Lessons from AI: stakeholder engagement

White Paper 1

Quantum Delta Nederland

Inhoudsopgave

Lessons from AI: stakeholder engagement

List of Abbreviations

Abbreviation	Definition
BDVA	The Big Data Value Association
ELSA	Ethical, Legal, Societal Aspects
AI	Artificial Intelligence
AGI	Artificial General Intelligence
UN	United Nations
ITU	International TeleCommunication Union
EU	European Union
GDPR	General Data Protection Regulation
AI HLEG	High-level Expert Group on Artificial Intelligence
OECD	The Organisation for Economic Cooperation and Development
WEF	World Economic Forum

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.[1]

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 economu, 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 history of AI regulation in Europe

As early as 2010, numerous AI initiatives and actors cropped up, initially engaged in research. Soon, they also took up initiatives that explored related regulatory issues. DeepMind is a prominent example, which was established that year in the UK with the goal of developing 'general and capable problem-solving systems, known as artificial general intelligence (AGI)'[2]. Fast forward to 2014, the European Commission (EC) created the Big Data Value Association (BDVA) as a public-private partnership with a focus on Big Data. This later changed into DAIRO (Data, AI and Robotics), reflecting the move away from Big Data towards AI as the dominant framing used for policy discussions. At the same time, negotiations over data privacy gained traction over the years - and culminated in the adoption of the General Data Protection Regulation (GDPR) in 2018.

With the GDPR, the EU presented a new blueprint for regulatory activity for emerging technologies, and shifted the attention of non-governmental and industry stakeholders from stimulating the use of AI (and growing the workforce) towards regulatory issues. In 2015 OpenAI was established as a private, non-profit initiative to develop AI and to ensure that 'AGI benefited humanity'[3]. Its purpose to develop 'friendly AI' was a tacit recognition that AI needed to be regulated in some fashion to truly harness its benefits. The announcement and establishment of the Partnership on AI in 2016[4] was yet another expression from private industry that AI should be regulated. The scope and extent of such regulation, however, remained contested. Private industry's drive to establish AI standards to regulate its use culminated in the Asilomar AI Principles in 2017, which included 23 ethical guidelines for AI research and development.[5]

From this point onwards, the focus on AI grew steadily across EU institutions and the Brussels landscape of stakeholders. On 10 April 2018, 25 European countries signed a Declaration of Cooperation[6] to collaboratively work on three main pillars for AI: industry capacity, socio-economic issues, and legal and ethical concerns. With the goal of establishing a Digital Single Market, work on these three pillars commenced immediately after the announcement of the Declaration of Cooperation. Also in 2018, the European Group on Ethics in Science and New Technologies called for a common and internationally recognised ethical and legal framework on AI.[7] In June 2018, both the AI High Level Expert Group (HLEG) and EU AI Alliance were set up to pool insights from leading AI experts and prepare documents that would later provide the groundwork for new AI regulation. National governments also got involved. A notable example is the German Data Ethics Commission established in 2018 with the aim to develop ethical and legal frameworks for AI. The UN and the OECD also started to develop AI regulation, for instance with the 2017 ITU AI Summit for Good.

The EC released a White Paper on AI in 2020, formally documenting its regulatory vision and strategy for AI. Its publication included a public consultation which solicited public feedback on its regulatory intentions. A year later, in 2021, the proposed Artificial Intelligence Act was released. The proposed Act, which is expected to enter into force after 2025, regulates the use of AI across sectors. Its regulatory scope, application and other details are still being debated and amended by the European Parliament as of 202

2 Lessons from AI: stakeholder engagement

#1 Dare to lead on global inclusion: the EU can guide a global discussion about human-centric use of technology

Every stage of technological development depends on choices; and these choices are shaped by context, (access to) resources and past experience. For AI systems, this basic fact has received ample attention from scholars, programmers, and policy makers. One of the most prominent examples: the data used to train and feed some AI systems have been found to replicate societal inequalities and biases[8]. For the EU, equality is a central tenet. In a White Paper on Artificial Intelligence, the Commission recognises that 'bias when present in AI could have a much larger effect, affecting and discriminating many people without the social control mechanisms that govern human behaviour'.[9] This fundamental concern also guided the work of the AI HLEG for a human-centric approach to AI regulation.

The concern is that the human-centric approach adopted by the EU may not be accepted as a universal principle. With other technological innovations, diverging visions among important regional blocs may not pose significant problems. However, for AI systems, the stakes are high: AI systems are being developed for healthcare and defence systems, for self-driving cars or climate change technology. This means that it matters how large training datasets can be used, and whether this differs per region. If EU rules are too restrictive, companies might choose to develop software elsewhere. In other words: globally consistent rules are required for effective AI regulation.

The EU has a clear normative vision for the regulation of emerging technologies. But the EU must navigate an international space shared by authoritarian countries and important market-first countries such as the US. Provisions to ban the use of some AI systems, for instance, may cede their control and direction to authoritarian and other regimes that do not share a human-centric approach to AI. On the other hand, the EU cedes its own political authority if it allows technologies that it would otherwise not accept on the basis of contradictory regulations elsewhere.

Fragmentation and exclusion may frustrate the ability to regulate effectively. Consider the case of China. The lack of meaningful participation by China in international collaborations, such as the Global Partnership on AI, diminishes the global efficacy of such initiatives. Accordingly, it risks reducing the EU's and collaborating partners' ability to influence the development of emerging technologies globally. This is a dilemma prominent for, but certainly not limited to, AI regulation.

Inspiration #1: For regulation to be meaningful, it must be a global effort. Try to develop a consistent vision on tech regulation, by safeguarding multiple voices and giving a voice to diverging views. This can be done on an expert level, ahead of policy discussions that follow. In practice, this could take shape in the form of informal hubs around the globe for best practice sharing.

#2 Leverage the EU's regulatory track record: make sure your rules become central to international fora

The so-called 'Brussels Effect' stands for the EU's ability to exercise regulatory power globally by the development of rules and standards applicable to European markets, affecting anyone who wants to access them. Scholars posit that the Brussels Effect is 'Europe's unilateral power to regulate global markets'[10], with the GDPR being a prominent example.

There is agreement that the Brussels Effect gives the EU a unique opportunity to shape regulation of emerging technologies to comport with its governance model. In a highly competitive geopolitical race to develop world-class technology, the EU has in the past focused on developing legislation, leaving much of the (early) investments to other market players, notably the US and China.

While losing investments to other geopolitical players may harm the short term interests of the EU, the decision to develop robust legislation has enabled the block to play an influential role in the global development of AI and other emerging technologies. This is most evident in the international impact of Europe's legislative activity regarding AI. The frameworks adopted by the OECD, Council of Europe and World Economic Forum share similar focus, vision and language as those developed in EU legislation. Additionally, the choice to develop robust legislation also has a gravitational pull that offsets some of the issues around lack of investment: the Brussels Effect means that tech companies also keep up their presence in Europe.

While doing so, Brussels should be paying critical attention to whether or not the larger (global) industry players benefit most from suggested regulations. During the negotiations around GDPR and the proposed AI Act it became clear that the tech industry favoured regulation out of efficiency rather than merit: the aim was to create a de-facto world-wide *API*[11]*, so that* businesses only have to comply with one regulation rather than multiple. The downside: important local and regional differences may be the casualty of such universal regulation.

Inspiration #2: Make use of the EU's reputation as a credible first-mover on tech regulation, de-facto engaging tech innovators by saying "if you don't make the tech human-centred, it will not be allowed to reach EU customers."

#3 Invest in early bridge-building across sectors – with clear mandates

The EU AI regulatory cycle saw numerous examples of policy collaboration between civil society and businesses. Most prominently, the EU AI Alliance[12] provided a platform for public consultation connecting business and civil society, which resulted in valuable feedback on the draft of the Ethics Guidelines for Trustworthy AI.[13] The Alliance allowed the EC and the HLEG on AI to solicit public feedback on their progress; including from civil society and businesses. Another example of collaboration between the business community and civil society is the work of the HLEG on AI itself.[14] The group was tasked with the specific mandate by the Commission to provide advice on its AI strategy. The experts chosen to be on the HLEG represented different disciplines, mixed industries, civil societies, geographies and efforts to maintain a gender balance[15]. The HLEG

produced four documents providing advice on the EU ethical framework for AI, its policy and investment strategy, and a self-assessment list for incorporating its Ethical Guidelines.

Notably, the work of both the HLEG on AI and the EU AI Alliance resulted in the first European AI Assembly[16], where different stakeholders, including businesses and civil society discussed the latest achievements in AI policy developments. The HLEG on AI also organised a unique consultation, with more than 350 stakeholders who contributed to operationalising the Ethics Guidelines for Trustworthy AI. This led to the so-called Assessment List for Trustworthy AI[17] – a self assessment tool which offers guidance to implement the requirements of the Ethics Guidelines.

In addition to these efforts, the BDVA provided another platform where business and industry representatives could engage with the regulatory activity of the EU. This engagement resulted in the Strategic Research, Innovation and Deployment Agenda[18] – a joint initiative between different industry players, which included a shared vision on AI research, policy framework and market preparedness for AI.

However, there is a danger that mandates are not clear, and civil society interests are confused with business interests as a consequence. Many of the civil society driven efforts were supported by large tech firms, and could not claim a neutral role. Civil society participation is ultimately only representative of those groups involved in the process. They are not accountable, and experts warn against giving non-governmental stakeholders too great a role in creating regulatory building blocks. An over-reliance on civil society risks the over-representation of some viewpoints and risks unwanted dominance of corporate lobbying efforts.

For quantum technology regulation to truly harness the benefits of engagement between civil society and businesses, such engagement should start early, should include clearly defined mandates and roles for the participating groups, and should be as inclusive as possible.

Inspiration #3: Engage early with civil society, and ensure the mandates are well-defined. Use civil society for information collection, and-interpretation, but leave decision-making to other, officially mandated bodies.

#4 Look out for relevant national governmental initiatives

On 18 July 2018, the German government launched the national initiative 'German Data Ethics Commission', with the goal to develop 'ethical benchmarks and guidelines as well as specific recommendations for action' on Artificial Intelligence.[19] The Commission was established in the context of a flurry of activities on AI with a broad range of stakeholders, including businesses, NGOs and governments. All of them had the intention to develop a regulatory framework for AI systems.

Throughout its existence, this commission had a major influence on the activities of the EU on AI. Notably, in April 2019, the EU High-Level Expert Group on AI echoed the need for 'human-centric ' AI systems – a principle featuring in its official report. Most importantly, the commission's framework for a risk-based approach was adopted in the EU proposed regulation on Artificial Intelligence in April 2021.

The work of this particular body stands out for two reasons. First, the initiative was less encumbered by the bureaucratic style of EU institutions, so it was able to get up to speed and reach impactful conclusions relatively quickly. Second, it carried the full German political weight within the EU. And, since the work of the commission was focused on developing ethical benchmarks and guidelines, its recommendations were seen as widely applicable. The combined effect of a specific mandate and a German government-backed initiative worked in its favour and gave it an influential role on the EU AI regulatory cycle and beyond. The OECD also adopted a 'human-centric' approach in its official statement on AI, leading to a de-facto universal approach to ethical principles for AI.

The role of a prominent EU member state is an interesting case for technology regulation. For the upcoming discussions about quantum technology and its societal impact, national commissions can carry weight, leverage resources and are more flexible than EU-wide initiatives. In addition, national efforts in other nation states with a particular standing in the technology should also be considered when looking for partnerships. This can be the case because quantum technology is being developed in a particular nation, or because the technology is expected to have a significant local impact (e.g. on national security). Experts recommend looking even beyond Europe, to include national efforts in the Global South.

Inspiration #4: national initiatives can drive regulation on tech; get involved with them early. Make sure there is an timely start with investing in outreach and engagement with countries that might not be in the driver's seat but affected by the introduction of the technology.

#5 Balance between non-binding instruments and full-fledged regulation

Over the years, there were a flurry of non-binding documents in the EU and elsewhere. For the EU, the most consequential work came out of the HLEG on AI. It produced the Ethics Guidelines for Trustworthy AI – developing and enshrining concepts such as 'Trustworthy AI', the 'human-centric approach' and 'risk-based model' (see also above). These concepts were further developed and incorporated in the proposed regulation on AI. Moreover, these terms also found their way in other, international, non-binding documents. The Ethically Aligned Design by the IEEE[20], the OECD AI Principles[21] and Recommendation on the Ethics of Artificial Intelligence by UNESCO[22] are examples of world-leading international bodies that utilised the ethical concepts developed by the HLEG on AI such as 'human-centred approach' and 'Trustworthy AI'.

The prevalence of such ethical concepts underscore their usefulness. First, they provide a straightforward way to encapsulate complex concepts. The idea of a 'human-centric' approach to AI is attractive and is able to provide an ethical vision that is easily tangible and vivid. Second, non-binding documents allow for consensus building. Getting everyone to agree on regulation is an arduous task, especially if the regulation tackles a complex technology such as AI. The process of developing non-binding documents can make the regulatory process manageable and less intimidating. The relatively quick turnaround of the HLEG on AI attest to this. The HLEG was established by the Commission in June 2018 and on 8 April 2019, it published its Ethics Guidelines for Trustworthy AI. Experts agree that the

relatively quick process by which such an ambitious project was published, is mainly due to the non-binding character of the document.

The rise of non-binding instruments for regulatory purposes, however, has not been uncontroversial. Consensus-building exercises around non-binding instruments tend to invite the dilemmas encountered concerning the selective nature of civil society participation and industry lobbying. Non-binding instruments carry the risk of ceding political authority to unaccountable stakeholders. Accordingly, reliance on non-binding instruments should not stand in the way of other important regulatory practices. Moreover, it is an often-heard comment even from those involved that the non-binding exercises produced too many best practices – which in turn meant that similar work was done in parallel through different channels.[23] Observers note that it was often unclear if certain voluntary principles could, should or would develop into regulatory language.

Inspiration #5: be mindful of non-binding agreements becoming the main focus. They should be used instrumentally, not as ends in themselves, and there should be a clear pathway towards regulations.

#6 Allow for experimentation when designing a regulatory process

Regulators cannot afford to respond constantly to a fast-changing technology landscape. Regulatory processes are by definition slow-moving and should not be modified with each latest innovation. When AI and related machine learning applications gained traction throughout various industries, it was difficult to respond adequately to the evolving technology, adapt the regulatory language and to anticipate the next wave of innovations and their impact on potential regulation.

The concept of 'regulatory dynamism' allows for such a balance. The risk-based model adopted by the proposed AI Act is a good example of how this can be applied. It is able to scale based on the threat AI technologies pose to fundamental rights. This offers an element of dynamism that allows for the technology to be developed without its innovation cycles being stifled by an overly stringent regulatory environment.

Taken to the next level, regulatory processes could try to adopt 'agile' principles, as recommended by the UK strategy on AI.[24] The rationale for agile regulation stems from the need to create the most conducive space for technology and innovation, while safeguarding that its application meets high legal and ethical standards.

That being said, there are also potential drawbacks to dynamic and agile regulatory practices, especially diminished certainty in advance. Clear and consistent rules allow actors to plan ahead. In the context of the Cambridge Analytica scandal[25] and the proliferation of disinformation[26], the benefit of regulation as a safeguard against abuse should not be underplayed. Too much flexibility carries the risk of sacrificing the primary benefits of regulation.

Inspiration #6: Invest in flexibility. Explore if it is possible to apply agile regulatory principles, which allows for trial and error. For example, define a phase of 2 years for deliberation and experimentation, before discussions about binding regulation starts.

#7 Explore the benefits of sector-specific legislation

Many experts involved in the regulatory process for AI in Europe point out that the legislative process lacked an effective feedback loop. Policy advisors, consultants and representatives from industry were often left in the dark with respect to the ways their feedback impacted the trajectory and outcome of the regulatory process. Although they felt assured that their concerns and general feedback were heard through the appropriate means, they couldn't assess the degree to which their feedback was taken into account.

Practically, it is an arduous task to include and take into account every feedback offered and this is perhaps not an achievable goal in complex deliberations. However, the perception that the EU decision-making is opaque or not responsive lowers confidence among stakeholders and could potentially damage the entire regulatory process by stifling and discouraging relevant feedback.

One suggested way to avoid this criticism and perception was to adopt more tailored, sectoral legislation. For example, cases from the healthcare industry suggest that the efforts to draft and adopt horizontal, cross-sectoral, legislation caused unclarity, which in turn resulted in uncertainties. For this sector, the proposed AI Act binds them to parallel legal obligations – a situation where similar rules exist in different legislations. Moreover, the proposed legislation is cumbersome and contains contradictory provisions, especially in the context of compliance under the GDPR.

The idea of sectoral legislation sounds attractive and has been implemented for data protection and privacy regulation[27].

AI is currently deployed differently by each business, sectors and industries, as a one-size-fits-all legislation is seen as counterproductive[28]. The downside, however, is a complex landscape of proliferating, narrowly-tailored regulatory regimes. For future regulatory discussions in quantum technology, the choice and strategy for legislative efforts should keep these advantages and drawbacks in mind.

Inspiration #7: work with specific sectors, where there is a wish to contribute to the deliberations, explore what could work in each vertical and then start horizontal discussion with wider policy fora.

#8 Invest in forecasting

AI technology is immensely complex and the regulation of AI and emerging technology reflects this complexity. 'In the same way that an architect from the outset of designing a house has to think of compliance with the building code, AI software engineers will have to think from the outset about how their future software tool could affect democracy, fundamental rights and the rule of law[29]'. In other words, those developing AI must be aware of the regulatory environment which is applicable to their products. But complexity increases over time. When it comes to future use of certain technologies, it is difficult to have a meaningful conversation about regulatory principles, and the impact on economy, politics and society as a whole.

The good news is that a deeper understanding of the technology and its possible market and societal implications can help. The question is whether and if so how extrapolations can be used for regulatory decisions.

The concept of technology forecasting offers a tool that can be used throughout the regulatory cycle but – if used – should feature prominently at the initial stages of the regulatory process. Some experts involved in the AI regulatory cycle suggested that technology forecasting is the weakest component of the regulatory process. At present, forecasts for AI use cases are vague and often aspirational without providing enough concrete information for the regulatory process.

Looking ahead, regulatory practices may benefit from a more prominent role for technology forecasting, be it prescriptive or as a guidance for deliberations. Robust technology forecasting requires resources and relevant expertise. This is not an easy feat but might have been a worthwhile investment in AI – and a missed opportunity until today.

Inspiration #8: explore how to integrate technology forecasting into the regulatory process – in close collaboration with academic institutions – as a means to drive the conversation meaningfully into the future (application) of the technology.

#9 Create easy-to-use tools for those working on and with the technology

In the case of AI regulation, many argued that the use of impact assessments was an important factor. These assessments, whether obligatory or not, can increase awareness and accountability. They offer a manageable way for businesses and other relevant stakeholders to test, apply and possibly adapt regulatory principles – even at a time when regulation is not yet in place. Later on, when used in combination with regulation, impact assessments help foster a wide ownership of emerging regulation by allowing relevant stakeholders to be actively involved in the regulatory process. This enhances regulatory compliance and also ensures that there is a steady stream of relevant feedback from stakeholders who are bound by the regulation.

Inspiration #9: use impact assessments. This is part of the toolbox to help those who are working on and with the technology – and those who are responsible for regulating it down the road.

#10 Embrace public participation

There is no doubt that the regulatory process benefits from active public participation. However, the manner in which such participation should occur is contested. The EU AI Alliance is notable for the extent to which public consultation was encouraged in the regulatory process for AI. The AI Alliance has 4000+ members. For the complexity of AI regulation and the high level calibre of discussion on the Alliance website, as attested by experts, this level of participation from the broader public is encouraging. This is easier said than done, given the complexity of technologies like AI (and quantum technologies). The technology does not lend itself to quick explanations, and target groups differ. Experts recommend avoiding the use of jargon, and it might help to seek multidisciplinary backgrounds for those experts that are involved.

Inspiration #10: Organise larger audiences. When you do, keep it simple when possible, use jargon when useful. At any given time, be mindful of your target group.

3 Way Forward

Looking ahead, the emerging quantum landscape will have to invest heavily in a meaningful pathway towards principles, guidelines and, ultimately, regulations. This is a rocky road, and we are fully aware of the fact the bar is **very** high.

The AI regulatory process throughout the EU has been a difficult one – as some of the examples discussed in this white paper have shown. It will not be any easier for the complexities that come with quantum technology – in any of its application areas such as computing, communications or sensing.

At Quantum Delta NL we are committed to invest 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 inspirations preseted 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 with our partners in the EU and elsewhere.

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. Available from 2022 onwards.

Acknowledging that it is certainly is a rocky road ahead, and we welcome any support along the way from our partners in the growing quantum economy.

4 Endnotes

1. See, 'Why Quantum Computing Is Even More Dangerous Than Artificial Intelligence'. 21 August 2022. <https://foreignpolicy.com/2022/08/21/quantum-computing-artificial-intelligence-ai-technology-regulation/>
2. See, DeepMind 'About' Web Page. <https://www.deepmind.com/about>
3. See, OpenAI Website. <https://openai.com/>
4. Partnership on AI is a coalition consisting of academics, industry, civil society and media organisations working to address ethical approaches to AI.
5. See, Asilomar AI Principles. <https://futureoflife.org/2017/08/11/ai-principles/>
6. See, Declaration of Cooperation on AI. 10 April 2018. <https://ec.europa.eu/jrc/communities/sites/jrccties/files/2018aideclarationatdigitaldaydocxpdf.pdf>
7. See, Statement on artificial intelligence, robotics and 'autonomous' systems. 9 March 2018. <https://op.europa.eu/en/publication-detail/-/publication/dfebe62e-4ce9-11e8-be1d-01aa75ed71a1/language-en/format-PDF/source-78120382>
8. See, 'Understanding algorithmic bias and how to build trust in AI.' PWC. 2021. <https://www.pwc.com/us/en/tech-effect/ai-analytics/algorithmic-bias-and-trust-in-ai.html>
9. See, WHITE PAPER On Artificial Intelligence - A European approach to excellence and trust, COM(2020) 65 final. <https://ec.europa.eu/info/sites/default/files/commission-white-paper-artificial-intelligence-feb2020_en.pdf>
10. See, Anu Bradford, The Brussels Effect, 107 NW. U. L. REV. 1 (2012). Available at: https://scholarship.law.columbia.edu/faculty_scholarship/271
11. Interview with an AI scholar.
12. See, EU AI Alliance, <https://futurium.ec.europa.eu/en/european-ai-alliance/pages/about>
13. See, Ethics guidelines for trustworthy AI. 2019. <https://digital-strategy.ec.europa.eu/en/library/ethics-guidelines-trustworthy-ai>
14. See, High-level expert group on artificial intelligence. <https://digital-strategy.ec.europa.eu/en/policies/expert-group-ai>
15. Interview with EU tech policy advisor and expert.
16. See, The first European AI Alliance Assembly. <https://digital-strategy.ec.europa.eu/en/events/first-european-ai-alliance-assembly>
17. See, Pilot the Assessment List of the Ethics Guidelines for Trustworthy AI. <https://ec.europa.eu/futurium/en/ethics-guidelines-trustworthy-ai/pilot-assessment-list-ethics-guidelines-trustworthy-ai.html>
18. See, 'Strategic Research, Innovation and Deployment Agenda for the AI, Data and Robotics Partnership' 2020. <https://www.bdva.eu/sites/default/files/AI-Data-Robotics-Partnership-SRIDA%20V3.0_0.pdf>

19. See, ‘Opinion of Data Ethics Commission.’ 2019. Page 5. <https://www.bmjv.de/SharedDocs/Downloads/DE/Themen/Fokusthemen/Gutachten_DEK_EN.pdf?__blob=publicationFile&v=2>
20. See, The IEEE Global Initiative on Ethics of Autonomous and Intelligent Systems. Ethically Aligned Design: A Vision for Prioritizing Human Well-being with Autonomous and Intelligent Systems, First Edition. IEEE, 2019. <https://standards.ieee.org/content/ieee-standards/en/industry-connections/ec/ autonomous-systems.html>
21. See, The OECD AI Principles. 2019. <https://oecd.ai/en/ai-principles>
22. See, Recommendation on the ethics of artificial intelligence by UNESCO. 2021. <https://en.unesco.org/artificial-intelligence/ethics>
23. Interview with an academic and AI expert.
24. See, National AI Strategy. 2021. Page 56. <https://assets.publishing.service.gov.uk/government/uploads/system/uploads/attachment_data/file/1020402/National_AI_Strategy_-_PDF_version.pdf >
25. See, Facebook, Cambridge Analytica and data mining: What you need to know. BY: Ian Sherr. 18 April 2018. <https://www.cnet.com/news/politics/facebook-cambridge-analytica-data-mining-and-trump-what-you-need-to-know/>
26. See, Gelfert, Axel. Fake news: A definition. (2018) <https://www.erudit.org/en/journals/informallogic/2018-v38-n1-informallogic04379/1057034ar.pdf>
27. For example: Proposal for an ePrivacy Regulation
28. Interview with experts in industry.
29. Constitutional democracy and technology in the age of artificial intelligence. BY: Paul Nemitz. 2018. <https://royalsocietypublishing.org/doi/10.1098/rsta.2018.0089>
30. Interview with an AI expert.

Quantum Delta NL

Lorentzweg 1

2628 CJ Delft

E info@quantumdelta.nl

Quantum Delta
Nederland