



Breaking Artificial Intelligence's Hidden Glass Ceiling

Amy Hodler



Amy Hodler
Director, Graph Analytics and AI Programs
Neo4j

Biography

Amy Hodler is Director, Graph Analytics and AI Programs at Neo4j (<https://www.neo4j.com>) where she manages the Neo4j graph analytics programs and marketing.

Amy has consistently helped teams break into new markets at startups and large companies including EDS, Microsoft and Hewlett-Packard (HP). She most recently comes from Cray Inc., where she was the analytics and artificial intelligence market manager.

As a network science fan, Amy promotes the use of graph analytics to reveal structures within real-world networks and predict dynamic behavior. She is also co-author of 'Graph Algorithms: Practical Examples in Apache Spark and Neo4j' published by O'Reilly Media.

The book can be downloaded at <https://neo4j.com/graph-algorithms-book/>

Keywords Artificial Intelligence (AI), Explainable AI, Graph technology, Graph databases, Machine learning (ML) Neo4j
Paper type Research

Abstract

In a world driven by connections, graph algorithms provide one of the most potent approaches to analyzing connected data because they are specifically built to operate on relationships. Graph analytics can uncover the workings of intricate systems and networks at massive scales – for any organization. So could graph technology build our trust in the algorithms being used to shape our future?

Introduction

There are too many examples of Artificial Intelligence (AI) and Machine Learning (ML) providing incorrect answers – answers that can make people really resent your brand¹. Is there a way to stop this from happening?

Anyone working in the commercial application of AI knows this is a big issue, but it is sometimes a significant challenge to understand what led an AI solution to make a particular decision.

This problem comes down to three aspects of a core new concept, 'explainability':

- **Data** – What data was used to train the model, and why?
- **Predictions** – Which features and weights were used for this particular prediction?



Analysis

- **Algorithms** – What are the individual layers and the thresholds used for a prediction?


The problem is that I could have the best AI system in the world, but if the data has been manipulated, why would I rely on it? We need to know where the data's been and who's touched it. We need to know exactly when it was changed, what the chain of relationships are, and how that data may be used somewhere else.

Clearly, we need to make AI predictions easier to trace and a lot easier to explain. It's crucial for long-term AI adoption. In many use cases where we would like to employ automation to improve outcomes and the customer experience, such as healthcare, credit risk scoring and even criminal justice, we must be able to defend and justify our proposed smart system's decisions.

Graph tech often used to meet data compliance regulations

Graph database technology could offer immediate help. When we store data as a graph database, it's easier to track how that data is changed, where data is used, and who used what data. For example, graph technology is often used for data lineage to meet data compliance regulations such as Europe's GDPR, or the California Consumer Privacy Act (CCPA). A data lineage approach is also used in NASA's knowledge graph to find past 'lessons learned' that are applicable to new missions.


If we started doing the same for our AI applications, graph technology could tackle the explainable data issue using data lineage methods. That's because graph technology incorporates context and connections that make AI more broadly applicable.





Understanding and monitoring data lineage also guards against the manipulation of input data. For example, corporate fraud research has shown that when the significance of input data is common knowledge, people will manipulate information to avoid detection. Imagine an energy supply system or voting application where we might be confident in our monitoring software, but could not rely on the input data. It wouldn't work as the whole system would become immediately untrustworthy.

Having that explainable data would mean that we know what data was used to train our model and why. This requires storing your data as a graph database. That's a worthwhile step since it provides the ability to track how data is changed, where data is used, and who used what data. Hence, we would always get 'explainable data' on the three axes of data, predictions and algorithms presented above.



Infer an explanation from the surrounding data

Fortunately, graph databases are really good at allowing us to track the chain of data change and subsequent ripple effects. Another area with significant potential is research into explainable predictions. Here we want to know what features and weights were used for a particular prediction. For example, if we associate nodes in a neural network to a labelled knowledge graph, when a neural network uses a node we will have insight into all the node's associated data from the knowledge graph. This would allow us to traverse through the activated nodes and infer an explanation from the surrounding data.



Analysis

In conclusion

The use of graph algorithms could enable us to understand which individual layers and thresholds lead to a prediction. I'm the first to put my hands up to say we are years away from solutions in this area. There is promising first wave research, however, that includes constructing a tensor in a graph database with weighted linear relationships. Promising early signs also indicate we may be able to determine explanations and coefficients at each layer, too.

To increase public trust in AI, the predictions and decisions AIs make must be more easily interpretable by experts and explainable to the public. Without this, people will reject recommendations that don't feel right, when they could be useful and interesting.

AI holds great potential, and using graph technology to help unlock that potential makes pragmatic sense. If you are building an AI solution, look at graph technology to give it the contextual power to push it through the 'hidden glass ceiling' of explainability that's holding it back.

Reference

- ¹ Morrison, S. (27 May 2021), "A disturbing, viral Twitter thread reveals how AI-powered insurance can go wrong", Recode. Available at: <https://www.vox.com/recode/22455140/lemonade-insurance-ai-twitter>