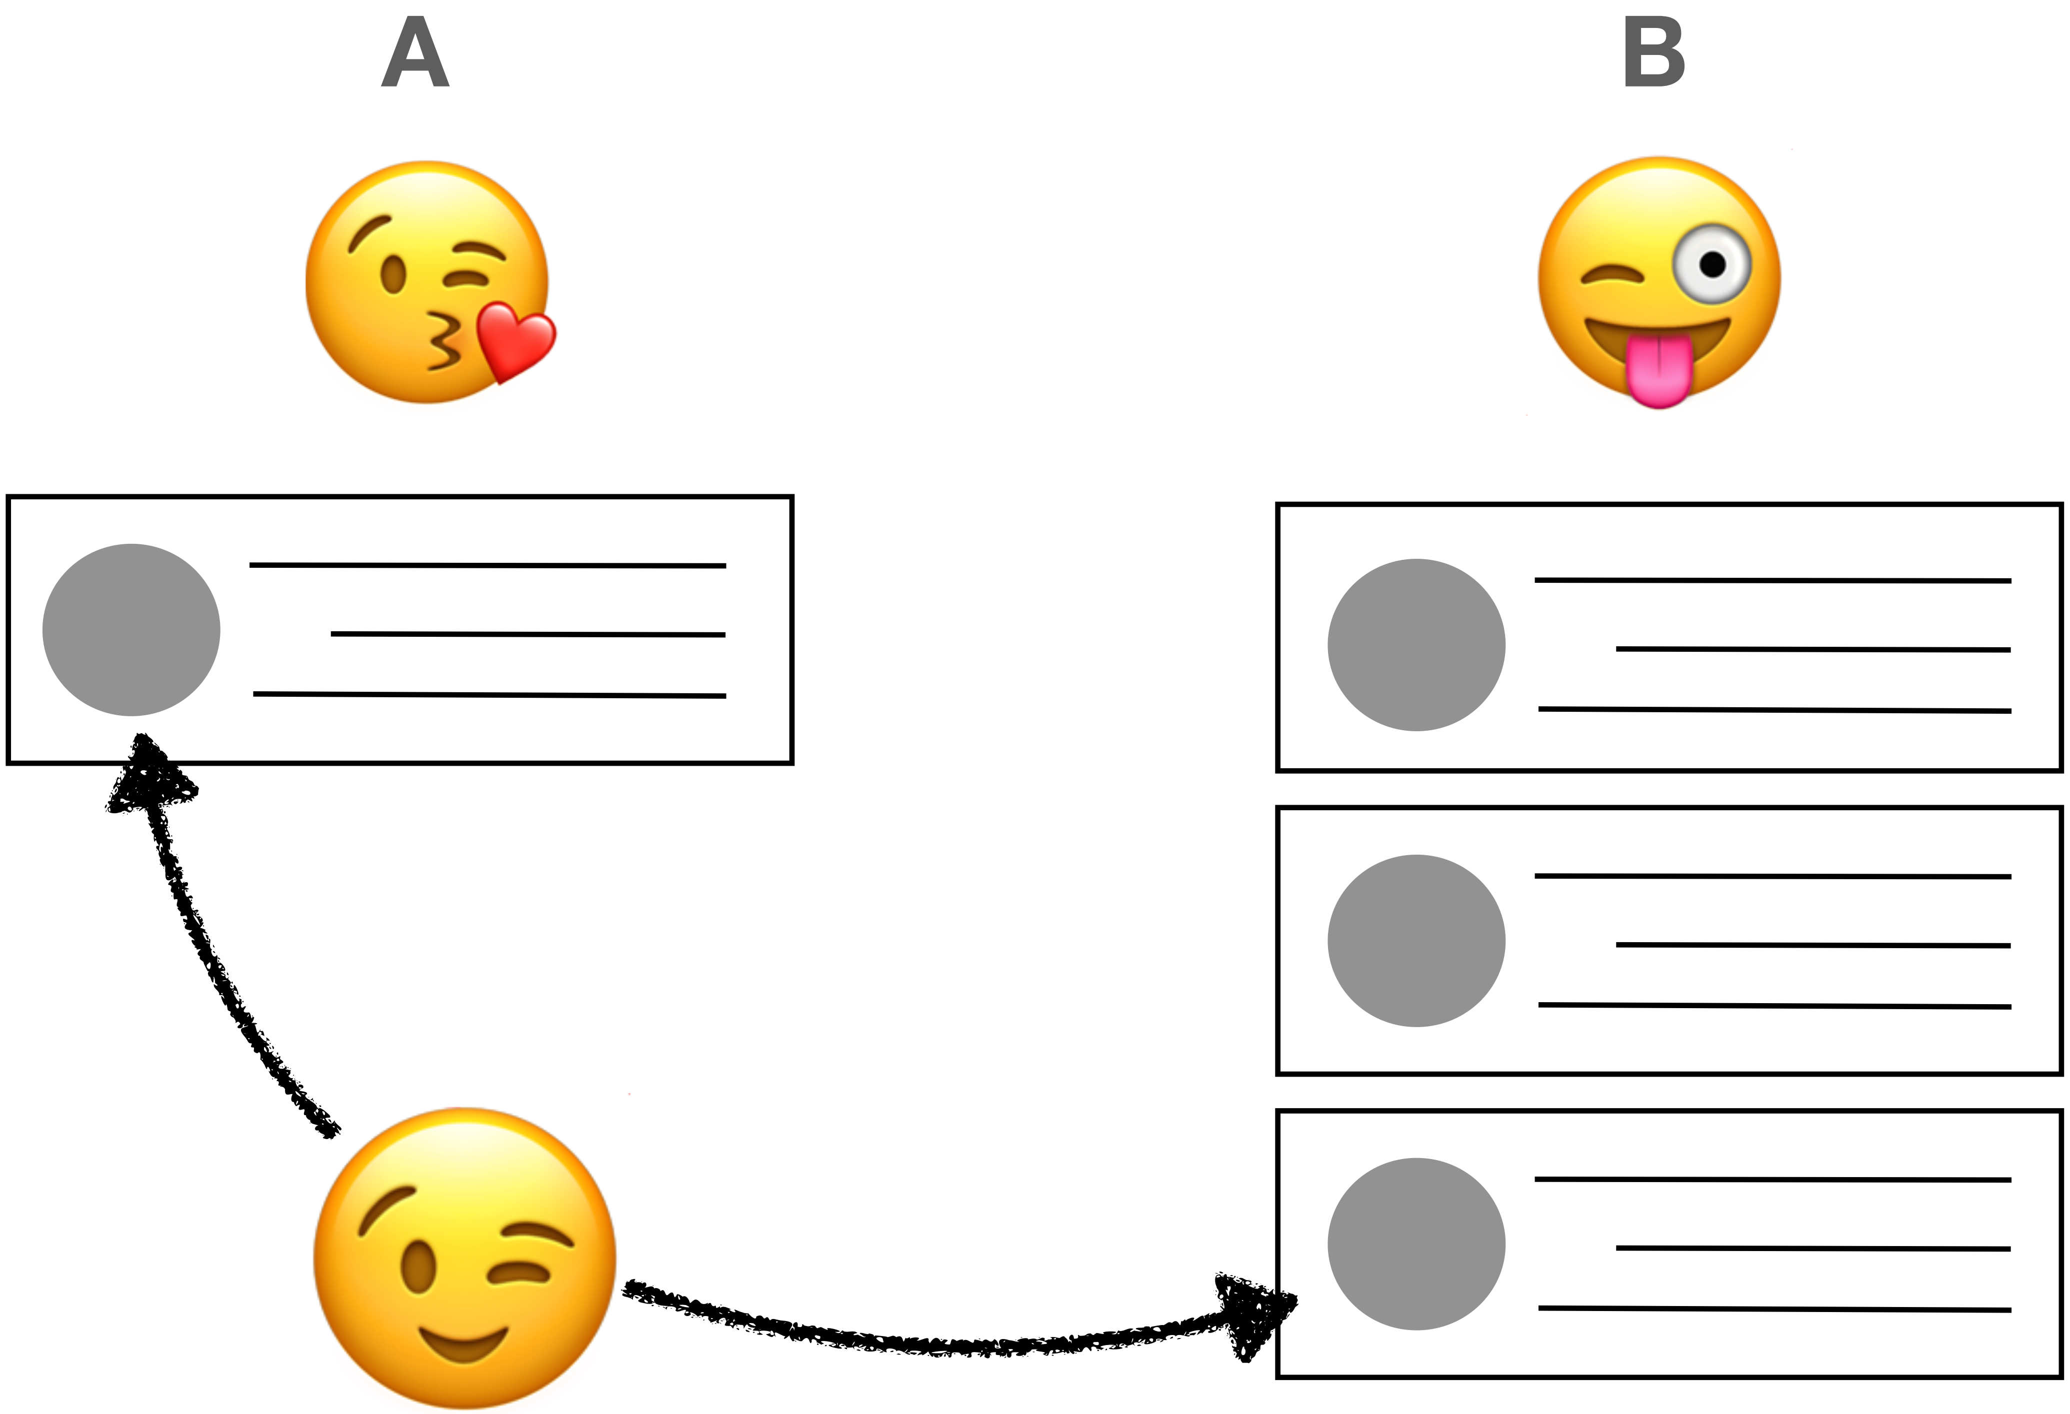

In which friend list am I likely to appear?

Sadamori Kojaku, Jisung Yoon, Isabel Constantino, and Yong-Yeol Ahn. Residual2Vec: Debiasing graph embedding with random graphs. NeurIPS (2021)

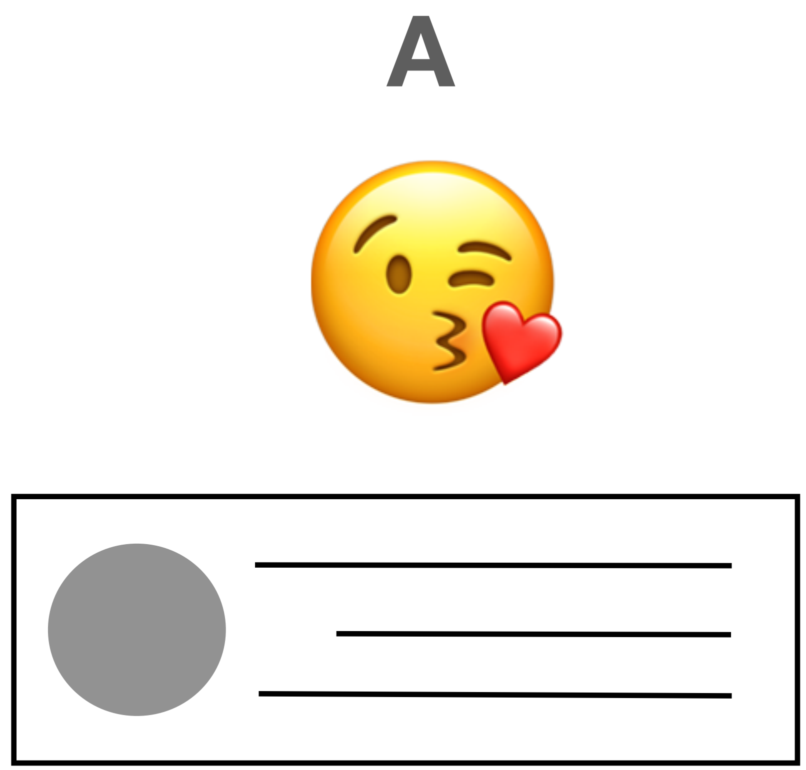

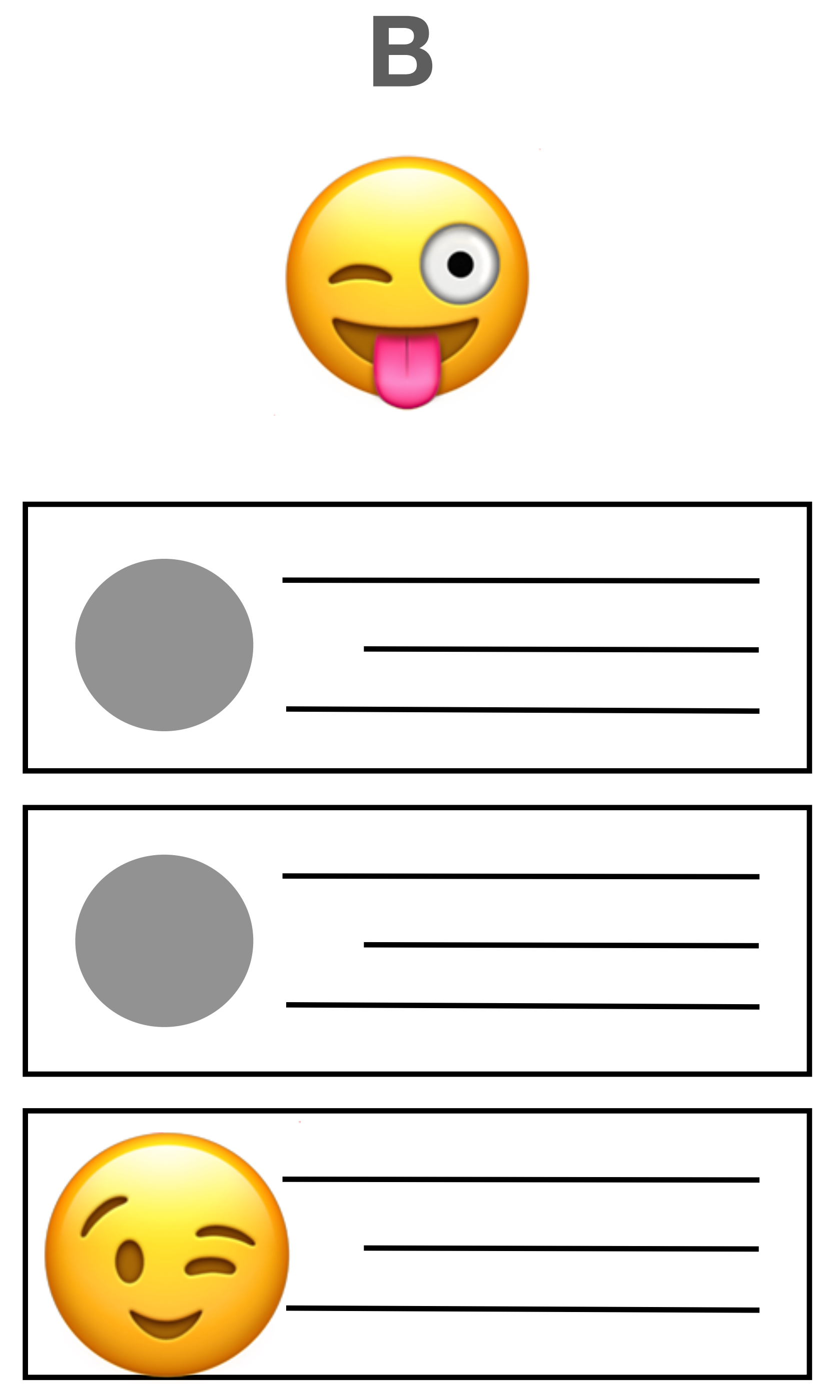

3 times more likely to appear in the B's friend list

Sadamori Kojaku, Jisung Yoon, Isabel Constantino, and Yong-Yeol Ahn. Residual2Vec: Debiasing graph embedding with random graphs. NeurIPS (2021)

Friendship paradox

Your friends have more friends than you do (on average)

...because more friends someone has,
more likely the someone appears in your friend list.

Can introduce *biases* in graph embeddings

Sadamori Kojaku, Jisung Yoon, Isabel Constantino, and Yong-Yeol Ahn. Residual2Vec: Debiasing graph embedding with random graphs. NeurIPS (2021)

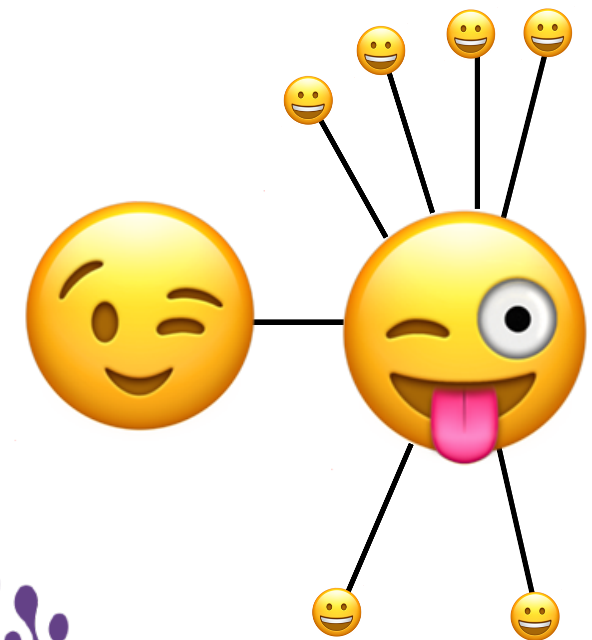

Residual2Vec: Debiasing graph embedding using random graphs

arXiv.org https://arxiv.org/abs/2110.07654

Sadamori Kojaku | Jisung Yoon | Isabel Constantino | Yong-Yeol Ahn

Sadamori Kojaku, Jisung Yoon, Isabel Constantino, and Yong-Yeol Ahn. Residual2Vec: Debiasing graph embedding with random graphs. NeurIPS (2021)

Graph embedding

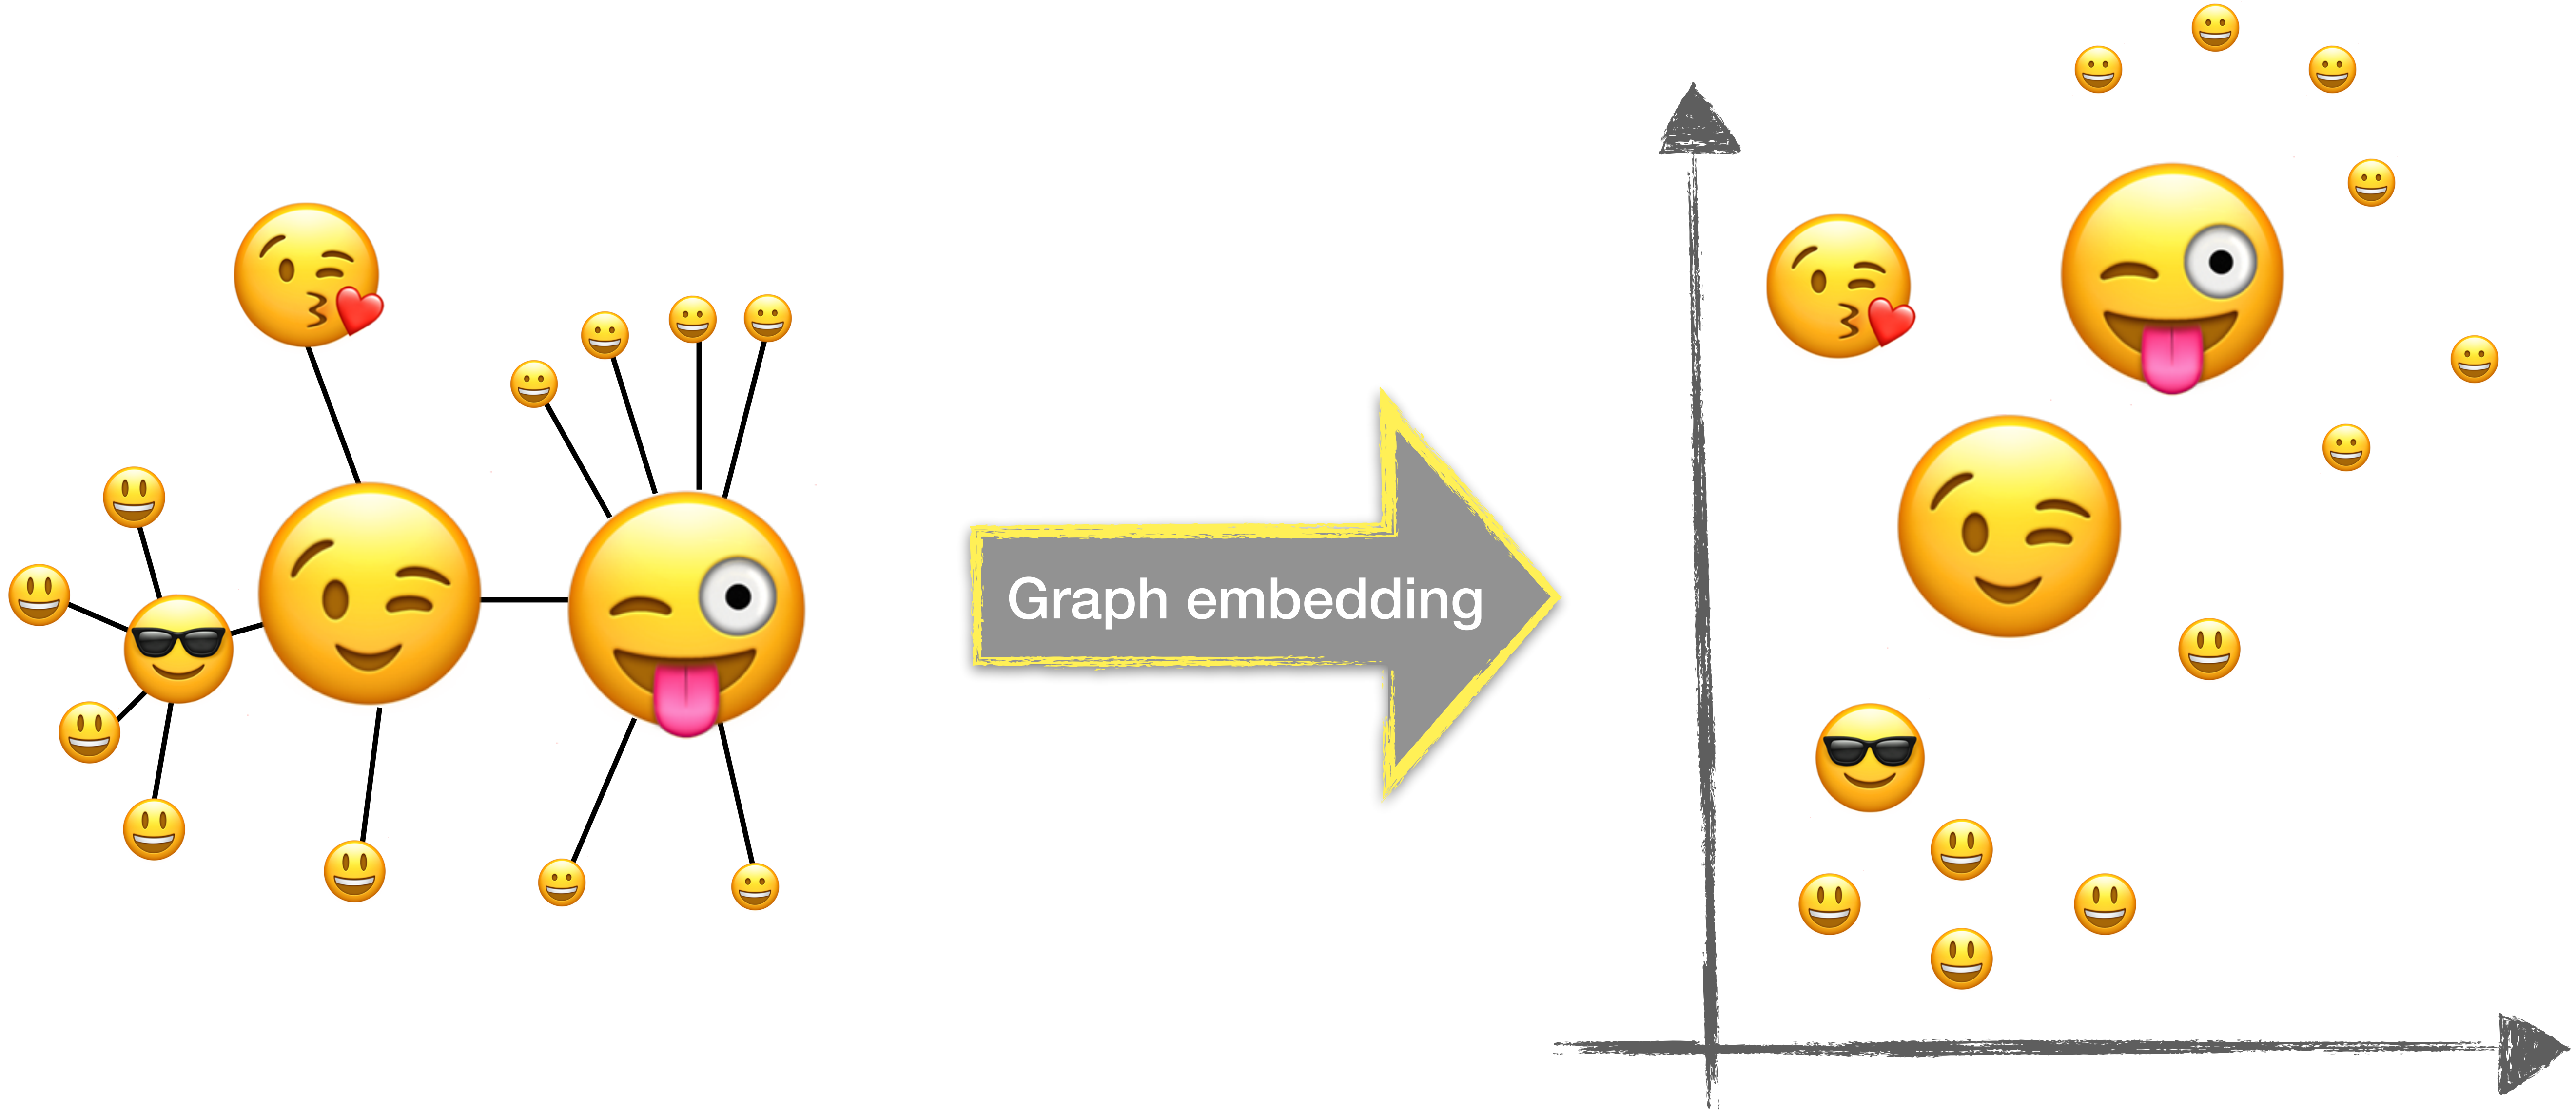

Sadamori Kojaku, Jisung Yoon, Isabel Constantino, and Yong-Yeol Ahn. Residual2Vec: Debiasing graph embedding with random graphs. NeurIPS (2021)

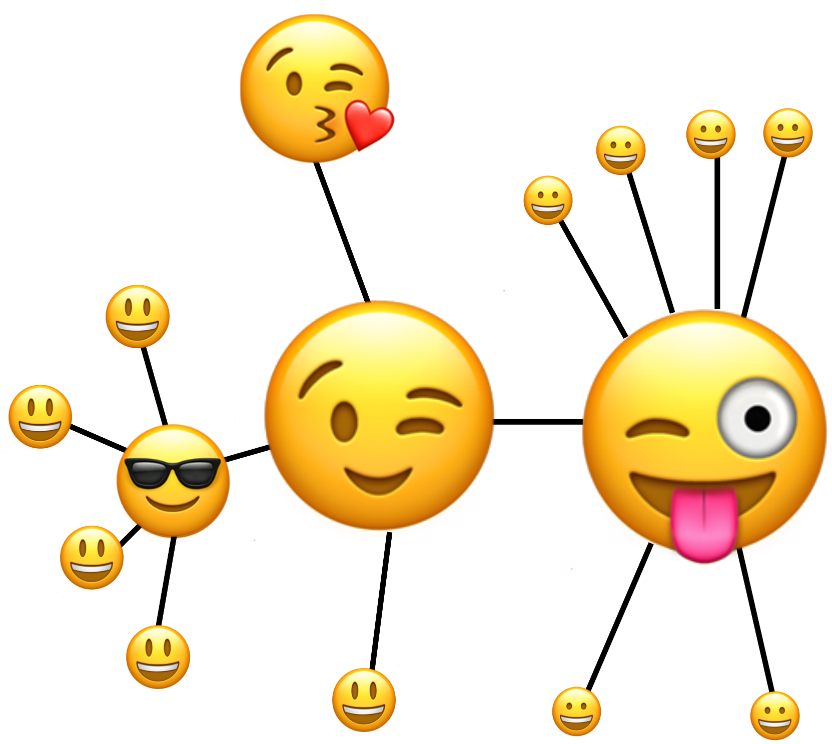

Sequence of nodes

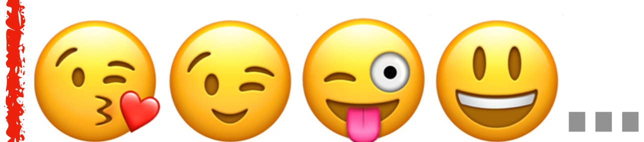

Embed by word2vec

Word2Vec*

Word embedding

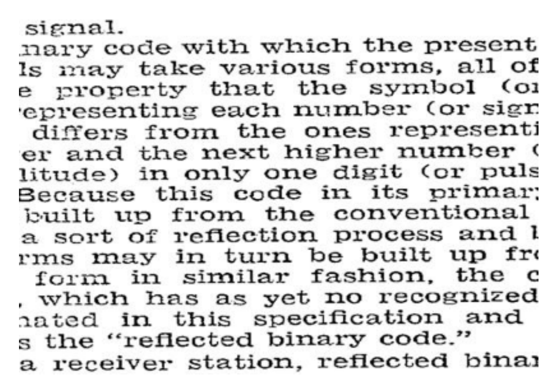

signal.
nary code with which the present
ls may take various forms, all of
e property that the symbol (o
epresenting each number (or sign
differs from the ones representi
er and the next higher number (
litude) in only one digit (or puls
Because this code in its primar
built up from the conventional
a sort of reflection process and l
rms may in turn be built up fr
form in similar fashion, the c
, which has as yet no recognized
nated in this specification and
s the "reflected binary code."
a receiver station, reflected bina

Sadamori Kojaku, Jisung Yoon, Isabel Constantino, and Yong-Yeol Ahn. Residual2Vec: Debiasing graph embedding with random graphs. NeurIPS (2021)

How to generate the sequence from a graph?

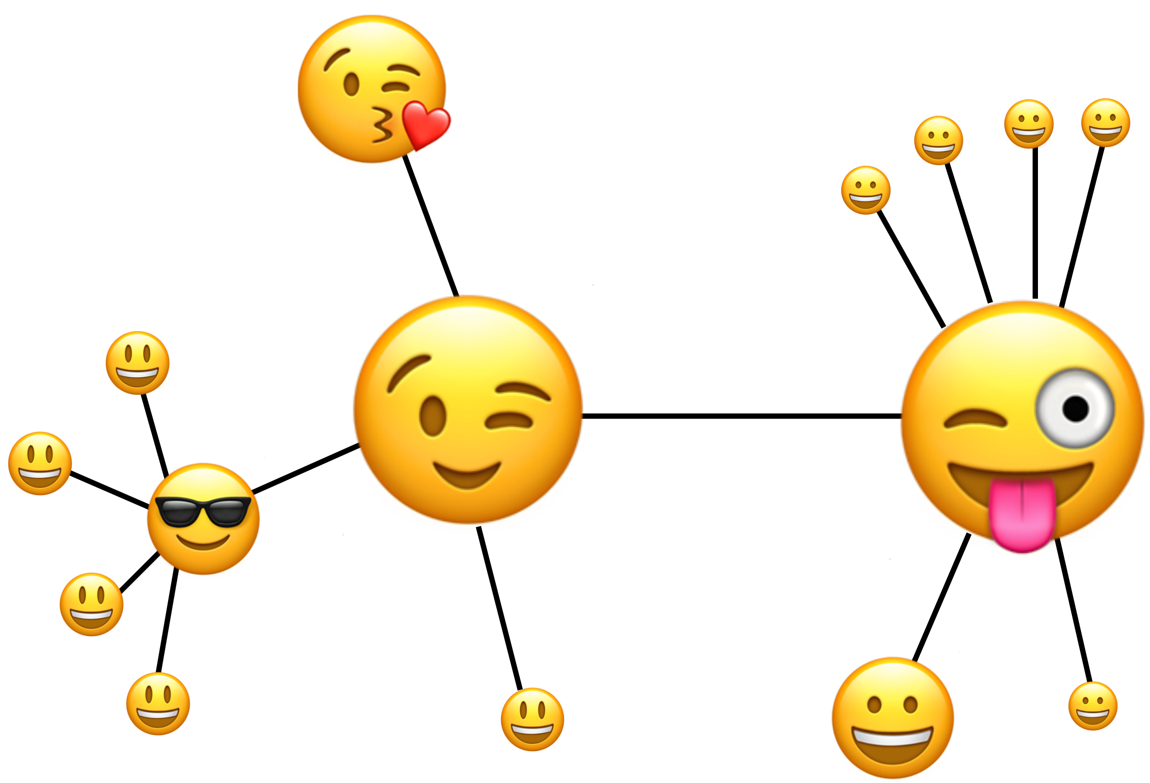

Sequence of nodes

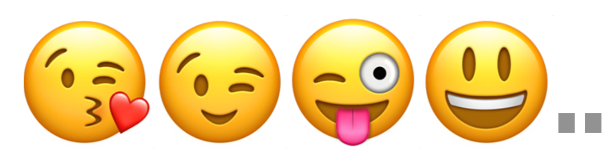

Sadamori Kojaku, Jisung Yoon, Isabel Constantino, and Yong-Yeol Ahn. Residual2Vec: Debiasing graph embedding with random graphs. NeurIPS (2021)

Random walker

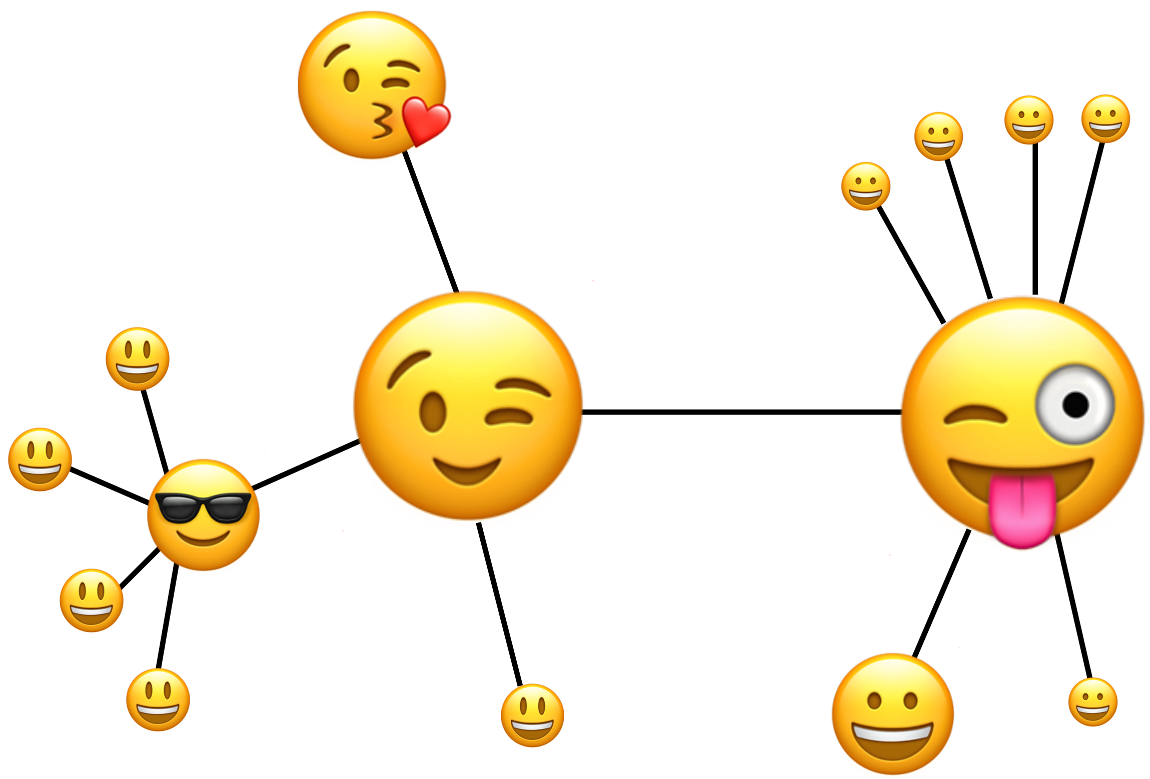

Sadamori Kojaku, Jisung Yoon, Isabel Constantino, and Yong-Yeol Ahn. Residual2Vec: Debiasing graph embedding with random graphs. NeurIPS (2021)

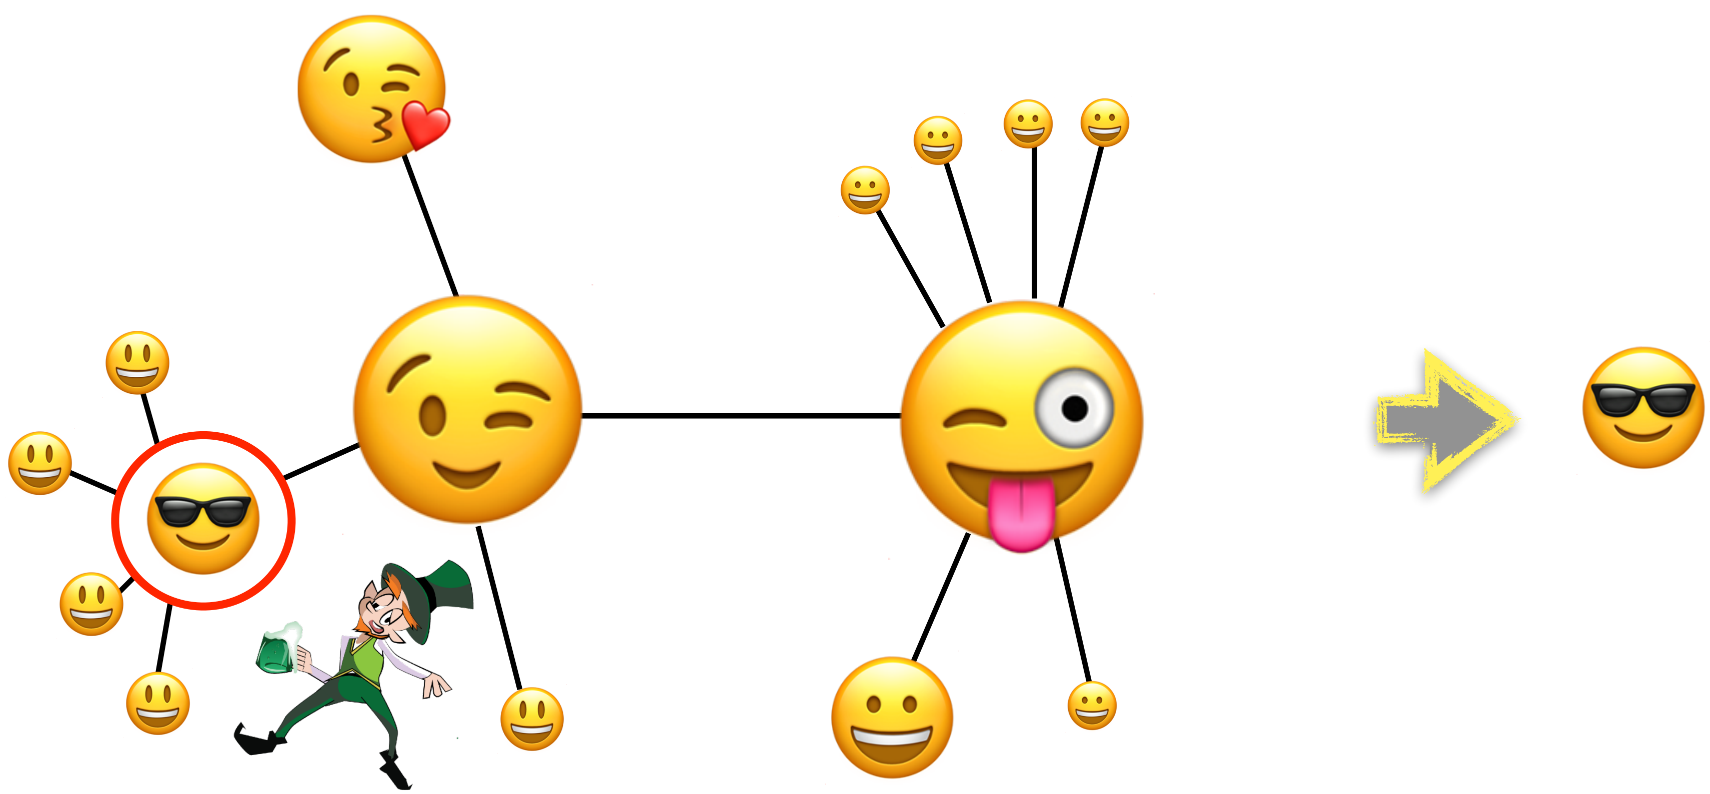

Sadamori Kojaku, Jisung Yoon, Isabel Constantino, and Yong-Yeol Ahn. Residual2Vec: Debiasing graph embedding with random graphs. NeurIPS (2021)

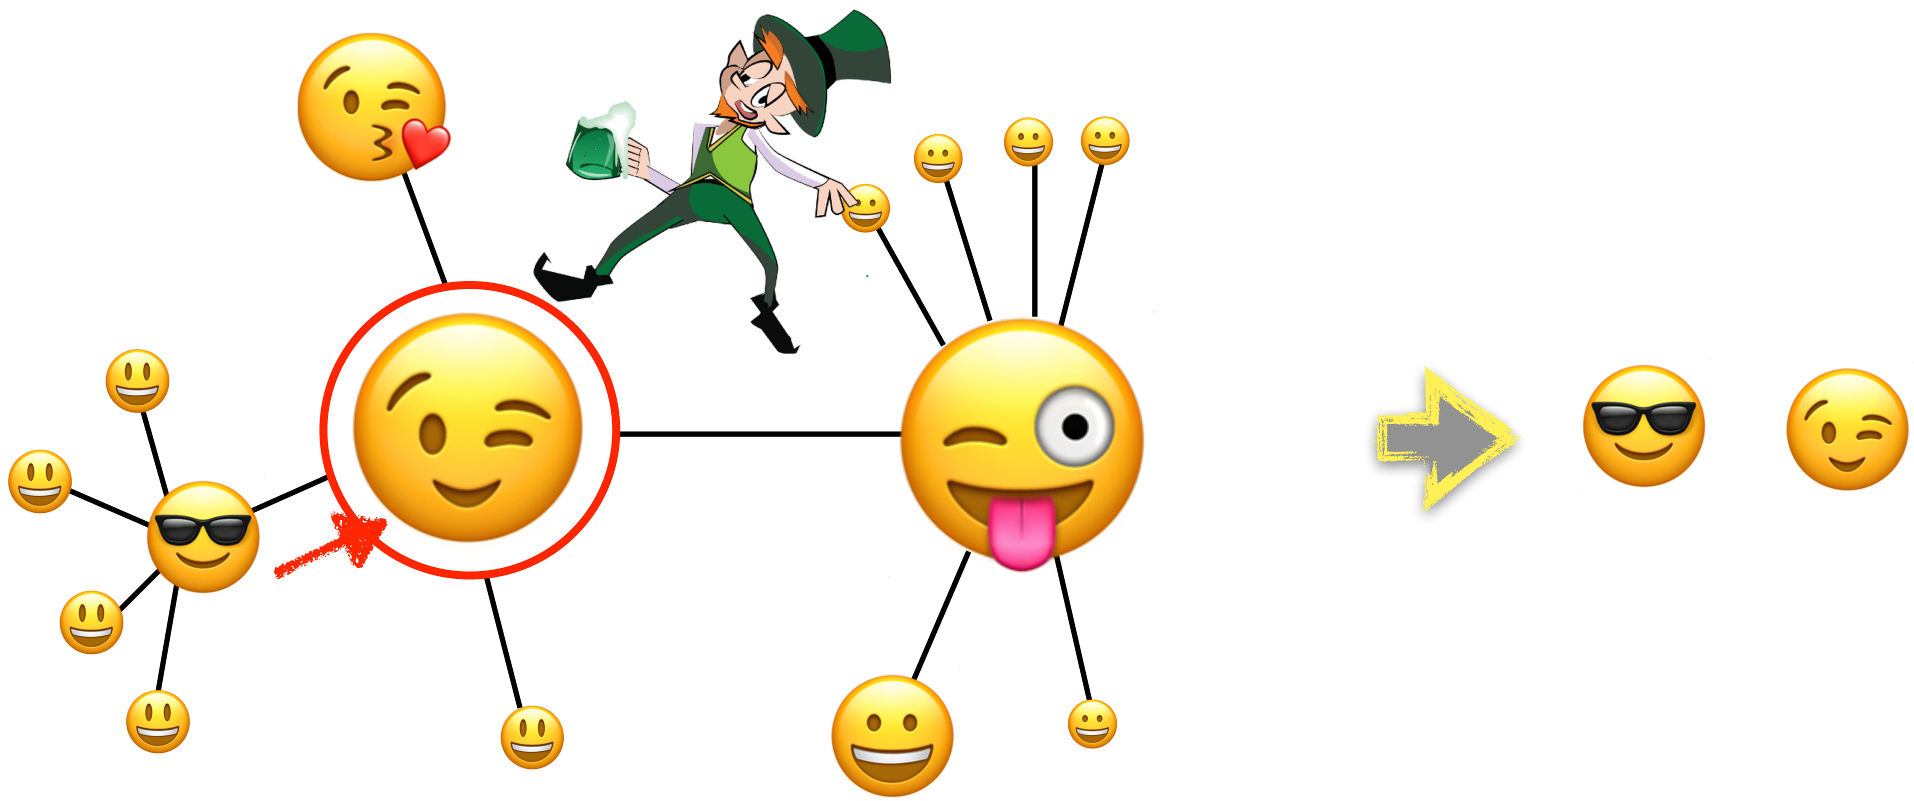

Sadamori Kojaku, Jisung Yoon, Isabel Constantino, and Yong-Yeol Ahn. Residual2Vec: Debiasing graph embedding with random graphs. NeurIPS (2021)

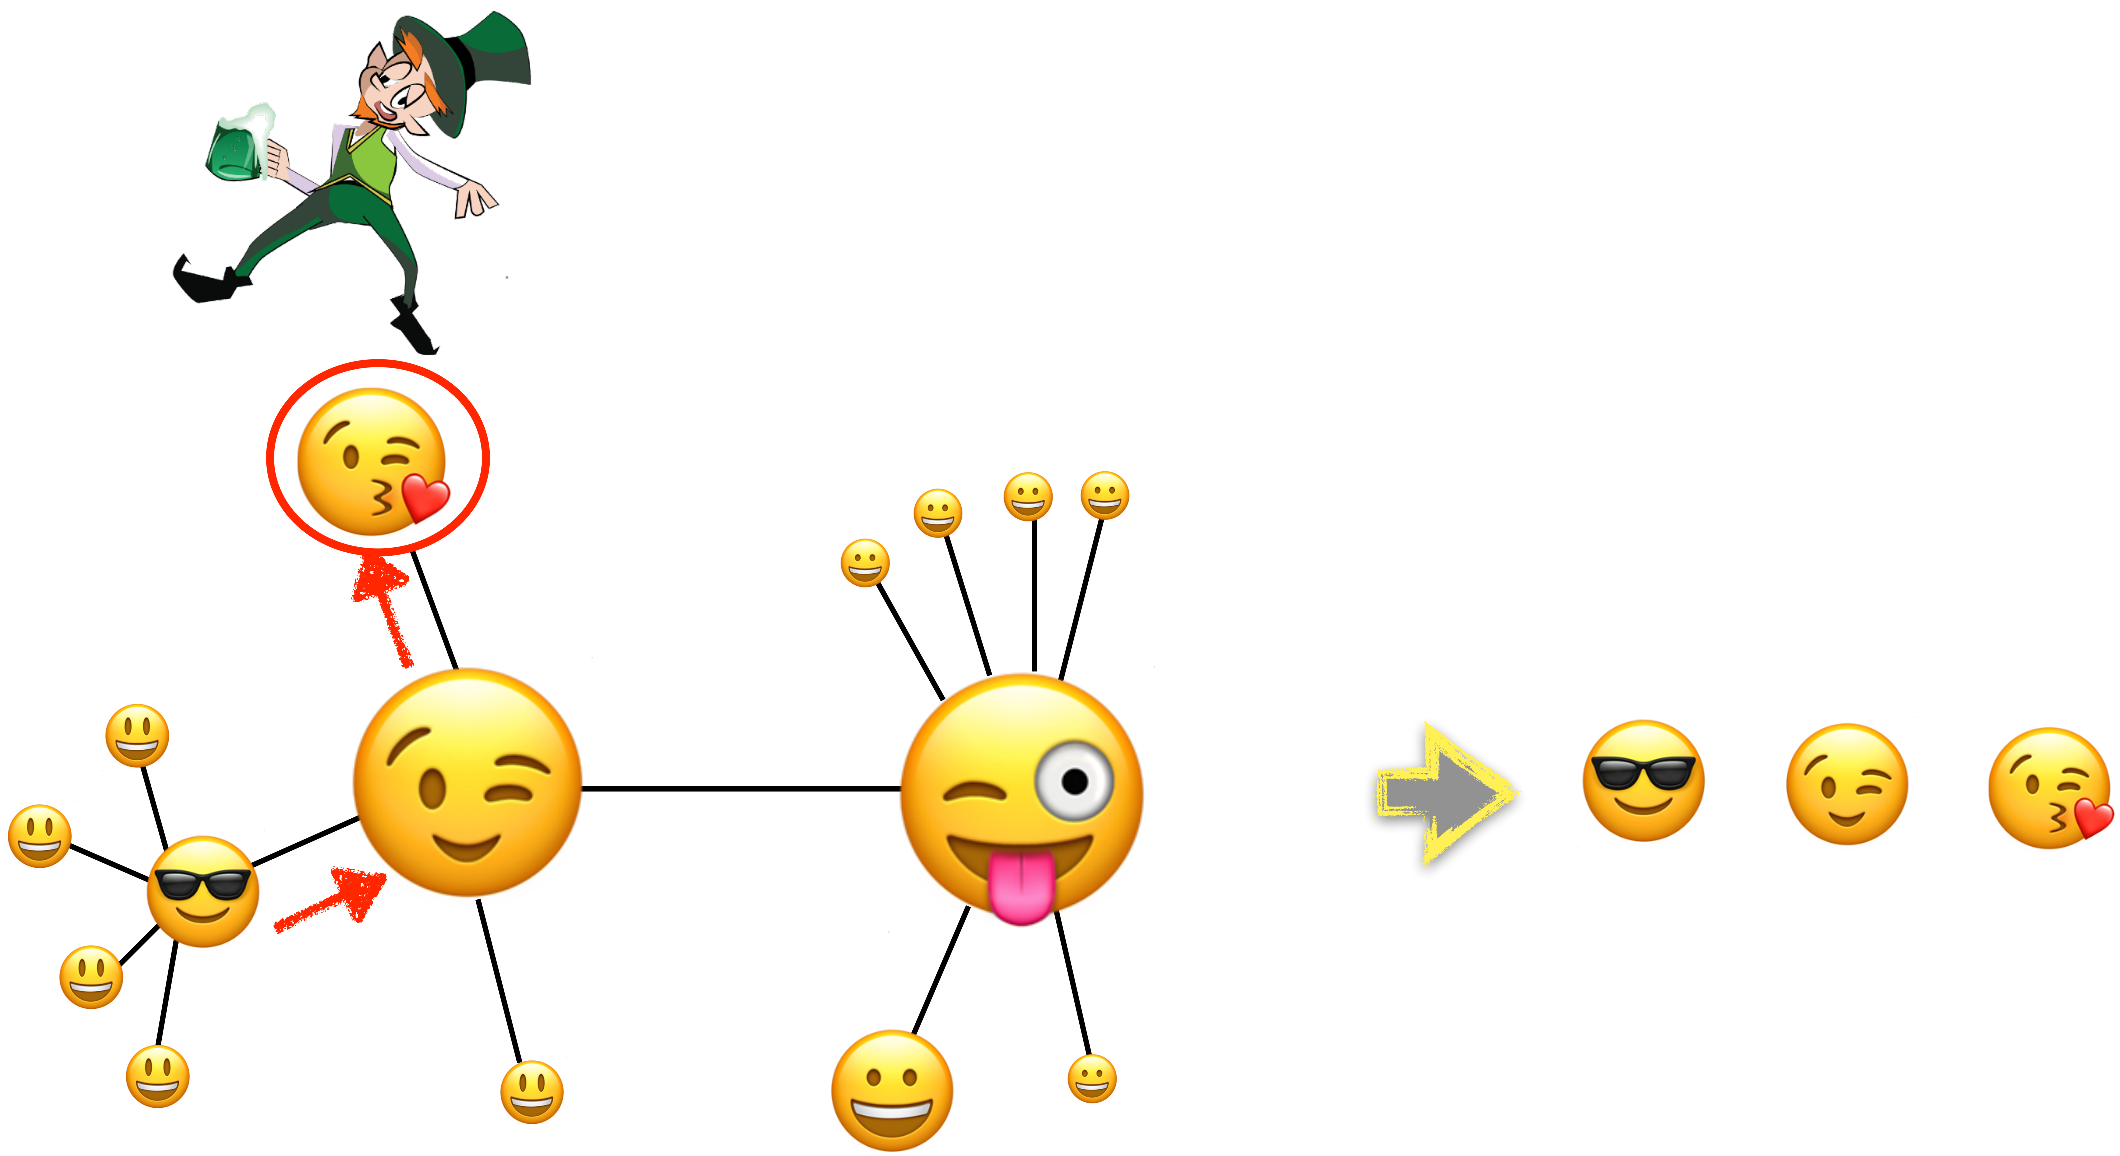

Sadamori Kojaku, Jisung Yoon, Isabel Constantino, and Yong-Yeol Ahn. Residual2Vec: Debiasing graph embedding with random graphs. NeurIPS (2021)

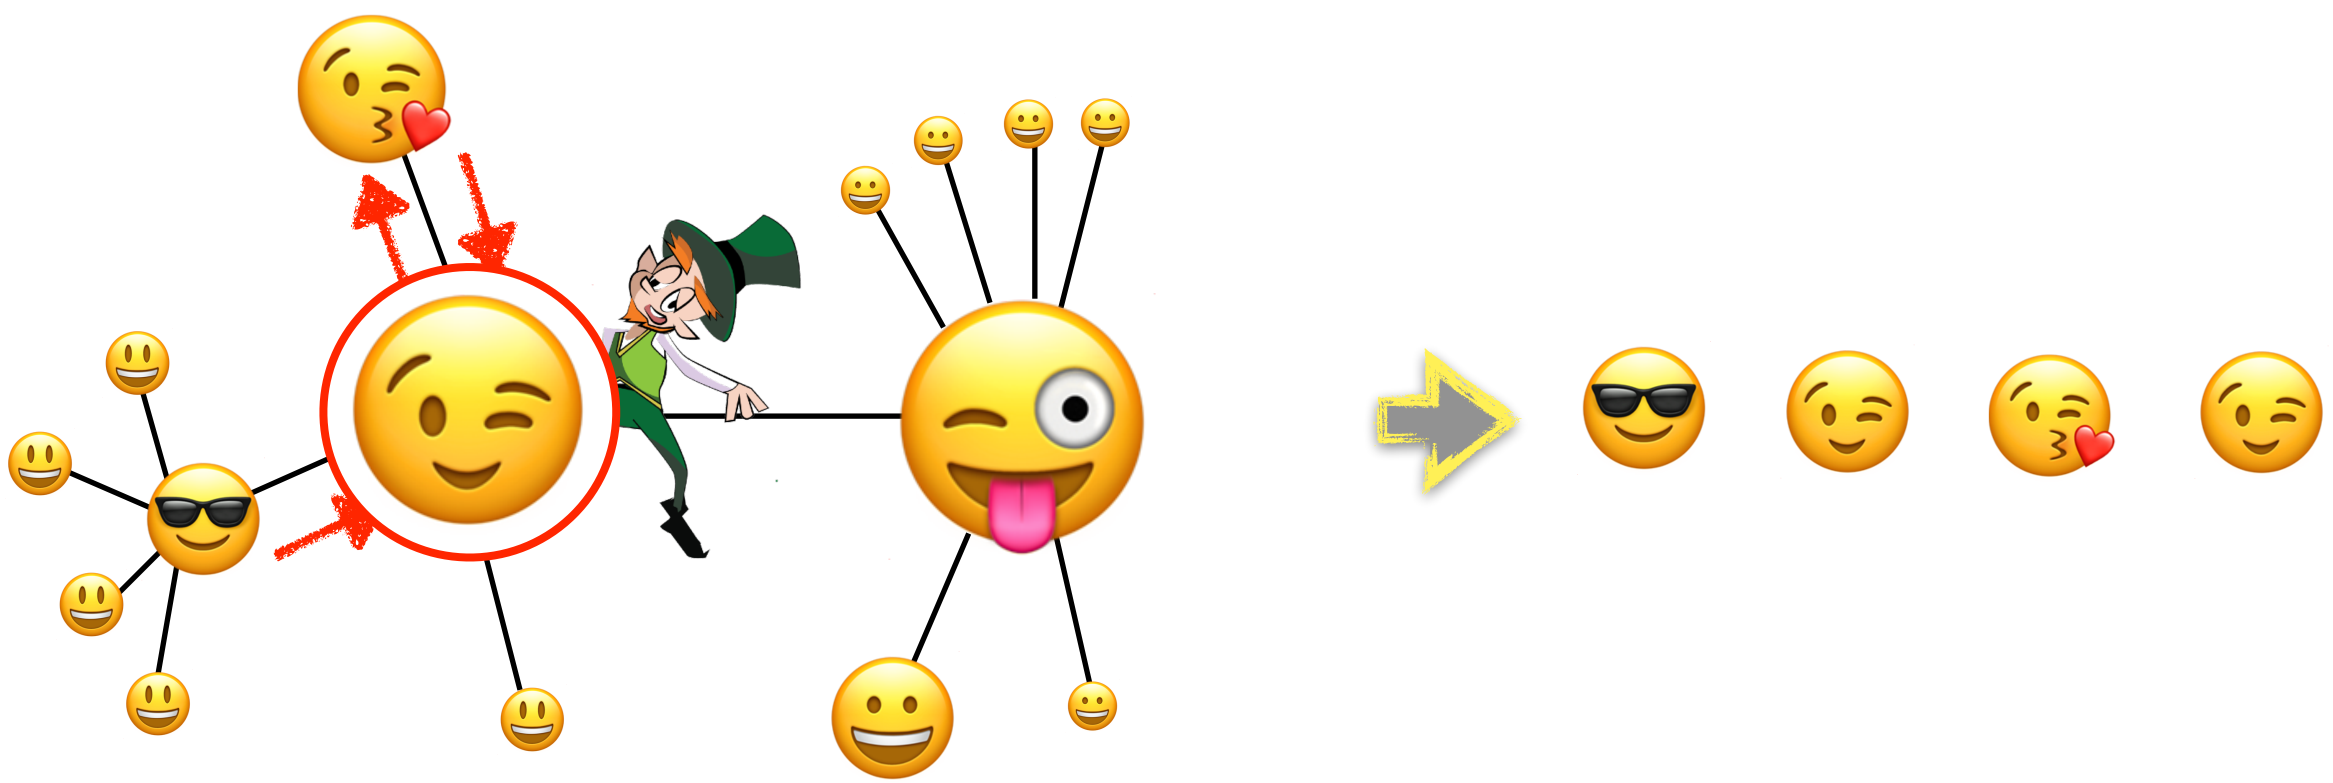

Sadamori Kojaku, Jisung Yoon, Isabel Constantino, and Yong-Yeol Ahn. Residual2Vec: Debiasing graph embedding with random graphs. NeurIPS (2021)

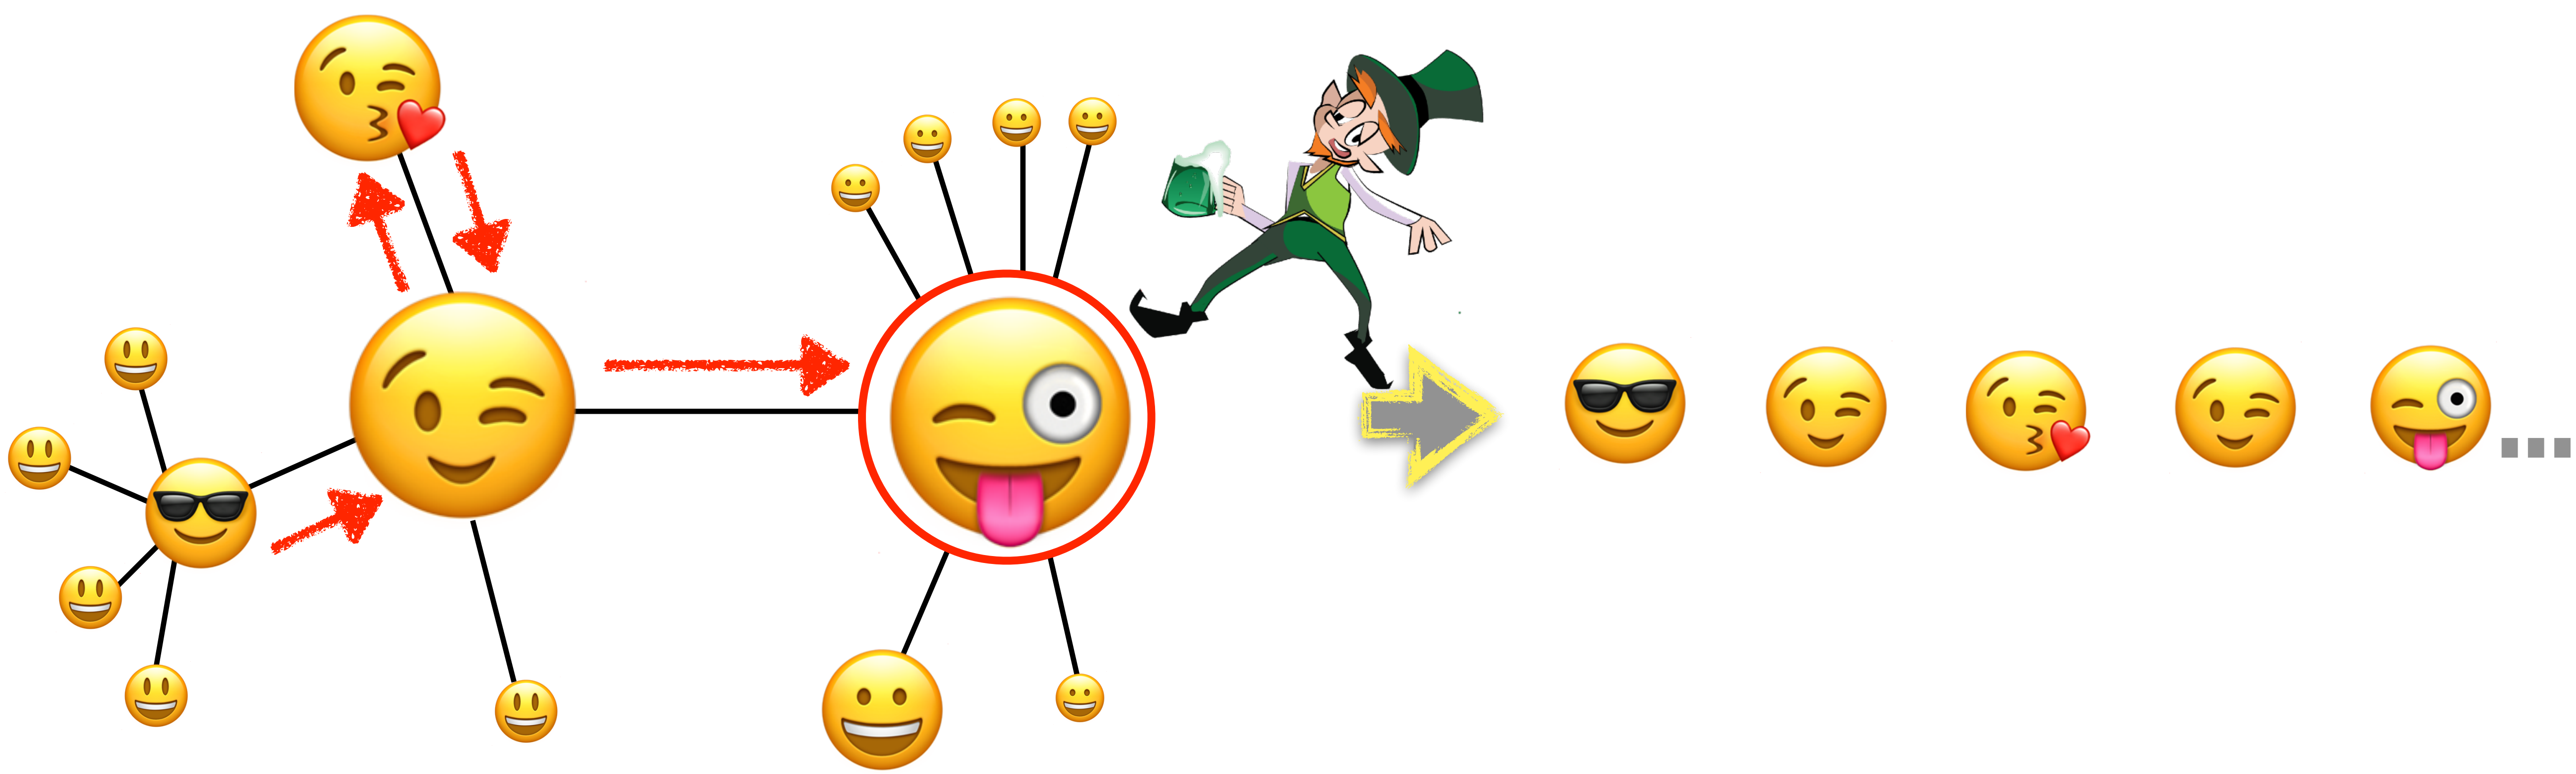

Sadamori Kojaku, Jisung Yoon, Isabel Constantino, and Yong-Yeol Ahn. Residual2Vec: Debiasing graph embedding with random graphs. NeurIPS (2021)

Friendship paradox

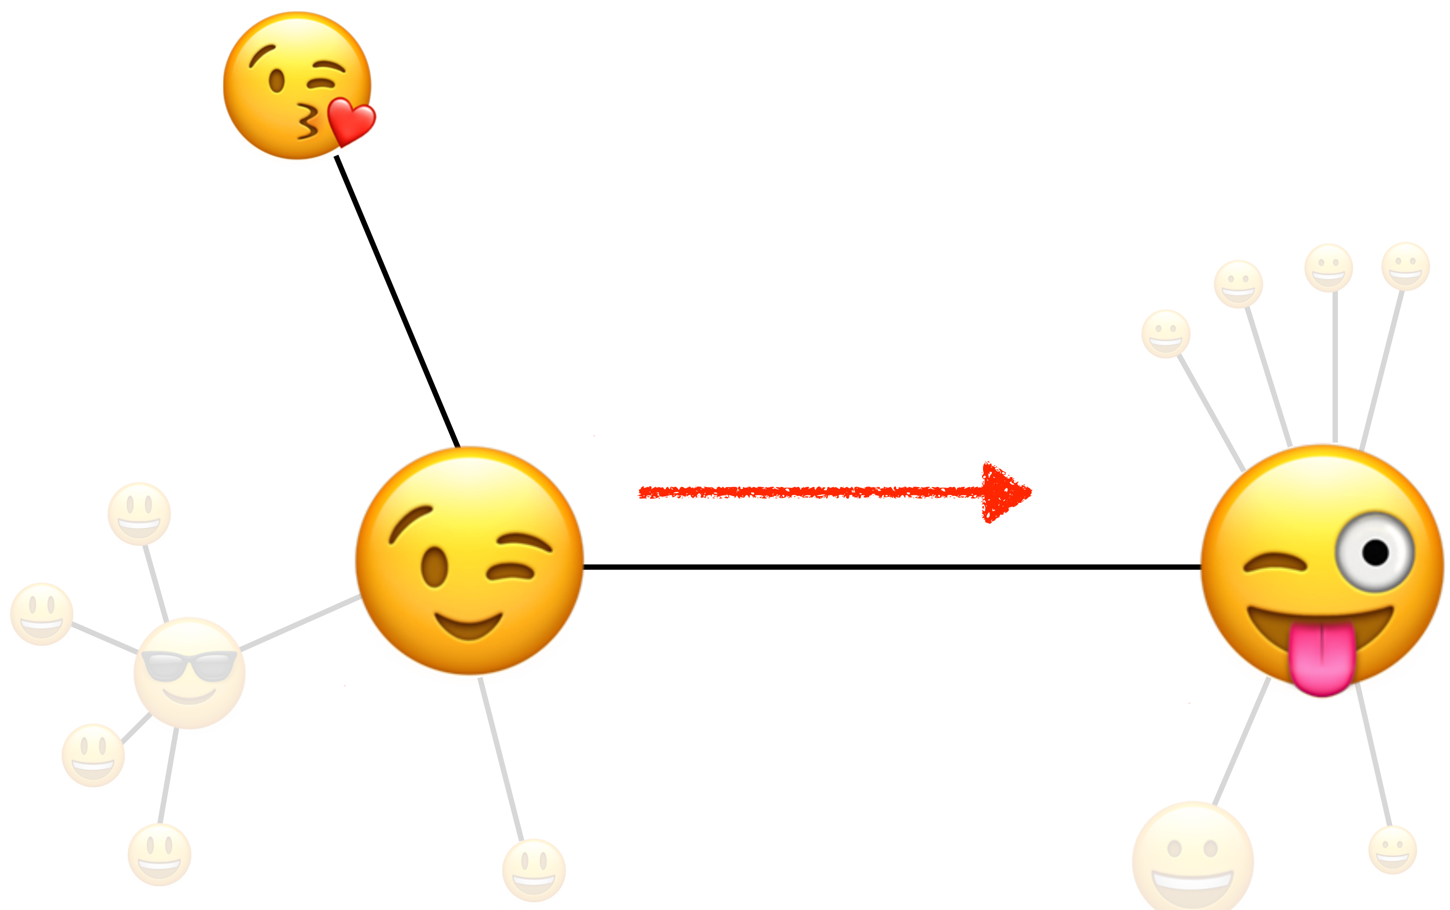

The walker is likely to visit ***a node with many friends.***

Following edges is a ***biased*** sampling that preferentially leads the walker to nodes with many neighbors.

Sadamori Kojaku, Jisung Yoon, Isabel Constantino, and Yong-Yeol Ahn. Residual2Vec: Debiasing graph embedding with random graphs. NeurIPS (2021)

Toy example

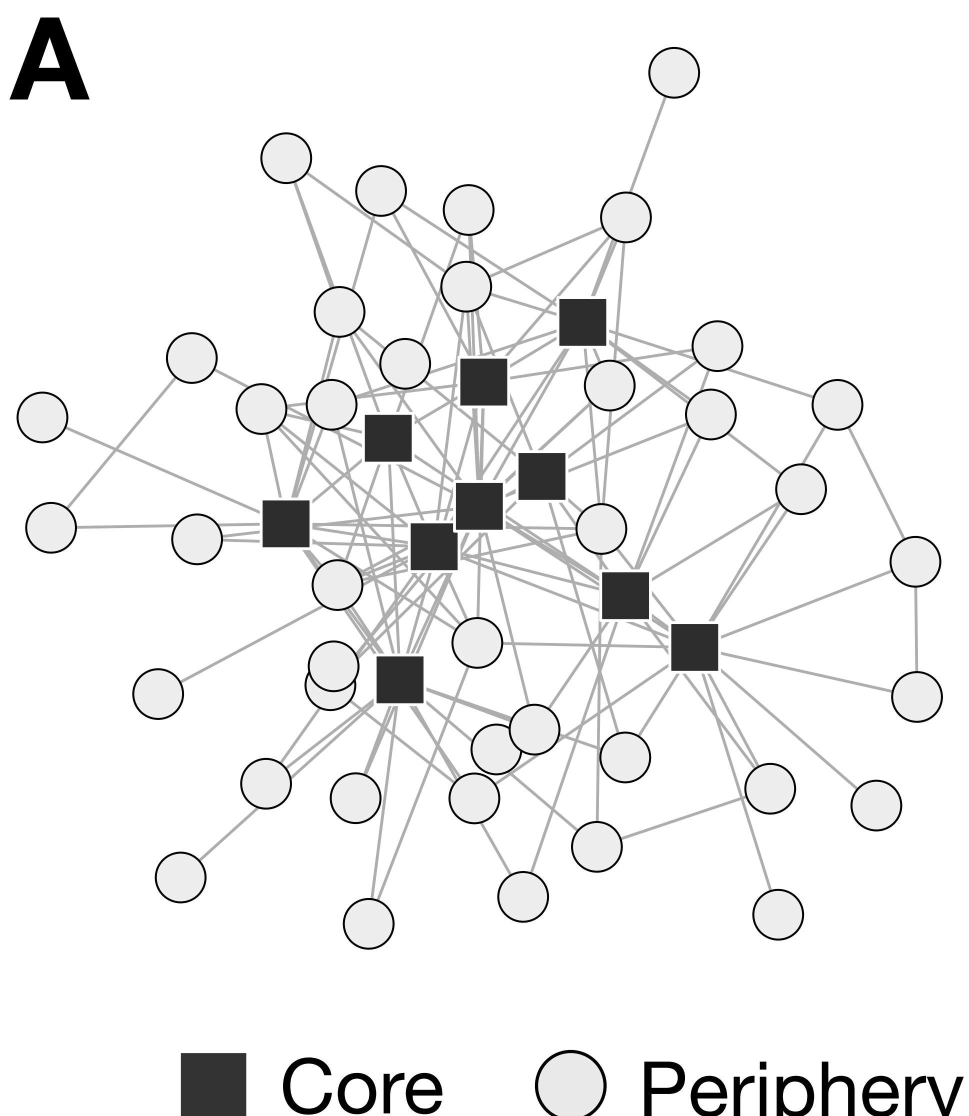

Sadamori Kojaku, Jisung Yoon, Isabel Constantino, and Yong-Yeol Ahn. Residual2Vec: Debiasing graph embedding with random graphs. NeurIPS (2021)

Toy example

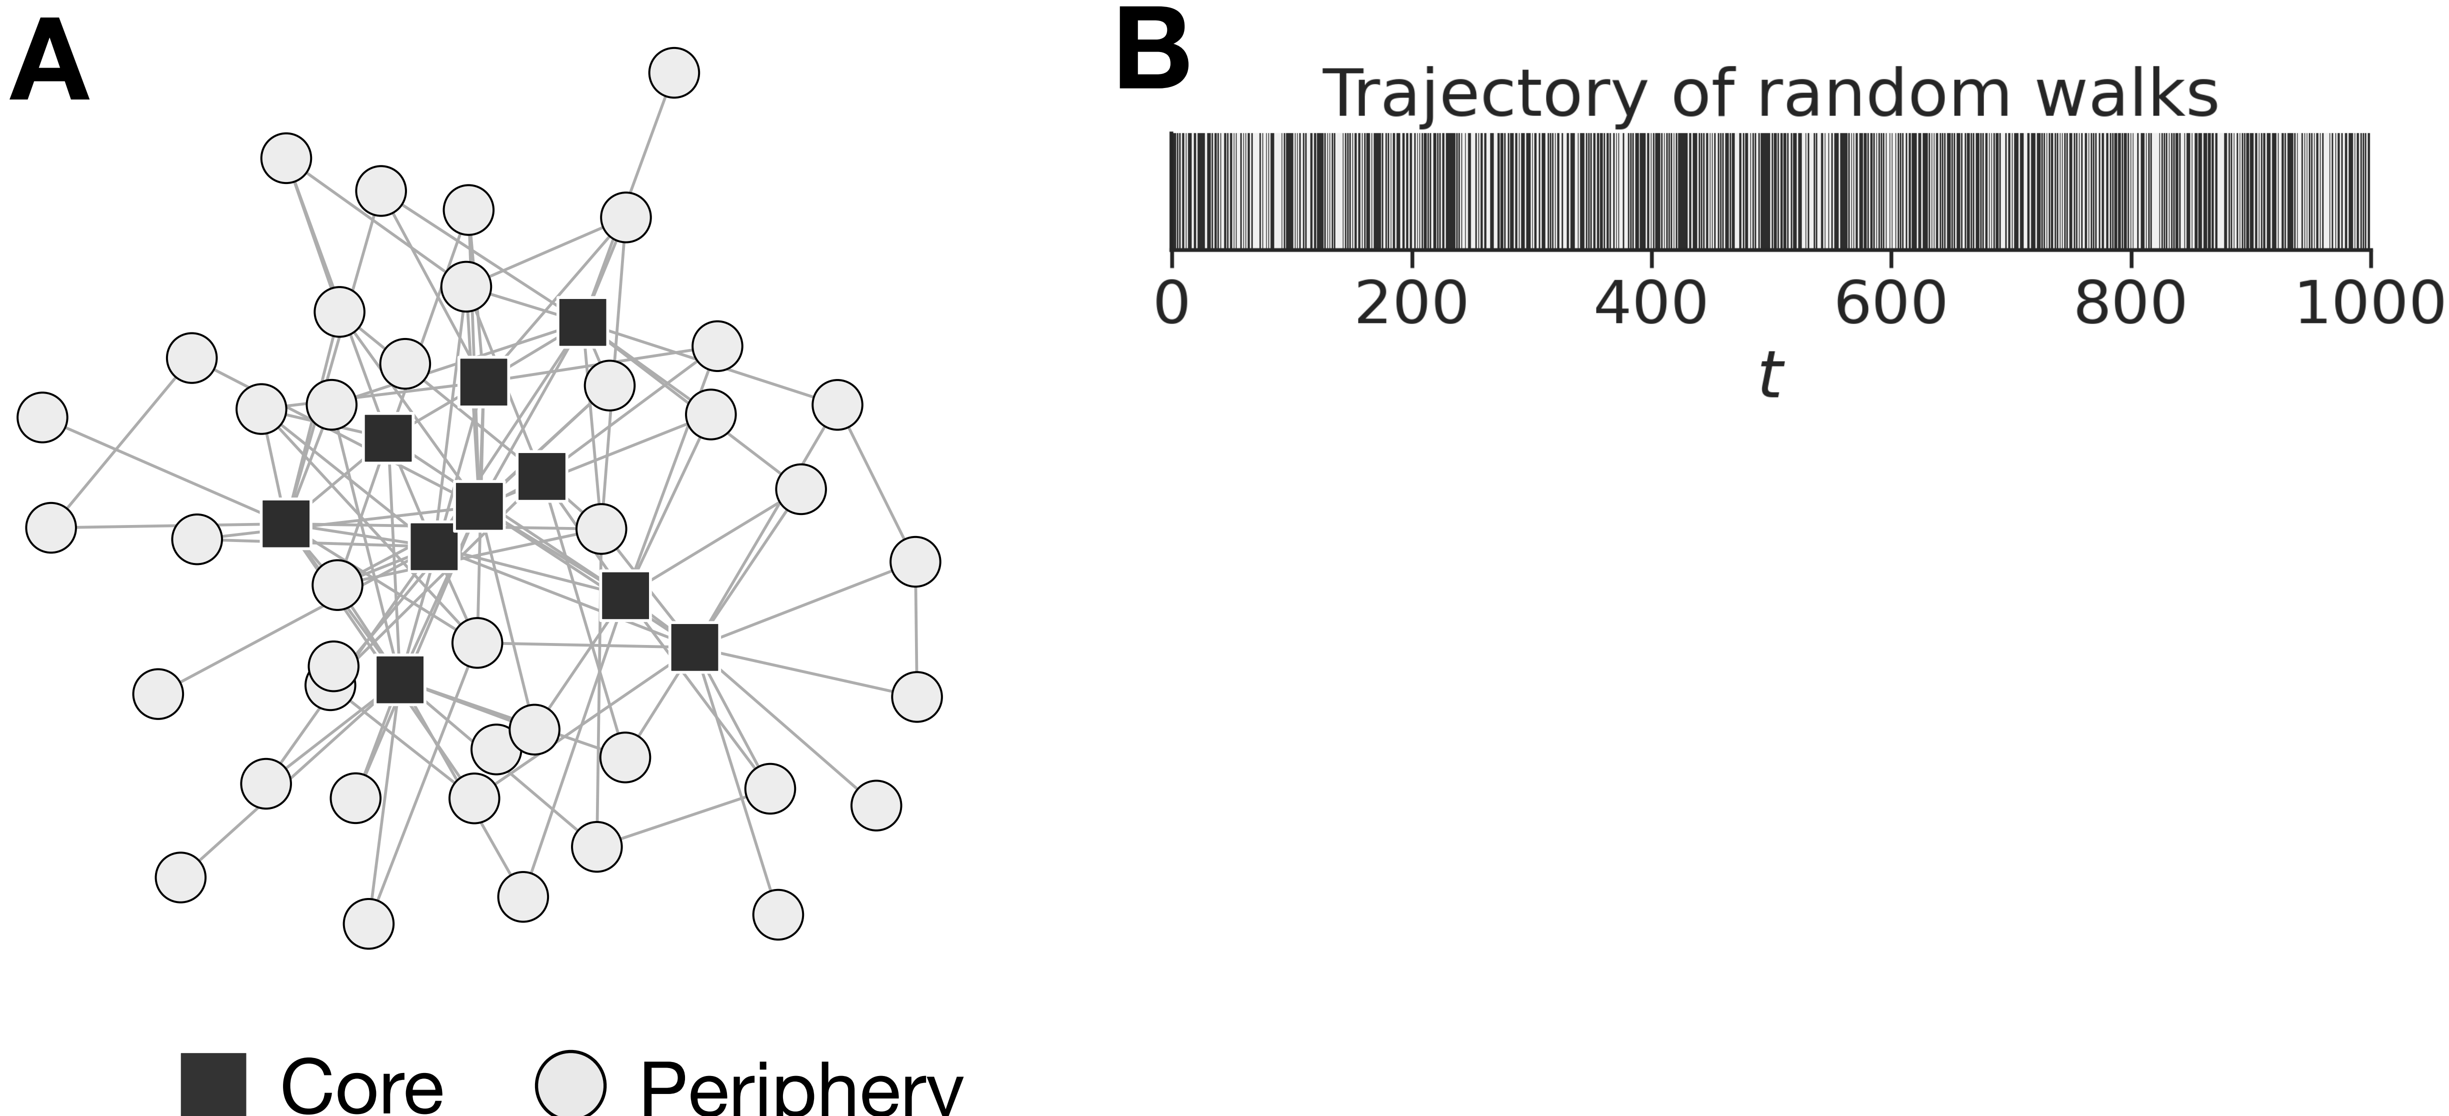

Sadamori Kojaku, Jisung Yoon, Isabel Constantino, and Yong-Yeol Ahn. Residual2Vec: Debiasing graph embedding with random graphs. NeurIPS (2021)

Toy example

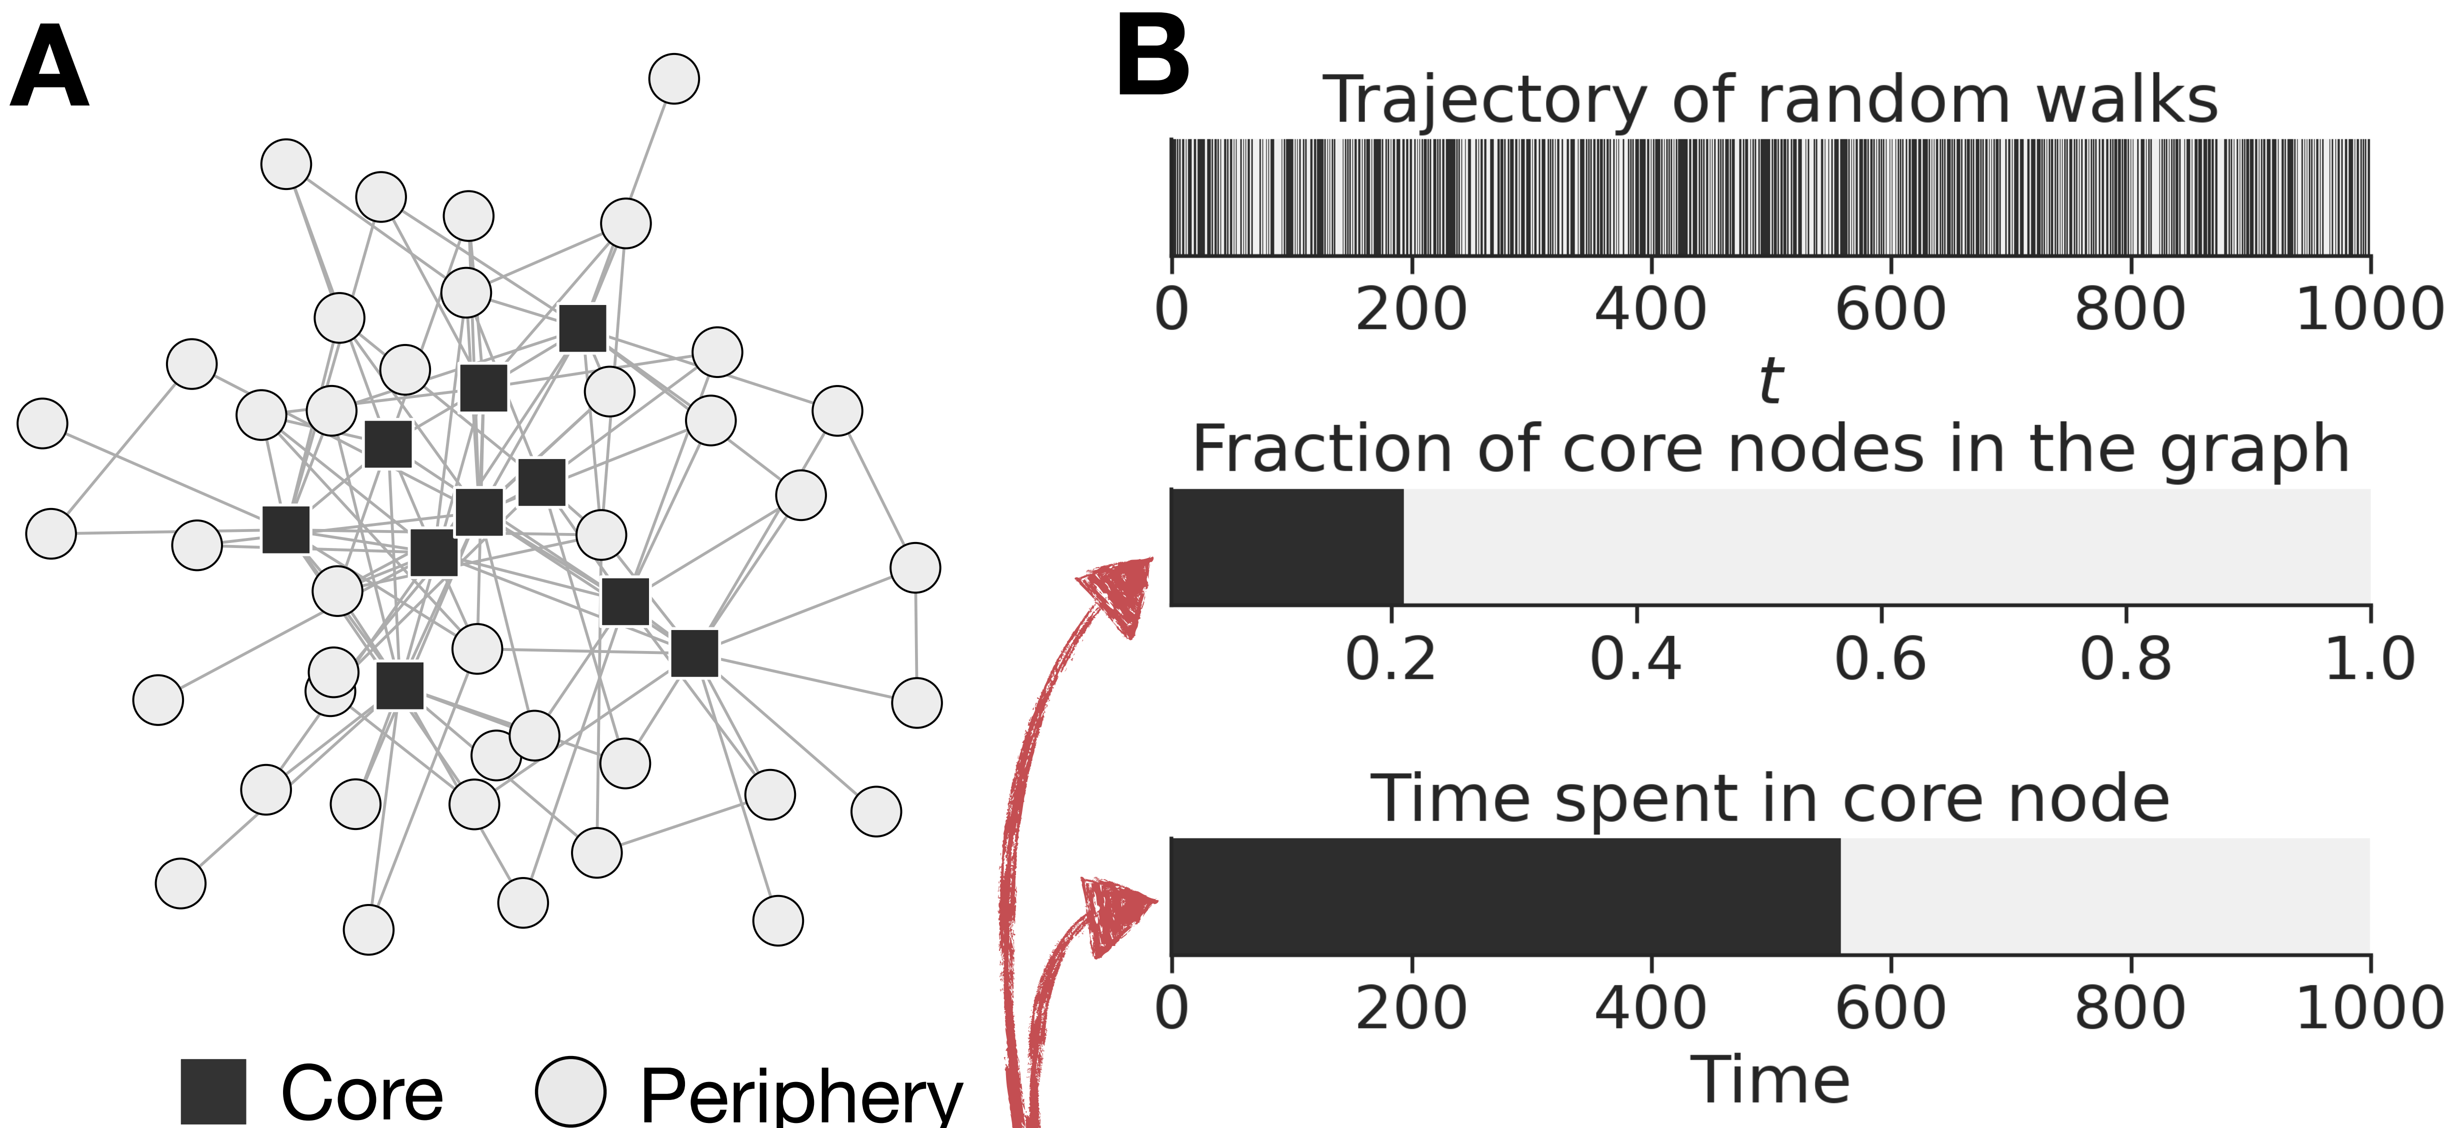

Core nodes are overrepresented due to *the friendship paradox!*

Sadamori Kojaku, Jisung Yoon, Isabel Constantino, and Yong-Yeol Ahn. Residual2Vec: Debiasing graph embedding with random graphs. NeurIPS (2021)

Does the sampling bias have negative impact?

word2vec trained with negative sampling has an overlooked built-in debiasing feature!

We demonstrate how to leverage this feature to debias other biases!

Sadamori Kojaku, Jisung Yoon, Isabel Constantino, and Yong-Yeol Ahn. Residual2Vec: Debiasing graph embedding with random graphs. NeurIPS (2021)

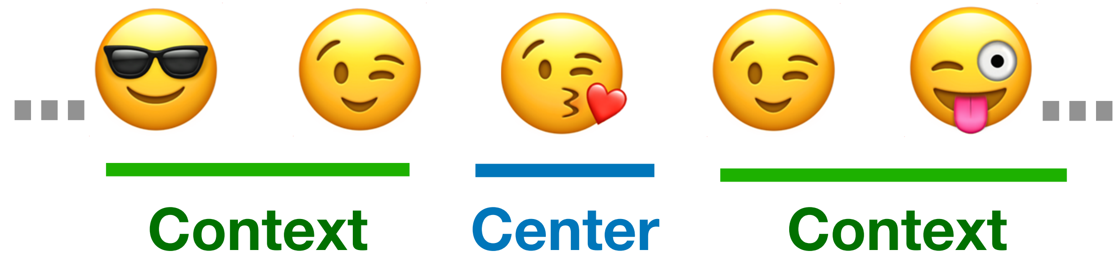

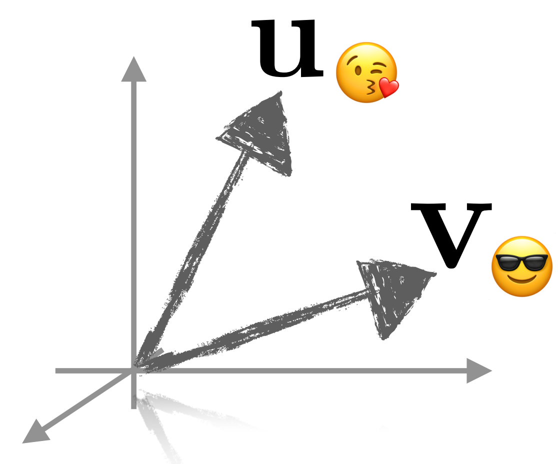

$$P_{\mathrm{w2v}}\left(\underbrace{\text{😎}}_{\text{Context}} \,\middle|\, \underbrace{\text{😘}}_{\text{Center}}\right) = \frac{1}{Z}\exp\left(\mathbf{u}_{\text{😘}}^{\top}\mathbf{v}_{\text{😎}}\right)$$

$$Z = \sum_j \exp\left(\mathbf{u}^{\top}\mathbf{v_j}\right)$$

Negative Sampling is used for training word2vec*

Simplified

💪**Noise Contrastive Estimation (NCE)**💪

*Mikolov, Tomas, Ilya Sutskever, Kai Chen, Greg Corrado, and Jeffrey Dean. 2013. "InDistributed Representations of Words and Phrases and Their Compositionality." In *Advances in Neural Information Processing Systems*, 1389–99.

Sadamori Kojaku, Jisung Yoon, Isabel Constantino, and Yong-Yeol Ahn. Residual2Vec: Debiasing graph embedding with random graphs. NeurIPS (2021)

💪Noise Contrastive Estimation (NCE)💪

🤓 🥸 🧐

k random context

😎 😘

Context Center

Noise distribution

$$\sim P_0(\text{🤓}) \propto \{\text{Frequency of 🤓}\}^{\gamma}$$

*Gutmann, Michael, and Aapo Hyvärinen. 2010. “Noise-Contrastive Estimation: A New Estimation Principle for Unnormalized Statistical Models.” In *Proceedings of the Thirteenth International Conference on Artificial Intelligence and Statistics*, edited by Yee Whye Teh and Mike Titterington, 9:297–304. Proceedings of Machine Learning Research. Chia Laguna Resort, Sardinia, Italy: PMLR.

Sadamori Kojaku, Jisung Yoon, Isabel Constantino, and Yong-Yeol Ahn. Residual2Vec: Debiasing graph embedding with random graphs. NeurIPS (2021)

💪Noise Contrastive Estimation (NCE)💪

🤓 🥸 🧐

k random context

😎 Context

😘 Center

$$Y_{\text{🤓}} = Y_{\text{🥸}} = Y_{\text{🧐}} = 0 \qquad Y_{\text{😎}} = 1$$

$$P(Y_{\text{😎}} = 1 | \text{😎}) \propto \frac{P(\text{😎} | Y_{\text{😎}} = 1)}{P(\text{😎})} P(Y_{\text{😎}} = 1)$$

Posterior

Likelihood

Prior

$$P(\text{😎} | Y = 1) = \begin{cases} P_{\text{w2v}}(\text{😎} | \text{😘}) & (Y = 1) \\ P_0(\text{😎}) & (Y = 0) \end{cases}$$

$$P(Y = 1) = \frac{1}{k+1}$$

*Gutmann, Michael, and Aapo Hyvärinen. 2010. "Noise-Contrastive Estimation: A New Estimation Principle for Unnormalized Statistical Models." In *Proceedings of the Thirteenth International Conference on Artificial Intelligence and Statistics*, edited by Yee Whye Teh and Mike Titterington, 9:297–304. Proceedings of Machine Learning Research. Chia Laguna Resort, Sardinia, Italy: PMLR.

Sadamori Kojaku, Jisung Yoon, Isabel Constantino, and Yong-Yeol Ahn. Residual2Vec: Debiasing graph embedding with random graphs. NeurIPS (2021)

💪 Noise Contrastive Estimation (NCE) 💪

Sigmoid

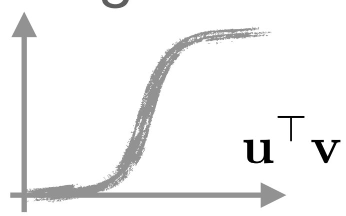

$$P(Y_{\text{😎}} = 1 \,|\, \text{😎}) = \frac{1}{1 + \exp\left(-\mathbf{u}_{\text{😘}}^\top \mathbf{v}_{\text{😎}} + \ln P_0(\text{😎}) + c\right)}$$

NCE is asymptomatically unbiased for

$$P_{\mathrm{w2v}}\left(\text{😎} \,|\, \text{😘}\right) = \frac{1}{Z} \exp\left(\mathbf{u}_{\text{😘}}^\top \mathbf{v}_{\text{😎}}\right)$$

Context Center

*Gutmann, Michael, and Aapo Hyvärinen. 2010. "Noise-Contrastive Estimation: A New Estimation Principle for Unnormalized Statistical Models." In *Proceedings of the Thirteenth International Conference on Artificial Intelligence and Statistics*, edited by Yee Whye Teh and Mike Titterington, 9:297–304. Proceedings of Machine Learning Research. Chia Laguna Resort, Sardinia, Italy: PMLR.

Sadamori Kojaku, Jisung Yoon, Isabel Constantino, and Yong-Yeol Ahn. Residual2Vec: Debiasing graph embedding with random graphs. NeurIPS (2021)

Noise Contrastive Estimation (NCE)

$$P(Y_{\text{😎}} = 1 | \text{😎}) = \frac{1}{1 + \exp\left(-\mathbf{u}_{\text{😘}}^\top \mathbf{v}_{\text{😎}} + \ln P_0(\text{😎}) + c\right)}$$

asymptomatically unbiased for

$$P_{\mathrm{w2v}}\left(\text{😎} | \text{😘}\right) = \frac{1}{Z} \exp\left(\mathbf{u}_{\text{😘}}^\top \mathbf{v}_{\text{😎}}\right)$$

Also unbiased for a general model

$$P(Y = 1|x) = \frac{1}{1 + \exp(-f(x) + \ln P_0(x) + c)}$$

asymptomatically unbiased for

$$P(x) = \frac{1}{Z} \exp(f(x))$$

*Gutmann, Michael, and Aapo Hyvärinen. 2010. "Noise-Contrastive Estimation: A New Estimation Principle for Unnormalized Statistical Models." In *Proceedings of the Thirteenth International Conference on Artificial Intelligence and Statistics*, edited by Yee Whye Teh and Mike Titterington, 9:297–304. Proceedings of Machine Learning Research. Chia Laguna Resort, Sardinia, Italy: PMLR.

Sadamori Kojaku, Jisung Yoon, Isabel Constantino, and Yong-Yeol Ahn. Residual2Vec: Debiasing graph embedding with random graphs. NeurIPS (2021)

Noise Contrastive Estimation (NCE)

$$P(Y=1|x) = \frac{1}{1+\exp(-f(x)+\ln P_0(x)+c)}$$

asymptomatically unbiased for

$$P(x) = \frac{1}{Z}\exp(f(x))$$

Negative sampling*

$$P(Y=1|x) = \frac{1}{1+\exp(-f(x)+\cancel{\ln P_0(x)}+c)}$$

$$= \frac{1}{1+\exp\left[-\underbrace{(f(x)+\ln P_0(x))}_{f'(x)}+\ln P_0(x)+c\right]}$$

asymptomatically unbiased for

$$P(x) = \frac{1}{Z'}\exp(f'(x))$$

$$P(x) = \frac{1}{Z'}P_0(x)\exp(f(x))$$

Sadamori Kojaku, Jisung Yoon, Isabel Constantino, and Yong-Yeol Ahn. Residual2Vec: Debiasing graph embedding with random graphs. NeurIPS (2021)

word2vec trained with Skip-gram **Negative sampling** (SGNS) is asymptomatically unbiased for

$$P_{\mathrm{w2v}}(\text{😎} \mid \text{😘}) = \frac{1}{Z} \underbrace{P_0(\text{😎})}_{\textit{Baseline probability}} \exp(\underbrace{\mathbf{u}_{\text{😎}}^\top \mathbf{v}_{\text{😘}}}_{\textit{Residual from the baseline}})$$

$$P(x) = \frac{1}{Z'} P_0(x) f(x)$$

Baseline probability

$\sim P_0(\text{😎}) \propto \{\text{Frequency of 😎}\}^\gamma$

Residual from the baseline

Information other than frequency

Built-in debiasing feature for frequency bias!

😳

Sadamori Kojaku, Jisung Yoon, Isabel Constantino, and Yong-Yeol Ahn. Residual2Vec: Debiasing graph embedding with random graphs. NeurIPS (2021)

$$P_{\text{w2v}}\left(\text{😎} \,|\, \text{😘}\right) = \frac{1}{Z} \underline{P_0(\text{😎})} \exp\left(\mathbf{u}_{\text{😎}}^\top \mathbf{v}_{\text{😘}}\right)$$

$$\sim P_0(\text{😎}) \propto \{\text{Frequency of 😎}\}^\gamma$$

Sampling bias due to *the friendship paradox***

$$\left\{ \begin{array}{c} \text{Frequency of 😎} \\ \underline{\text{in the sentence}} \end{array} \right\} \propto \left\{ \begin{array}{c} \text{\# neighbors of 😎} \\ \text{in the graph} \end{array} \right\}$$

The friendship paradox has ***no effect*** thanks to *the built-in debiasing feature* of SGNS word2vec if $\gamma = 1$.

🤯

*Masuda, Naoki, Mason A. Porter, and Renaud Lambiotte. "Random walks and diffusion on networks." *Physics reports* 716 (2017): 1-58.

Sadamori Kojaku, Jisung Yoon, Isabel Constantino, and Yong-Yeol Ahn. Residual2Vec: Debiasing graph embedding with random graphs. NeurIPS (2021)

Residual2Vec

$$P_{\mathrm{r2v}}(\text{😎} \mid \text{😘}) = \frac{1}{Z} \underline{P_0(\text{😎} \mid \text{😘})} \exp(\mathbf{u}_{\text{😘}}^\top \mathbf{v}_{\text{😎}})$$

Model the baseline explicitly to control bias

Sadamori Kojaku, Jisung Yoon, Isabel Constantino, and Yong-Yeol Ahn. Residual2Vec: Debiasing graph embedding with random graphs. NeurIPS (2021)

Residual2Vec

How can we model the baseline?

$$P_{\mathrm{r2v}}(\text{😎} \mid \text{😘}) = \frac{1}{Z} \underline{P_0(\text{😎} \mid \text{😘})} \exp(\mathbf{u}_{\text{😘}}^\top \mathbf{v}_{\text{😎}})$$

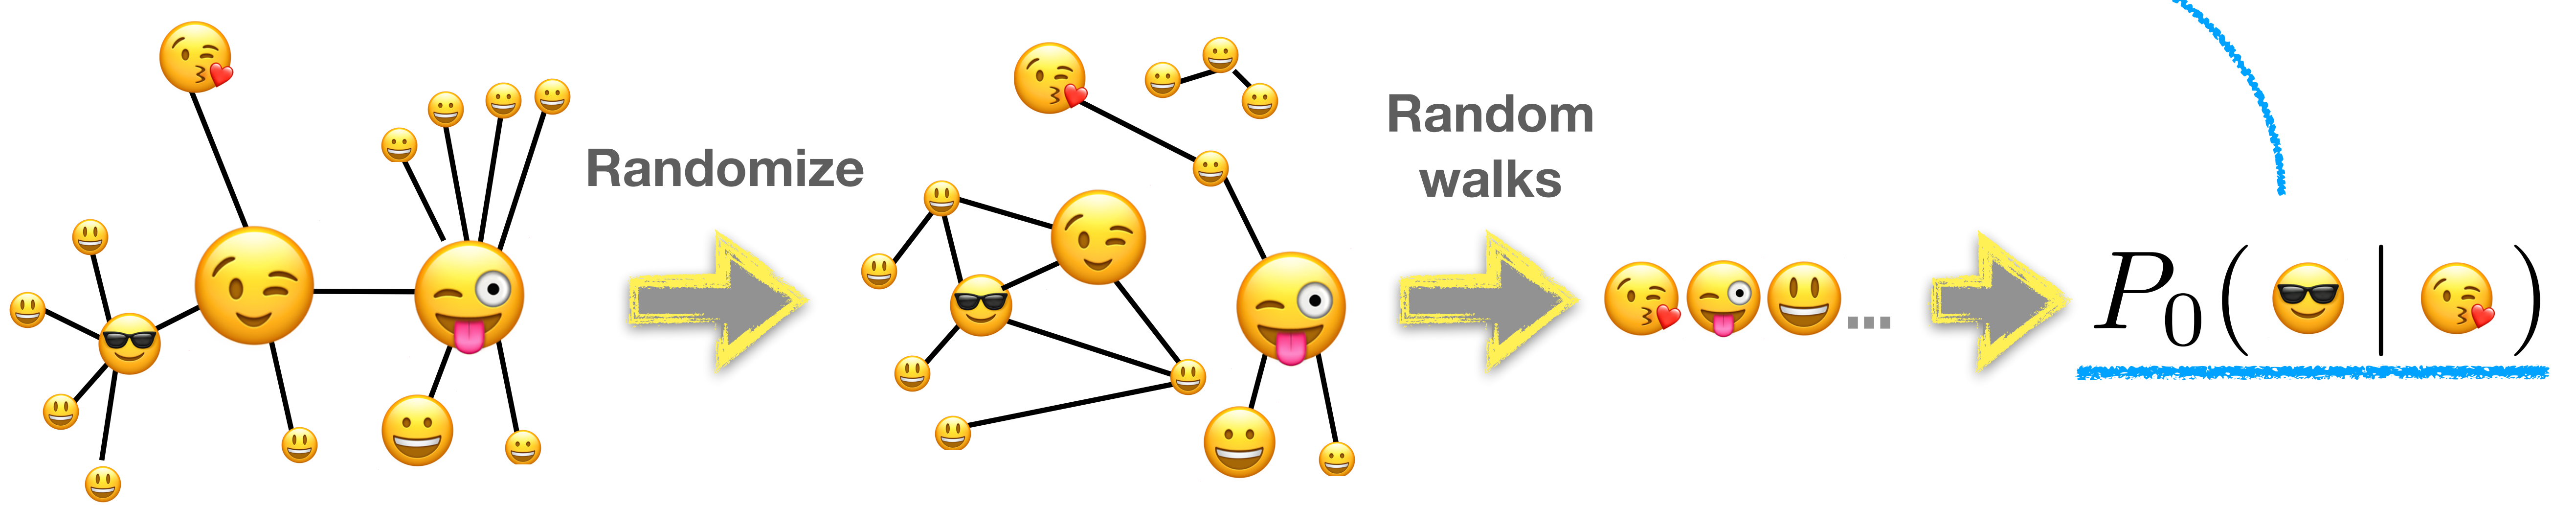

Can be done analytically without simulations!

Sadamori Kojaku, Jisung Yoon, Isabel Constantino, and Yong-Yeol Ahn. Residual2Vec: Debiasing graph embedding with random graphs. NeurIPS (2021)

Performed the best or nearly the best for all the six graphs of different domains

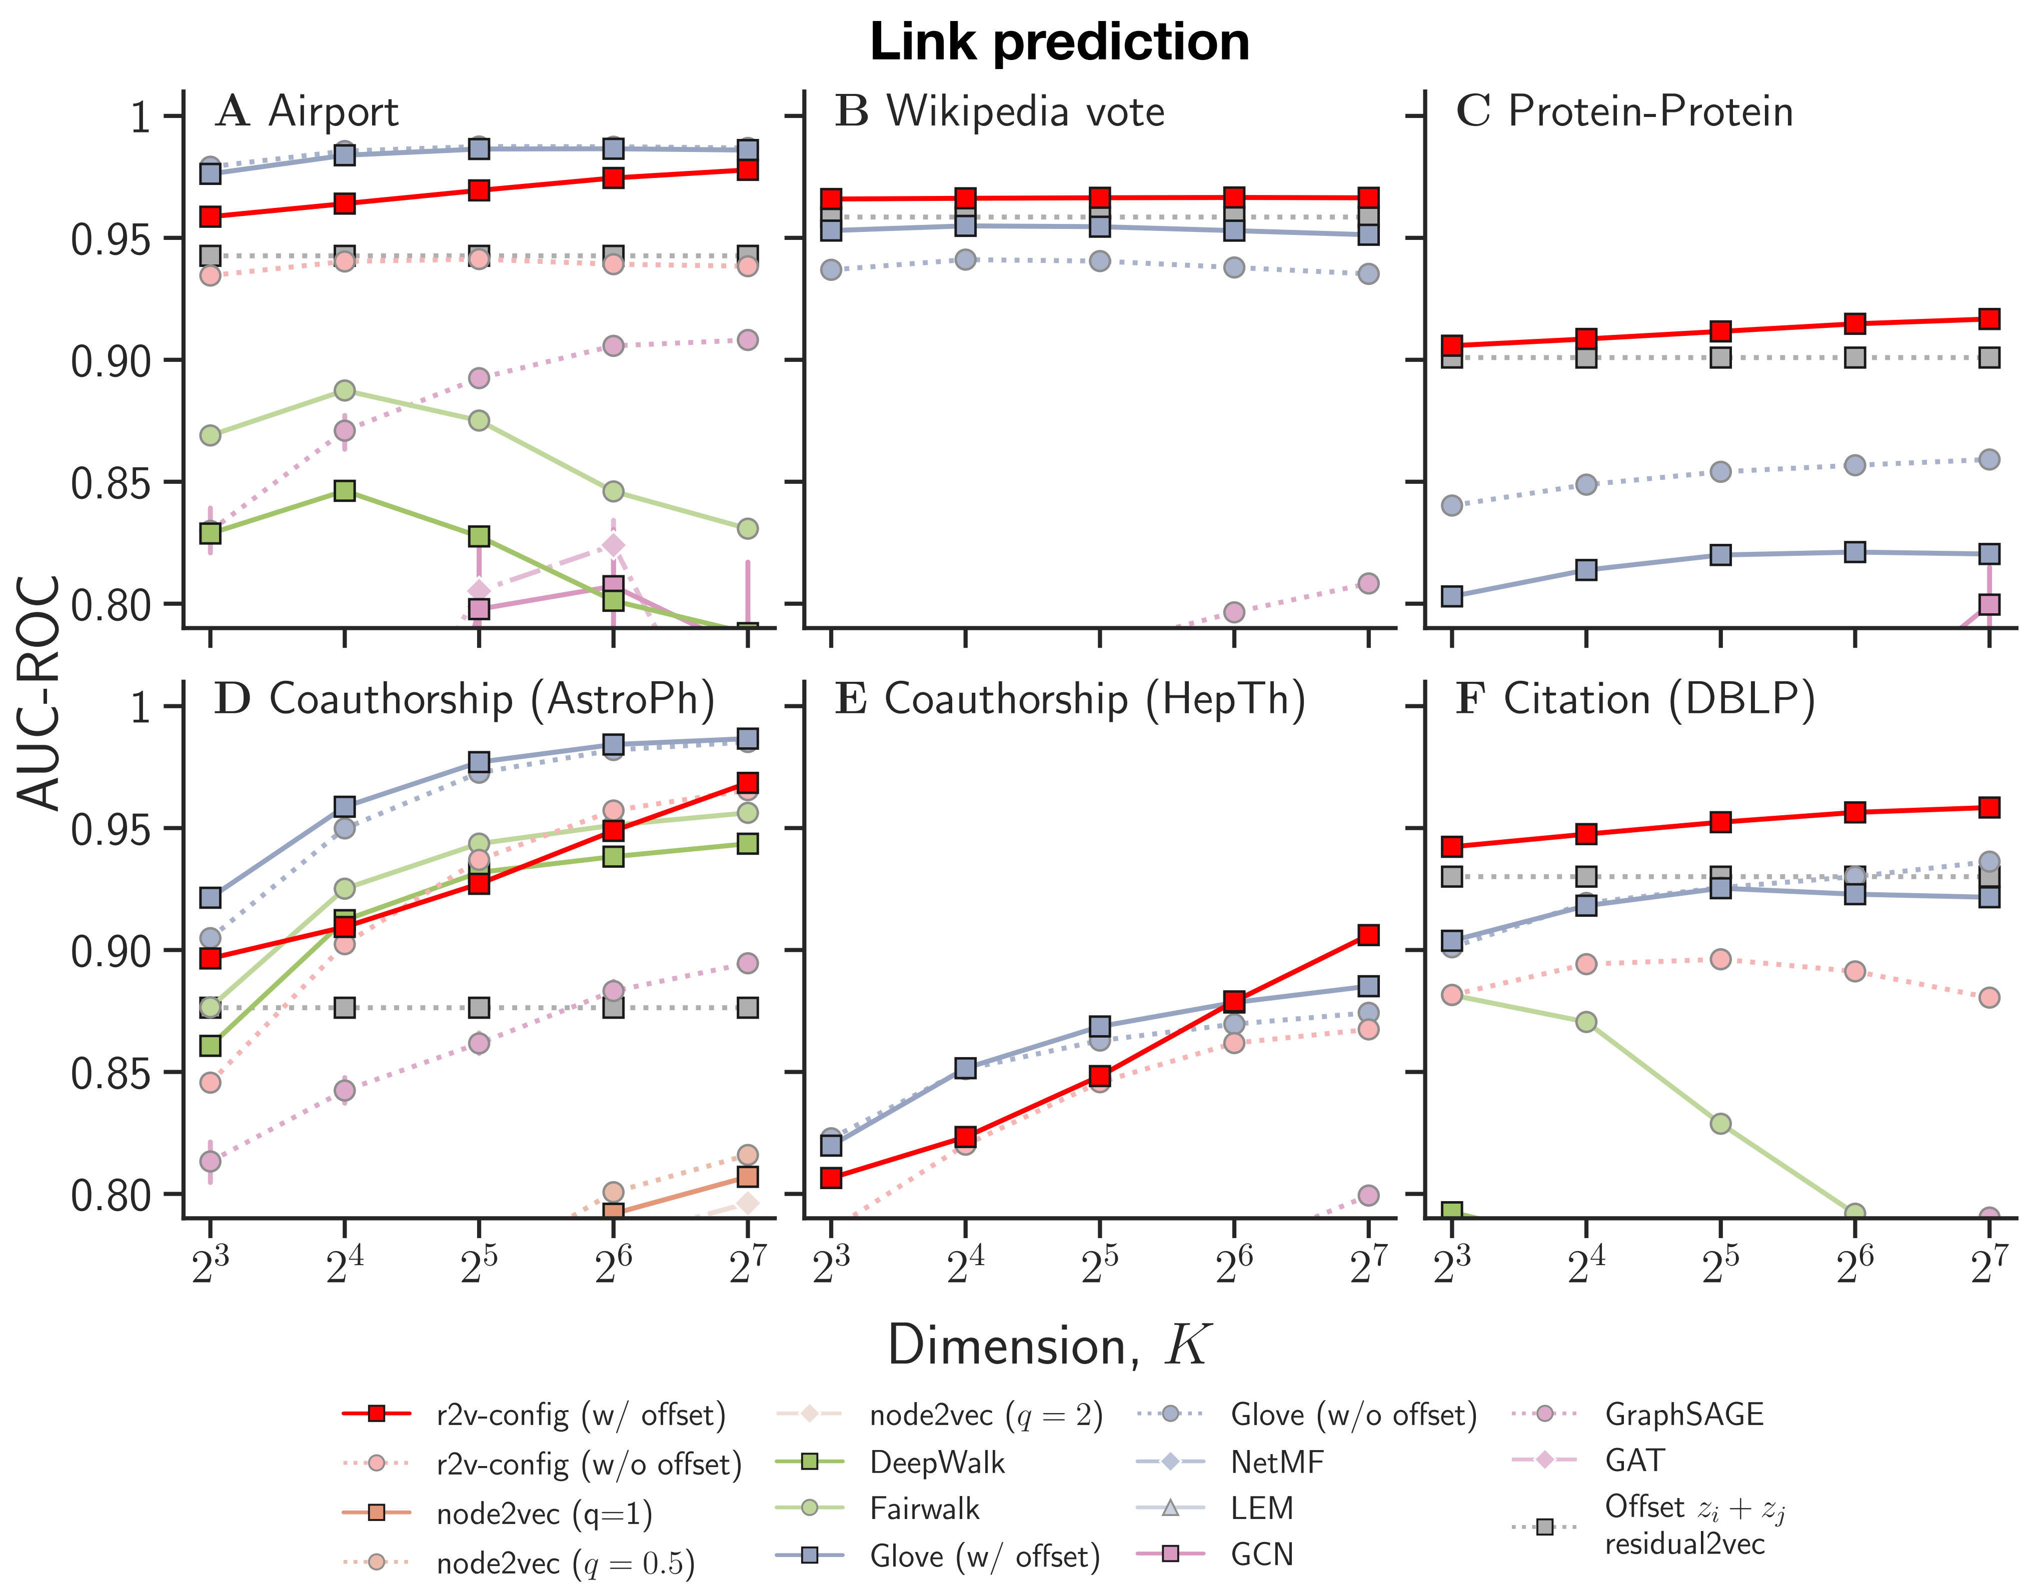

Sadamori Kojaku, Jisung Yoon, Isabel Constantino, and Yong-Yeol Ahn. Residual2Vec: Debiasing graph embedding with random graphs. NeurIPS (2021)

The best or the second best performer for a community detection benchmark

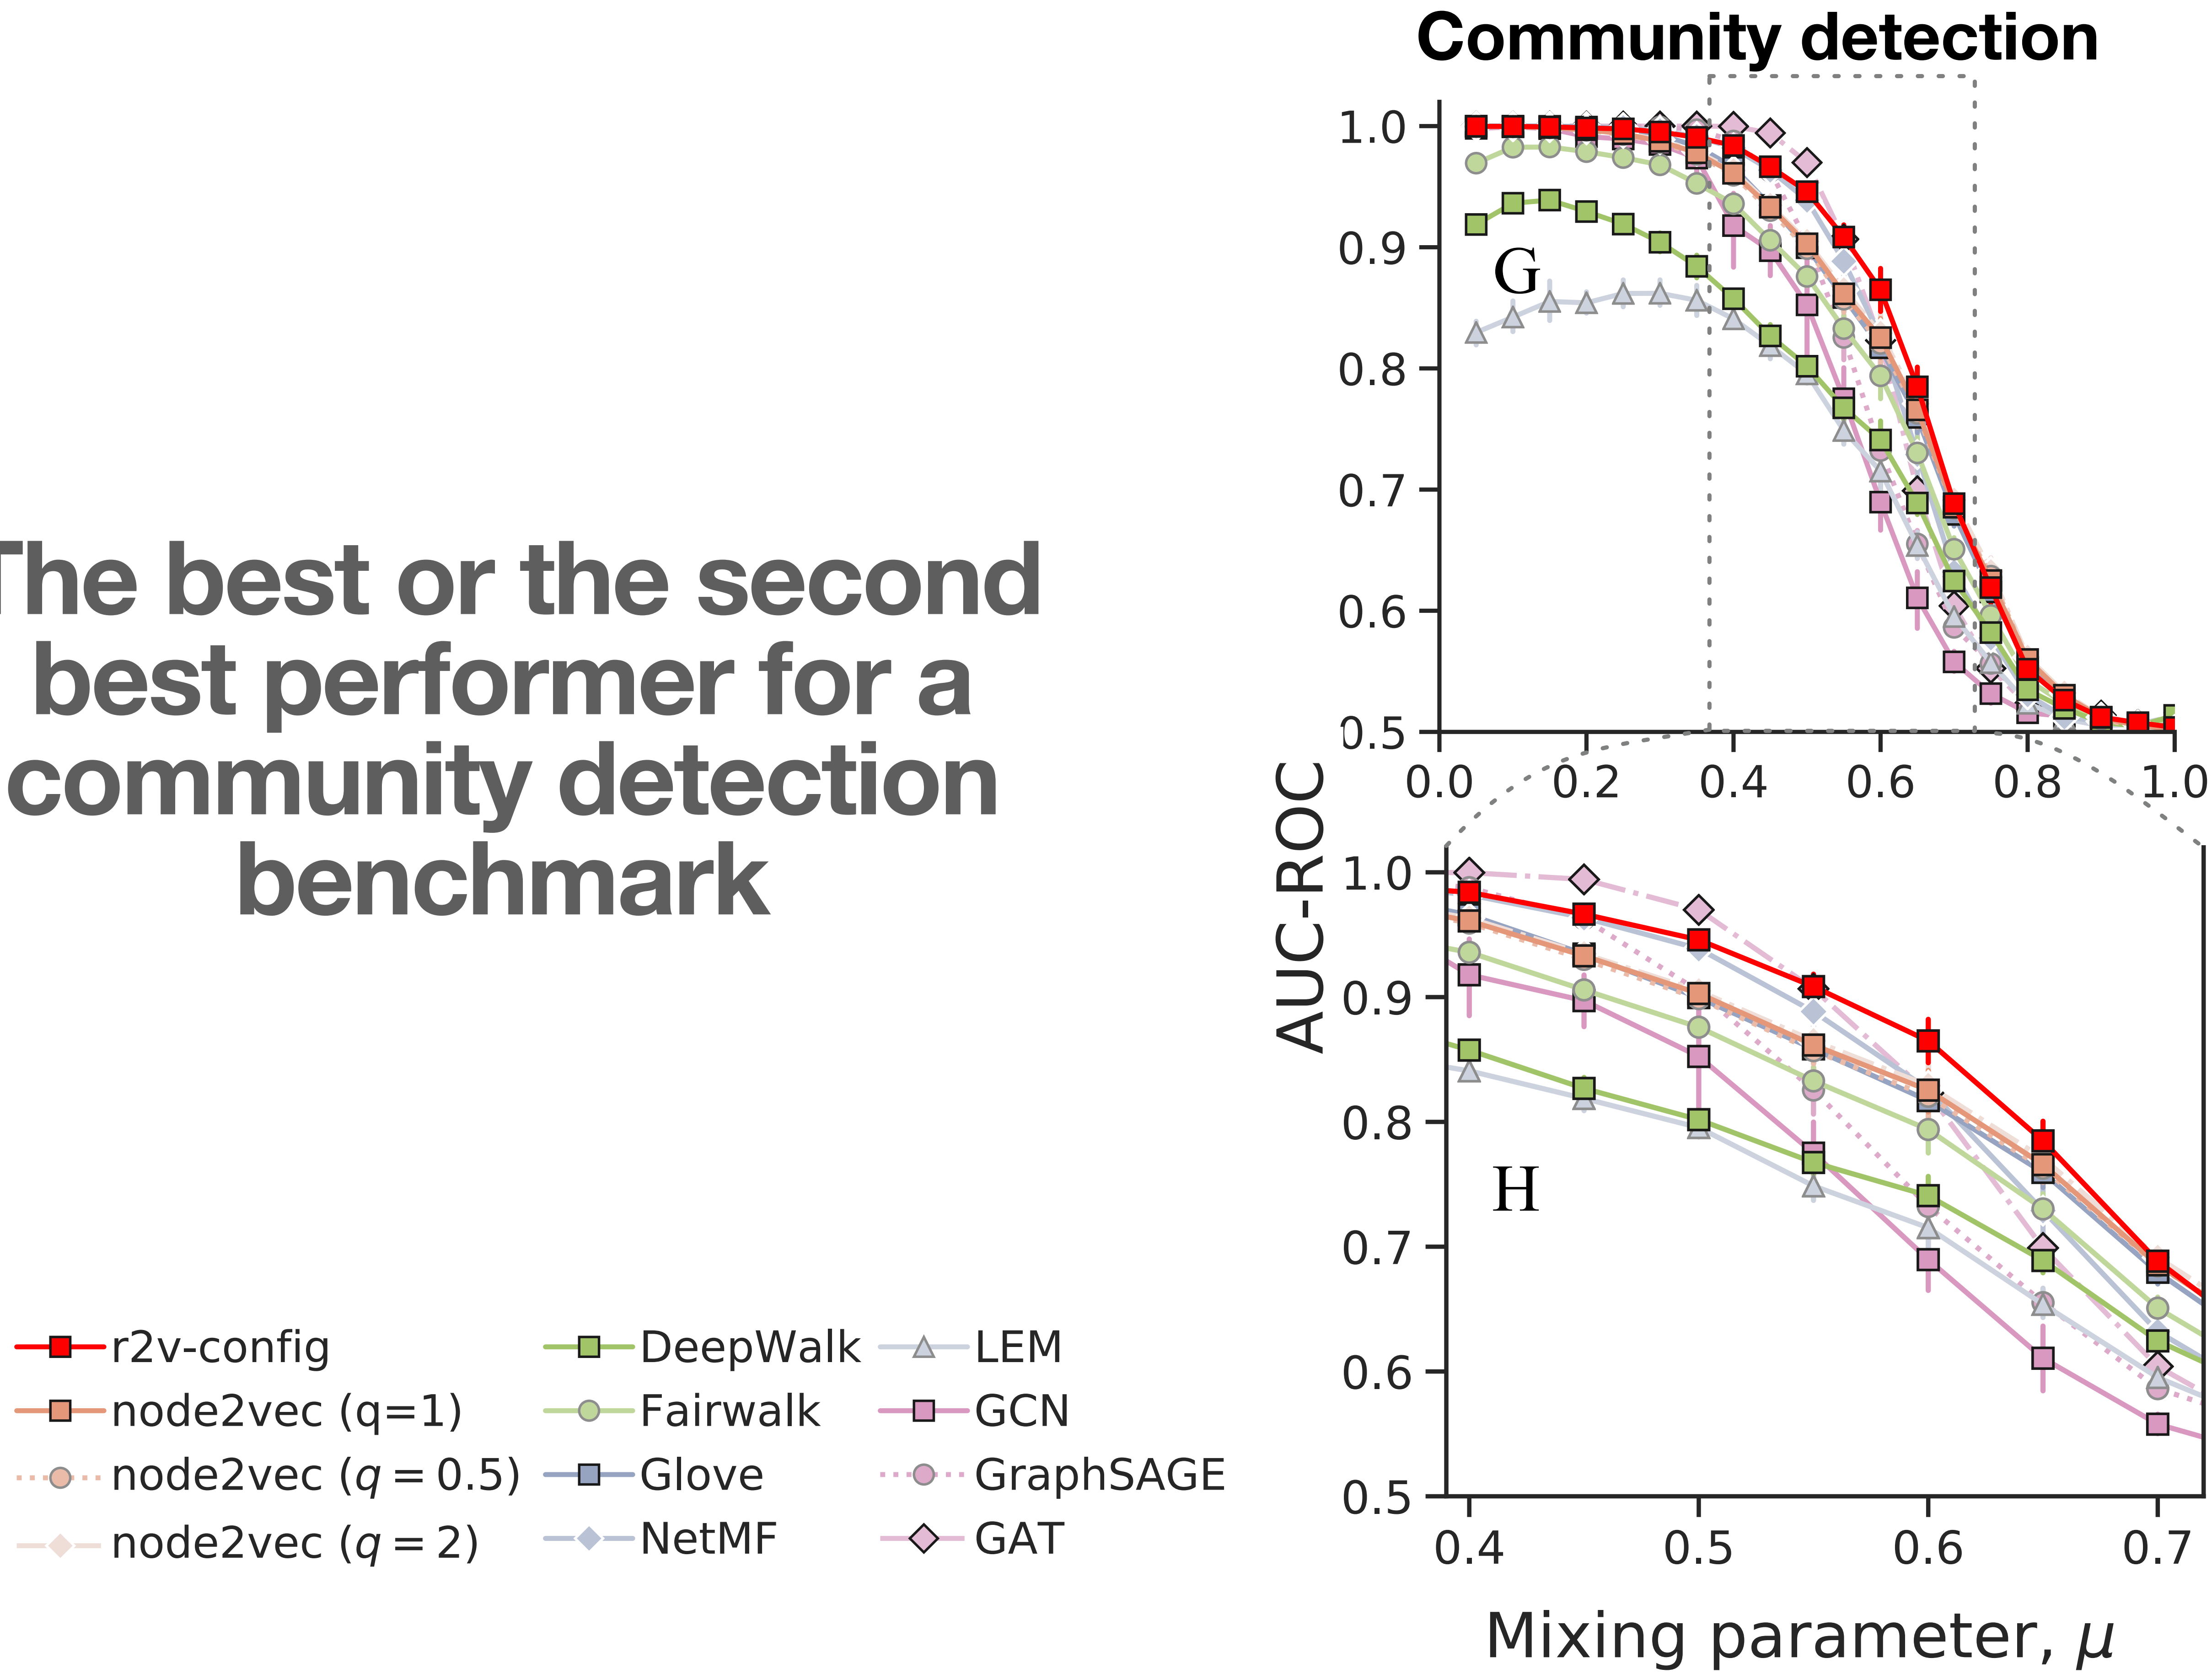

Sadamori Kojaku, Jisung Yoon, Isabel Constantino, and Yong-Yeol Ahn. Residual2Vec: Debiasing graph embedding with random graphs. NeurIPS (2021)

Case studies

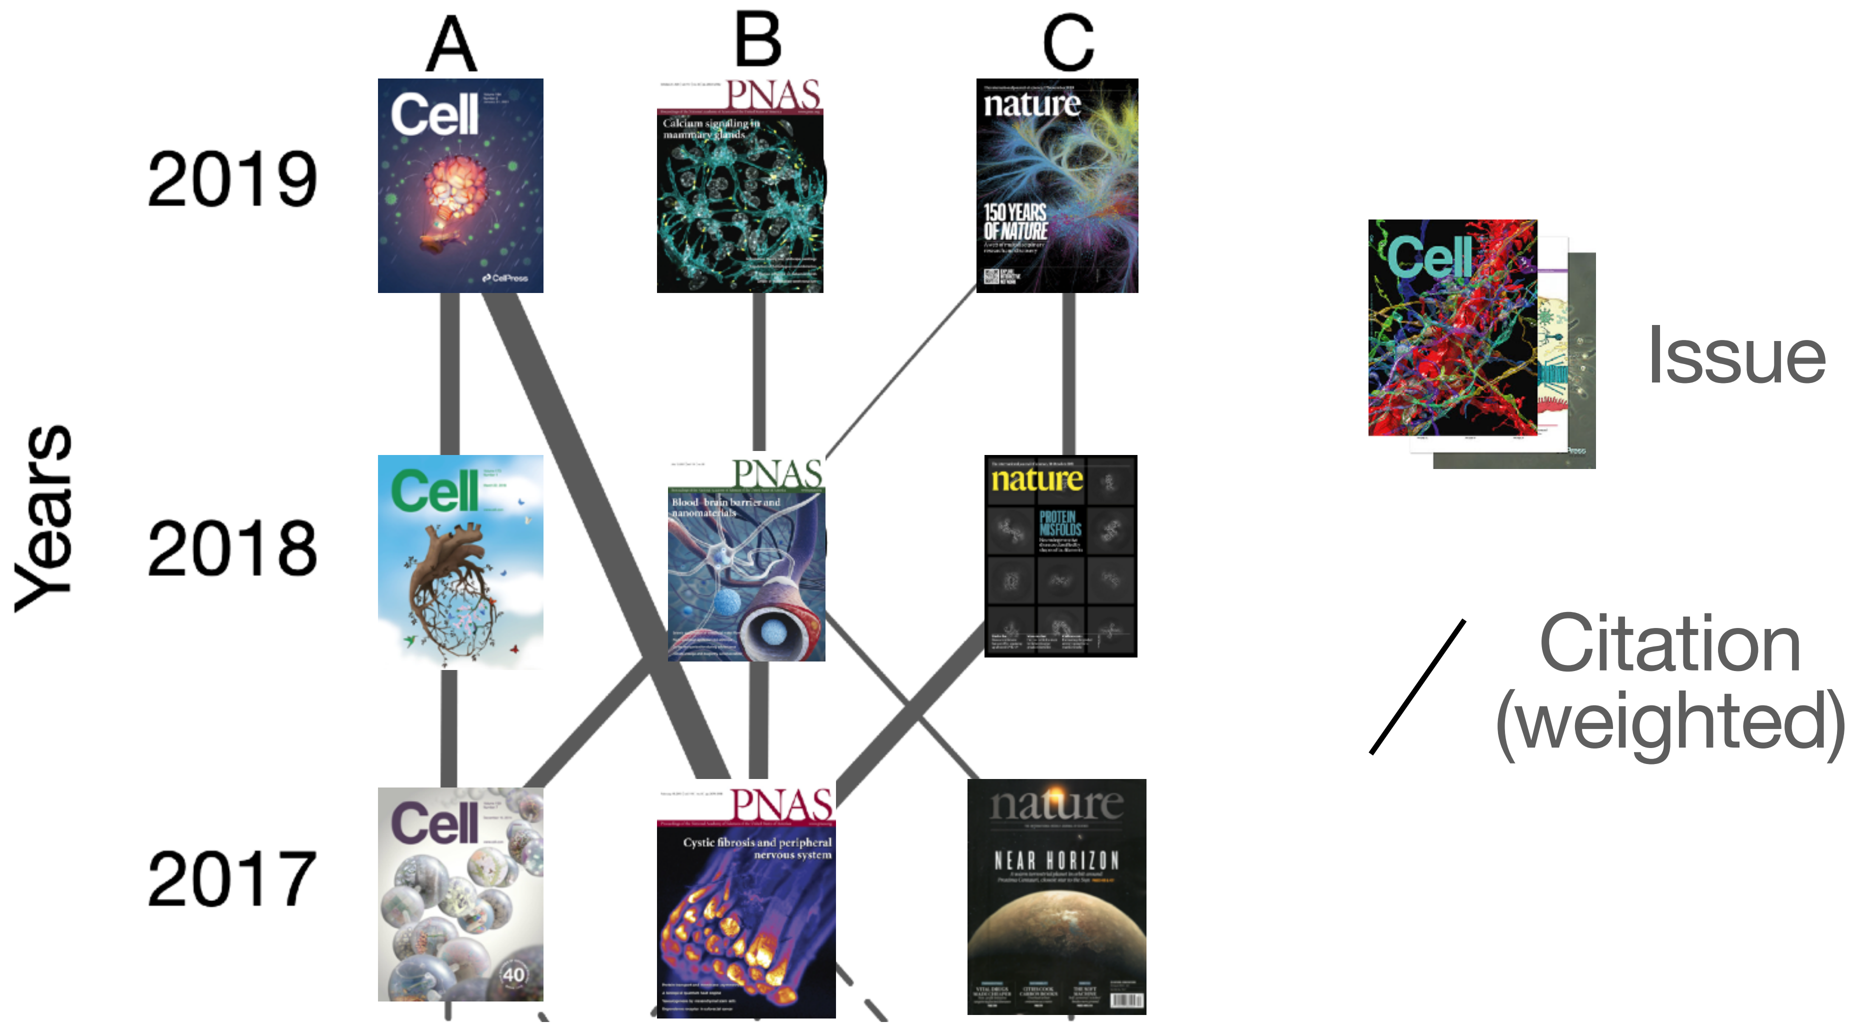

Citation graph of academic issues

Constructed using the Web of Science

240K issues and 250M citations

Sadamori Kojaku, Jisung Yoon, Isabel Constantino, and Yong-Yeol Ahn. Residual2Vec: Debiasing graph embedding with random graphs. NeurIPS (2021)

Glove*

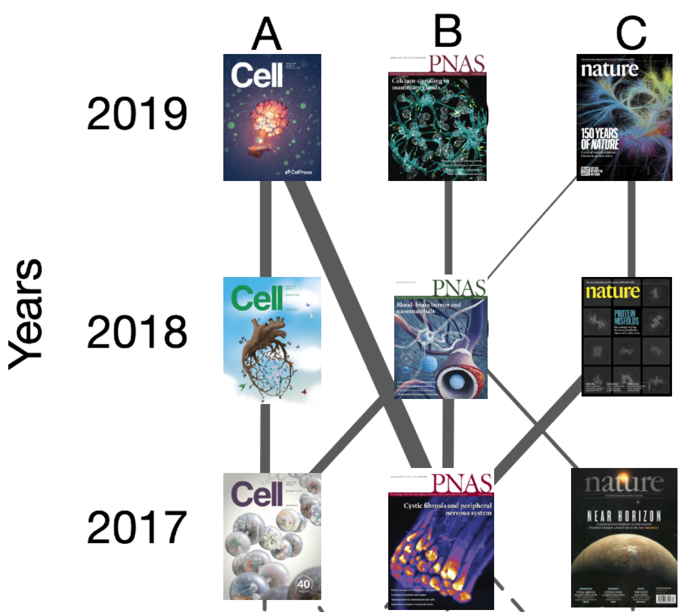

Issue

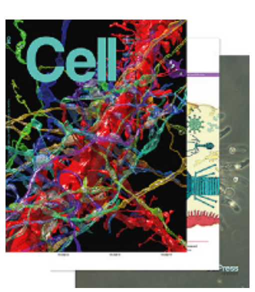

Citation (weighted)

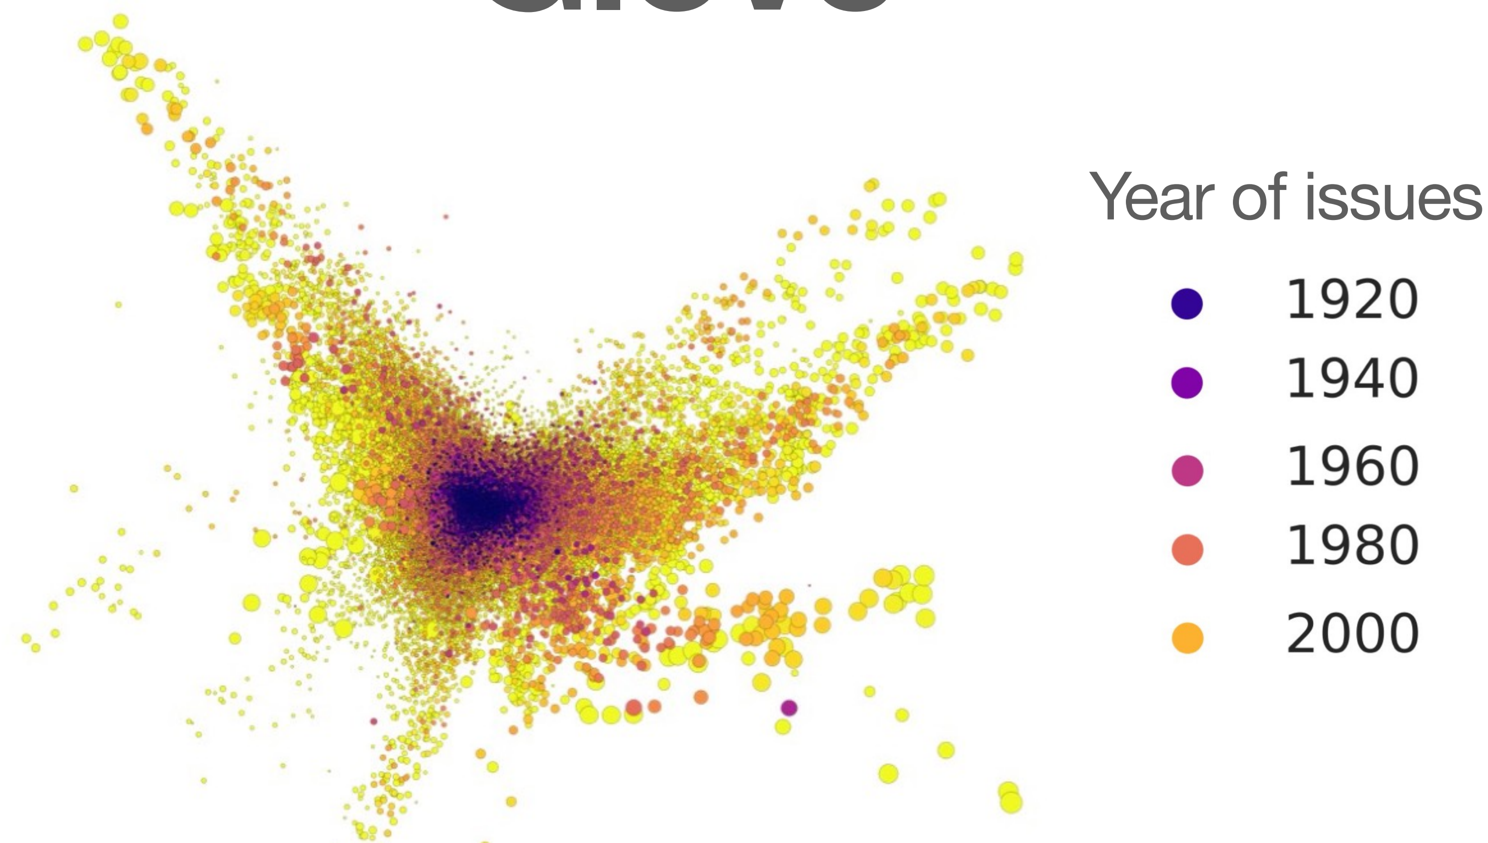

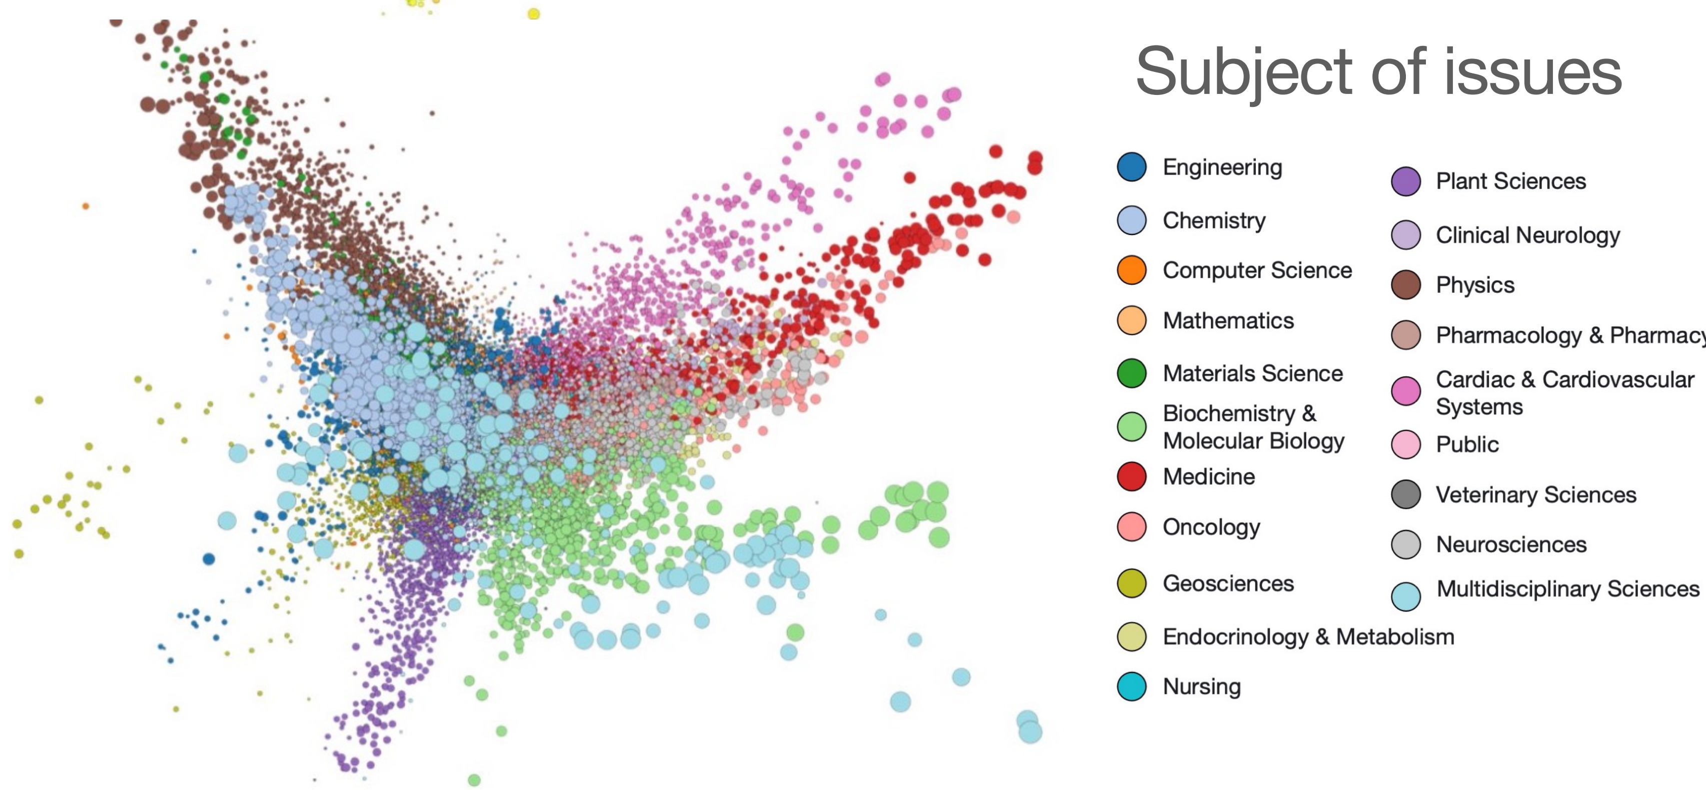

*Jeffrey Pennington, Richard Socher, and Christopher D. Manning. "Glove: Global vectors for word representation." *Proceedings of the 2014 conference on empirical methods in natural language processing (EMNLP)*. 2014.

Sadamori Kojaku, Jisung Yoon, Isabel Constantino, and Yong-Yeol Ahn. Residual2Vec: Debiasing graph embedding with random graphs. NeurIPS (2021)

Residual2Vec
(temporal biases are debiased)

Glove

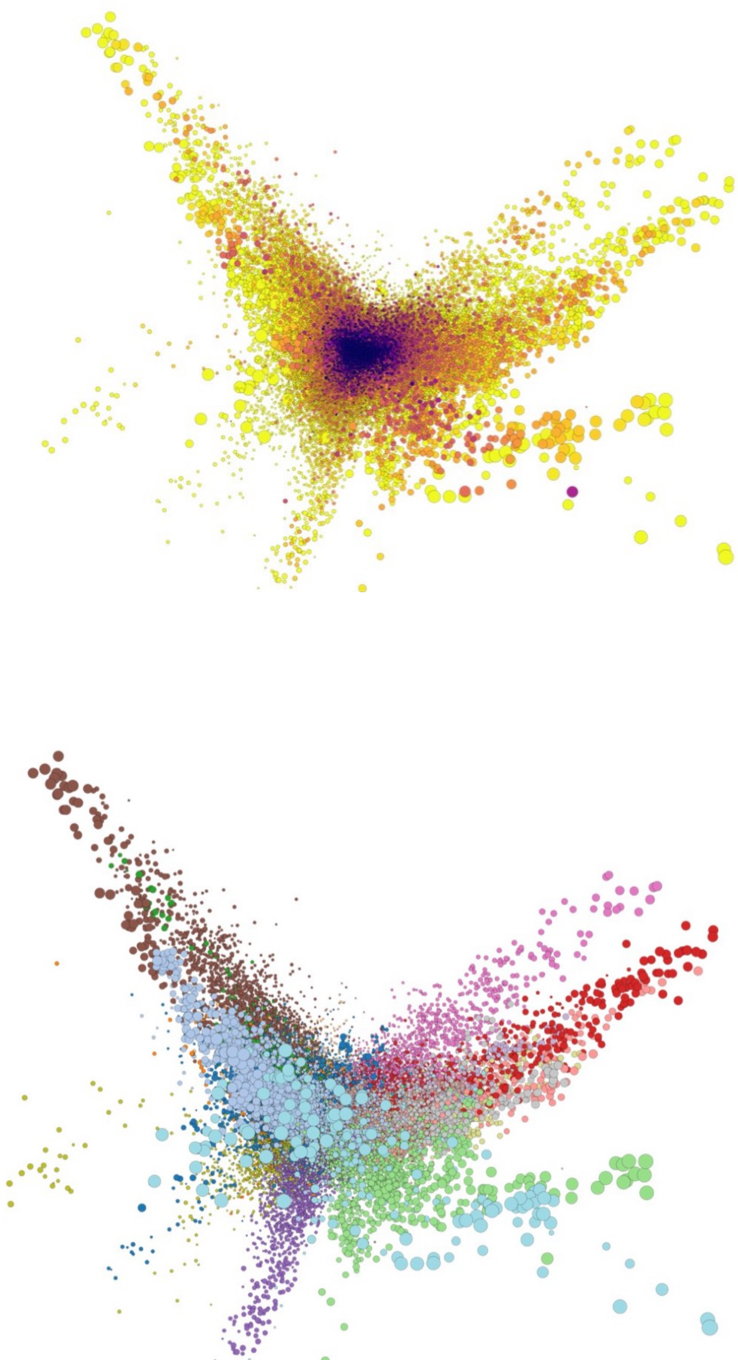

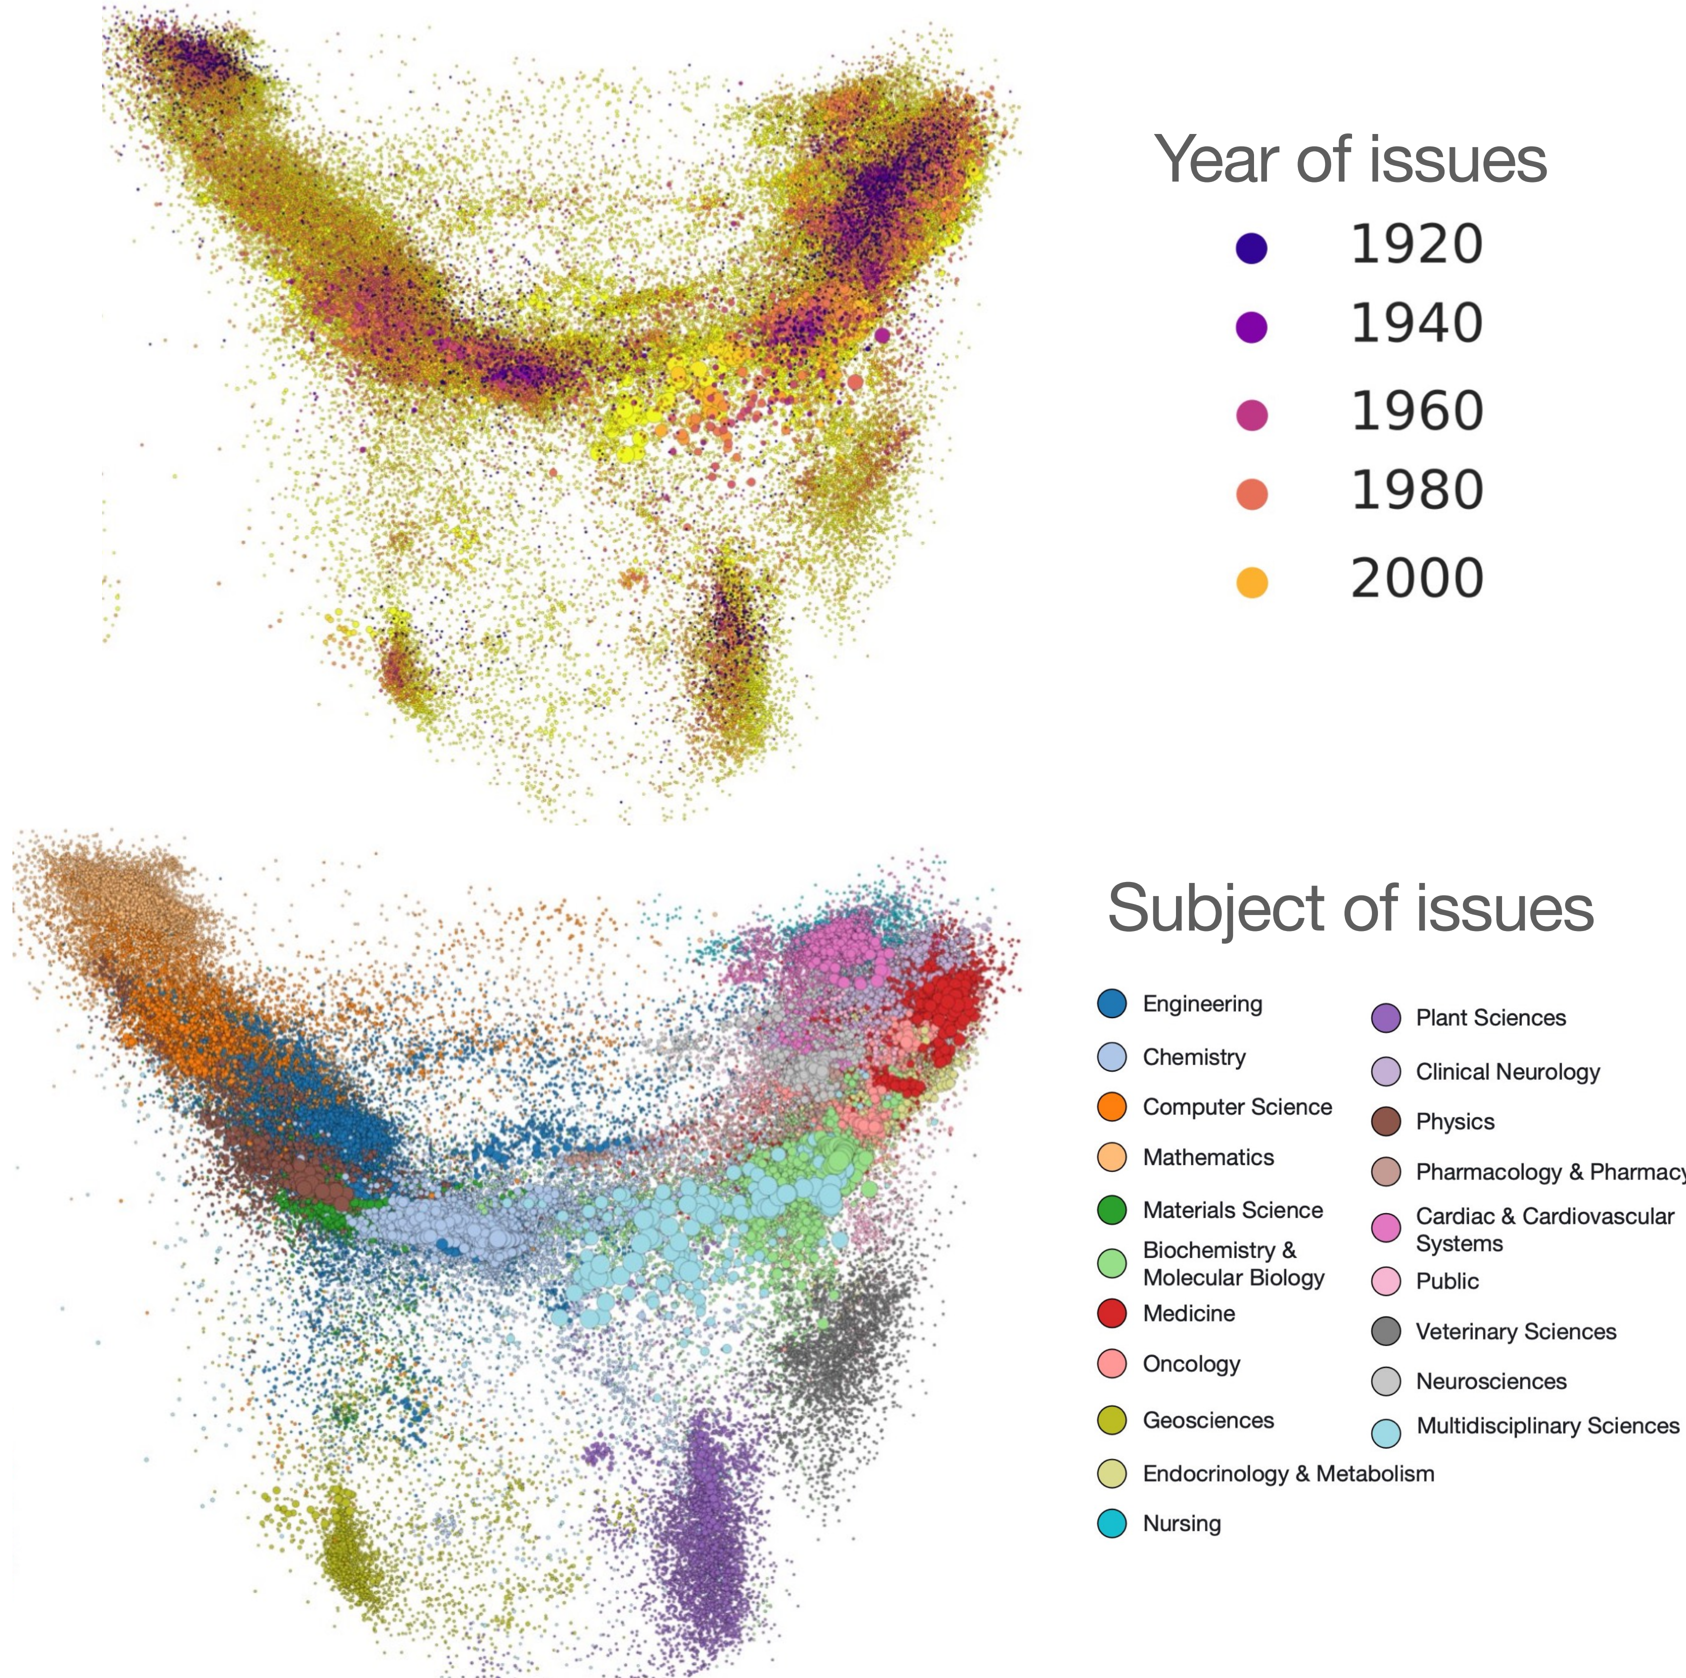

Predicted the impact factor and the subject of journals well

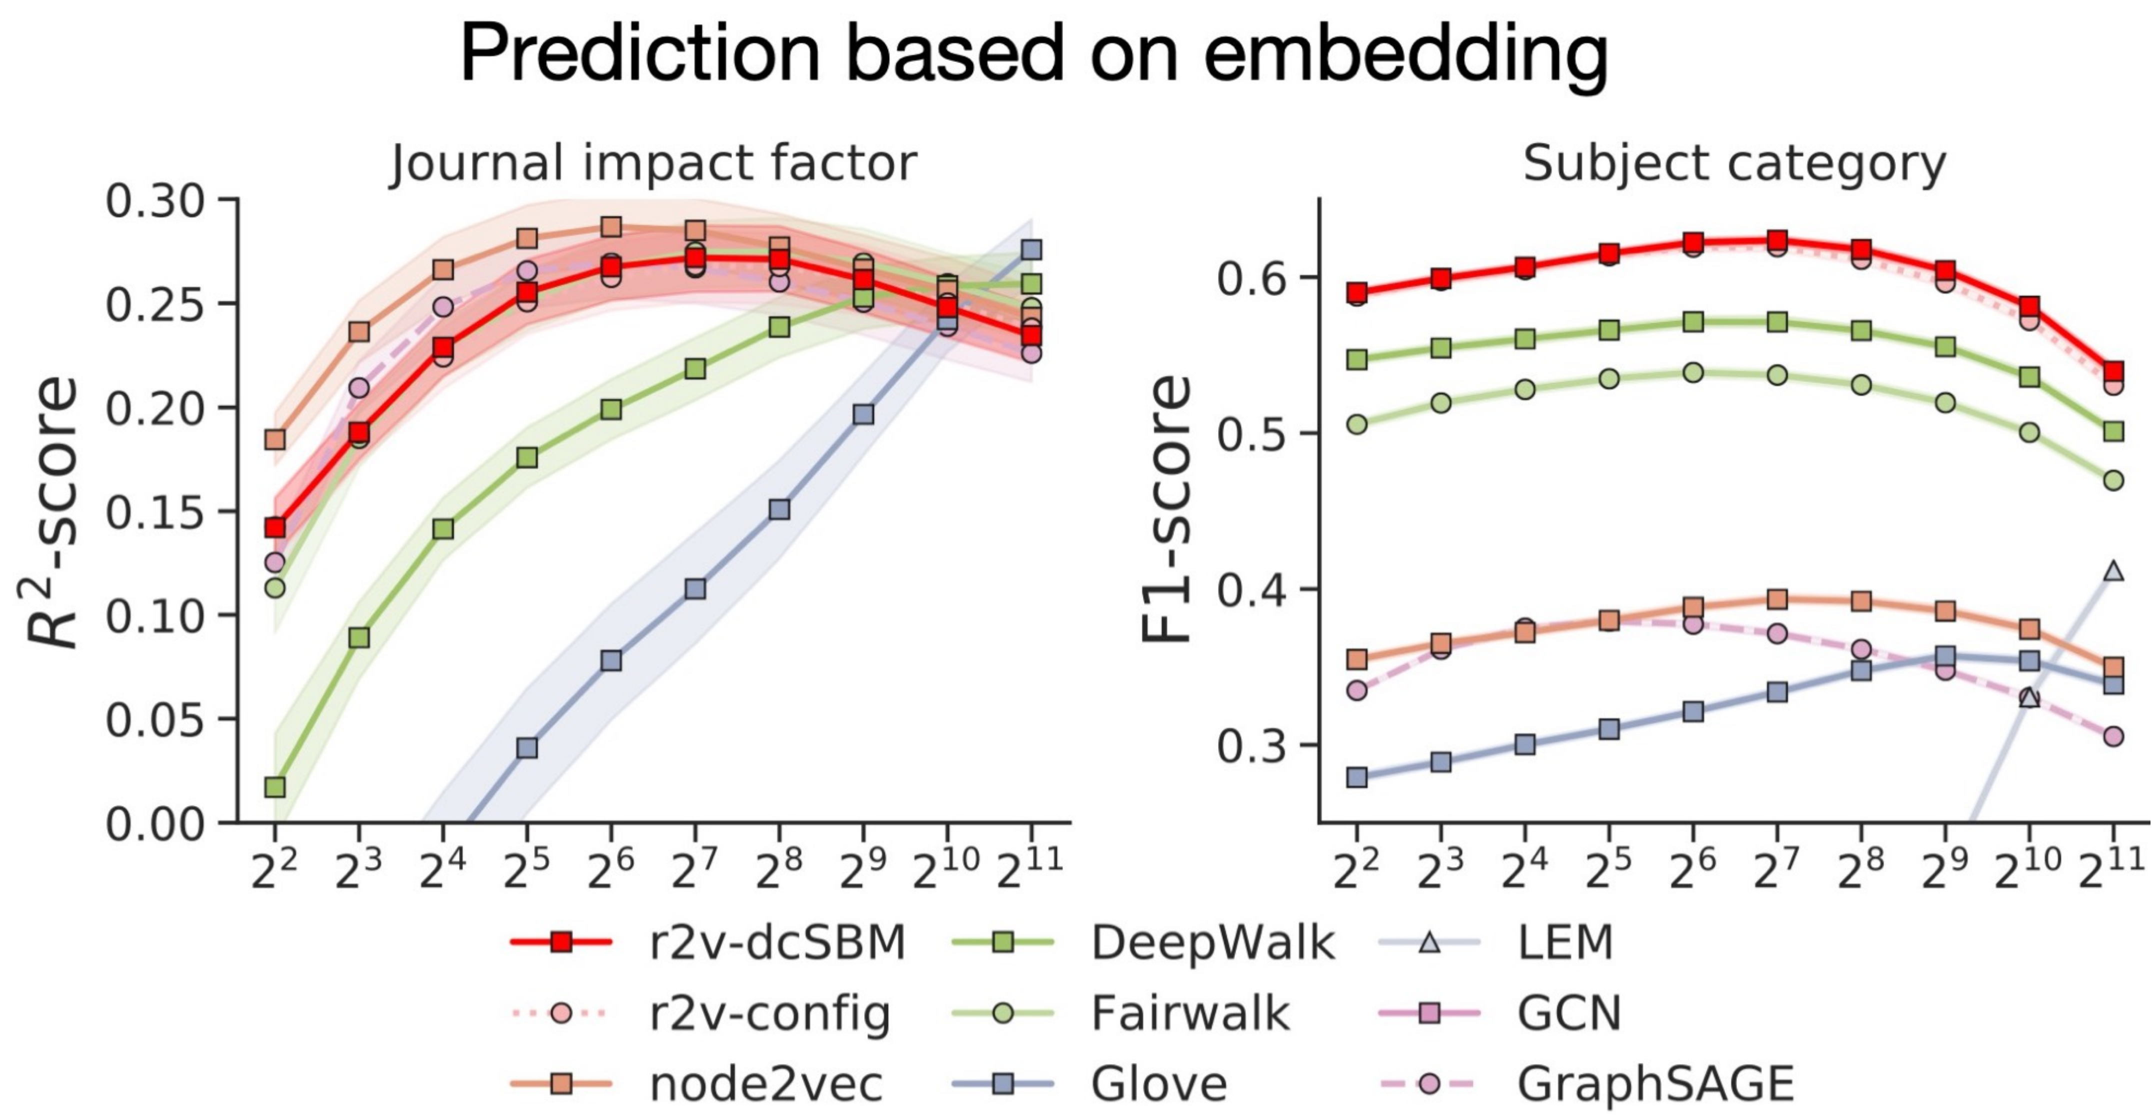

Sadamori Kojaku, Jisung Yoon, Isabel Constantino, and Yong-Yeol Ahn. Residual2Vec: Debiasing graph embedding with random graphs. NeurIPS (2021)

Summary

word2vec has a built-in debiasing feature attributed to negative sampling

Inspired by this finding, propose *residual2vec* that can negate other types of structural biases

Good performance and enable more control on the biases in embedding.

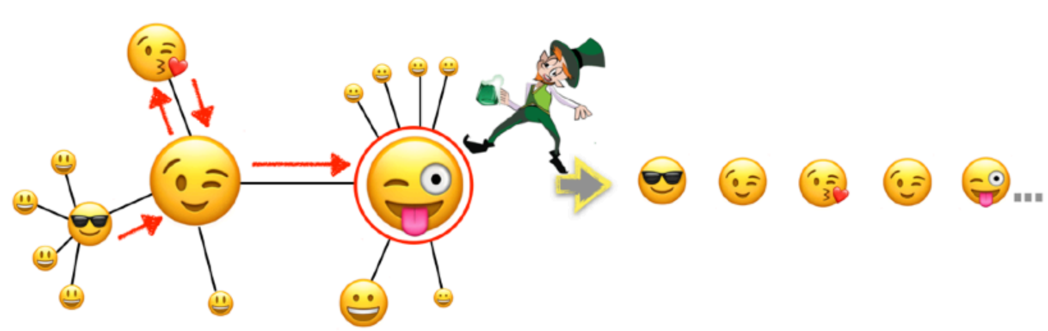

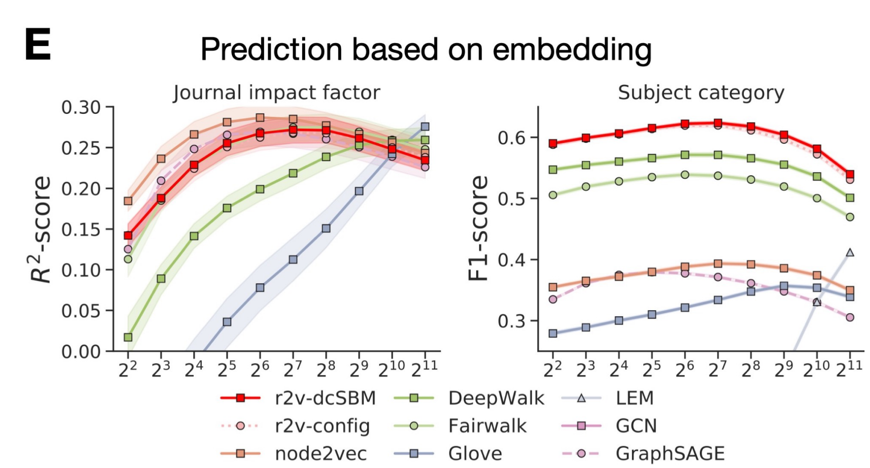

Sadamori Kojaku, Jisung Yoon, Isabel Constantino, and Yong-Yeol Ahn. Residual2Vec: Debiasing graph embedding with random graphs. NeurIPS (2021)

Discussion

We demonstrated a new potential of negative sampling as a way to mitigate bias in representations

This figure is taken from https://www.bbc.co.uk/programmes/w3ct1ls5

Sadamori Kojaku, Jisung Yoon, Isabel Constantino, and Yong-Yeol Ahn. Residual2Vec: Debiasing graph embedding with random graphs. NeurIPS (2021)

INDIANA UNIVERSITY

POSTECH
POHANG UNIVERSITY OF SCIENCE AND TECHNOLOGY
포항공과대학교

INDIANA UNIVERSITY NETWORK INSTITUTE

MIT Connection Science
the technology of innovation

Residual2Vec: Debiasing graph embedding using random graphs

arXiv.org https://arxiv.org/abs/2110.07654

github.com/skojaku/residual2vec

skojaku@iu.edu

Sadamori Kojaku

Jisung Yoon

Isabel Constantino

Yong-Yeol Ahn

Sadamori Kojaku, Jisung Yoon, Isabel Constantino, and Yong-Yeol Ahn. Residual2Vec: Debiasing graph embedding with random graphs. NeurIPS (2021)