



AI and Ethics in the Health Innovation Community

EIT Health pilot of the AI & ethics guidelines of the
European Commission's (EC) High-Level Group



Contents

Foreword	3
Executive summary	4
Setting the scene	5
EIT Health	5
AI in EIT Health	6
The European Commission's AI and ethics guidelines.....	6
Process	7
Survey	
Case studies	
World Health Summit	
Gastein	
Findings	8
Survey results	8
Human agency and oversight.....	10
Technical robustness and safety	12
Privacy and data governance	13
Transparency.....	14
Diversity, non-discrimination and fairness	16
Societal and environmental well-being	17
Accountability	18
Case studies	20
Analysis and interpretation	22
General	
1. Human agency and oversight	
2. Technical robustness and safety	
3. Privacy and data governance	
4. Transparency	
5. Diversity, non-discrimination and fairness	
6. Societal and environmental well-being	
7. Accountability	
Suggestions for next steps	24
For the European Commission	
For EIT Health	
Framework for trusty AI	25
References	26

Foreword

There is probably no technology that has received more attention in the last two years than artificial intelligence (AI).

AI holds great promise for health and ageing innovation and, moreover, is already delivering impressive results. AI is considered to be equivalent to world-leading doctors in determining the correct treatment for over 50 diseases of the eye¹. 80% of health executives believe that AI will work alongside humans in their organisation within the next two years².

AI is of great relevance to EIT Health too. I expect that we will see AI appearing more and more in each and every EIT Health activity: innovation projects, business creation and acceleration, professional and citizen education, as well as regional outreach.

At EIT Health, we hold a unique position from which to contribute to the broader debate on AI, in Europe as well as globally. We have the ability to put AI and its ethics into practice, across our vast European network who share our mission of bringing cutting-edge innovation in health and ageing to those who may benefit.

Therefore, in 2019, we have undertaken three related activities:

- A. **EIT Health is piloting the recent 'Ethics Guidelines for Trustworthy AI'**³ of the European Commission's High-Level Group on AI, by means of a pan-European survey amongst innovators within our community; including start-ups, entrepreneurs and EIT Health partners from industry, academia and research organisations. The results are presented in this report and were showcased at the 2019 World Health Summit (WHS). This is a unique and early testing of the practical applicability of the EC's AI and ethics guidelines in one of the most relevant areas; health and ageing.
- B. **EIT Health delivered a Forum session at European Health Forum Gastein 2019** entitled "AI: It's not (just) about the technology!". You will find a short description of the discussions in this report.
- C. **EIT Health held a panel session at the WHS on 27 October 2019** in Berlin on AI and ethics in health innovation with high-level participants from policy, industry (small and large), and academia. We will report on the debate in an update to this report.

I believe that EIT Health should be a leader in the responsible use of AI in health and ageing innovation. We look forward to supporting innovators and policy makers in implementing guidance into real-world scenarios. Enjoy the reading!

Jan-Philipp Beck, Chief Executive Officer, EIT Health

Executive Summary

Important and rapid progress is being made in AI for health and care innovation.

In parallel, the issue of responsible, ethical use of AI has risen to the top of political agendas. A High-Level Group established by the European Commission issued guidelines for AI and ethics in April 2019. Incoming European Commission President, Ursula von der Leyen, has committed to putting forward AI legislation in the first 100 days of her mandate.

Ethical use of AI is of the utmost importance for EIT Health, a network of best-in-class health innovators backed by the European Union.

Responding to the European Commission's invitation to pilot the guidelines, EIT Health ran an AI and ethics survey to which 82 innovators, start-ups, and other stakeholders from 21 countries responded. The survey, though limited in size, is complemented by case studies in this report. Combined, they provide a relevant early indication of how the health innovation community deals with the ethics of AI in their field. It is still early days for the AI and ethics guidelines: only 22% of respondents were already aware of them. The survey itself was therefore an important awareness-raising tool.

The survey shows that the highest priority for ethics of AI in health innovation is given to privacy and data governance (which includes data protection and access to data), technical robustness and safety (which includes cyber-resilience and reproducibility of the AI), followed by traceability and human agency and oversight. Generally, respondents feel that AI should assist rather than be relied upon completely.

Lower ranked though still relevant were the ethics of diversity, non-discrimination and fairness; accountability (established methods appear to be considered adequate); and societal and environmental well-being. Finally, a clear majority (60%) of respondents expect that their AI solution will require regulatory approval. The survey suggests recommendations for further work to the EC and to EIT Health.

EIT Health also organised a panel on AI and ethics in health innovation at the WHS 2019 in Berlin on 27 October, with high-level participants from policy, industry, and academia. Finally, EIT Health organised a session at European Health Forum Gastein 2019 entitled "AI: It's not (just) about the technology!". This addressed the impact of AI on clinical roles and organisational structure as well as on the education and skill set needs of healthcare professionals.

Setting the Scene

EIT Health

EIT Health is a network of nearly 150 expert partners, made up of leading organisations and institutions from academia, business and healthcare delivery to answer some of the biggest healthcare and ageing challenges facing our society today. Supported by the European Institute of Innovation and Technology (EIT), a body of the European Union, it collaborates across borders to nurture and deliver new solutions and initiatives that can enable European citizens to live longer, healthier lives.

Headquartered in Munich, EIT Health has a pan-EU representation via seven regional innovation hubs (Co-Location Centres), which operate as strong clusters of relevant actors collaborating as a thriving ecosystem. The InnoStars cluster brings together organisations from more progressive regions where the overall pace of innovation is more moderate. In addition, they have expanded the community's presence with further hubs in the United States and Israel, which connect innovators across Europe to other key thriving ecosystems beyond the EU.

EIT Health powers up innovation; connecting industry, academia, and research so that the right players can work together to identify and bring promising health innovation to fruition.

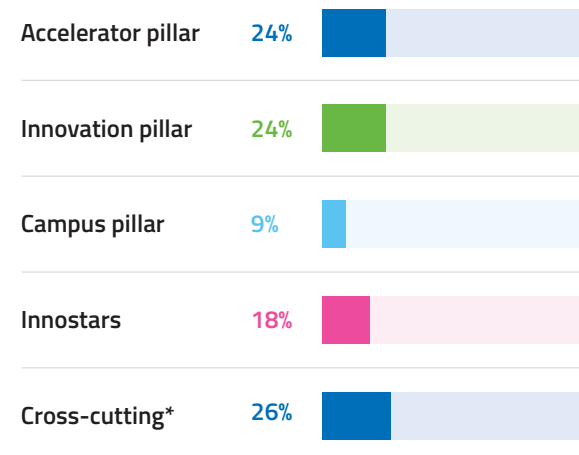
This is operationalised via three core mechanisms (or pillars):

- The Innovation programme provides comprehensive support for innovators developing a product or service with a high potential for societal and economic impact. The most promising ideas are developed into commercially-viable products through a multi-disciplinary approach, involving business, medicine, IT and other fields of knowledge.
- The Accelerator is a business creation programme, set up to support the best and brightest health entrepreneurs, providing skills and services to get promising business ideas to market. Support for healthcare entrepreneurs is provided at every stage of the process.
- Campus is the education incubator of EIT Health, providing up-to-date knowledge, skills and attitudes to help turn the brightest learners into healthcare leaders and entrepreneurs, who will shape the future of Europe's health. We are also committed to empowering patients and citizens to manage their own health more effectively through targeted learning programmes.



AI in EIT Health

Within the pillars of EIT Health, AI is already being utilised in concrete ways: as a technology within the Innovation projects, as a theme in entrepreneurship (Accelerator), in training and education (Campus), and in regional outreach, for example within our InnoStars hub. The table below provides an indicative distribution of AI-related activities across the organisation.



Some examples: Notch Interfaces is in the market with an app that uses AI to analyse and compare data from sensors that register the yoga positions of users. EIT Health's OncoWatch project improves prostate cancer diagnostics and treatment with AI and neural networks for image analysis based on access to 80,000 needle biopsies. RGS@HOME develops a Rehabilitation Gaming System combining brain theory, AI, virtual reality and cloud computing, to assist patients with motor and cognitive recovery after a stroke.

AI is rapidly growing in significance and prominence within EIT Health, as well as in health and ageing innovation in general. Today, an estimated 20% of EIT Health activities utilise AI in some form, and in the next few years a conservative estimate is that this is expected to double.

In Europe, as in other regions in the world, there is a high interest in responsible use of AI in health innovation. The world is closely following Europe's path into AI, in terms of innovation, ethical guidance and law. EIT Health can be a credible, respected and necessary contributor to the debate on AI in Europe, with a **focus on the practice of AI in health innovation**.

*Activities for EIT Health as a whole.

The European Commission's AI & ethics guidelines

In April 2019, the High-Level Group on Artificial Intelligence set up by the European Commission published its "ETHICS GUIDELINES FOR TRUSTWORTHY AI" and stated that trustworthy AI should be:

1. **Lawful - respecting all applicable laws and regulations**
2. **Ethical - respecting ethical principles and values**
3. **Robust - both from a technical perspective while taking into account its social environment**

These guidelines address seven areas: human agency and oversight; technical robustness and safety; privacy and data governance; transparency; diversity, non-discrimination and fairness; societal and environmental well-being; and accountability.

The guidelines are currently being piloted by interested parties from various sectors of society and economy including EIT Health, given its specific standing in the healthcare landscape as one of the largest innovation networks across Europe.

Process

Survey

EIT Health has conducted an online survey to measure the health innovation community's awareness of the guidelines on AI and ethics of the European Commission's High-Level Group and determine the relevance of the guidelines to their project and/or start-up. The target group of health innovators was reached through EIT Health's network and social media channels such as Twitter, LinkedIn and Facebook. The survey addresses, in a one-to-one way, the seven key requirements for trustworthy AI as defined by the EC High-Level Group.

Case Studies

We provide a number of case studies in order to highlight the real-life experiences of innovators with ethical issues when applying AI in health or ageing innovation. The intention is to illustrate the practice of AI and ethics; in terms of solutions, open issues and possibly 'hard-to-crack' dilemmas. The case studies are taken from the field of brain lesions, anti-microbial resistance, malaria and dementia. They demonstrate, without attempting to be complete, the diversity of AI and ethics considerations in this type of innovation.

World Health Summit

EIT Health has engaged in a partnership with the WHS, a leading international forum for global health. At the WHS on 27 October 2019 in Berlin, EIT Health delivered a panel on AI and ethics with high-level participants: Roberto Viola, DG European Commission; Peter Dabrock, Chair, German Ethics Council; Hans Hofstraat, Vice President, Philips Research; Hila Azadzoy, Managing Director, Ada; Susan Dehmel, BITKOM, Jan-Philipp Beck and Paul Timmers (moderator), EIT Health. The aim is to contribute to the broader debate on AI, and to Europe's policy agenda. A panel report will be included in an update to the present report.

Gastein

EIT Health delivered a Forum session at the European Health Forum Gastein 2019 entitled "AI: It's not (just) about the technology!". This session aimed to assess the implication of AI in healthcare beyond the technological advancements. It involved two panel discussions – one focused on the impact of clinical roles and organisational structure reviewed by Drs. Dominik Pfürringer, Medical Specialist for orthopaedics and trauma surgery, University Hospital Munich Rechts der Isar, Germany, and Steven Petit, Medical Physicist and Assistant Professor, Erasmus University Medical Centre. The second panel discussion explored the impact of AI on the changing education and skillset needs of healthcare professionals. Participants included Antanas Montvila, Doctor and Board Member, European Junior Doctors Association, Kristaps Krafte, CEO of Vigo.



Findings

Survey Results

In total, we collected 82 responses of which we will show the results in the graphs below. The survey was relatively extensive and required about 40-60 minutes to answer. As a result, not all respondents filled out every question, and therefore the overall number will vary per question.



The companies and start-ups in this survey have an average of **6 employees.**




The respondents represent the following countries: Poland, Spain, UK, The Netherlands, France, Belgium, Romania, Germany, Italy, Austria, Hungary, Finland, Cyprus, Bulgaria, Denmark, Ireland, Portugal, Sweden, Lithuania, Croatia and Switzerland.



34 out of 58

respondents expect their AI solution will need regulatory approval, (e.g. by the European Medicines Agency). However, to many, the pathway to receive approval is unclear.




13 out of 58


respondents suggest more topics to be addressed in the guidelines on ethics and AI. For example: the quality of network infrastructure; post-marketing surveillance of AI products; contribution of AI to society through tax or dividend and workforce displacement and transition.

Before going into the AI and Ethics guidelines, we provided several questions to receive clarification on the maturity of the health innovation, the type of AI they are using, and their knowledge of the guidelines.

For startups: Is your product already on the market?




How are you using AI in your projects and/or organization?



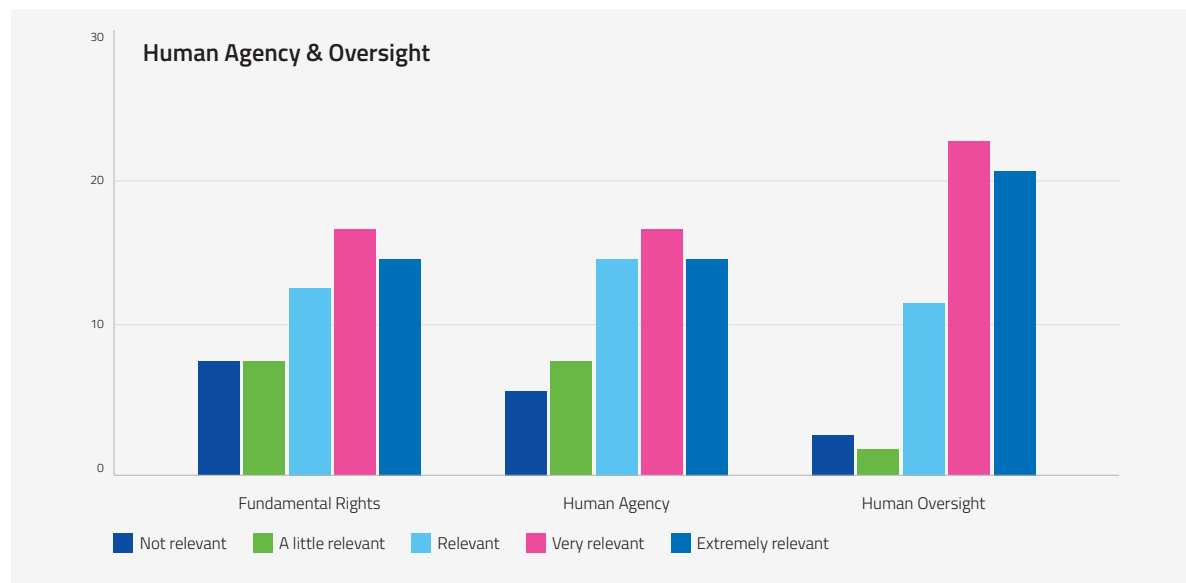
'Other' uses of AI included; patient risk stratification, discovery of drug candidates for unmet medical needs, EEG analysis, foot gesture recognition, virtual therapist, and automated patient training.

Did you hear about the European Commission's Ethics guidelines for trustworthy AI before this survey?



The following provides an overview of the relevance of the seven guidelines on ethics and AI on a scale of not relevant to extremely relevant. The seven guidelines include human agency and oversight, technical robustness

and safety; privacy and data governance; transparency; diversity, non-discrimination and fairness; societal and environmental well-being and accountability.



Human Agency and Oversight

European Commission High-Level Group on human agency and oversight

AI systems should empower human beings, allowing them to make informed decisions and fostering their fundamental rights. At the same time, proper oversight mechanisms need to be ensured, which can be achieved through human-in-the-loop (HITL), human-on-the-loop (HOTL), and human-in-command (HIC) approaches.

Respondents were asked to give scores on three elements of human agency and oversight as formulated by the High-Level Group:

Fundamental rights: AI systems could negatively affect fundamental rights. In situations where such risk exists, a fundamental rights impact assessment should be undertaken and include an evaluation of whether those risks can be reduced or justified as necessary in a democratic society, in order to respect the rights and freedoms of others.

Human agency: AI systems can sometimes be deployed to shape and influence human behaviour through mechanisms that may be difficult to detect, which may threaten individual autonomy. The overall principle of user autonomy must be central to the system's functionality.

Human oversight: human oversight helps to ensure that an AI system does not undermine human autonomy or cause other adverse effects. When using AI, it must be ensured that public enforcers have the ability to exercise oversight in line with their mandate. Oversight may be achieved through governance mechanisms such as HITL, HOTL, or HIC.

WHAT RESPONDENTS SAY:[†]

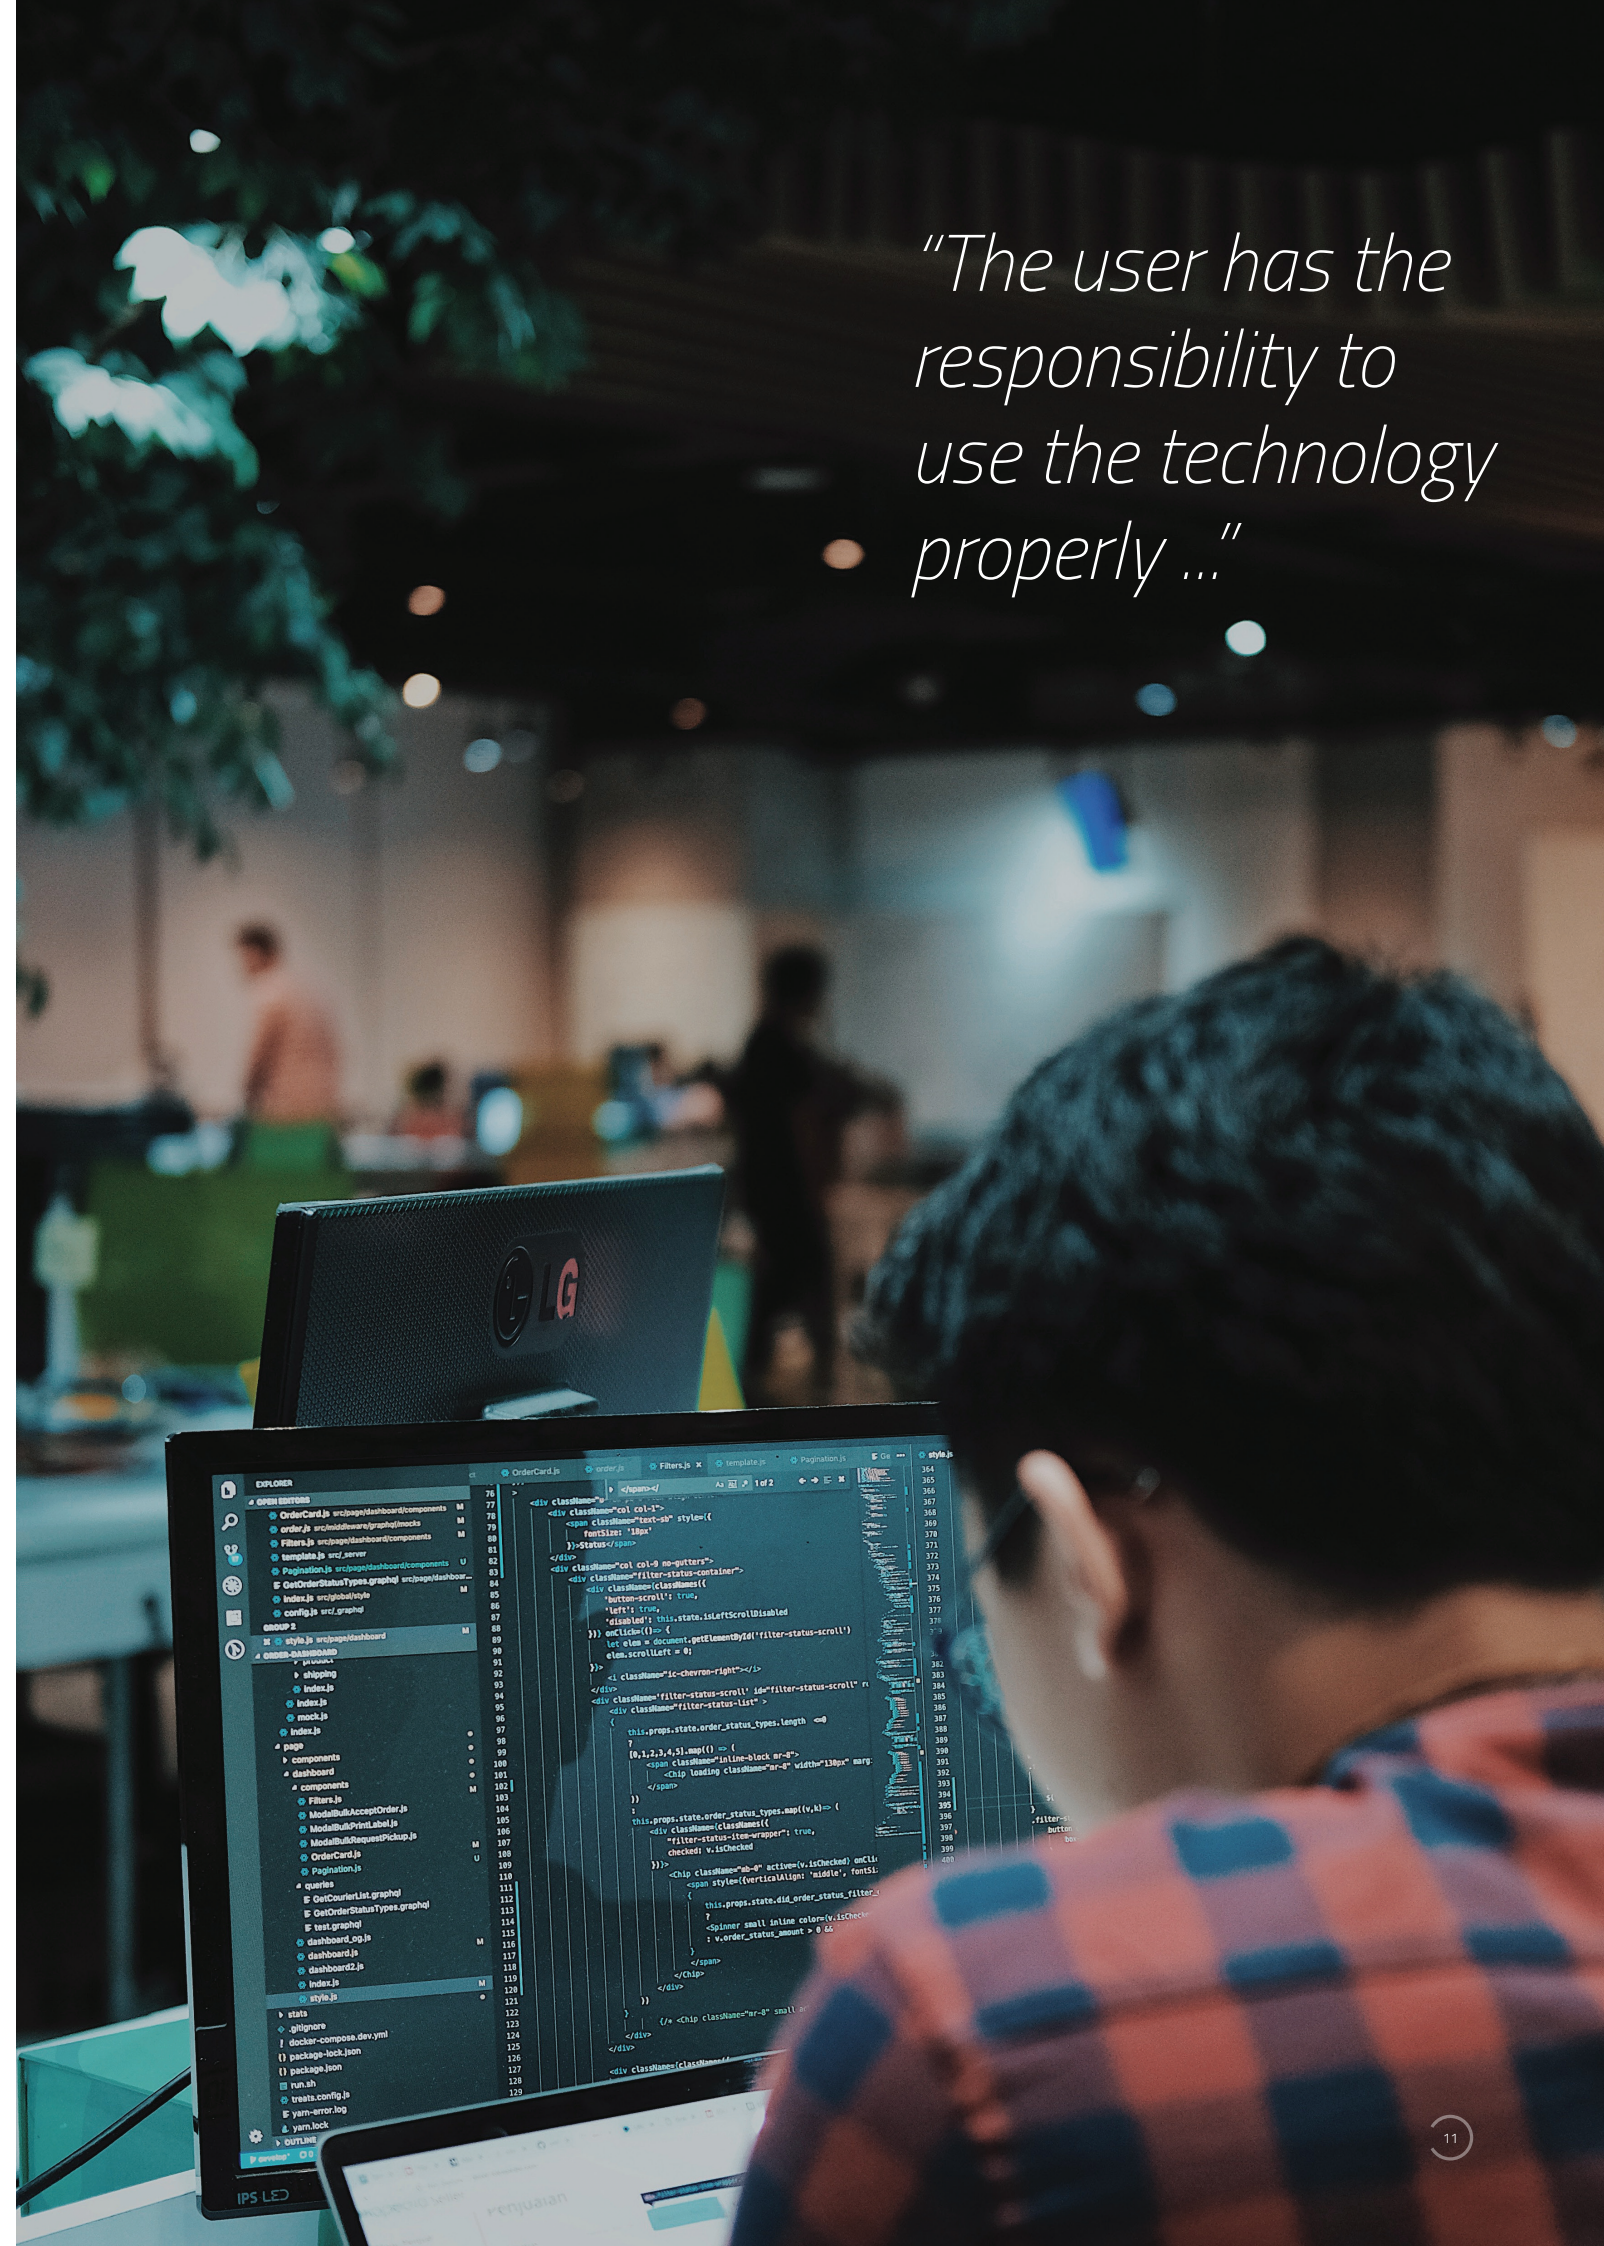
“It is very important to consider human agency and oversight when using AI tools, as a principle aim of developing or using AI is to assist, rather than replace, human brain. For example, we are trying to develop an AI platform which will provide doctors with necessary information regarding patients' current condition and best possible medications. However, it is between the patient and doctor to decide which medication best suits the patient. In this case AI will not be able to measure the emotions of a patients.” (A company from Germany)

“As developers, we only ensure that the AI system works according to specifications. The user has the responsibility to use the technology properly. [...] We can train human users to properly use our technology.” (A company from Spain)

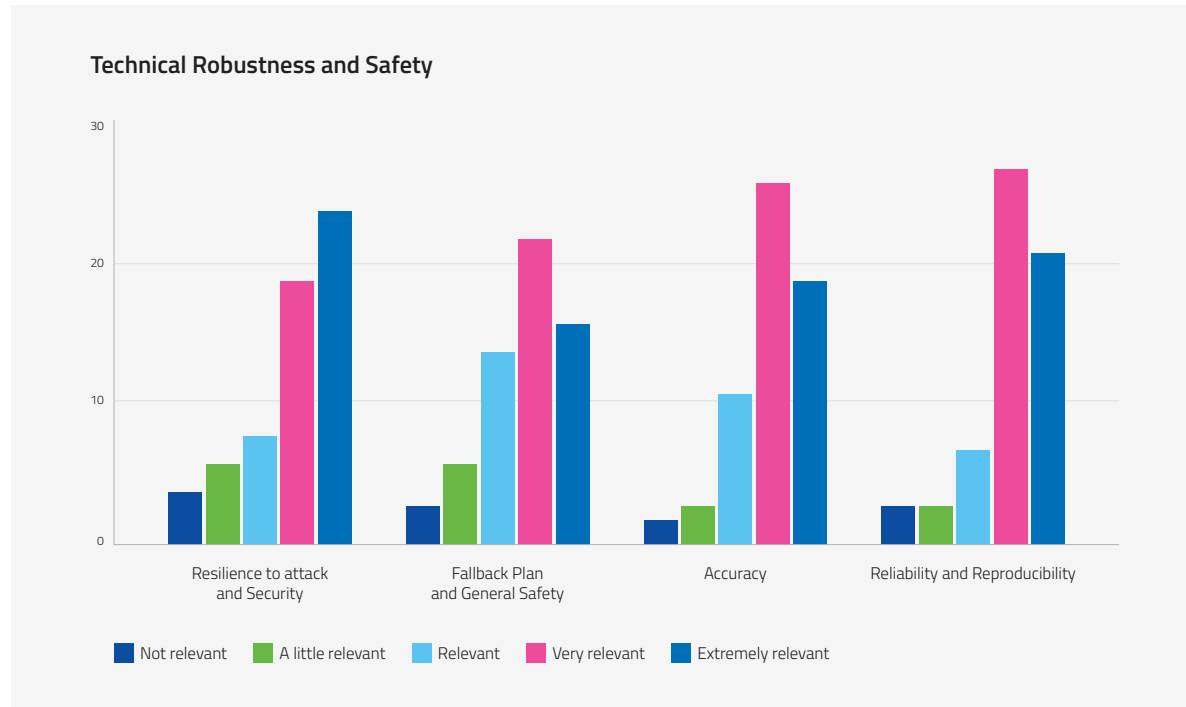
“The results of the automatic patient image analysis are presented in an easy way to understand. It is more than a simple 'Yes/No' detection or classification. We put effort into the user-experience (UX) to show the decision maker - in this case the radiologist - why the automatic evaluation process got to that conclusion. Then the human oversight has the responsibility to check the results and verify if they are correct. Only then do we look towards recommendations of treatment based on this detection or pathology classification. Every other AI tool should have this, regardless of whether it is a medical product or not. UX and clarification of why a certain AI result is presented is necessary.” (A company from Germany)

“We use AI to help process and make sense out of data. The outcomes are, for example, an advised setting or advice in taking medicine. Without AI, the outcomes are created by human processing. The advantage of AI is the processing power and offering insight e.g. in form of relations in data that a human could not come up with. The end result, or outcome, can and will be evaluated by a human, while the human tries to make sense of the advised outcome. It still is a human decision to follow the AI advised outcome or not.” (A start-up from The Netherlands)

[†] Minor spelling mistakes in respondents' answers have been corrected.



“The user has the responsibility to use the technology properly ...”



Technical Robustness and Safety

European Commission High-Level Group on technical robustness and safety


AI systems need to be resilient and secure. They need to be safe, ensuring a fall back plan in case something goes wrong, as well as being accurate, reliable and reproducible. That is the only way to ensure that unintentional harm can be mitigated and prevented.

WHAT RESPONDENTS SAY:

“Standardisation is a prerequisite to enhance interoperability, quality, security and data analysis. Also, recommendations for data quality, validation, regulatory and acceptability are key in a healthcare AI project. Data governance process needs to be challenged at the origin of the concept to prepare criticism, audit and regulatory constraints (for future registration).” (A start-up from The Netherlands)

“Whether an application or device is providing medical diagnosis or not, it is very important that the historical audit trails are stored that describe what decisions the AI made. Right now, it can be very difficult to understand how an AI engine arrived at a particular answer. In our development process, we have implemented quality management processes that require any data point that is inferred by an AI algorithm to be logged by the system as it is executing. As the system executes, we build a dependency model that describes which (AI) software functions were executed. For example, if we report that during a given day, a patient has completed ‘Food Intake’ ADL, the administrator can view a chart that shows the data items used to arrive at that decision and which software modules they depend on.” (A start-up from Ireland)

“The main check for our development is to try and understand the outcomes, ‘can we understand why the AI advises us an outcome?’. If not, we are very cautious in applying it. Reproducibility is also a very easy check that we use often.” (A start-up from The Netherlands)



Privacy and Data Governance

European Commission High-Level Group on privacy and data governance


Besides ensuring full respect for privacy and data protection, adequate data governance mechanisms must also be ensured, taking into account the quality and integrity of the data, and ensuring legitimised access to data.

WHAT RESPONDENTS SAY:

“Each echocardiography on which our algorithm is learning is coded, thus apart from the time when physicians use our software for analysis, patients are untraceable. Quality of the images, based on what the algorithm is learning, is ensured by a professional team of clinicians, which annotates these images and then feed them into the algorithm.” (A researcher from Lithuania)

“Privacy by nature should be at the origin of the project. A privacy officer should be the ethical guardian for legal accountability/responsibility to design the privacy by nature at the origin of the project.” (A company from The Netherlands)

“[...] the important thing to note is that there cannot be the same restrictions for data privacy and security for different types of AI. Narrow AI and prototyping/testing AI solutions should not have many restrictions, if any at all. Otherwise, this will increase the barrier for the development of any kind of AI so drastically that no AI innovation will be created in EU ...” (A company from Poland)



Transparency

European Commission High-Level Group on transparency

The data, system and AI business models should be transparent. Traceability mechanisms can help achieving this. Moreover, AI systems and their decisions, should be explained in a manner adapted to the stakeholder concerned. Humans need to be aware that they are interacting with an AI system, and must be informed of the system's capabilities and limitations.

WHAT RESPONDENTS SAY:

“All our data is traceable, and we try to show the algorithms in simple logic to patients so they can understand what's 'under the bonnet' when a decision is made. This mirrors what doctors do when they explain to patients how they arrived at diagnoses.” (A company from the UK)

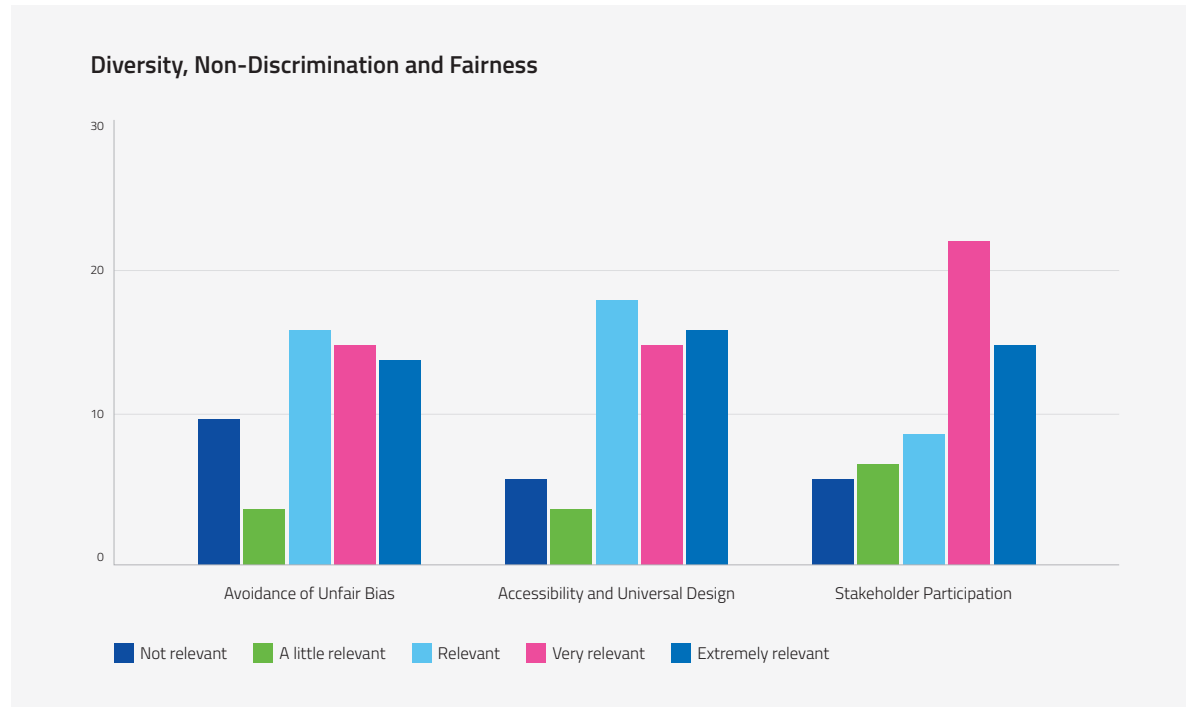
“Transparency and traceability will be most important when errors are made. We provided a dependency model that links a decision made to the data items and software algorithms used to arrive at that decision.” (A company from Ireland)

“For us, explainability is key. We use AI to give us insight, these insights gain as we can explain the outcomes better. If explainability is needed by rule, this can be a big help for other AI developments. True power of AI comes if we find a way to 'trust' the outcomes without the explicit need of explanation by human analysis.” (A start-up from The Netherlands)

“Transparency, for us, means to directly involve caregivers in the development of our solutions - to show them how it works and why they shouldn't be afraid to leverage its power during their everyday activities.” (A company from Italy)

“Transparency is relevant, yet I think it is much less relevant than outcome. Just like people don't know how Google Maps finds the best road, or how an e-mail is sent. Full transparency is very hard to achieve, and in general people are not really interested. And again, it matters for much more than just AI algorithms, including topics from election systems, the process behind medical reimbursement (or any other decision made by a government or big company) to food processing.” (A start-up from Belgium)

“All our data is traceable and we try to show the algorithms in simple logic to patients so they can understand what's 'under the bonnet' when a decision is made ...”



Diversity, Non-discrimination and Fairness

European Commission High-Level Group on diversity, non-discrimination and fairness


Unfair bias must be avoided, as it could have multiple negative implications, from the marginalisation of vulnerable groups, to the exacerbation of prejudice and discrimination. Fostering diversity, AI systems should be accessible to all, regardless of any disability, and involve relevant stakeholders throughout their entire life circle.

WHAT RESPONDENTS SAY:

“Our AI systems use both video and audio. For our video system, we have carried out testing of a range of mechanisms that might defeat the AI recognition - with special focus on both grooming and clothing styles that vary due to geography and religious beliefs.”
(A company from Ireland)

“It is also important, but not a priority for us. Data analysis in AI must have diversity, however, sometimes is not easy to have a big representation of the data we want to analyse. Data sometimes is private, and to have access to them we must work closely with stakeholders. In our case, we are doing this with hospitals and researchers.”
(A company from Spain)

“We have leveraged Inclusive Design Principles since the beginning. Our team is worldwide, and some developers are impacted by the condition themselves. We have an AI ethic board in place. Stakeholders (parents) have been included in every step of the design/development.” (An NGO from Italy)



Societal and Environmental Well-being

European Commission High-Level Group on societal and environmental well-being


AI systems should benefit all human beings, including future generations. It must hence be ensured that they are sustainable and environmentally friendly. Moreover, they should take into account the environment, including other living beings, and their social and societal impact should be carefully considered.

WHAT RESPONDENTS SAY:

“As a start-up, frankly, it is costly for us to consider all environmental challenges (sustainable materials etc. for our MedTech device). When cash is stretched, we are unfortunately still using plastics etc., and although we do want to move to a more eco-friendly process, it is not the most financially savvy way and may kill our business. If, however, there were legal penalties or financial incentives to switch to these we would strongly consider that.”
(A start-up from the UK)

“To my knowledge, there has not been much research on the environmental impacts of widescale deployment of AI in healthcare. It is known that the power consumption of GPUs [Graphics Processing Units] is significant from a cost perspective, but little is known about the long-term impact on the environment should AI reach the level of market penetration that analysts expect.”
(A company from Ireland)

“Training of algorithms is a great burden in consumption of energy. In order to minimise it we are using cloud providers to store and process our data.”
(A company from Portugal)



Accountability

European Commission High-Level Group on accountability

Mechanisms should be put in place to ensure responsibility and accountability for AI systems and their outcomes. Auditability, which enables the assessment of algorithms, data and design processes plays a key role therein, especially in critical applications. Moreover, adequate and accessible redress should be ensured.

WHAT RESPONDENTS SAY:

“Liability issues need to be clarified for the users of medical devices making AI-based treatment decisions (such as automated insulin delivery or artificial pancreas systems). Clinical trials are mandatory in order to measure the risks of AI-based tools against the non-AI-based devices.” (An innovation project from Germany)

“For our software to reach the market, we will need to get a CE certification. ISO certified quality management system is also a prerequisite. Naturally, we will only go to the market when we have clear evidence from clinical trials on the effectiveness and accuracy of the algorithm.” (A start-up from Lithuania)

“Redress is really important. We have an internal audit process, for example, to ensure that any adverse events are fully investigated as part of a monthly review to ensure quality/safety standards.” (A company from the UK)

“We always document our model development as well as performance and possible pitfalls, but I feel like there is a lot more we can do in this area, especially in terms of making sure the models’ users are aware of how the models are built and that they understand edge cases and best practices.” (A researcher from Italy)

“We will continuously test our AI via human analysis in clinical settings, and signal patients at risk for a re-screening at the hospital, thereby confirming the accountability of the data collected and the AI output.” (A start-up from Germany)



“We will continuously test our AI via human analysis in clinical settings, and signal patients at risk for a re-screening at the hospital ...”

Case studies

Cases studies aim to illustrate the practice of AI and ethics, without claiming to be fully representative.

Abtrace

Abtrace builds an AI-based solution that advises clinicians about the best decision for their individual patient based on historical data from many patients, focusing on the prescription of antibiotics and tackling unwarranted variation. Getting antibiotics right is essential, since overuse is the main driver of antimicrobial resistance and a third of antibiotics prescriptions are inappropriate. Abtrace uses machine learning to recognise patterns in large datasets and Natural Language Processing to extract clinical meaning from medical records generating comprehensive insights. These insights are then fed back to clinicians in the format of integrated clinical recommendations. Semi-supervised learning is used, meaning that a limited labelled dataset is used and that a team of clinicians checks the output of each iteration of the algorithm improving the accuracy and robustness of the process.

With respect to the AI and Ethics guidelines, the translation of each requirement into the context of Abtrace's solution is considered to be fairly straightforward. Bias can be tackled by increasing the training data assuring diversity and using representative samples of the population in which the AI product will be deployed and revalidating regularly and at each new deployment. Further, the interpretability of the recommendation is considered to be key to convince clinicians, data holders and regulatory bodies about the transparency and safety of adoption of such systems, whereby these are provided with clear information about the data (and the source) that are used to train the algorithms. Abtrace considers all seven requirements of the AI and Ethics guidelines to be relevant by (early) design. One ethical concern, which may need further exploration, is how to ensure ethical compliance, while simultaneously ensuring that patients are able to benefit from new technological developments in a reasonably quick manner.

AiDx Medical

AiDx Medical has developed a device for image-based detection of malaria. AiDx's device automatically screens malaria blood samples, combining an optical system with AI to automate the process. An algorithm replaces the human factor and screens the images on the presence of malaria parasites, thereby reducing the expert dependency of malaria diagnosis and increasing the speed significantly in comparison to traditional diagnosis through human operated microscopy on blood samples. The aim is to increase diagnostic capacity in African hospitals.

With respect to the AI and Ethics guidelines, the developers analysed the requirements together with stakeholders through an extended context analysis. For instance, in relation to both the requirements of technical robustness and safety and human agency and oversight, AiDx Medical cannot provide a 100% reliable outcome, whereby the solution is positioned as a "decision support system" and not a diagnostic

device. The developers aim to improve the solution by gathering larger amounts of samples to train and further optimise the algorithms using machine learning. Once the screening is done, the device will show the images of what is detected. A medical expert can judge whether these are indeed parasites in order to confirm the result. The AI solution can reduce human error (analysing images with high precision), increase diagnostic quality (looking at larger amounts of blood than a human can do) and increase speed (save up to 40 minutes per diagnostic test), however all should be checked and confirmed by experts before a final diagnosis can be made. The developers consider the guidelines to be extremely useful, but they can imagine that not all product developers have time to do such an elaborate exploration as they did. They consider the guidelines to be clear in formulation but propose that they can be extended with examples of existing cases to provide more guidance.

RGS@Home

The Rehabilitation Gaming System (RGS) is an eHealth technology for post-stroke functional recovery (neurorehabilitation), which uses virtual reality and machine learning enhanced individualisation to deliver patient-specific rehabilitation. AI techniques for the individualisation of the training protocols include neural network data compression techniques, reinforcement learning, and deep learning. The types of data include kinematics, performance data from the training protocols, explicit responses to questionnaires and data from wearables including electrocardiogram, GPS and gyroscope. The RGS has been in development for the last 15 years and has proven to significantly reduce costs, while enhancing the quality of care and is in daily use in all leading hospitals in the Barcelona area and a further 50 clinics and hospitals in Spain, Germany, The Netherlands, Sweden, France, USA, and Brazil. The RGS solution generates data on the patient's recovery path which allows the doctor at the rehab centre to remotely see progress, adjust therapy, and give feedback. It also enables better monitoring of actual outcomes and integrates well in a care chain that spans the neuro centre at the hospital, the rehab

centre and the patient's home. Studies are also running to look at the relevance of RGS in other pathologies such as cerebral palsy, oncology, addiction, and psychosis under the assumption that the same brain principles for diagnostics and recovery apply here.

With respect to the AI and Ethics guidelines, one of the creators of the RGS stresses the importance of various principles, but also notes that there is some redundancy and overlap of various principles, and that the guidelines should not complicate or hinder innovation in this context that is in such great need of technological solutions: it is argued that given the strong growth in demand for health services (and thus rising costs) we are forced to become more reliant on technology. Therefore, it is suggested to carefully consider what real problems are faced when AI enters healthcare on a large scale, and how general principles can be fine-tuned to real-world contingencies and specific AI-applications. For RGS, technical robustness and safety is, for instance, important because the developers expect to deliver increasingly more autonomy in the decision making of the RGS system.

Sensara Lifestyle Monitoring

The Sensara Senior Lifestyle Monitoring System is an Internet of Things (IoT) service for (senior) people to live at home longer, and provide peace of mind to them, their loved ones and their (in)formal carers. Sensara does not use intrusive cameras or wearables, but rather uses a limited number of small infrared sensors in the home with a few motion (PIR) sensors and a few open/close sensors. Data is uploaded to an analytics engine to recognise living patterns: at what time someone gets up, how long the bathroom is used, with what frequency the kitchen is used, how long one is out of the house for shopping, when are they are sleeping etc. When the behavioural patterns are known, exceptions can be detected and analysed, for instance, minor exceptions like skipping a meal, or major ones like not getting out of bed in the morning, or even suspicious inactivity in the afternoon. Based on detection of exceptions and anomalies, it sends a notification via the smartphone app of a carer, family member or friend. The service is based on the user-in-the-loop principle – by providing information

to involved formal and informal caregivers that they can use in their decision-making. Sensara does not execute any actions on its own or make any (medical) diagnosis.

With respect to the AI and Ethics guidelines, it is considered that Sensara must be safe, ensuring a fall back plan if the system does not function properly (i.e. sensor battery empty, internet connection offline), by means of push-notifications on the smartphone app. Technical robustness and safety, privacy and data governance and transparency are considered to be the most important and best applicable principles to Sensara, whereas diversity, non-discrimination and fairness and societal and environmental well-being are considered to be non-applicable. The developers are familiar with the guidelines, but not in detail, since the system was developed long before these guidelines were defined. It is suggested that the guidelines can be helpful when designing a new service / product or when a service provider would like to bring its service to the 'next level'.

Analysis and interpretation

General

The fact that only 22% of respondents were already aware of the EC's guidelines can be explained in several ways, such as the extent of visibility of the European Commission's publication of the guidelines and the extent to which such information is followed by the respondents. From the comments, however, it is evident that the survey was well appreciated.

Several respondents expressed that the survey was important in raising their own awareness of the issues involved in AI and ethics. This confirms the expectation that it is still early days in the adoption of AI and ethics guidance. Not every respondent will already have a highly differentiated approach with regards to the relative importance of the seven areas of AI and ethics.

The fact that 60% of respondents answered 'yes' to the question regarding whether they expect that their AI application will need regulatory approval is in line with surveys of the general public and CEOs from other sectors of the economy.

When calculating the number of responses given to 'very relevant' and 'extremely relevant', the overall picture is that the highest priority is given to privacy and data governance; technical robustness and safety; followed by traceability; and human agency and oversight. Lower ranked, though still relevant, were the ethics of diversity, non-discrimination and fairness; accountability; and societal and environmental well-being.

1. Human agency and oversight

A large majority of respondents regard fundamental rights, human agency and human oversight as very to extremely relevant. Notably human oversight scores high (44 out of 61 respondents). There is a high level of recognition of potential risks of AI. Generally, respondents feel that AI should assist people rather than be relied upon completely. One reason is that AI may be missing out on important human factors such as emotions and social nuances, which are very important in healthcare, but difficult to capture with current AI technology. Finally, human agency is considered as relevant during the design stage and should also be a reality during implementation and use (e.g. in users' routines).

2. Technical robustness and safety

There is high level of recognition of the need for technical robustness and safety, with well over two-thirds considering this very or extremely relevant. Within this, security and reliability are highly rated with 48 of 61 scoring this as very or extremely relevant. Many respondents mention the importance of safety when working with patient data (very sensitive), and stress that this principle is more important when there is greater (negative) impact of incorrect outcomes. Some respondents provide both practical guidance on how to address technical robustness and safety, while others do not yet know how to approach it.

3. Privacy and data governance

Likewise, respondents reported a high level of recognition of the relevance of these topics. Obviously, privacy is considered fundamental when handling patient data, such as health records. Comments by respondents range from practical suggestions; such as about anonymisation to stating concerns. One respondent raised concerns about the risk of too restrictive rules for narrow AI or prototyping, while another expressed concern that the public debate (and fear) around privacy and the legal difficulties and costs of gathering personal data could lead to the risk of Europe falling behind in AI.

4. Transparency

Here there is a larger spread of views. The degree of transparency that is needed according to respondents seems to vary, which could potentially be explained by the nature of the algorithms and the extent to which developers are able to explain the algorithmic outcomes. Comments by respondents show that there is little 'hard' data-based, or algorithm-based, tooling for transparency (one respondent mentions that creating a transparency system might be more difficult than the actual AI application). Suggested measures tend to be around 'soft' measures, such as good communication and awareness, as well as involving care givers, i.e. a more human-focused rather than technology approach. Opinions are divided on whether we can afford AI algorithms in health to be a black box - for some respondents the result matters more than transparency.

5. Diversity, non-discrimination and fairness

This is reported as less of a priority than the other guidelines above. However, respondents do recognise that it is relevant to be inclusive and are, to some extent, working on it. For example: some include a diverse group of stakeholders, others are trained to cover a representative sample of the population. Also, some consider it important to tackle issues related to this principle in an iterative manner, e.g. step-by-step; improving datasets and learning about or removing potential biases. Further, some respondents seem confident about their ability to detect bias within their data(sets) based on their (domain) expertise. From the current results it is not possible to determine to which extent this takes place.

6. Societal and environmental well-being

On average, while rated as relevant, it is not seen as highly relevant. Most respondents do make a link with societal well-being (some consider that self-evident as they are working on health solutions), but environmental impact is less mentioned or not really taken into account, or environmental standards are considered expensive for a start-up.

7. Accountability

This is considered moderately relevant - perhaps surprisingly so, given that here we have a still early-stage technology which is already applied in a highly sensitive area, namely health. Traditional methods appear to be relied upon such as auditing, post-market surveillance and procedures for redress. Comments show that this is something start-ups wish to make more of a priority in the future, such as during post-market surveillance. A specific barrier that is mentioned is proprietary intellectual property on the AI algorithm such that in case of contested accountability, the algorithm could only be revealed in a court case.

Suggestions for next steps

For the European Commission

- 1) The EIT Health survey was a clear way to raise awareness of the AI and ethics guidelines of the EC's High-Level Group. The survey was running in August and September 2019, which was approximately four months after the guidelines were issued, and when awareness was still low. The European Commission may wish to investigate and understand the evolution of awareness, and define targets, as well as provide more examples of the application of the guidelines.
- 2) A clear majority of respondents to the survey expect that regulatory approval will be required of their AI in this field. This is relevant input for the preparation of a legislative proposal. The related impact assessment could address efficiency and effectiveness of regulatory options taking into account existing regulation in the field on health technologies and medicines.
- 3) There is a need to break down the wide field of AI. The EC could consider a further detailing of the guidelines for example by sector (health and ageing being one of those), with weights for the individual areas of guidance, reflecting their importance for the sector concerned.
- 4) On a number of specific issues that were raised, the EC is recommended to do further study:
 - a) How to deal with accountability and transparency of AI when intellectual property is at stake
 - b) Whether established approaches to accountability suit AI-accountability (this is particularly relevant for the medical/pharma sector but likely also for other sectors)
 - c) Whether there is a need for more illustration of issues that may arise regarding 'diversity, non-discrimination and fairness'
 - d) Which tooling can be made available for technology-based transparency of algorithms as required by EU rules
 - e) Whether further guidance and tools should be made available (beyond anonymisation) to address data protection in relation to AI solutions

For EIT Health

- 1) Running this AI and ethics survey was clearly positively appreciated by respondents. Such, or other, forms of policy engagement merit further discussion within EIT Health.
- 2) EIT Health can add more value which is clearly necessary and would be much appreciated by:
 - a) Helping its innovators to exchange practical ways for approaching AI and ethics in health innovation
 - b) Collecting and building up more experience with AI and ethics in EIT Health
- 3) EIT Health should consider issuing specific guidance and/or training on AI and ethics in health innovation, within the frame of the general AI and ethics guidelines and addressing specifically sensitive issues for health and ageing such as human oversight and transparency / accountability, privacy / data governance and technical robustness / safety.

Framework for Trustworthy AI

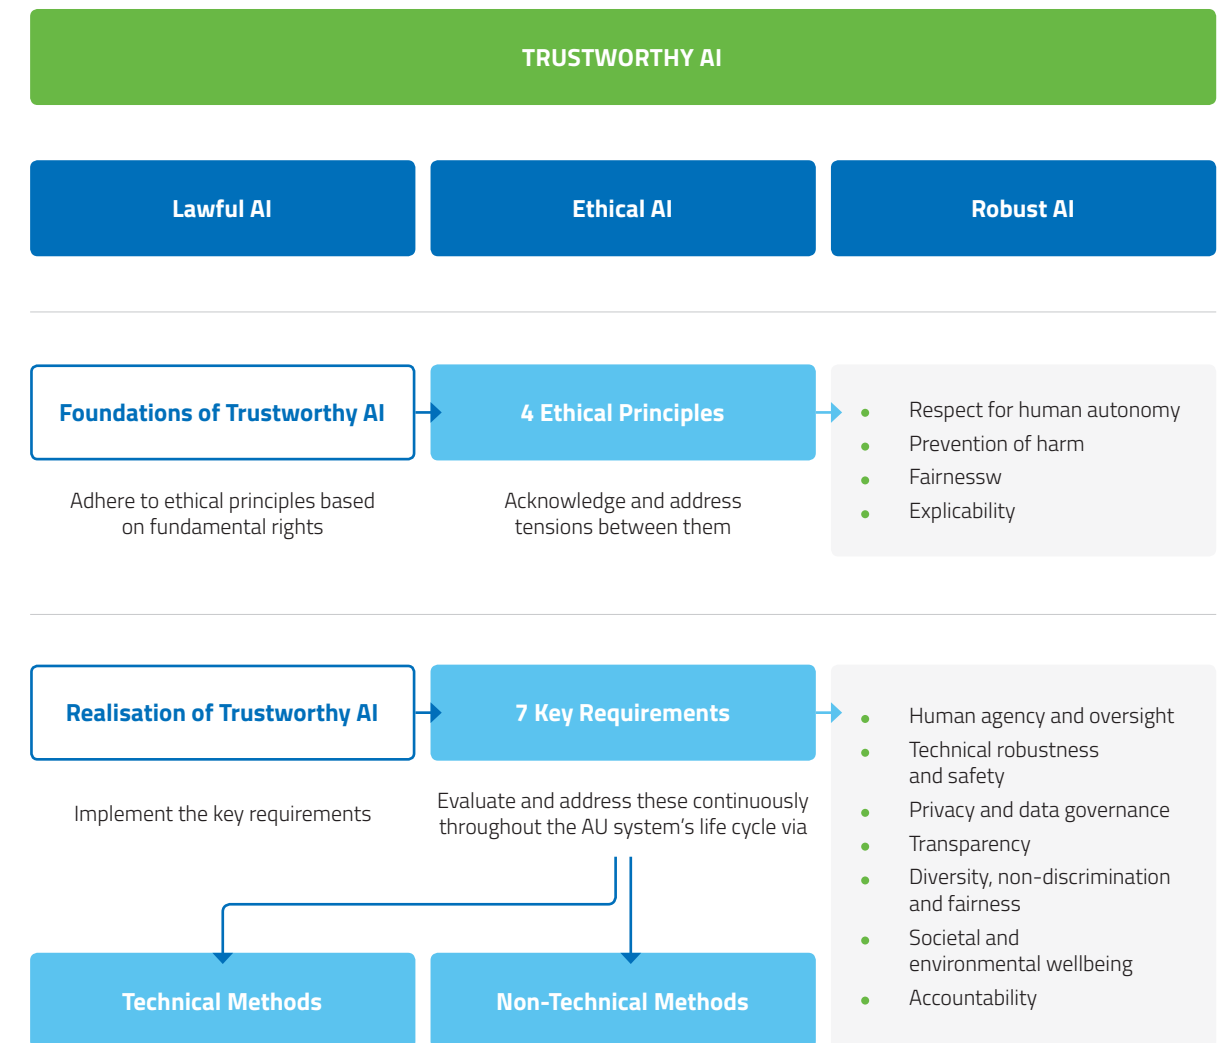


Figure 2 Guidelines for Trustworthy AI (after: European Commission, April 2019)

References

1. Nat Med 2018; 24(9), 1342-1350. doi:10.1038/s41591-018-0107-6 Available at: <https://www.ncbi.nlm.nih.gov/pubmed/30104768>
2. Digital Health Technology Vision 2018. Accenture Consulting, 2018. Available at: https://www.accenture.com/t20180518t100409z__w__/_us-en/_acnmedia/pdf-78/accenture-digital-health-tech-vision-2018.pdf#zoom=50
3. EC (2019), Ethics guidelines for trustworthy AI. Available at: <https://ec.europa.eu/digital-single-market/en/news/ethics-guidelines-trustworthy-ai>



EIT Health e.V.
Mies-van-der-Rohe-Str. 1C
80807 Munich
Germany

EIT Health is supported by the EIT,
a body of the European Union

