Explainable Link Prediction for Privacy-Preserving Contact Tracing

Balaji Ganesan, Hima Patel, Sameep Mehta

IBM Research, India

{bganesa1,himapatel,sameepmehta}@in.ibm.com

NEURAL INFORMATION PROCESSING SYSTEMS
SpicyFL 2020

International Workshop on Scalability, Privacy, and Security in Federated Learning (SpicyFL 2020).

People are not willing to share data on contract tracing apps. What can we do?

Contact Tracing has been used to identify people who were in close proximity to those infected with SARS-Cov2 coronavirus. A number of digital contract tracing applications have been introduced to facilitate or complement physical contact tracing. However, there are a number of privacy issues in the implementation of contract tracing applications, which make people reluctant to install or update their infection status on these applications. In this concept paper, we present ideas from Graph Neural Networks and explainability, that could improve trust in these applications, and encourage adoption by people.

Link Prediction using GNN models can help fill the gaps in contact tracing[2].

Link Prediction is the task of finding missing links in a graph. Given a property graph where nodes are people, and their physical contacts are links, Graph Neural Network (GNN) models can be trained to predict additional *exposure links.* These links can happen even when there is no recorded physical proximity event (because of apps being switched off or people not carrying the phone on them during a chat in office and other similar potential exposure events). A typical off-the-shelf GNN model gives a reasonable performance, but this can be improved with better features.

Dataset	Model	ROC AUC	Std. Dev.
UDBMS	GCN	0.4689	0.0280
	P-GNN	0.6456	0.0185

Comparison of GNN model performance on the UDBMS Dataset

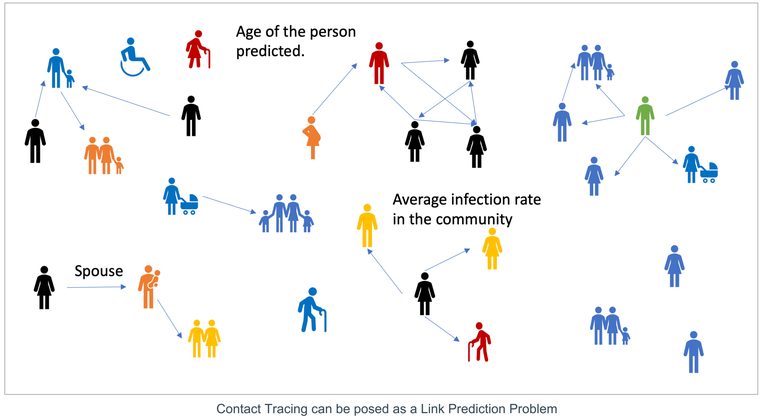

Contact Tracing can be posed as a Link Prediction Problem

Explainability[1] is key to gain users' trust

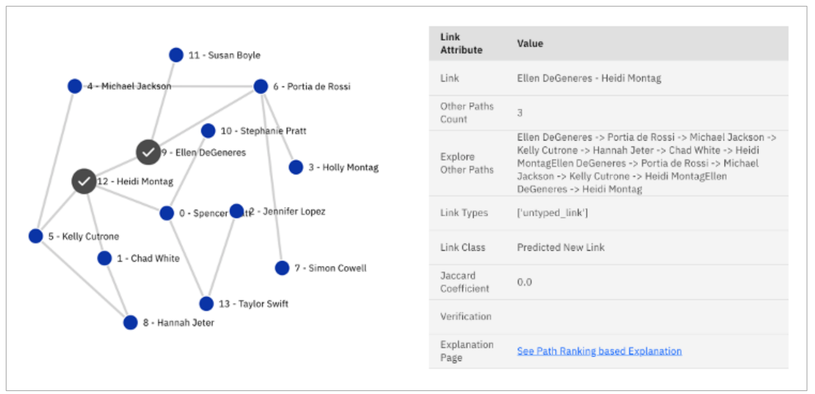

Explainable Link Prediction Demo

Explaining the neural model predictions is naturally very important to build trust in the digital contact tracing apps. But making the explanations more human understandable is particularly important in applications aimed at the general population. We present two ideas that draw from state of the art research to explain exposure link predictions.

Anchors as Explanations[3] can explain features

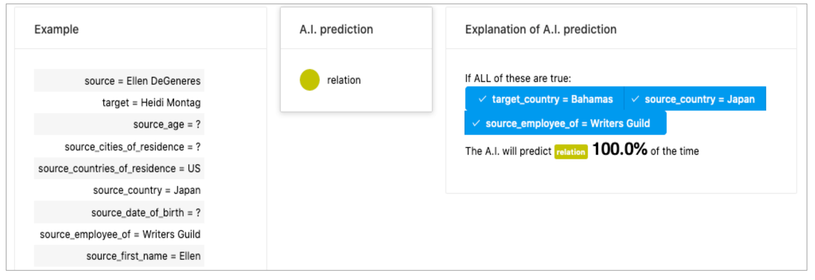

Anchors can be used to prune feature-based explanations and show important features

Path Ranking Algorithms[4] can explain exposure links

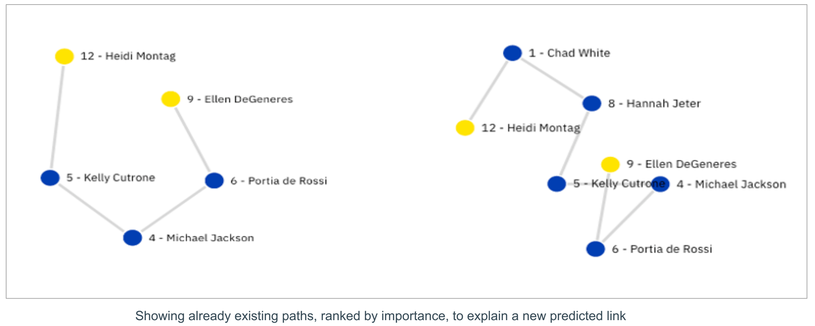

Showing already existing paths, ranked by importance, to explain a new predicted link

We can nudge[5] users to share data by pointing to explanations on feature importance

Age:
(99% of users share this. Age is a significant factor in assessing risk).

Email Id:
(1% of users share this. Email is only for COVID19 related advise. Email Id is not useful for assessing your risk).

Location:
(3% of users share this. Our app does not use your location for contact tracing. But you can use this to automatically call healthcare staff, in case of an emergency like shortness of breath)

Nudge users to share personal data explaining why it's needed for contact tracing.

Graphsheets[2] with details of the model, data and the system can add transparency to apps

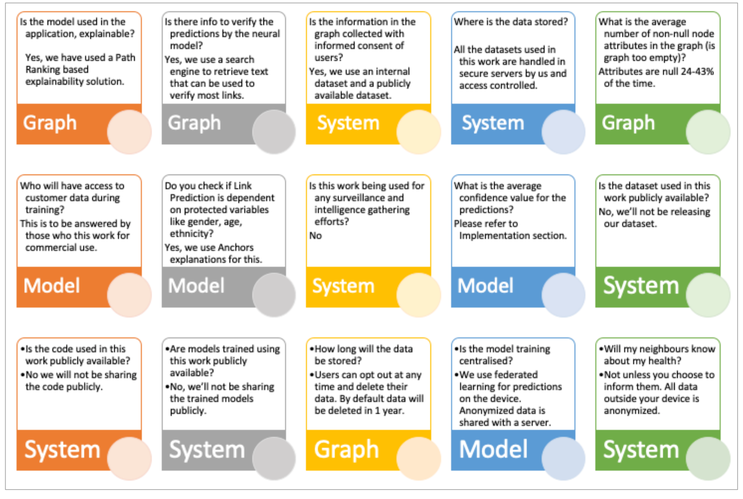

Graph Sheets answer questions users of the system may have.

References

[1] Ganesan, Balaji, Matheen Ahmed Pasha, Sumit Bhatia, Neeraj R Singh, Srinivasa Parkala, Gayatri Mishra, Hima Patel, Somashekar Naganna, Sameep Mehta. "xLP: Explainable Link Prediction Demo". 34th Conference on Neural Information Processing Systems, 2020.
[2] Ganesan, Balaji, Gayatri Mishra, Srinivas Parkala, Neeraj R. Singh, Hima Patel, and Somashekar Naganna. "Link Prediction using Graph Neural Networks for Master Data Management." arxiv preprint arXiv:2003.04732 (2020)
[3] Ribeiro, Marco Tulio, Sameer Singh, and Carlos Guestrin. "Anchors: High-Precision Model-Agnostic Explanations." In AAAI, vol. 18, pp. 1527-1535. 2018.
[4] Melo, André, and Heiko Paulheim. "Detection of relation assertion errors in knowledge graphs." In Proceedings of the Knowledge Capture Conference, pp. 1-8. 2017.
[5] Thaler, Richard H., and Cass R. Sunstein. *Nudge: Improving decisions about health, wealth, and happiness.* Penguin, 2009.