



# Lessons from AI: Communication

White Paper 2

# Lessons from AI: Communication

## Auteurs

MSc. Linshan Feng

Dr.ir. Julia Cramer

Dr. Ulrich Mans



# Inhoudsopgave

## Lessons from AI: Communication

<b>1</b>	<b>Introduction</b>	<b>4</b>
1.1	About the series	4
1.2	About Quantum Delta NL	4
1.3	A short story of AI communication	5
<b>2</b>	<b>Lessons from AI: Communication</b>	<b>6</b>
<b>3</b>	<b>Conclusions</b>	<b>11</b>
<b>4</b>	<b>Endnotes</b>	<b>13</b>



# 1 Introduction

## 1.1 About the series

Technology is not developed in a vacuum. As such, its Ethical, Legal, Societal Aspects (ELSA) must be carefully considered. ELSA posits that technological innovations are not independent of our current society and must be developed to cohesively integrate and enhance ethical, legal and societal values we hold to be paramount.

Quantum technology is no different. Even though quantum technologies still are at their nascent stages, research advances are moving fast and the field is gradually moving from science to application. It is the right time to look at how best to consider the ethical, legal and societal aspects of quantum technologies. In recent publications, various groups have argued that quantum stakeholders should ensure that mistakes made in the field of AI should not be repeated, and that there is a need to work out guidelines ahead of fully functional quantum systems.<sup>1</sup>

Quantum Delta NL's white paper series attempts to do just this: take a step back and search for lessons from the development and regulation of Artificial Intelligence. This is part of Quantum Delta NL's mission to study and facilitate societal impact of quantum technology. Its Centre for Quantum & Society – the first of its kind – is the place where this work comes together. This series consists of three white papers about Lessons from AI. One paper focuses on stakeholder engagement across policy, science,

industry and civil society in the European context. Who was involved, when, and with what result? The two others zoom in on risk management and communication channels. Each white paper draws from in-depth interviews with experts from a varied number of organisations, and presents a select number of “inspirations”, which we think are worth taking on board on our journey to make quantum technology a positive force for science, business and society.

## 1.2 About Quantum Delta NL

Quantum Delta NL is the keystone of the Netherlands' national ecosystem for excellence in quantum innovation; the foundation that connects the most important knowledge institutions in the field of quantum technology in the Netherlands. Our starting position is excellent: Dutch universities and knowledge institutes are leaders in the field of quantum technology research, our startup and industrial ecosystem is growing rapidly, and our national policy is strongly developed. With the allocation of 615 million euros from the National Growth Fund in 2021, we have the mandate to execute the Netherlands' National Agenda Quantum Technology (NAQT) over the next seven years. Our mission is to further build the quantum economy, and to become the most relevant ecosystem for Europe. Quantum Delta NL consists of five major quantum hubs and several universities and research centres. The hubs collaborate on innovation by bringing together top-



quality scientists, engineers, students and entrepreneurs, working together on the frontier of quantum technology.

### 1.3 A short story of AI communication

The formal research field of Artificial Intelligence began in 1956, but its arguably first widespread application only came about 30 years later, with the application of Machine Learning for software services and mobile devices<sup>1,2</sup>. Since then, AI applications have become more commonplace in daily life. With this increase in applications we experience increased coverage and public interest<sup>3</sup>.

More recently the press has been involved in writing about AI Technologies, with media coverage generally increasing over time<sup>1</sup>. Since 2009, media attention has become markedly more frequent, which coincided with front-page stories of an AAAI Presidential Panel on long-term AI Futures and advances in AI research that applied deep learning to natural language and perception tasks. Over time and with increasing applications of AI Technologies, topics of focus have shifted from concerns about a lack of progress in AI research to concerns about potentially dangerous AI innovations being on the horizon.

Yet, a report from 2021 found that the communication efforts of AI have not been sufficient and there are growing concerns that public trust in AI and understanding of AI technologies have not been established and secured

sufficiently<sup>3</sup>. These worries are underlined by findings of the Royal Society in the UK, which show that only 9% of the survey participants had heard of Machine Learning<sup>4</sup>. Other studies highlight that current narratives surrounding AI are often problematic<sup>5</sup>. For instance, that AI is often illustrated through an extremely utopian or dystopian lens. Another report also finds that there is too much public focus on the positive and negative impacts of future breakthroughs, while there is too little awareness about the benefits of existing AI applications<sup>3</sup>.

A misinformed public can have many consequences for the development and implementation of AI technologies. For instance, concerns over hyped-up issues such as humanoid robots could lead to currently important topics surrounding data use and equal access to AI technologies receiving less attention than they deserve<sup>4</sup>. On the flip side, over-hyping a technology could lead to disappointment from investors, hurting the development of said technology in the long term. This has in fact happened twice during the development of AI<sup>2</sup>.

With the acknowledgement of shortcomings in public understanding, effective communication of AI technologies has received much more recognition in recent years. For instance, MIT Technology Review started the podcast “In Machines we Trust” as recent as 2020. The podcast assesses the impact that existing and potential AI Technologies have or could have on people’s daily lives. Scientists are also starting to reach out to the public directly, for instance through staged public debates<sup>3</sup> or giving TED talks. In the last 5



years, there have been several projects in for example The Netherlands that aim to educate the general public about AI Technologies from technical, ethical and social perspectives, such as “Inzicht in AI” or The National AI Course.

About 20 years into the transition of AI technologies from research to application, effective communication to and with the public is finally receiving more attention and efforts are being made to inform the public in an accurate, engaging manner that encourages participation of the public in the implementation process of the technology. Overall, however, experts agree that AI has been communicated poorly to the public, especially in the early days of the technology. As we look ahead and seek inspiration for other technology fields, we can both look at “what worked well” and “what should have been done”.

For this white paper, we interviewed 15 professionals from the broad field of AI communication. The field of communication is vastly diverse and we aimed to capture a range of experiences from people whose work is related to AI communication. Among our interviewees were researchers who were thinking of starting with science communication and ones who have been doing outreach for a long time. We also talked to journalists who report on advances in technology and their impacts on society. Others voiced their experiences with outreach coming from a university, non-profit or industry perspective. Thus, the range of perspectives covers many of the common actors in the communication of AI<sup>3</sup>. Most of the interviewees are based in the Netherlands and other European

countries, two interviewees are based in the US.

To schedule interviews with relevant experts, we emailed people with relevant backgrounds or asked for references. The semi-structured interviews were conducted online via Zoom and took about 30 minutes each. Each interview was transcribed and anonymized and then analysed in ATLAS.ti. After the first and the second interview were coded, the codes of the second interview were compared to the codes of the first interview to see if any new codes could apply or whether some codes could be merged. This was repeated for every subsequent document with all previous documents. Finally, the codes were grouped into larger themes where applicable. The final lessons were formulated based on those findings and a round of feedback with those involved in Quantum Delta NL’s communication and outreach activities.

We thank all the participants for their time and collaboration.

## 2 Lessons from AI: Communication

### #1 Show the people behind the technology

Digital innovation is driven forward by many different stakeholders. The direction that the technology takes is not necessarily dictated by the technology itself. It is merely a tool, and it is people who are developing and using it. This distinction should be clear to people in order to avoid technological determinism<sup>3</sup>. The people working on the technology can change how it will



develop. As a consequence, when communicating AI, it is key to give room to public scrutiny towards developers and help keep them accountable: engineers have a responsibility to create applications that are fair, inclusive and safe.

This also holds true for the use of AI. Headlines in mainstream media tend to portray AI – as a technology product – as responsible for the problem it caused. This framing can be misleading and generates a wrong perception about the ‘power’ of AI. Instead, it is important to communicate about the influence any given AI user has: even the best-intended technology can be misused by people with bad intentions.

*Inspiration #1: Use storytelling to convey how people are in the driving seat when developing and using technology*

## **#2 Invest in image libraries**

We need to be aware of how people ‘view’ new technology. Looking back, the words and images used for communicating AI are not always carefully chosen. This can create problems when they have problematic connotations. Think about the word “intelligence” for example. Even though professionals in the field of AI know how different AI is from human intelligence, for most lay people, the word underscores the impression that AI is similar to human intelligence. Experts acknowledge that this might have contributed to the popular fear of intelligent robots taking over the world.

Another problem for technology communicators: AI is notoriously

difficult to accurately depict in a single item. This is because AI is not a single unit, but is made of a collection of computer algorithms. A google search on AI results in various hits ranging from high tech robots to blue numbers with a black background. Neither of those depictions meaningfully portray what AI is or can do<sup>4,6</sup>.

With AI much more established than 5-10 years ago, people seem to be more aware of what AI technologies really entail now<sup>3</sup>. However, the images portraying the technology are still the same and are often not helpful for conveying the written messages they accompany. There are efforts, such as Better Images of AI, that now aim to create repositories for images that better represent what AI technologies entail.

Consequently, they contribute to a bias in the public debate – away from pressing issues of ethics, fairness in AI benefits or data bias towards a gloom-and-doom scenario with killer robots. Using (more) accurate images for technology communication is especially important in today’s increasingly image dominated (social) media landscape.

*Inspiration #2: Create repositories of images or words that can help professionals in the fast-paced realm of communication to use readily available - yet adequate - visuals.*

## **#3 Keep in mind who your target audience is**

There are several ways in which the information can be tailored towards the chosen target audience. First, it is important to consider what the target



audience might want to know or should know from a communicator's perspective. For instance, talking about the risks of data bias might be more relevant in a presentation for industry representatives or federal agencies, while everyday applications and their impact might be more relevant for a lay audience.

Second, it is important to modify the language that is used for every specific audience. When talking to non-experts, for instance, it is important to avoid jargon and choose a language register that most people can relate to<sup>4</sup>.

Third, having a clear goal makes it easier to effectively reach the audience. Often, a speaker would target the “general public”, but this definition is more problematic than one might think. For example, if you are trying to reach people who are not yet interested in AI or tech, then simply publishing material on social media or the web might not be the best way to do it: social media algorithms are designed in a way that people read about the topics they already liked in the past. More creative, in-person strategies might be better to create awareness to new groups of people. This could be done through street surveys, popular events, or present your work on platforms that have a distinctively different audience and create a ‘hook’ for a specific group to engage with the topic.

*Inspiration #3: Decide whom you want to reach with your communication and tailor your information to that audience. Choose the content, language, and channels based on your audience.*

#### #4 Engage your audience

As mentioned in the history section, popular framing of AI tends to be either extremely utopian or dystopian, which is not helpful for public understanding of AI<sup>5</sup>. Yet, underlying these portrayals are constraints based on the current “attention economy”. Media outlets have to compete for attention and have to use elements that triggers people's interest. Since each person's willingness to pay attention is limited, so-called engaging narratives are often the most effective<sup>5</sup>. Although AI communication should not lean into sensationalist headlines or imagined conflicts, it should acknowledge that engaging the audience is crucial. The easiest way to do so is to relate to the topics that are close to a person's personal life. For instance, many education efforts highlight the current widespread presence and applications of AI in medical applications, which are topics that are relevant to all of us<sup>5</sup>. The information can also be made more personal and interesting if it is connected to real or fictional people. For instance, when it is a human-interest story of how AI applications influenced someone's life<sup>4</sup>.

Another way is to make the learning experience interactive and show people what AI technologies can do. For instance, one project showcasing how AI can process large data sets gives the audience a task where they need to sort an increasing amount of data. As the dataset becomes larger, it becomes increasingly difficult for humans to sort these immense files, but the AI is able to handle it. Interactivity can also be achieved in a museum, e.g. a game that asks the visitor to distinguish between AI-generated and human-generated





pictures. This encourages visitors to actively think and debate about what AI can and cannot achieve.

*Inspiration #4: Find ways to engage your audience. Draw inspiration from the wider field of science communication and science education.*

## **#5 Apply your resources smartly**

Every communication project depends on the available skills, budget, and time. These parameters should be the basis for deciding how much and what information the project can include, which medium is the most suitable, and which target audience the project can realistically reach.

For instance, if public speaking is a known strength in a given team, then choosing public events might be a good idea. Experts mentioned public speaking or poster sessions as examples of how to reach audiences with limited budgets. Another way of leveraging scarce resources is to partner with commercial and institutional partners whom you can task with outreach activities through their own, existing channels.

This effect can also be achieved in the social media domain. Instead of building your own following, which takes time and resources, it can help to make use of tech enthusiasts or other types of social media influencers who can help spread the message.

*Inspiration #5: At the beginning of communication projects, think about how you can make the most of the resources available.*

## **#6 Prototype, Iterate, and Measure Impact**

To ensure that efforts communicating AI are conveying the intended message, it is important to continuously receive and incorporate feedback. If possible, start with a quick prototype and develop further iterations based on initial feedback. This way, you gain an early understanding of what works and what does not, and whether the chosen language is appropriate. In the very early stages of a project, even feedback from relatives or friends can be very insightful (and for free).

This approach can be very effective for example when building an informative website. A first sketch delivers the look and feel, and first user feedback can help finetune how users can orient themselves through the website and to change navigation options. Important in this: don't stop after the website is 'live'. Feedback is a continuous process and AI-related content means different things to different users. While feedback from initial user groups might be positive from tech experts, university students might struggle to understand everything on the site<sup>7</sup>.

No feedback loop is final without measuring the impact. While it is tempting to focus on 'sending' it is advisable to spend resources on the actual reach. This is particularly true for long-term impact of communication projects, which can be challenging. Yet, it is the only way to ensure that the efforts are going in the right direction. When a project relies on direct interaction, it can be easier to assess the type of people reached, as well as their background. When working with schools for example,



teachers can give helpful insights into what their observed effectiveness of the programme is<sup>8</sup>.

*Inspiration #6: Create feedback loops right at the start of your project and continue to do so.*

### **#7 Combat misconceptions before they stick**

AI research and development went through two hype cycles at the end of the last century, during which media attention led to exciting promises for AI applications that were not attained. This led to investors temporarily losing faith in AI technologies and to a decrease of funding<sup>2,4</sup>. Hype can also lead to general misconceptions about technology in society, which could impact policy-level decisions<sup>3</sup>.

Although the current developments in AI are different from previous hype cycles because it has already found widespread application in society, some hyped-up voices from different actors still exist – think about the advertisements for self-driving cars. This type of ‘mis’-information can be tackled by pointing out what it is inaccurate and instead offer alternative expert insights, which highlight the current scientific consensus on the topic. Another way is to cooperate with media outlets and help them create more accurate reports<sup>4</sup>.

This is not to say that hype is by definition problematic. Hype creates attention, and as a result people become aware of the topics associated with the hype. While it should not be encouraged, it is possible to harness a hype cycle in a way that is beneficial to increase more

realistic communication alongside, for example when used as a point to connect to previous knowledge that the audience has. This can make communicating AI potentially more engaging.

*Inspiration 7# Tackle hypes head-on. Make sure you know when and where hypes emerge and prepare alternative communications when needed.*

### **#8 Be prudent about relevant technical details**

It is important to communicate relevant technical background information so that people can get a good understanding about what AI technologies really entail<sup>4</sup>. This should not be overdone, however. Conveying technical information to a lay audience is challenging and will require a lot of time and often concentrated learning. Specific technical details might be necessary but be clear when this is the case – and when it distracts from the main message. The amount of technical information included should be very important for the message you are trying to convey.

One way to convey technical information is to use descriptive language instead of technical terms and to take your time when explaining those. To explain how a Natural Language Processing works for instance, one can create a sentence word by word together with the audience and then explain that an AI would do the same, but just choose each word based on how probable it is to come next. It is intuitive to understand that for each next word, there are only certain words that would fit there.



*Inspiration #8: Limit the use of specific technical information – and take your time explaining it.*

### **#9 Start ASAP**

Even 20 years after the entrance of the first AI-related technologies, communicating it in an appropriate manner is still a luxury. If done well at the start, communication can be an effective means to inform the public and build trust in a technology before it becomes at sale. Several experts voiced their concern that communication about AI started too late. The public was largely unaware of AI technology, and there were too few efforts to start an early engagement on AI. Instead, early focus rested on technical advances rather than explaining what those milestones entailed, or whether or not the potential applications are in fact worth pursuing. Over time, this sentiment changed, and awareness about AI grew with the onset of life-changing algorithms such as navigation systems on the positive end of the spectrum, and automated systems designed to identify fraud with in-built biases against minorities at the other end. It should not come as a surprise that in the young years of a new technology, resources flow into research, testing and creating business models rather than communicating the potential effects over the years to future consumers. Yet, the last 10+ years have confirmed that investments in communication have come too late, and without the needed funds. The future consumers of the next generation (digital) technologies have the right to be informed in a timely, adequate manner as technical experts make headway.

*Inspiration #9: Build trust by starting on day 1. Even when applications are not yet clear, it is worth exploring communication means to relay the power of the technology and its impact.*

## **3 Conclusions**

Communicating AI has received much attention in recent years, with many stakeholders becoming increasingly involved in reinforcing public understanding of AI technologies. This effort resulted in a shifting public understanding of AI – from fears over gloom-and-doom scenarios to a more realistic perception about the social and economic impacts of AI applications. Still, much remains to be done and communicating AI is in constant change, with recent efforts focusing on talking with the public instead of talking to the public.

Looking at the current challenges in communicating AI, it becomes clear that efforts on future technologies such as quantum should start as soon as possible, adopt the best practices from AI, and should start with addressing potential pitfalls in public understanding today. At Quantum Delta NL we are committed to investing in best practice sharing. We believe that we should learn from the past, and take inspiration from those who have taken similar pathways before. The nine inspirations presented here have been drawn from a select number of expert interviews and are merely a beginning, a conversation starter. We hope they can serve as a starting point for a fruitful discussion within the quantum communication community.



At the same time, we are investing in concrete follow-ups. Where inspirations lead to concrete suggestions on how to do things in quantum, we will take these up. For example, we have started working on impact assessments back in 2021, along the lines of AI impact assessments. We are also working on public awareness campaigns and a quantum course, for free

and without any jargon. Both are available from 2023 onwards.

Acknowledging that there is much work to be done in communicating quantum, we welcome any support along the way from our partners in the growing quantum ecosystem.



## 4 Endnotes

- <sup>1</sup>. Fast, E., & Horvitz, E. (2017). Long-Term Trends in the Public Perception of Artificial Intelligence. *Proceedings of the Thirty-First AAAI Conference on Artificial Intelligence*, 31(1), 963–969.
- <sup>2</sup>. Peter Stone, E. (2016). Artificial Intelligence and Life in 2030. One Hundred Year Study on Artificial Intelligence: Report of the 2015-2016 Study Panel, September, 52. [https://ai100.stanford.edu/sites/default/files/ai\\_100\\_report\\_0901fnlc\\_single.pdf](https://ai100.stanford.edu/sites/default/files/ai_100_report_0901fnlc_single.pdf)
- <sup>3</sup>. Littman, M. L., Ajunwa, I., Berger, G., Boutilier, C., Currie, M., Doshi-Velez, F., Hadfield, G., Horowitz, M. C., Isbell, C., Kitano, H., Levy, K., Lyons, T., Mitchell, M., Shah, J., Sloman, S., Vallor, S., & Walsh, T. (2021). *Gathering Strength, Gathering Storms*: Stanford University, September.
- <sup>4</sup>. The Royal Society. (2020). Communicating AI: the role of researchers in enabling public conversations about AI. 2015(May), 1–5. [www.royalsociety.org/machine-learning](http://www.royalsociety.org/machine-learning)
- <sup>5</sup>. Cave, S., Craig, C., Dihal, K., Dillon, S., Montgomery, J., Singler, B., & Taylor, L. (2018). Portrayals and perceptions of AI and why they matter. The Royal Society, 1–28. <https://royalsociety.org/-/media/policy/projects/ai-narratives/AI-narratives-workshop-findings.pdf>
- <sup>6</sup>. Chubb, J., & Çetin, R. B. (2022). We need better AI imagery for better science communication. *Impact of Social Sciences Blog*. <https://blogs.lse.ac.uk/impactofsocialsciences/2022/07/25/we-need-better-ai-imagery-for-better-science-communication/>
- <sup>7</sup>. Interview with communication expert October 24th 2022
- <sup>8</sup>. Interview with communication expert November 16th 2022



Quantum Delta

Lorentzweg 1

2628 CJ Delft

E [info@quantumdelta.nl](mailto:info@quantumdelta.nl)