

# AI-GR Pod 38 Marinka Zitnik

## 01.15.26

[00:00:00] But the challenge becomes, well, how we can build a matrix or build a model that will be able to go to all the thousands of drugs that are available on the market and many thousands of disease phenotypes and diseases that affect humans, and build a model that can fill in that matrix and be able to predict a score.

How likely it would be that each and every drug among those that are on the market would be able to have some positive effect and be still be safe on patients diagnosed with a disease. And so, TxGNN is a model of that kind. So, it's a, it's, it's a foundation style model in the sense that it has been trained over 17,000 disease phenotypes.

Hi, and welcome to another episode of *NEJM AI Grand Rounds*. I'm your co-host, Raj Manrai, and I'm here with Andy Beam and today we are absolutely delighted to bring you our conversation with Dr. Marinka Zitnik. [00:01:00] Marinka is an associate professor of biomedical informatics at Harvard Medical School. My department, the Department of Biomedical Informatics, and she is an absolutely amazing and amazingly productive scholar working on AI for drug discovery, AI for science, building AI scientists and co-scientists. And essentially working, I think, Andy, you'll agree, on pretty much every area of applying AI to both biology and medicine.

This was a fascinating conversation. She both told us a lot about her research, but I think she also gave us pretty cool productivity hacks and stories behind the stories with some of her big papers. I had a lot of fun talking to her.

Yeah. Uh, Marinka's research output can be summarized as everything, everywhere, all at once. Like, she is an absolute beast when it comes to productivity. And, you know, what I learned in this conversation today is that she really is like a bioinformatics OG almost. We learned in the conversation today that her interest in biomedical AI and computational biology goes back to high school. [00:02:00]

I've known Marinka to be technically. very deep. Mathematically, very sophisticated. But I was really struck by, like, computational biology and bioinformatics was really kind of her first love. And so, she's carried that through. She did a postdoc at a very famous lab at Stanford where she was the

person who introduced biology to this lab, and again, has been just an amazing pioneer at the intersection of biology and machine learning.

And so, this conversation was really fascinating and learned a lot about someone who I thought already knew quite a bit about.

The *NEJM AI Grand Rounds* podcast is brought to you by Microsoft, Viz.ai, Lyric, and Elevance Health. We thank them for their support.

And with that, we bring you our conversation with Marinka Zitnik on *AI Grand Rounds*. Well, Marinka, thanks for joining us today on *AI Grand Rounds*. We're super excited to have you.

Excited to be here and looking forward to this discussion.

Marinka. Lemme just echo Andy. Great to have you [00:03:00] on *AI Grand Rounds*.

So, this is a question that we always get started with. Can you please tell us about the training procedure for your own neural network? How did you get interested in artificial intelligence? And what data and experiences led you to where you are today?

Yeah, that's such a great question. I was always a bit of a nerd in mathematics and kind of software and statistics.

And back in the days when I was in primary school and then high school, I attended all these various international Olympiads in mathematics and computing and won awards there. And then became very interested in the opportunity to make a real-world impact and real-world impact for me meant at that time to be a doctor or to help people.

And then I quickly realized that doing that directly by being in contact with patients might really not be the best for me. Given that I was not the best in environments around, there would be lots of blood or, or other types of things happening. But then I realized in end of high school that there is an emerging [00:04:00] field of bioinformatics or computational biology where someone with interest in mathematics and computer science can have an impact in biology and medicine.

And that was so compelling to me because I could not have imagined how running and knowing how to code and knowing the algorithms can help

discover new things. And so, I had the chance early on to be part of a project in collaboration with Baylor College of Medicine in Houston.

The group there was very interested in finding and understanding how the small creatures, amoebas, which are social amoebas, these creatures that form communities and develop the resistance mechanisms to different types of bacteria. They were interested in understanding what are genes that are responsible for the mechanisms of resistance to bacteria that these social amoebas develop.

And so, at that time, I was able to run some of the clustering algorithms on the [00:05:00] genomes of the, of social amoebas to and produce a ranked—.

Marinka. Yes.

Sorry to interrupt. So, this is, this is in high school or this is in college?

Uh, this is end of my high school.

Okay.

Yes.

Okay, keep going. So, clustering algorithms.

This is end of my high school.

I was able to run clustering algorithms comparing the sequence similarity of the genes of that amoebae. It had 12,000 genes and they knew about five genes that were involved in the resistance mechanisms of social amoebae and the expectation was that there are probably dozens of more genes that are part of the full mechanism that they wanted to identify.

And so, I ran the clustering algorithm. I implemented, ran it, and then, and then produced a ranked list of genes that collaborators from Baylor College of Medicine used to run mutagenesis screens. And so, what came out of those experiments were that they were able to identify nine additional genes out of top 12 genes that my algorithm [00:06:00] recommended at that time.

And so, to me, that seems so powerful because at that time, which was around 2012, it was impossible for them to just exhaustively try out and generate a

knockout for each gene in this organism. And so, instead of having to do that exhaustively, which would take so many, many years, they were just able to run experiments based on the priority lists produced by my algorithm.

And so, that was the first time where I saw real-world impact of algorithms that I designed and implemented where there was new insights, new discoveries enabled by the outputs of AI models. And, and so, uh, since then, I just became hooked to this idea of coupling AI models with interesting multimodal, multiomics data, and produce and generate AI based hypothesis that could then be validated in different types [00:07:00] of experimental, biological, or medical systems.

Amazing Marinka. So, what I'll first remark is that, I don't know if this is like a standard thing for bioinformatics, but we've known each other for, for several years at this point, and we've had lots of discussions about research and early career stuff, and I didn't know until your comments just now and taking us sort of through your background that we both got our start in bioinformatic through clustering of gene expression data.

Although I was not working on amoebas. But I think it's a sort of a standard and a well-accepted point of clustering algorithms being a compelling way to get started in, specifically with gene expression data, to get started in bioinformatics.

So, very, very cool. And then before we sort of dig into your work on AI for drug discovery, AI for science, all the many projects that you're doing now, maybe you could help us fill in the time from, let's say, grad school and postdoc, before the work that you're leading [00:08:00] now as a professor. How did you sort of decide like where and what you wanted to study in graduate school and then in the rest of your formal training before starting your independent lab?

Here is my lengthy journey of going from high school where I started with clustering gene expression values, gene expression profiles, all the way to Harvard where I am now. So, I did my undergrad in computer science and mathematics at University of Ljubljana in Slovenia, which is where I'm from.

And then the journey took me to Baylor College of Medicine based on that initial interesting study that we have done with them in social amoeba, there were many groups that became very interested in AI and the use of my model. So, I spent a year or so there in Houston working with a number of experimental labs, putting algorithms to use for the analysis of genomes and for predicting

protein functions. I then spent a bit of time at University of Toronto where [00:09:00] I worked with Gary Bader and Charlie Boone on studying the systems biology of yeast. And afterwards I went back to Slovenia where I completed my PhD. I did my PhD in two years and 10 months, and then I wrapped it up. It was in computer science. And then I moved to Stanford in 2016 where I started my postdoctoral fellowship there. I worked with Jure Leskovec in computer science department, and at that time I do remember I was the first postdoc in the group working on computational biology.

So, my first meeting with my PI was very interesting because I was talking with my PI about biological functions and protein functions. So, the kind of roles and activities that these important molecules, proteins play in human body, in human cells. And the question that I got back from the PI was, well, can I just write out on, on my whiteboard what is that function and what does the form that that function has?

Because the PI was very much used to mathematical [00:10:00] functions, not necessarily studying biological function. And so, I was very fortunate to be able to introduce computational biology to the group of the PI that I worked with, and at the same time benefit from just tremendous amount of core machine learning and AI research that the group has done before.

And so, there were several very productive, exciting years I spent at Stanford, merging and blending together machine learning models, specifically graph machine learning models and geometric deep learning models, the type of models that operate on networks and kind of relational data sets with problems in biology, increasingly focusing on questions that relate to understanding of drugs and mechanisms of drugs, and what patient populations could benefit from a specific treatment and whether the treatments are safe and effective.

So, those are the kind of questions that I became increasingly interested in during the time between 2016 and 2019, after which I moved to Harvard and [00:11:00] in 2020 and started my lab. Focusing on really developing AI models that can help transform our understanding of existing medicines and also help us with developing better medicines and better drugs for many patient populations, including those that currently do not have good treatments for their diseases.

That's a great transition point, Marinka, thank you for that. And that's a great transition point to where we wanna dig into next, which is your work on AI for drug discovery. And so, I was trying to select a few papers to sort of organize our discussion, right, and dig into, and honestly, you've been, we're gonna ask you about this later, too, so I'm just forewarning you now.

But you've been so productive over the last few years that it's hard for me to pick just one or two papers that encapsulate all the work that you've done on AI for drug discovery, but nonetheless, I tried. And so, the first one that I wanna start with, and I think these both sort of show off two different sides of the work that your lab is really [00:12:00] pioneering and leading.

The first one is your TxGNN paper, which I believe was published in *Nature Medicine* in 2024. The title of the paper is "A foundation model for clinician-centered drug repurposing." Maybe you could take us through the sort of background, like why did you want to write that paper and what that paper's about, and also reflect on the response by the community to the paper after you guys published it.

Yeah, so that TxGNN paper is one of my favorite papers from the last few years because it is an example of how the development of ML models and advances we have seen in the last few years can directly address a previously open problem in the development of treatments specifically for rare diseases, which was the focus of this paper.

So, for a bit broader context, so in the context of drug discovery, we can think about developing new drugs fundamentally in a few different ways. So, one way is to think about developing a new drug from scratch. And here this process is very [00:13:00] long and expensive and resource intensive. It can take on average 12 to 13 years.

It can cost several billions of dollars to develop a new drug from scratch. And it means going to many stages of the development, wanting to develop a new, identify new disease target, then design new molecules in chemistry that would bind and modulate the function of that disease target and do all the preclinical and then clinical work.

And so, it's, you can imagine, that's a very expensive and, and very lengthy process. An alternative, an enticing alternative to drug discovery in developing a drug from scratch is drug repurposing or where the idea is to see whether we can take drugs that are already on the market, that the patients are already taken that are available in pharmacies, and we can expand the use of those drugs for new indications.

So, if we can effectively find new tricks for existing drugs, if we can see that a drug that might have been initially developed for one specific disease might also [00:14:00] show beneficial effect for patients that have a different disease. So, why would we want to do that? We would want to do that because the

timeframe for the development as well as the cost associated with the development of such a drug is much, much shorter, right?

One starts from an existing drug, which presumably has acceptable safety profile, and so it would be shorter and faster to expand its use for a new disease. But the challenge becomes, well, how we can build a matrix or build a model that will be able to go through all the thousands of drugs that are available on the market and many thousands of disease phenotypes and diseases that affect humans, and build a model that can fill in that matrix and be able to predict a score, how likely it would be that each and every drug among those that are on the market, would be able to have some positive effect and be still be safe on patients diagnosed with a disease.

And so, TxGNN is a model of that kind. So [00:15:00] it's, it's a foundational-style model in the sense that it has been trained over 17,000 disease phenotypes. Including 7,000 rare diseases that have no existing treatments, where drug repurposing is an interesting therapeutic approach because it would allow us to find a drug for patients with rare diseases in a more cost-effective manner and also in a faster manner.

Since many populations of patients with rare diseases are relatively, as by the name implies, small populations of patients. And therefore, when these, their diseases can be incredibly heterogeneous and because of the small patient population, complexity of the diseases, poor understanding of the diseases, developing new drugs from scratch for rare disease is very challenging.

And often also the financial incentive for the large companies to develop new drugs from scratch for those patient populations is rather low. And so, drug [00:16:00] repurposing was just an opportunity to help those many patient populations across 7,000 rare diseases, of which only 5% have any FDA-approved drugs.

Meaning for 95% of patients with rare diseases, after they're diagnosed, they really, there is no treatment available for them. And so, how can we handle and tackle this big challenge? That was the main problem and the main challenge that we were after, which TxGNN model that, and then we trained TxGNN model across thousands of diseases, including many common diseases, to be able to learn a metric of similarity, phenotypic similarity, and then molecular similarity between common diseases or diseases with treatments and many diseases with no existing treatments.

And so, after we developed that, that, that detection and model, we were able to then show that it works well. We did it rigorously with standard machine learning. Models and machine learning [00:17:00] evaluations and different types of complicated and challenging splits to convince ourselves and collaborators that predictions about potential drugs that would be effective for a disease, would generalize in zero-shot manner to disease that has no existing treatments.

And after we have done that, we were able to publish the work. We supplemented the model with an interactive open website so that anyone from the community, including clinical researchers or even family members of patients with diseases, or biomedical researchers, can directly interact with the model without having to train it from scratch.

And so, that was the end of publication. Since then, we have seen lots of interest and have several follow-ups where we have used the model to nominate drugs for new diseases. Most recently, we've been working with the center for X-linked parkinsonism at, uh, Mass General Hospital, [00:18:00] um, which is a center focused on a specific, rare subtypes of parkinsonian disease that affects males in southeastern, Southeast Asian populations.

And this is an example of a rare disease that is devastating to patients and their family members but has no existing treatment. Over the last decade, the center has produced lots of insight into understanding of the disease, and they were able to phenotype disease, derive disease signatures, and we were able to run TxGNN model on their description of this rare disease in order to predict the drugs that would likely have the effect on the disease, and that was recently predictions were validated through an in vitro chemical screen, and now we are proceeding to the next stage of evaluation.

Can I hop in here and ask a question, Marinka?

Of course. So yeah, maybe you could give us some intuition for like what drug repurposing is doing. So, I think there's at least two explanations I can think of. So, [00:19:00] if a drug, let's just say aspirin or something you know, you take it for a, a headache, but let's say secretly it's also a cancer drug. Is the, is the right intuition that secretly cancers and migraines, probably in this case not, have the same disease mechanism and we don't know it.

Or that aspirin binds to something or activate some other biological pathway that we didn't understand. Is this a statement about common disease

mechanisms that aren't appreciated or that this drug can act on lots of different biological pathways?

Yeah, so both mechanisms are possible in the context of drug repurposing. So, the rationale for drug repurposing can be, uh, defined in the context of a drug or in the context of a disease. So, from the drug perspective, it is really that a drug typically affects, has its direct target, meaning that when it's ingested by human, it would bind the utmost, often the specific disease target a specific protein, and that protein is typically a protein that [00:20:00] is associated with the disease that we want to treat. But what happens in reality is that many drugs really do not affect only their nominal targets.

Not they don't affect solely their disease proteins, but they might bind also to other proteins and other parts in the human genome. And so, these are known in the literature as off-target effects. Where the drug might influence also the function of other proteins in the human body beyond those proteins that are associated with the disease the drug was initially approved for.

So, that can create problems. It also creates opportunities. So that can create problems because often drugs that bind to many proteins in the human body beyond those that they were designed and to bind to, they are associated with higher toxicity profile. The increased likelihood of safety events.

You can imagine that, well, if the drug should bind to protein A because protein A is the disease pro associated protein. But in reality it also [00:21:00] occasionally, sometimes binds to three other or 10 other proteins that might be so involved in biological processes unrelated to that disease. Then that can lead to unwanted effects in human body.

It can lead to safety issues. It can lead to adverse events. And so, drugs that have these very broad effects on the human body, they often call the drugs that are promiscuous in their binding profile. They can sometimes have challenges in the preclinical and clinical studies in terms of acceptable safety profiles.

When that is not the case, however, we can take those off-target effects of a drug where it binds to other proteins and see whether the drug that might have been developed to target protein pathways associated with lowering cholesterol levels would also target proteins involved with the DNA repair mechanisms.

And if so, we might see and use potential use of this drug in diseases [00:22:00] where the DNA repair mechanisms are perturbed and altered or do not function properly. So, this is one rational for how or why we can even expect that a drug

that was developed for one disease might have effect on diseases that phenotypically, symptomatically look very different from each other.

Right? The alternative approach, the alternative rationale for drug repurposing is disease centric. So, we often look at diseases that have similar phenotypic representations, and then we can look for an opportunity whether these drugs that worked for disease A could also work for disease B based on different measures of similarity between these two diseases.

And here is where the ML comes into play, where we can do self-supervised learning, learn effectively mapping functions of similarities between diseases. That would be where this similarity process of learning similarities or embeddings or signatures of diseases would be [00:23:00] guided based on the similar phenotypes, similar molecular signatures of diseases.

And once we learn that, we can then use that similarity profiling as a transfer learning mechanism. And so that we can say, well, if we have a drug that works for disease A, and now I know this is similar to disease B, symptomatically under molecular signatures, so different levels of data. Then perhaps this drug would also work for disease B.

And so, depending on the strength of similarity, you can imagine that drug would have a higher, get a higher or lower score. So, the two rationales for drug repurposing then are either drug centric based off drug centric approach. It is often based on characterizing the binding profile of what are the molecules, the drug binds to human body or disease centric approach that looks at similarity and profiling of similarity of the diseases.

Cool. Thanks Marinka. That was a really great explanation. So, we're gonna move on to some of the—. Oh wait, actually Andy. Yeah. I [00:24:00] actually do have one more question. Sorry about that. One, one more question on this topic. So just quickly, Marinka, there's another paper that I just wanted to highlight and I think we can maybe just spend a minute or two on it before we, 'cause I really wanna dig into the next topic with Andy.

I just wanna highlight it because I think it showcases another aspect of your group that I think has really inspired a lot of the community and I think also led to a lot of traction for you on interesting biological problems. And we're always interested in sort of like how collaborations start and also what we do as bioinformatics, biomedical informatics folks.

And this was a paper that is called the “Therapeutic Data Commons,” and it’s a machine learning data sets and tasks paper for drug discovery and development that I think you published in NeurIPS a few years ago. And the data sets and benchmarks track back in maybe 2021. And first of all, is my understanding correct or did I hallucinate any of those new?

No. Your understanding is correct. TDC or Therapeutic Data Commons has been published in several different [00:25:00] venues, but one of them is the NeurIPS data sets and benchmark tracks. And I think this speaks broadly to our interests in providing resources for the research community as well as identifying what are machine learning tasks that are real bottlenecks for drug discovery and development.

So, as someone who has been formally trained in computer science, I have always been a bit cautious whether the problems and the task and the challenges that we are going after and that we are doing research on are relevant, can really move the needle, are relevant for drug discovery development. And I’ve seen that especially until last few years, the computer science community or machine learning community developing models for biology and then biology drug, drug hunters, drug developers, those who study how drugs work, were not talking to each other much and there were not many [00:26:00] venues of how they would connect and how they would interact with each other.

That has very much changed in the last two to three years, but around 2020 to 2021, it was still fairly, is these two communities was still fairly isolated from each other. So, our idea was can we really start creating bridges between, on one end, those scientists who are very interested in developing new models.

And on the other side, drug designers, drug developers who really know what are the bottlenecks that needs to be solved, where machine learning could help. And so, the interface of that, we created Therapeutic Data Commons, which is an open ecosystem of data sets of machine learning models, machine learning tasks. For each task we provide, a life leaderboard where others from the community who are developing models for those tasks can make their submissions.

Those submissions are scored on a standard unified [00:27:00] test set so that we know what are the models that perform best for specific task and where are the gaps that require new innovation and new model development.

And this also helps us understand what type of algorithms work particularly well for specific areas of drug discovery development. And so, Therapeutic

Data Commons has been really a substrate catalyst of many research projects in my group. What started as a small summer project for a masters student then became a major initiative where we have currently between 60 and 70,000 active users of Therapeutic Data Commons every month, over 800,000s of downloads of AI models and tasks and leaderboards over the last two years.

And it has led to many collaborations between my group and other groups, both in academia who are developing new drugs as well as industry. So, it's [00:28:00] been really exciting and interesting to see how Therapeutic Data Commons, I think, one of the many initiatives nowadays can help contribute to building better bridges and connecting the diverse interdisciplinary communities together in order to have an impact of drug discovery and development.

One thing, however, that I noticed perhaps in the last few months is a sudden jump in the use of Therapeutic Data Commons. And now what I can say is that the users of Therapeutic Data Commons are not just human scientists anymore, but we have seen the data sets and resources of TDC being incorporated in variety of toolboxes for agentic AI as part of the MCP servers.

And so, it's interesting to see how what started as an initiative that was targeting human scientists and with human readable descriptions of tasks and data sets now became a really fruitful resource for AI agents to [00:29:00] query and incorporate, as well as other machine learning groups to train their foundation models on the many tasks of Therapeutic Data Commons.

So, most recently, Google Health Team has trained two models. One called TxGemma and the other Tx-LLM, that were models that started from base Gemma models, but then trained on data sets and tasks from Therapeutic Data Commons because it's designed in such a way that one can evaluate the utility of pre-trained foundation models across several dozens of tasks that we have covering adverse event prediction, drug repurposing, design of molecules, and binding activity prediction.

To then modeling different types of molecular modalities, and it then becomes really useful resource to benchmark models at scale across many, many tasks. And these models over the last months have become agents which we race to the sudden spike in the use of TDC.

That, that, that's a great, uh, [00:30:00] transition point.

And in full disclosure, if you track some of the IP's that have been pulling down data from the Therapeutic Data Commons, you probably would trace some of

them to the office that I'm sitting in now because we've definitely been using that as a resource here at Lila. So, I think that's a great transition.

So, I think what I'd love to discuss now is some of the work that you've done in AI scientists or agents or sort of whatever the right, uh, noun is to describe this body of work where large language models are using tools to accomplish tasks. So, first maybe could you give us your definition of what an agent is and how these agents are gonna have impact on scientific discovery?

Yeah, so my definition of agent is that agent is fundamentally an entity. Entity, an AI agent. Fundamentally, an entity that can take actions in a specific environment and then get feedback or results of those actions that are performed by the agent in environment [00:31:00] in order to then solve more complex tasks.

And then that involve taking a series of actions over a longer period of time each being informed by what has happened in the past and eventually leading towards a solution for a more complex task or complex analysis than an agent has to perform. Nowadays, AI agents for us are large language models that serve as orchestrators of tools or other machine learning models that then can be invoked and executed in order to perform analysis, to retrieve data in real time, to analyze data in real time, and then help scientists and biomedical researchers perform their everyday research activities and increasingly also not only improve the efficiency of their work, but also potentially [00:32:00] provide insights that might not have been possible with just standard large language models that are reading literature or reading the Internet alone.

And the key functionality, for me, that is so important, the capability, key capability that these AI agents have for science and biomedical discovery is the ability of those AI agents to use tools. I would target, especially for biomedical research. There are lots of insights that exist in scientific literature that one can get and reason about them by reading many, many scientific papers.

But there is just tremendous amount of data that are not described in PDFs of our papers. So, you can imagine if there is a large perturb perturbs data set or some other experimental data set generated, only a few insights or a few, a few examples. A few samples will be discussed in the PDF of the paper, but there is just tremendous amount of [00:33:00] background data or the actual experimental data that might be deposited in some repository that does not naturally come in the form of text form, and that in order to understand the meaning of the data, one really needs to analyze it.

And practically that means that we can imagine, and we hypothesize that we're able to unlock new capabilities of agents by being able not only to read scientific literature and reason about the literature and understand everything that is discussed in the PDF documents of our papers that we publish, but by also going and per and retrieving the experimental data and performing analysis on those data.

And couple that with analysis of literature with generating scientific hypothesis with retrieving information from various curated databases of medical and bio and scientific knowledge, and that that can be done by these models. Now, leveraging tools to analyze data, which [00:34:00] can be bioinformatics workflows these tools can be machine learning models.

They can be API endpoints where the models can retrieve information in real time. So, what I'm most excited now, one of our agent model called TxAgent, which is a follow-up to TxGNN, but it's now an agent model for therapeutic reasoning. This model can reason about all drugs developed since 1939.

How can it do that? It does that by actively sending out API requests to FDA through open FDA mechanism and it can access information about regulatory documents, adverse event reporting profile, information about the use of drugs in specific patient populations for all drugs since 1939. And so, the benefit of that is twofold.

First is that these agents can retrieve information in real time, so it's never outdated. Practically, what it means is that [00:35:00] for TxGNN, which we discussed earlier, we trained the model in one point in time, and so one year later the model is outdated to the extent that we would need to retrain it again, to be able to know what has learned, what has happened, and what was data was published in the year since it was initially trained.

For TxAgent, in real time, the agent will go and get information from FDA. So, during the process of training, we train the agent to know how to use tools, so if it knows how to use tools and the API at FDA endpoint. The FDA has not changed. Then this agent knows how to retrieve information for a drug that was approved last month, even though the TxAgent has been trained six months ago.

Assuming the API signature does not, has not changed. The model can always retrieve relevant information and its information for trusted, validated sources. Which we specify, but why by the defining [00:36:00] which tools the agent can use. So, the tool use capability is, to me, the most critical capability that differentiates AI agents for science from—

Mm-hmm.

—large language models in the traditional sense.

Yeah, so that was a wonderful overview, Marinka, and I think it's obvious that these are gonna be huge productivity boosts for scientists doing research. Maybe if I could ask you, like, a slightly more philosophical question, because this is obviously a very hot area.

Lots of people are talking about AI scientists. You'll see Sam Altman say things like, you know, GPT-6 is gonna cure cancer. And that gets much more into, like, the AI scientist vein. And scientists do more than information retrieval. They do experiments. They reason. So, what's your sense of these as extremely good information retrieval tools and mechanisms versus doing something that feels, like, much more, like, of what a scientist would do, which involves a lot more [00:37:00] cognitive abilities.

Yeah, that's such a great point, Andy. Especially, especially nowadays where we are very, there are many in the field, from big tech to large pharma to academic researchers, excited about this vision of creating co-pilots for scientists so that everyone will eventually be able to have their own research assistants to collaborate with them, or these models even perform certain tasks and analysis by themselves.

So, I think that's a fantastic vision to have. I share that vision. The reality, the current reality, I would say is that the current systems and models that we are building, and that communities building also at large, are being demonstrated to be helpful for identifying interesting hypothesis. And I would think of them, those hypotheses as connecting the dots type of hypothesis.

And there's lots of discussions in the field. What is a novel scientific hypothesis? Which is at the core of [00:38:00] of this problem of scientists are generating new scientific hypothesis. That the novelty means pushing forward the frontier of knowledge. And so, if AI scientists, this kind of compound agent systems can generate new hypothesis, are there's really novel hypothesis.

I would say the current generation of AI agents for science are connecting the dots in the sense that they would be able by reading literature from different subfields or sub areas of research and analyzing different types of data, realize that if we take, like, insights from 10 different subfields that were.

We would be able to put together an interesting hypothesis with a research plan that can then translate into an experimental validation of that hypothesis. And so, there is a value already in the being able to generate this type of hypothesis when somebody can connect the dots. I think recently there [00:39:00] was a fantastic example of that in mathematics in particular where GTP-5 Pro was initially thought of us having solved on its own a few open problems from the list of open problems in mathematics that you see Berkeley—.

Is this what Sebastien, Sebastien Bubeck tweeted out? These are, these are the Erdos problems. The Erdos problems that

Erdos.

Erdos, yes. Erdos. I think there were two or three problems that initially the field, we were very excited about those models, kind of those insights, because the expectation was that the model has independently solved them.

It turned out that what happened really is that there were solutions of those problems in the field, but they were published in obscure journals, or they were published in technical reports, not in peer reviewed venues somewhere back in the day, several decades ago. And so, other scientists were just not aware of them.

And so, this model was able to elevate, find those proofs, find those find [00:40:00] for these open problems and present them to human scientists and human mathematicians. And, and so this is an example of being able to just do very deep thorough literature research and in a nontrivial manner and connecting the dots from different subfields and from different areas of science to be able to generate interesting novel hypothesis.

And I do hope there will be more and more of that in in the future. We now have several anecdotes of early that are early examples of success, whereas scientist models were used to generate interesting hypothesis, to some extent, novel, that then were followed up in experimental screens to validate and evaluate them.

I do hope that the future will bring two things: A. To do not just one round of this AI and then experiment, experimental campaign, but have multiple rounds of human-AI experimental discovery loops. So, AI scientists makes prediction, it [00:41:00] going to validating that prediction in experiments, the data being generated, and then we repeat that process, taking the generated data, inform the

model for the next round of experimental campaign. And so that would lead to some continual loop of discovery.

And second, there is an interest in taking these models and embody them in some ways so they would be able to more directly interact with experimental systems, which I think it's also interesting to explore further.

Cool. Awesome. Marinka, that, that was awesome. Um, I think we're ready to move to the lightning round. Are you, are you ready, Raj? I'm ready. Let's do it. Are you ready, Marinka?

I'm ready to, uh, for the lightning, lightning round. Okay. Let's do it.

Alright. So, the first question in the lightning round is, what was your first job?

My first [00:42:00] job, um. Oh yes, I know my first, oh, my first job was actually, was actually offering, um, I was, I was when I was, when I was, when I was that, yeah. So that was end, uh, first year of my high school, first year of my high school, which is Slovenia with me.

I was around, I was 16 years old. Uh, so at that time I was. Offering. I was tutoring, uh, uh, students from primary school in mathematics and computer science and logic in particular. So, there was, there is strong emphasis in Slovenia in kind of solving these logical puzzles and playing chess and do this kind of work.

So, I was very good at that. So, I was preparing primary school students for the national competition in mathematics and in logic. And then those who won at the national competition went to the International Olympiad in mathematics. And so I was, um, tutoring those students and coaching them.

[00:43:00] Wow. Super cool. So, Marinka, should we, this is not a lightning round question, this is just like a piggyback to that one. Yes. But should we start a DBMI chess league or a, uh, a puzzle league or, I think there's a fair number of people who are interested.

We, we should, we should do that. And I think we should also have an AI be one of the players in the league.

Oh, I, it's gonna, it's gonna make me feel insecure. It's gonna make me feel insecure. Okay, next one. Marinka, what is your ultimate productivity hack?

Um, my ultimate productivity hack is the following. So, every day I write on a blue post-it note, a list of things that I need to accomplish by the end of the day.

So, these are like urgent To-Do's. And so. My goal is that the day ends when the list is empty. It doesn't often happen nowadays that the list is empty when, when we get to the evening of the day, but it just really helps me ground so I don't get [00:44:00] distracted too much by a flood of emails or various requests so that I still get focused on what are just really the key items or the key activities that I wanted to complete in a given day.

So, yeah. That's been really helpful for me. Keeping the main things, the main thing. Right. And then what—,

Exactly.

—when do you, just as a follow to that one, when do you write out that list? The night before or the morning of?

The night before.

Okay. Nice. Yeah.

Yeah.

Cool.

Awesome. Um, what is the, sorry. The next lightning round question is, what is your most favorite place that you have ever visited?

My most favorite place that I've ever visited is a lake called Bled, B-l-e-d. So, it's a really, really nice, fantastic spot in Slovenia. Slovenia is a small country. There's region of it that's close to Alps that overlaps then with Italy and Austria. So, there is a place there called Bled, and it's a lake and there's an island in the middle of the lake.

[00:45:00] And it's incredibly, incredibly, nice. And so, around that lake, uh, and the island, there are mountains and so you can go there during summer you can, one can swim. In the winter, there's some snow. But you can still go to that lake, to the island in the middle of the lake. And it's really beautiful.

You should, you should Google for it.

Sounds really incredible. Yes.

Amazing. Okay, Marinka, the next lightning round question, who is your scientific hero or, uh, stated another way, other than your PhD or postdoc advisors, whose lab would you be in if you were in graduate school today?

Oh, wow. That's an interesting question.

Whose lab would I be in if I were a grad, graduate student today? Um hmm. So, that means that my hero needs to be someone who's alive today.

Uh, it could be any, any era. It doesn't have to be alive today.

Okay. So, I, I, I to, you know, broadly speaking, I would say [00:46:00] my, my hero would certainly be, um. The would be, uh, famous, uh, famous phy, the physicists.

I just think the way, uh, then, uh, the, the way the philosophy of how science is, is done is resonates very closely to me. I also closely with me, I also appreciate the idea of thinking about complex systems and how complex systems can be modularized into simpler systems that we can address and tackle with AI models or different modules or different agents nowadays.

Um, and, and to, to build. To build up solutions, uh, for, from, for, for, for bigger challenges. Um, whose lab would I be now? I would say someone from DBMI department, DBMI Lab in at Harvard Medical School.

Amazing. Marika. So, in, I'm pleased to offer you any, any uh, any appointment you would like in my lab. I wish.

I wish. I'd join your lab. I wish, Marinka. I wish. Alright. Andy, next to you.

Cool. Thanks Marika. So, we are recording this on [00:47:00] October 31st, which means it's Halloween. And so, the Halloween lightning round question is, what is your favorite kind of candy?

My favorite candy, uh, is, uh, Twix.

Oh, nice. Mine too. It's classic.

Wow. We have a lot of, there's a lot of interesting still there is there.

I, I'm, I'm a sour, I'm a sour patch guy.

Oh, that's, that's Rachna's. It's my, that's my wife's favorite by far. So, alright, here's, here's the last one. Marinka. Will biology be driven more by biologists or computer scientists?

Hm hmm. Uh, um, perhaps by. Uh, the way things are developing now, I would say increasingly more by computer scientists who know which are the problems to solve. I, I, I would say conceptually I've seen these changes in the field that to me were, are just incredible. When I was a PhD student, I do remember I was knocking on the doors of experimental biology.

And I on. I was knocking on the doors and I said, well, I [00:48:00] have this exciting algorithm. I think it works really well. They have some interesting problem or interesting data I can apply this algorithm to. And nowadays I feel the last few years.

It's inverted? The rules has have, have, have changed. Yeah. And so it's often just we're developing different types of models and are being approached by, by fantastic researchers who have interesting biological systems or biological questions that they would like to answer.

And so, it's almost like, I feel that, I mean, this store with lots of candies and it becomes a question of which is a candy in terms of biological question or an opportunity that would be of most interest to us or being that we could engage with and collaborate with that group.

I do think that the, the challenge however, is that we sometimes, as computer scientists, do, especially those that are newcomers after a few years to get humble. But initially, we do not necessarily appreciate the significance and complexity of the human biology. [00:49:00] And so that might, that can lead to some of very broad terms about how we will be solving many complex diseases in just a year or two, perhaps that might not happen.

I do appreciate the enthusiasm and lots of resources put to use in AI towards medicine and towards biology. So, I think computer scientists would have play a major role in driving the biology further, but in order to make real impact, we need to think about what are the key bottlenecks where AI can productively help and specifically the AI that we have today as of, as opposed to the one we had a few years ago and the one that we might have a year from now.

Excellent. And, uh, also, uh, extra points for using, reference the candy again on Halloween in that answer. So, Marinka, amazing job with the lightning round. You've survived, you've passed with flying colors. Congratulations. We have just one last [00:50:00] question. It's sort of a, kind of concluding, you know, big picture question that we want you to reflect on.

And I think you're probably the best guest that I can imagine to reflect on this. Um, and actually, Andy, I might ask you to sort of opine a little bit if you're, if you're comfortable with, as well, 'cause you both are two of the most qualified people to really talk about this.

So, this is a question about sort of the moment we're in and where the best perch is to be working on interesting problems at the interface of artificial intelligence and let's say biomedicine. So, biology and medicine, broadly. Marinka, could you give us the strongest case for staying in academia during the current AI boom?

And then as a follow on to that. Uh, what is the strongest case for not staying in academia right now?

Okay, let me start with the strongest case to sustain academia right now. I'm sure that Andy can, can add to that substantially. So, I think one, perhaps the strongest case is the, is really [00:51:00] complete freedom to chase hard questions, potentially questions that need multi-year arcs to be answered and to be solved.

I often think, together with students, when we are starting new research projects. The question that we need to answer before we start a new research project is always, why are we the best group to address this research question? So, why cannot be solved by many others? And if we can come up with a meaningful argument that we are, we have an ideal set of insights, unique set of insights, unique set of expertise, unique combination of data sets and models, or unique combination of data sets, models, and collaborators. Then we go and work on that research project, and I often think about projects that are challenging. The two, three require multi-year arc of development, so I think that's one very compelling case for why one who is interested in these types of questions can go about these very challenging problems in academia.

The second one is [00:52:00] perhaps the ability to cross different modalities. And so, by that I mean ability to cross wet lab with clinical context with computation. And Harvard, for example, is such an ecosystem with so many

different schools, many different departments, many different health care systems.

That allows you researchers to collaborate with a broad range of groups with very different expertise. I do acknowledge that there are now several very interesting, exciting ecosystems in industry that are replicating or building debt, including the kind, some of the, ecosystems that Andy is part of.

But I still find this intercross disciplinary access, very important for modern science. Then I would say those who are interested about training people who will be the next generation of science, then that's, that's what academia is for and, and ultimately being able to have access to some of the data.

Either data from consortia or [00:53:00] patients that shape what we built and would be challenging to access from us, from industry. And so, all these four things together. So, the freedom to chase very hard problems, the ability to connect different modalities and different disciplines. The desire to train people who will shape the frontier and the future of AI and the future of biomedical AI and access to interesting new data sets or instruments form this case for why being in academia is still great and fantastic thing to do.

That was very well said, Marika. I think the breadth of things that you can work on in academia is great, and you and Raj are both, like, shining examples of that done well.

I think I agree with what you said. In startup it's increasingly less true that there's that kind of breadth in industry and increasingly like industry is where you can make a concentrated research bet [00:54:00] and you can, like, if you need a lot of resources behind one big bet then industry is a good place to do that.

But I think you, you and Raj both have done a great job at showing what you can do given academic freedom to work on many different things.

Yeah. One thing to add here is also scale. And so, some innovations really require scale, and it can be scaling data generation scale in this size of AI models that we are training in terms of the number of model parameters that an AI model has or scaling the number of use cases or evaluations or tests or molecules to design that, that we would need, that we would want to achieve.

And so, so the scale piece is some is, is uh, I think where industry really shines. I'm mostly excited also about the opportunities to collaborate with industry to

scale some of the prototypes or initial models that we are developing in academia and scaling them up towards real [00:55:00] world use by scaling either the data or the model or the downstream applications or analysis that models are involved in.

And that's kind of the balance between the scaling the data model or evaluations on one end, and then just the ability to ask in a very unconstrained manner, interesting, exciting research questions without being necessarily tied to a specific model or tied to a specific use case is where there is an interesting synergy between academic research and researching industry.

Amazing. Well, Marinka, I think that is a great note to end on. Thank you so much for being on *AI Grand Rounds*. This was wonderful.

Thank you for having me. It was really a fantastic conversation. Uh, and, um, yeah, I'm, I'm very happy to be here and, and look forward to just seeing what the future brings for us in biomedical AI and AI agents and, and the [00:56:00] transformation of scientific discovery.

Amazing. Thank you, Marinka. Cool. Thanks, Marinka.

This copyrighted podcast from the Massachusetts Medical Society may not be reproduced, distributed, or used for commercial purposes without prior written permission of the Massachusetts Medical Society. For information on reusing NEJM Group podcasts, please visit the permissions and licensing page at the *NEJM* website.