



Utrecht  
University

Sharing science,  
*shaping tomorrow*



INAUGURATION

Daring to Trust AI

Prof. dr. Pinar Yolum

July 2, 2024

# Daring to Trust AI

**Pinar Yolum**

Department of Information and Computing Sciences  
Utrecht University

*Inaugural Lecture*  
*Delivered on July 2<sup>nd</sup>, 2024*

Dear Mevrouw Rector Magnificus, Dear Colleagues, Friends, Family;

As I look at the audience, I see many familiar faces from various stages of my life. Thank you all for being with me here today.

In the next 45 minutes, I will tell you a bit about AI, my take on it, my journey---if you will---to dare to trust it.

But, I also hope to tell you a parallel story of daring and that is of daring to be an AI academic, as a woman, as a minority, and as a dissident.

Let's begin.

Early 1990's, I am in high school, did a course on programming, and completely fell in love with it. Binary trees, stacks, and recursion are my new cool. I am spending hours and hours trying to write programs on everything I see around me. One day I am working on a program that can play backgammon, another day I am thinking if it is possible to automatically generate fortune-telling from coffee cups? You get the picture.

Following year, I do national exams on programming and I am selected to be trained for the Turkish team for International Olympiad in Informatics. We do weeks-long camps on AI algorithms, complexity, programming. I am the only girl in a group of around 40. I am scared to death to ask questions, I am scared to death to speak up. With every new problem, every new algorithm, I think that the boys probably already know the answer. A feeling, I am sure, many computer science women in the audience recognize.

Spending long days, long nights, and all holidays programming, I make it to the team. Incredibly exciting times at a personal level, which, officially marks my journey into becoming an AI researcher.

## **AI around us**

Fast forward 30 years to now, This time being an AI researcher is exciting in a different way. Pick any of your favorite area of interest---from business to media to sports to history---you will come across AI systems that are doing incredibly sophisticated tasks. Sometimes solving problems that have been open for years; sometimes easing tasks that have been tedious before.

This is also a time when it is extremely draining to be an AI researcher as we try to understand, internalize, and keep up with all the AI buzz out there, trying to figure out constantly whether that new best AI tool---whatever it is---can actually do any of the things that it claims to be doing.

This is also a time when being an AI researcher is somewhat embarrassing (or at times even disgraceful).

As AI is rapidly moving from university and industry to society at large, we are witnessing bias, unfairness, and lack of accountability in what AI systems do as well as how they are used.

From chatbots that reproduce hate speech to HR support systems that favor men over women to military support systems that “intelligently” identify targets to hit, AI—as is used now—is not always representative of the values many of us stand for.

Interestingly, the particular AI techniques that create a lot of excitement are the exact same ones that create embarrassment.

Take ChatGPT as an AI system that we have all played with recently in some form. We ask ChatGPT questions and it answers in ways that we can't tell apart from human answers. Often we are impressed by its functioning. But, at the same time, the same technology generates text in which we find numerous reasoning errors, from calculations that do not add up to facts that do not check out and we are shocked by these mistakes.

Take AI systems that recognize our face when we use our phones. We don't think twice in using them. Many of our passwords are saved behind our face and we are grateful for the ease. But, the same great technology is not that good in detecting all kinds of faces. You have to be part of the majority to enjoy the benefits. What is worse, sometimes you even suffer from the same technology. Say, when used by law, to detect if someone has been involved in a crime. Imagine all the mistakes that can come about.

Or, take recommender systems that help us pick content on Netflix based on what we have liked before. It works great when it finds the perfect movie for a Friday night. At the same time, the same technology can keep track of what choices we are torn between and try to persuade us towards one direction.

Remember the Cambridge Analytica scandal when the recommender system on Facebook showed content to nudge users to vote in a certain way, essentially changing the outcome of an election, maybe even the flow of history. Or, better yet remember the Dutch childcare benefits scandal, where many parents with immigration backgrounds were falsely accused of doing fraud. Allegedly a similar AI system was the mastermind.

## So, what is going wrong?

- Are we using the AI techniques in *wrong applications*?

Sometimes. Just because a technology we have **can** do something doesn't mean we have to use it for every purpose we can think of.

Choosing which applications to use AI for is a political decision. An analogy that I find useful is with cars. There are parts of the world where you would have to get into your car to get bread. You are car-dependent. In other parts of the world, say here, you can live by for weeks without getting in a car. How and when we choose to use technology is a choice we make as a society.

EU AI Act, which passed earlier this year, aims to regulate when and where AI can be used and what measures need to be taken for safe usages. However, even after regulations are in place, there is still a lot of room for things to go bad.

- Are the people that use or develop AI not trained well?

Most of the current AI systems are developed, without necessarily thinking who will use them or what skills they would need to have.

Take cars again. You can't drive a car without a driving license. When you obtain a driving license, you don't only learn how to drive a car but also how it works, what can go wrong, how you can understand and attend to these things. However, so far this has not been a common practice in the deployment of AI systems.

- Are we using the wrong data?

It is fashionable these days to blame the "training" data for anything that goes wrong.

Some AI systems need to look at examples of how a task was done before, to learn how they can redo a similar task. More examples, the better. Intuitively, if we provide enough, well-balanced examples, the AI systems would learn just the right things and for example, not have bias. While to some extent true, this doesn't do justice to the situation. Think of yourself as an example. You are exposed to all kinds of hate speech, bias, unfairness every single day, especially those of you that hail from my neck of woods. Yet, we don't learn to be unfair or we don't learn to discriminate. Apparently, there is something else that stops us that does not stop the AI systems.

Which then leads to my final question:

- Are the AI *techniques* that we develop insufficient?

I would say largely.

We are interested in developing AI systems that do tasks that are useful for humans but at the same time we are not necessarily trying to replicate how humans attempt to do these tasks. On one side, this is useful because AI systems can compute things much faster than us and can leverage on this. But at the same time it is tricky because they can't apply all the checks and balances we do as humans.

We need to improve our existing AI techniques to prioritize working with humans.

## **But how?**

When answering this, there are many concepts that are rolled on the table. Many say: AI should be accountable for its actions, transparent in its reasoning, fair in its decisions, explainable in its processes, ethical overall. You name it. Often the discussion focusses on the definition of these terms, how they relate to each other, whether one subsumes the other, and so on. For some reason, we are OK with institutions that are not transparent, people who are not ethical, a world that is not fair but when it comes to AI, it needs to fulfill all of these golden properties simultaneously. This, I find, is an unrealistic standard to live up to. Instead, I am going to pick parts of these properties and opt for AI that we can trust.

Trust is a concept that I have been working on since my PhD times. At the time, I looked at how trust comes about based on the interactions of network of AI systems.

So, when I say I am going to opt for AI that we can trust, I don't mean blindly to trust them. Just to contrary, we need to be developing AI techniques that exhibit certain desired properties that will enable us to trust them.

Similar to how other humans make themselves trustworthy by taking certain actions, showing certain behaviors, or even not showing certain behaviors, AI systems should be built in ways to be trustworthy.

At the same time, we need to remember:

Trust is not a choice between right and wrong; trust is relative.

Trust is not confidence in external events; trust requires taking an action.

And, maybe more importantly trust is not a guarantee; on the contrary, trust is taking a risk for the benefit that it will bring.

## **But, what then is trustworthy AI?**

I am going to take a pragmatic view and move away from dense terms. If the AI systems that we build are trustworthy, they, the AI systems, should be able to say the following five phrases in their interactions with us and with each other.

**“I don’t know”**: This is a sentence that I wish men would use more often and I hope AI systems will pick up soon. Most current AI systems are not equipped to judge their own expertise or limitations when doing a task. When they are torn between two different things, they can pick one somewhat arbitrarily, rather than stopping to say “I don’t know”. Making such arbitrary choices lead to mistakes.

However, reflecting accurately on how much we know of something essentially reflects competence and improves trustworthiness.

**“Here is why”**: We, as humans, do not only take actions or communicate decisions to others but we also expose our thought process to others. The steps taken can easily show what factors we considered, at what stages. AI systems do not always have a human-like thought process so it is difficult to say “Here is why” in a way that would be intuitive for an end user.

Yet, being able to say this to its user, and when needed to other AI systems, reflects willingness to collaborate and improves trustworthiness.

**“I got this”**: On a first glance, this might sound like a flip side of “I don’t know.” But, this sentence says more. Not only expertise to carry out a task, but also a strong autonomy to decide on what to consider, what to discard and to take the initiative. It requires the AI system to know the social conventions, norms appropriate in a situation. Using this autonomy responsibly improves trustworthiness

**“Fair enough”**: For humans, understanding others is a virtue. This means to stop thinking of ourselves alone, to stop pushing our thoughts to others, to stop and understand the other person also has a point. I think the same should hold for an AI system.

Rather than always pushing for its own utility, own benefit, an AI system should be able to understand and respect others’ choices and say “Fair enough”. Having and exhibiting such empathy improves trustworthiness.

The last phrase that I think AI systems should be able to say is:

**“With all due respect”**: Another way of saying “I disagree”. We sometimes assume that humans that work with the AI systems always know what is best. Of course, not always true. A user might ask an AI system to do something, which might not be appropriate in a situation, might have undesirable consequences for the user or for others. In that case the AI system should simply disagree and not take that action. Knowing that the AI system has this integrity increases trustworthiness.

I think when AI systems start to say these five phrases more often, they will come closer to being trustworthy.

Not surprisingly, these statements are also related many of the AI problems I have worked in throughout my career. I have been fascinated with values and rights that human societies run on. I have tried to understand these concepts from a social perspective, creating computational abstractions to capture them, building AI systems that could understand their intricacies and thus work better with humans.

One value, or even human right, that I worked with is privacy. Let me tell you more on privacy and then show how these five phrases play a role there.

## AI and privacy

When we hear the words AI and Privacy in the same sentence, it is generally not good news. But, could AI instead be helpful to us in managing our privacy?

Before, answering this, let's look at why we would need help in the first place.

- Humans are notoriously bad at managing their privacy. Take Websites. As soon as we enter a Website, we see a privacy window that asks us to configure things or to accept certain cookies. Most of these privacy configurations are filled with jargon, making it difficult for many of us to understand. We don't have the patience to deal with them, so we often just accept all without even reading them.
- Social media is another beast. A part of us wants to share pictures, thoughts, jokes with others. At the end of the day, why are we on social media if we are not going to share content? The other part of us wants to make sure that our private moments do not reach unwanted people. Thinking of what consequences sharing might have is again time-consuming and error-prone.
- That said, it is not only these pragmatic concerns that are the problem. Often, we also genuinely don't know what is private for us. When asked directly, people often express strong concerns about their privacy. However, when it comes to real-world actions, they tend to disclose more personal information than they previously state. A phenomenon known as Privacy Paradox.
- What is also challenging is that privacy is more than access-control to data.

What one person finds private, **the other** is happy sharing.

What one person finds private **now** does change later on.

What one person finds private **for outsiders**, they might not find private for their friends.

The dynamics of privacy is intricate, context-based, and personalized.

## AI for privacy

But, what can AI actually do for privacy? Over the past 10 years, with my students, I have been working on to answer this question. We have designed and developed many techniques that are used by AI privacy assistants. Let me tell you more on what these privacy assistants are and what they do.

Privacy assistants are AI systems. They work by your side to help you manage your privacy. It could be an app on your phone or as a plugin for your browser. When you are about to make a decision on

privacy, say you are sharing something on Instagram, your privacy assistant can analyze the implications and warn you if it foresees this to create privacy violations. Thus, the damage by these violations could be minimized or even avoided. Moreover, they can learn what is private for you and help recommend what sharing decision is in line with your preferences, thus in some way bridge the gap between your privacy preferences and your behavior. These assistants can tell verbally and visually why they perceive a certain image to be private.

In order to accomplish these, we use many existing AI techniques, that vary from computational argumentation to deep learning to model checking to knowledge representation and adapt many to fit our purposes.

## **What is private?**

One way of knowing what is private is to ask the owner of the picture to describe what is private for them. For example, a person can say “I find pictures of babies private”.

The privacy assistant can then take this information and flag any picture with a baby as private. But, more often, privacy preferences are abstract.

A person can say, “I find pictures that reveal my location to be private”. This is more difficult of a task for a privacy assistant as it cannot only look for objects, like a baby, but it is looking for any information from which location can be inferred.

This can be a landmark like Dom Tower; it can be others in the picture; it can even be a combination of say, local flowers and tiles that give away the location.

## **Detecting online privacy violations**

Led by Akin Günay, Özgür Kafalı<sup>1</sup>, and Nadin Kökciyan<sup>2</sup>, we designed a privacy assistant that has an understanding of the world, the concepts of the world and how they are related to each other.

The assistant first lists what information it can identify from a content, and then uses its world knowledge to infer new information. For example, would having a particular flower mean a certain city? Or, if having certain individuals in the picture mean a certain context?

In a systematic way, the assistant both accumulates and at the same time filters information to check if any of the newly derived information violates the user’s privacy preferences. If so, it warns the user to take an action.

But, if the privacy paradox is really true, then what people will say will not really match what they do. So, maybe a better way to go about it is not to ask people what they find private but to look at what they are really sharing (and if possible not sharing) and then learn what is private for the user.

---

<sup>1</sup> Kafalı, Ö., Günay, A., & Yolum, P. (2014). Detecting and predicting privacy violations in online social networks. *Distributed and Parallel Databases*, 32(1), 161-190.

<sup>2</sup> Kökciyan, N., & Yolum, P. (2016). P ri g uard: A semantic approach to detect privacy violations in online social networks. *IEEE Transactions on Knowledge and Data Engineering*, 28(10), 2724-2737.

## “I don’t know”

Led by Gönül Aycı<sup>3</sup>, we developed this idea into another privacy assistant. Interestingly, automatically learning what people find private by looking at their sharing behavior is difficult. For the privacy assistant to actually learn this, the user needs to share a lot of images first, which in principle is a bit unrealistic.

But, even then, there were many times the assistant would make mistakes. This brought us to the idea, can we make the assistant itself understand that it is likely to make a mistake in certain situations and say “I don’t know”?

What are some such situations this would come about? Say, the image has parts that would be understood as public (maybe, it is taken outside) as well as a private (maybe, it shows a baby). It is not easy to say which aspect should be taken into account to make a decision. Or, it could be that the contents of the picture are not like anything the assistant has seen before, say the first time it sees a picture of a party. Then, there is not much to base a decision on.

We capture these under an umbrella term of uncertainty. Using a technique that merges deep learning and subjective logic, the assistant assesses its uncertainty. We see that when the assistant does not answer such uncertain cases and hands the decision back to its user, it makes far fewer mistakes. Then, these uncertain cases are handled by the user.

It’s a great example of how humans and AI can work together, each bringing their strengths to the table, instead of having humans superficially “check” the AI’s decision.

## “Here is why”

Of course, an immediate follow-up question is what if the user wants to know why the assistant says a picture is private or public. For a user, it might not be important to know how the decision was made by the AI system itself. The details might even be too much to grasp. What is important is to make the decision understandable for an end user. And, what is generally most understandable for end users is the topics of a picture. Telling the user which topics were considered, how these topics relate to each other exposes a thought process.

Our privacy assistant<sup>4</sup> associates topics with pictures and evaluates the impact of each topic on the understandability. Interesting idea is that not all relevant topics are always useful for the explanation. Sometimes identifying a single influential topic, sometimes multiple topics with smaller contributions, or sometimes topics with opposing classifications are useful. We have further done experiments with real humans---as they are called these days in the AI community---and saw that the real humans also find these descriptions useful.

---

<sup>3</sup> Aycı, G., Sensoy, M., Özgür, A., & Yolum, P. (2023). Uncertainty-aware personal assistant for making personalized privacy decisions. *ACM Transactions on Internet Technology*, 23(1), 1-24.

<sup>4</sup> Aycı, G., Özgür, A., Şensoy, M., & Yolum, P. (2023). Can We Explain Privacy?. *IEEE Internet Computing*, 27(4), 75-80.

## “I got this”

But, privacy isn't just about labelling things as private or public or keeping things secret. It's about sharing the right information in the right context. For instance, when you're at the doctor's office, it's okay for the doctor to ask about your health, but probably not your salary. Both are private, but only one is appropriate in that situation.

Led by Nadin Kökciyan<sup>5</sup>, we created a privacy assistant that looks at the information in the environment, which could be anything from rules to knowledge to natural language text, and figures out which context the user is in.

Each piece of information could yield a different context. In some ways, this is like having different contexts fighting to win. When the assistant decides on one winning context, it looks at the norms of that context and takes an action that is appropriate.

This essentially requires the assistant to say “I got this” and have the autonomy to act as would be appropriate, even when you haven't specified what to do or even when this does not agree with what you might have instructed it in the first place.

## “Fair enough”

But what if the content to share or not to share does not belong to a single individual? Consider a group picture you have taken at a party.

You might have taken the picture and want to desperately share it on Instagram but maybe others in the picture would not like it to be there. Even though you took the picture, others “co-own” the content as much as you do and should ideally have a say in this. This makes the task of a privacy assistant much harder.

The privacy assistant cannot only consider your context, your preferences, and so on as before but also needs to agree with other people's privacy assistants to come to a conclusion.

We have worked on different mechanisms to enable this “agreement” among privacy assistants. Our main idea was that if different co-owners of a content want different outcomes---some want it to be shared, some not---then, the privacy assistant should be able to understand why the others want something different, be able to say “**Fair enough**” and go with their decision.

We approached this in various ways:

In one line of work, using logic-based methods, we enabled privacy assistants to carry out a dialogue to check if they could convince one another on share or not to share. What we tried to mimic with AI-based privacy assistants was a round-table discussion<sup>6</sup> on privacy.

In a second line of work, we thought what if the privacy assistants could come to a partial solution, say to share the content with common friends only? Maybe, explicitly blocking the location or time? Led

---

<sup>5</sup> Kökciyan, N., & Yolum, P. (2020). Turp: Managing trust for regulating privacy in internet of things. *IEEE Internet Computing*, 24(6), 9-16.

<sup>6</sup> Kökciyan, N., Yaglikci, N., & Yolum, P. (2017). An argumentation approach for resolving privacy disputes in online social networks. *ACM Transactions on Internet Technology (TOIT)*, 17(3), 1-22.

by Dilara Keküllüoğlu<sup>7</sup>, we had our privacy assistants negotiate different aspects of a content for privacy purposes and come to an agreement. A key idea is to imitate social responsibility such that the personal assistants sometimes give up their own benefits to respect each other's privacy. Essentially, privacy assistants take turns to help each other preserve privacy. Our experimental results show that when social responsibility is in place, the society as a whole can preserve privacy better than one where a user tries to do it on their own.

## **“With all due respect”**

One thing that is well-known from privacy studies is that some people are super knowledgeable in privacy, some are not. Some are motivated to learn about and change their privacy understanding. Some are not. Say we have a user who is neither knowledgeable nor motivated. Do we want to build a privacy assistant who is also neither knowledgeable nor motivated, because in the end the assistant represents the user? Or, do we want the assistant, contrary to its user, learn about privacy and act with knowledge and motivation?

In two separate lines, we looked at whether we can build such an assistant that can learn from other people's privacy assistants rather than its user, especially because their user has no clue on privacy to begin with. Led by Can Kurtan<sup>8</sup>, we enabled privacy assistants to imitate their friends in a social network and share only content that are similar to what others are sharing.

Led by Onuralp Ulusoy<sup>9</sup>, we extended this to learn from other assistants using reinforcement learning. The assistant can check if its user's choices are really the best and if other users are acting similarly. When the user's choices are not aligned with the society, the privacy assistant can say “With all due respect” and override the user to preserve their privacy.

Overall, with my research, I have strived to develop social AI techniques that can say these five phrases in the right places.

These have played a key role in developing privacy assistants that treat users fairly, can account for their lack of expertise, change their behavior for different contexts, and explain their reasoning to their users. And, thus help preserve their privacy.

## **Follow-up questions**

Where does these lead us? With my newly formed group on Responsible AI, I want to work on these in the next few years.

- How should we design AI systems to be trustworthy-by-design? If we want to have AI systems that say these five phrases, we don't only need to design AI systems differently, we also need to rethink how they will interact and work with their users. Our initial work on this particular

---

<sup>7</sup> Kekulluoglu, D., Kokciyan, N., & Yolum, P. (2018). Preserving privacy as social responsibility in online social networks. *ACM Transactions on Internet Technology (TOIT)*, 18(4), 1-22.

<sup>8</sup> Kurtan, A. C., & Yolum, P. (2021). Assisting humans in privacy management: an agent-based approach. *Autonomous Agents and Multi-Agent Systems*, 35(1), 7.

<sup>9</sup> Ulusoy, O., & Yolum, P. (2021). Panola: A personal assistant for supporting users in preserving privacy. *ACM Transactions on Internet Technology (TOIT)*, 22(1), 1-32.

question led by Davide Dell'Anna<sup>10</sup> shows that many aspects of human teams, such as interdependence or norm-awareness are actually also desired for human-AI teams.

- Can an AI system itself assess the trust humans put in itself and understand that a human doesn't trust it enough and maybe even change its behavior to create trust? We have started working on this problem, led by Emre Erdogan<sup>11</sup>, by providing computational theory of mind as a method to model others through everyday abstractions.
- How do interactions between users affect the trustworthiness of systems? Do people's perceptions change when in a group? Or, do people adjust their expectations based on others? Led by Huseyin Aydin and Onuralp Ulusoy, we are now looking at this from the perspective of privacy.
- How does having trustworthy AI systems affect concepts such as explainability or transparency? Would we still need AI systems to be explainable if we trust them to begin with? Could trust be a prerequisite for explanations to be accepted? In three different lines, led by Daan Di Scala, David van Kuijk, and Anastasia Apeiron, we are looking into assessing trustworthiness and realizing responsible autonomy.

## Thanks!

During my career, I have been inspired from and closely worked with many colleagues. Some are in the audience today. Many of my thoughts have been shaped by your work, input, and feedback. Thank you.

I particularly want to thank the members of the Hybrid Intelligence project, who made me part of their community, soon after I moved to the Netherlands and gave me ample opportunities to collaborate. I am very much looking forward to many more collaborations in the future.

I have placed a great importance in supervising and educating students. For the past four years, I have been the program coordinator for the AI Master's program and since the beginning of this year, serving as its leader. I am indebted to Mehdi Dastani for involving me in the AI Master's program here at Utrecht. Working close by him and the members of the Program Council of the AI Master's program over the years, namely with Rosalie Iemhoff, Rick Nowen, Tejaswini Deoskar, Lasha Abzianidze, Chris Janssen, Martijn Mulder, Floris Bex, Natasha Alechina, Brian Logan, Shihan Wang have been educational in itself.

I think having trustworthy AI systems is also a matter of education. We need to have an AI education that is diverse, inclusive, and socially impactful so that we can educate the next generation of AI researchers who can deliver AI systems that are on par with our values. In the coming years, I hope to work toward this goal, by increasing the diversity of the student body and extending our links to the societal partners, NGOs, with the hope to make a direct social impact.

I am very proud of the Global Justice Investigations Lab that we have formed with Brianne Mc Gonigle, Jessica Dorsey, Martijn Oosterbaan, and Imar de Vries, with support from Utrecht University Education

---

<sup>10</sup> Dell'Anna, D., Murukannaiah, P. K., Dudzik, B., Grossi, D., Jonker, C. M., Oertel, C., & Yolum, P. (2024). Toward a quality model for hybrid intelligence teams. In *23rd International Conference on Autonomous Agents and Multiagent Systems, AAMAS 2024* (pp. 434-443).

<sup>11</sup> Erdogan, E., Dignum, F., Verbrugge, R., & Yolum, P. (2025). TOMA: Computational Theory of Mind with Abstractions for Hybrid Intelligence. *Journal of Artificial Intelligence Research*, 82, 285-311.

Fund, where we have set up a new course to digitally study human rights and environmental abuses. My aim in the next few years is to create a similar, interdisciplinary course on Trustworthy AI.

## **“My students are my estate”**

With inspiration from Emily Dickinson, I will say “My students are my estate”.

Over the years, I have worked with many, incredibly talented students. Some of them, you have met through this talk. We managed to become a team, working together, worrying together, celebrating together, moving further together. But, mostly celebrating together. Many of my students pursued academic careers themselves. I am proud of every one of you. If I am standing here today, it is because of your brilliant work. Thank you.

Academic mentorship is akin to a family tradition. The way we are mentored shapes how we, in turn, mentor our students. My PhD supervisor, Munindar Singh, (who is here with us) worked with me relentlessly, explaining concepts, discussing results, correcting my writing. Here is a notes page from one of our discussions from 2001. I am sure some of my students will notice the similarity with the pages they leave from their meetings with me. Over the years, I and my students worked with my academic siblings from North Carolina State, visited each other, went to conferences together, becoming an academic family. If this is not privilege, I don't know what is.

Throughout this past 30 years, I have served on many program committees, organizing bodies, editorial boards, funding panels, departmental committees, university committees. Wherever I went, men; especially men from developed countries have been the majority. I am no longer afraid to speak up or ask questions. I no longer fear boys know the answer. But, I also know that this is a system that needs to change.

In hopes to do this, for the past 12 years, with my husband Ilker, we are running a platform called BolBilim to support young Turkish academics with their career. We wrote numerous blog entries and shot many YouTube videos. Through this platform, we came into contact with many new academics at early stages of their career. I have tried to pay special attention to women and prioritized coming together with them and mentoring them. I am very proud to see many of them pursuing academic careers throughout the world.

So this is to all the junior minority academics out there, be it gender-wise, ethnic-wise, or any other reason, please hang in there. You got this!

## **Bogazici University**

Before moving to the Netherlands, I worked roughly 15 years as assistant, associate, and then full professor at Bogazici University. In many ways, working at Bogazici has been life-defining. Shaping not only my identity as a researcher and educator but also instilling in me a deep awareness of societal, political, and civic matters. Today, I am honored to see colleagues from Bogazici in the audience. Your commitment to academic freedom, particularly over the past three years, serves as an inspiring example for many of us.

## **Utrecht...**

We, as a family, started over in Netherlands in 2018. Adjusting to a new country is difficult. Too western for Turkey, too eastern for the Netherlands, I spent a good deal of my life trying to understand: people, situations, conflicts.

Berna, Ceren, Claudia, Milly, Karazutra: I am incredibly grateful to have you. You make me tick.

Turkish Diaspora in the Netherlands: You have been an invaluable support. Thank you for each one of you for understanding what it means to be physically in one country and mentally in a different one.

Ne geçmiş tükendi, ne yarınlar  
Hayat yeniler bizleri  
Geçse de yolumuz bozkırlardan  
Denizlere çıkar sokaklar

My PhD thesis was dedicated to you Ilker, with the famous Metallica lyrics. Trust I seek and I find in you. I am still on the same page.

Güneşçim, Poyrazcim: You are the only reason life still makes sense. Thank you for making me smile every day.

Een Afrikaans spreekwoord luidt: 'Er is een dorp nodig om een kind groot te brengen.' Ik denk dat hetzelfde geldt voor een immigrant. Mijn goede vrienden, mijn Nederlands docent, mijn burens, leden van mijn boekenclub, collega's die in de loop van de tijd vrienden zijn geworden: Bedankt dat jullie zo goed voor me hebben gezorgd.

With that, my final thanks is to the Netherlands as a country for giving me a second chance to continue my journey as an AI academic.

Ik heb gezegd.