# Let's Talk AI with Matthias Fey

Matthias Fey[1] and Barbara Steffen[2]

[1] Kumo.AI, Department of Machine Learning,
matthias.fey@tu-dortmund.de
[2] METAFrame Technologies GmbH,
barbara.steffen@metaframe.de

***"Scale is hopefully not all you need."***

---

## The Interviewee - Matthias Fey

**My Personal AI Mission:**
Productionizing Graph Neural Networks - a set of models with strong inductive bias that require less data and are easier to interpret.

---

## My Takes on AI

**Artificial Intelligence:** A marketing term for machine learning :)

**Trust:** Fair and robust model outputs.

**Explainability:** Insights into the signals the machine learning model is picking up.

**Essential Elements of Human Capabilities:** Social and emotional intelligence.

## The Interview

**Barbara** *Today I have the pleasure of interviewing Dr. Matthias Fey. Please introduce yourself and your relationship to artificial intelligence.*

**Matthias** My name is Matthias and I completed my PhD in graph representation learning [3]. This involves developing neural networks that receive graphs as input rather than images or sequences. This approach allows us to view neural networks as a much broader architecture. Currently, I'm using this knowledge in a startup called Kumo AI, which aims to bring these ideas to the industry.

**Barbara** *What is your current focus? Are you focusing more on the research side or more on productionizing your technology?*

**Matthias** My work is quite diverse. Since this topic is still new and everyone is still researching it, I also do that as part of my daily work. However, there's also a lot of work in engineering and figuring out how to make things more efficient. I'd say it's a 50/50 split between research and engineering.

**Barbara** *Could you describe some of your current work in one or two AI-related research questions?*

**Matthias** One question, for example, is how we can scale graph neural networks to an industry-sized graph, which could mean like 10 million or 100 million nodes and billions of edges. There are also questions about explainability: how can we allow the user to gain insights from the model and build trust in the model to bring it to production. And then there are research questions on how we can train these models on all kinds of different tasks in a general way.

**Barbara** *What role does trust play in AI adoption?*

**Matthias** Personally, during my PhD, this wasn't a big focus for me because we were researching novel methods and models with better performance and new insights. However, during my time as a machine learning engineer in industry, I noticed customers asking for it, and I realized it's a valuable thing to have.

**Barbara** *And what about your personal trust in large language models as a developer or researcher in the field?*

**Matthias** I have to admit I'm biased. I train these models, I know how they work and how they are able to learn. As such, I am confident about them and generally which hints I need to give the models in order to trust them to learn the right things. But I understand that people see them as a kind of black box and that we need verification that they're doing the right thing. And I agree with that. We are adding a lot of explainability on top in our product in order to build this general trust.

**Barbara** *Do you see any essential measures for the ethical use of AI?*

**Matthias** I don't have any concrete metrics in mind. I think we always need to ensure that there's no unfairness involved, and what that means depends on

the business case. I believe we can achieve a lot of that through publishing and making things open-source, and allowing them to be confirmed by the community.

**Barbara** *On a scale of 1 to 10, what do you think is possible in terms of the future technical capabilities of artificial intelligence? Where 1 describes artificial intelligence systems like ChatGPT. And 10 describes artificial general intelligence that surpasses human capabilities.*

**Matthias** I'm not yet a believer in general artificial intelligence, so I think it's somewhere in between. I believe what we need to do is combine all these models into a larger learning system. This aligns well with my research, where you have all these kinds of agents or models, and you need to route information between them and let them learn from each other.

**Barbara** *Can you give a specific number on the scale from 1 to 10?*

**Matthias** I would say it's more of a 6 or a 7.

**Barbara** *And how does that make you feel? Looking into the future and thinking about all the possibilities from dystopia to utopia, what do you think awaits us?*

**Matthias** I'm not a pessimist, so I would say it's generally a good thing. AI will certainly change a lot of things, but that's not necessarily bad. For example, there are many jobs that I believe are not necessarily enjoyable, and letting AI assist and help here is definitely a welcome change. However, we also need to enforce stricter regulations on which models we're allowed to use and which we should disallow. I think the line should be drawn where we let AI make life-changing decisions for humanns.

**Barbara** *Taking a step back, what do we actually mean by artificial intelligence today?*

**Matthias** That's a really good question, and I think there's no consensus among communities on that. At this point, artificial intelligence is more of a buzzword or a hype term that you use when you want to attract a lot of attention. Personally, I would prefer to simply call it machine learning, where we learn an underlying data distribution and the model adapts to that data distribution. Intelligence would be more something where the model encounters a completely new situation and then tries to adapt to that situation. I think we're very far from that, but we're already claiming that we have artificial intelligence and all these systems already. And then there's this term of AGI, which is even more confusing.

"I have to admit I'm biased. I train these models, I know how they work and how they are able to learn. As such, I am confident about them and generally trust them to learn the right things."

**Barbara** *So when we talk about artificial intelligence and establishing a common understanding of when artificial intelligence is achieved, we first need to establish a common definition of intelligence, right?*

**Matthias** Yes. Currently, there are over 100 definitions of intelligence, and we can't even agree what intelligence actually means for a human being. So it gets even more confusing when you think about artificial intelligence.

**Barbara** *Coming back to machine learning. What are the cutting edge approaches today?*

**Matthias** Everyone is talking about Large Language Models (LLMs) [4], and I understand the hype. I see many of my previous researchers moving to that field. LLMs are trained on a massive amount of data, and there's this claim that compute is all you need [5], which I find kind of boring from a research perspective. These models are then trained to take in a list of sentences or tokens and produce the next word or token that the LLM should output. I come from the perspective of graph neural networks (GNNs), where you have graphs as input rather than sequences. This makes it more exciting for me because you have all this sparsity and all these different kinds of graphs. They don't follow a regular grid layout, which is common in most previous deep learning approaches. The question is how can we efficiently process this graph to make predictions for a certain set of entities. This is a completely different view of neural networks. It also requires less data and less compute, which I strongly favor [1].

"Personally, I would prefer to simply call it machine learning, where we learn an underlying data distribution and the model adapts to that data distribution."

**Barbara** *And today, LLMs and GNNs get applied to similar challenges?*

**Matthias** I would say they're used for completely different challenges. If you use something like GPT-3, it gives very precise answers from text it has already observed during training. But it can't really predict future scenarios. So if you ask GPT-3 who will be the next president in the United States, it will just say that it cannot predict that. But if you're thinking from a predictive AI or predictive machine learning standpoint, then you could use all the information you have in your database or the knowledge you've collected to make some kind of prediction for the future. That's the approach we take with GNNs.

"Intelligence would be more something where the model encounters a completely new situation and then tries to adapt to that situation. I think we're very far from that [...]."

**Barbara** *Could you elaborate on this? Is there an example that is also understandable for non-experts?*

**Matthias** Yes. One example would be that you're a food delivery service and you want to make recommendations for a certain set of users about which restaurants

they should try. If a new restaurant comes up and quickly gains popularity, an LLM couldn't give you the information to go to that restaurant. But a GNN could because it knows about all the recent facts and popularity measures and could make a prediction from the most recent data right away. This is where retrieval augmented generation comes in to strengthen's the LLM's predictive capabilities [2].

**Barbara** *And how do the GNNs know that?*

**Matthias** They don't necessarily know. The learn pattern from the past that are part of the input data.

**Barbara** *Where do you normally get the input data from?*

**Matthias** It's your own database, for example. And whenever that database is refreshed, you would have immediate access to that information and can drive the prediction based on that. An LLM can't necessarily do that.

**Barbara** *Do the models of the neural networks remain fixed, so that only the outputs change given the new inputs, or does the updating of the database also affect the model itself?*

**Matthias** What you see is that the model without retraining is quite stable over a period of time, but it's generally always recommended to retrain it every month or so. This is based on measurements you make on how well your model performs over time. But generally, it stays quite stable.

**Barbara** *And how do you measure or evaluate the model's performance?*

**Matthias** You would measure it in production in an A/B test. You would make predictions and then after a certain amount of time, you would re-evaluate how your model has performed. If you see a decrease in performance over time, then it's a good idea to retrain the model and refresh it with all the updated knowledge.

**Barbara** *So, you basically have a continuous way to evaluate the performance of the model because each output is a prediction that can be compared to what happened.*

**Matthias** It's a difficult problem in general, but overall, that's correct. The main problem is that your predictions may have implications about what your future graph may look like, and that may conflict with the way the model is making predictions in the first place.

**Barbara** *Can you go back to earlier versions of the model? For example, if new predictions have led you in the wrong direction, could you go back to an earlier version that was a better fit? Would that make sense?*

**Matthias** In general, retraining should always improve the model, but the conflict I was talking about is that you're using the neural network to make a prediction for the future, and based on that prediction, you take certain actions. For example, you would send a user a coupon or a notification to go to a restaurant, and that also changes the way the graph is evolving because otherwise, you wouldn't have sent the notification in the first place. And the model basically just knows about what the graph would look like without that notification.

"It will be a requirement for everyone to live with AI in the future and to understand these technologies. We will have to adapt to make that possible."

**Barbara** *Okay, interesting. Are there any specific tests or metrics to compare LLMs and GNNs? For example, are there metrics regarding the amount of data or computing power required, or tests to compare the quality of the output?*

**Matthias** I think the most promising research direction people are currently studying is how to combine these two models. So you could have an LLM and you can enhance it with information coming from a knowledge graph, for example, or you can have your graph that has rich text descriptions and you want to embed these text descriptions using an LLM. This combination of these two models is super interesting but faces major scalability issues, and that's why we haven't seen much progress yet.

**Barbara** *Do you have generic neural networks, or are they domain-specific, for example, for different groups of customers?*

**Matthias** The general idea is that we have one model trained on a certain task, but the model architecture is essentially fixed, and we want this model to be able to generalize to any kind of task.

**Barbara** *Could you give an example?*

**Matthias** The basic idea in our case is that a customer has a single database and wants to query that database to receive information. So, we would have one model that learns from this graph and then have all these smaller models on top where you can make quick decisions in a few shots.

**Barbara** *Now from your personal perspective, what should be the AI vision for the future?*

**Matthias** To be honest, I haven't given it much thought. There are some wild ideas out there about people wanting us to be replaced by robots that can explore the universe for us. But I believe AI should be a tool that assists us, not something that replaces us. If we can build tools that help us do our daily jobs better, that would be my vision for AI.

**Barbara** *Okay, but what if one person using AI applications becomes as efficient and productive as maybe ten people without AI support? I mean, we already saw it with ChatGPT, right? Today, you can write and generate text much faster than*

*if you wrote it fully by yourself. So, when you integrate AI tools into certain tasks, people become more productive, which means that companies need fewer people to do the same amount of work.*

**Matthias** Yes, I think it will be a requirement for everyone to live with AI in the future and to understand these technologies. We will have to adapt to make that possible. Of course, there will be changes in that regard, but it's not like robots will destroy us at some point. I don't believe in that.

**Barbara** *Thank you, Matthias, for your time and insights, especially for adding the perspective of going from AI research to AI product. Have a great day!*

## References

1. Matthias Fey: On the Power of Message Passing for Learning on Graph-Structured Data. PhD Thesis, TU Dortmund University, 2022
2. Yunfan Gao et al.: Retrieval-Augmented Generation for Large Language Models: A Survey, arXiv:2312.10997, 2023
3. William L. Hamilton: Graph Representation Learning, Synthesis Lectures on Artificial Intelligence and Machine Learning, 2
4. Shervin Minaee et al.: Large Language Models: A Survey, arXiv:2402.06196, 2024
5. Samuel L. Smith et al.: ConvNets Match Vision Transformers at Scale, CVPR 2023