



THE UNIVERSITY *of* EDINBURGH

Edinburgh Research Explorer

Legal provocations for HCI in the design and development of trustworthy autonomous systems

Citation for published version:

Urquhart, LD, McGarry, G & Crabtree, A 2022, Legal provocations for HCI in the design and development of trustworthy autonomous systems. in Proceedings of the 12th Nordic Conference on Human-Computer Interaction (NordiCHI'22): Aarhus, Denmark 8th – 12th October, 2022., 75, ACM, pp. 1-12.
<https://doi.org/10.1145/3546155.3546690>

Digital Object Identifier (DOI):

[10.1145/3546155.3546690](https://doi.org/10.1145/3546155.3546690)

Link:

[Link to publication record in Edinburgh Research Explorer](#)

Document Version:

Peer reviewed version

Published In:

Proceedings of the 12th Nordic Conference on Human-Computer Interaction (NordiCHI'22)

General rights

Copyright for the publications made accessible via the Edinburgh Research Explorer is retained by the author(s) and / or other copyright owners and it is a condition of accessing these publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy

The University of Edinburgh has made every reasonable effort to ensure that Edinburgh Research Explorer content complies with UK legislation. If you believe that the public display of this file breaches copyright please contact openaccess@ed.ac.uk providing details, and we will remove access to the work immediately and investigate your claim.



Legal Provocations for HCI in the Design and Development of Trustworthy Autonomous Systems

Lachlan D. Urquhart*

School of Law, University of Edinburgh. lachlan.urquhart@ed.ac.uk

Glenn McGarry & Andy Crabtree

School of Computer Science, University of Nottingham. glenn.mcgarry@nottingham.ac.uk,
andy.crabtree@nottingham.ac.uk

We propose a series of legal provocations emerging from the proposed European Union Artificial Intelligence Act 2021 (AIA) and explore how they open up new possibilities for HCI in the design and development of trustworthy autonomous systems. The AIA continues the ‘by design’ trend seen in recent EU regulation of emerging technologies. The AIA targets AI developments that pose risks to society and citizens’ fundamental rights, introducing mandatory design and development requirements for high-risk AI systems (HRAIS). These requirements regulate different stages of the AI development cycle including ensuring data quality and governance strategies, mandating testing of systems, ensuring appropriate risk management, designing for human oversight, and creating technical documentation. These requirements open up new opportunities for HCI that reach beyond established concerns with the ethics and explainability of AI and situate AI development in human-centered processes and methods of design to enable compliance with regulation and foster societal trust in AI.

CCS CONCEPTS • Human computer interaction (HCI) • Artificial intelligence • Applied computing → Law

Additional Keywords and Phrases: European Union AI Act (AIA), trustworthy AI, mandatory design and development requirements.

ACM Reference Format: This block will be automatically generated when manuscripts are processed after acceptance.

“The proposal is without prejudice and complements the General Data Protection Regulation and the Law Enforcement Directive with a set of harmonised rules applicable to the design, development and use of certain high-risk AI systems” (European Commission, Proposal for an AI Act).

1 INTRODUCTION

In this paper, we consider proposed European regulations targeted at autonomous systems, particularly the recent proposal for an AI Act (AIA) [23] and its implications for HCI. The AIA continues the bold move made by the General Data Protection Regulation (GDPR)¹ [21], which introduced legal requirements in 2016 to “implement appropriate technical

* Place the footnote text for the author (if applicable) here.

¹ Summary Key for all acronyms used throughout. **AIA** = European Union Proposed Artificial Intelligence Act 2021; **GDPR** = European Union General Data Protection Regulation 2018; **HCI** = human-computer interaction; **HDI** = human data interaction; **AI** = Artificial Intelligence; **HLEG** = European Union High Level Expert Group on Artificial Intelligence; **CE mark** = signifier of conformity with EU standards; **HRAIS** = high-risk artificial intelligence

and organisational measures” enabling “data protection by design and default” (GDPR, Article 25). This introduced a unique set of legally mandated *design requirements* to bridge between legal principles and the technologically-driven use of personal data in the EU and beyond. The AIA continues this trend towards ‘by design’ regulation, introducing “common mandatory requirements” for the “design and development” of AI systems (AIA, Explanatory Memorandum). The overarching aim in both cases is to create *trust*. For GDPR, the issue of trust lay in fostering consumer confidence in the digital economy (GDPR, Recital 7); for the AIA, the issue is to “ensure AI is developed in ways that respect people’s rights and earn their trust” (AIA, Explanatory Memorandum). Through design and development, the AIA thus seeks to “promote trustworthy AI” and create an “ecosystem of trust” (ibid.).

Our goal in this paper is to sensitise the HCI community to the mandatory design and development requirements of the AIA which are aimed at shaping legal “global norms and standards” in AI regulation (AIA, Explanatory Memorandum). In doing so, we argue the AIA creates a series of *legal provocations* for AI governance that open up new opportunities for HCI beyond the established design-oriented concerns with ethics, fairness, bias, model interpretability and explanation. We borrow the term *provocation* from design research where a large body of work has focused on critical design practices, the role of critique as a design resource, and the need to challenge existing norms and attitudes, often through creation of design artefacts [17; 56; 76]. Whilst the form of our provocations is textual, we draw on the spirit of many research through design approaches which encourage questioning, discussion, critique and challenging of established ideas. Brown et al [7] explored ethical practices in HCI research through provocative text based statements to encourage reconsideration of current assumptions around ethics. Similarly, our *legal provocations* present provisions of the law to invite further *reflection, critique and action* on how these implicate HCI. Many of the new legal requirements may be strange or unfamiliar to HCI designers and practitioners. The pathways for navigating and addressing these legal requirements also remain unsettled and contested as these are complex, multi-layered and interrelated design requirements. Responding to the legal provocations may prompt discomfort about ways of working in HCI, usher in shifts in existing design practices, provoke frustration at legal language, and require reflexivity about how HCI practitioners engage with law going forward. We seek to reframe the relationship between HCI and AI, where it is not just about designing better methods of explanation or interpretation around AI systems, but turning to *designing for action* between AI and humans. HCI has much to offer in moving towards more human centred AI and exploring the intersection between law and HCI is key for emergence trustworthy autonomous systems in everyday life.

In terms of structure, we begin by examining the scope, applicability, and operational context of the AIA. We then unpack the legal conception of trust the regulation trades on before turning to a selection of mandatory design and development requirements imposed on AI systems. These requirements seek to build societal trust in AI by mandating data governance practices, testing regimes, risk management, creation of technical documentation, and human oversight. We move on to consider how these legal provocations align with and open up new opportunities for HCI. This includes finding a greater role for HCI in the technical heartlands of AI development, exploring novel means of rendering AI systems accountable, developing human-machine interface tools for better oversight and action in monitoring HRAIS, and embedding AI system development in human-centred design methods and processes, as well as legal processes, to enable compliance. The legal provocations encapsulated in the AIA show that there is much more to the “human side” of AI [1] from a regulatory perspective than just focusing on the interpretability of machine learning systems and explainable AI and the commensurate need to make AI systems intelligible to ordinary users. This paper thus sets out a novel agenda that

systems; **RMS** = risk management system; **PMS** = post-market monitoring system; **QMS** = quality management system; **IML** = interpretability in machine learning; **XAI** = explainable AI.

aligns HCI with the regulation of AI and situates the design and development of AI systems in human-centred approaches by *fiat*.

2 INTRODUCING THE PROPOSED EU AI ACT 2021

The AIA defines AI systems as those that use machine learning approaches, including supervised, unsupervised, reinforcement, and deep learning approaches, or logic and knowledge-based approaches, or statistical including Bayesian approaches (Annex I, AIA). The AIA classifies AI systems in terms of the differentiated *risk* they present to citizens and society. At the base of the pyramid is *low risk* AI, systems that pose minimal risks and on which fewer compliance obligations are placed (e.g., transparency). In the middle are *high risk* AI systems, which face wide ranging design and development requirements. And at the top of the pyramid is *prohibited* AI, which covers systems that pose risks that are deemed unacceptable to health, safety, and citizens fundamental rights (AIA, Article 5).² We primarily focus on the regulation of high-risk AI systems (HRAIS) in this paper, as they are subject to comprehensive ‘by design’ regulation measures. An AI system is deemed high-risk in the AIA through one of **two mechanisms**. The **first** is if *both* of the following conditions are fulfilled: 1) an AI system is either a product itself or is intended to be used as a safety component of a product covered by the EU harmonisation legislation listed in the annexes to the AIA (Article 6) *and* 2) the product or AI safety component is required to undergo a third-party conformity assessment. The Annexes to the AIA [24] provide a long list of related legislation covering a broad range of riskier product types including toys, machinery, lifts, fuel burning appliances, protective equipment, radio equipment, medical devices, civil aviation, motorised vehicles, watercraft and marine equipment and rail in which AI could be used AIA (Annex II). The **second** is if an AI system is used in one of the contexts listed in Annex III. This includes AI use in critical national infrastructure; educational and vocational training; employment, including worker management and access to self-employment; access to and enjoyment of essential public or private services and benefits; law enforcement, migration, asylum and border control; administration of justice and democratic process; and biometric identification and categorisation uses (Annex III). This list covers some of the examples that have dominated public concerns about AI in recent years such as use in sentencing support systems for judges, for automated credit scoring by banks, use in hiring screening and performance evaluation by employers, use in education to assess students doing tests, and use by law enforcement in predictive policing or evaluating reoffending risks. As litigation and case law emerges once the law is passed, concrete examples will emerge. There is a process for updating this list over time by the European Commission, providing some degree of future proofing.

The regulation stipulates that a *quality management system* (QMS), already commonplace in business,³ must be put in place by the provider of a HRAIS to manage the compliance requirements and risks that a system poses across these product areas and application domains. The QMS may be subject to “internal control” (self-assessment) or third-party conformity assessment by a “notified body” designated by a member state to do conformity assessments [26].⁴

² For example, the AIA places prohibitions on certain uses of remote biometric identification systems (facial recognition). Thus, and with certain caveats, the “use of ‘real-time’ remote biometric identification systems in publicly accessible spaces for the purpose of law enforcement is prohibited, unless it is strictly necessary” (i.e., unless occasioned by a threat to public safety or the apprehension of persons involved in significant criminal offences). Prohibited too is the use of AI systems that deploy subliminal techniques to exploit the vulnerabilities of persons in order to materially distort behaviour in a manner that causes or is likely to cause physical or psychological harm, and the use of social scoring systems by public authorities to evaluate and/or classify the trustworthiness of people based on their social behaviour or known or predicted personal or personality characteristics (AIA, Article 5). For clarification on the caveats to facial recognition and the wider use of AI by law enforcement, see Urquhart and Miranda [74].

³ See, for example ISO 9001 [39].

⁴ The need for third-party conformity assessment turns on various considerations including whether or not harmonised legislation already provides mechanisms for assessing conformity, or “harmonised standards” published in the *Official Journal of the European Union* have been applied (AIA, Article 40) or “common specifications” adopted through future EU legislation have been applied (AIA, Article 41). If the answer is no, then third-party conformity applies unless a provider is eligible for “internal control” (AIA, Article 43) or can follow the compliance process within the harmonisation legislation (AIA, Article 43). The rules also differ for application domains covered by Annexes II and III, according to Article 43 (AIA) e.g., it differs for HRAIS conducting biometric identification and categorisation.

Demonstrating the product safety and market access focused approach of the AIA, it requires a declaration of EU conformity be drawn up by the provider (Article 48 and Annex V), and a CE mark of conformity affixed “visibly, legibly and indelibly” (Article 49) to packaging and documentation before a high-risk system is put on the EU market.

The QMS is key to the *management of risk* and must furnish an account of the provider’s compliance strategy, technical documentation, and the procedures put in place to ensure the efficacy of the QMS over a system’s lifetime [24]. When third-party conformity assessment is required, this may include providing the designated body with ‘full access to training and testing datasets’ through an application programming interface (API) or other means of remote access. Alternatively, the provider of the system may be obliged to provide ‘access to the premises where the design, development, testing of the AI systems is taking place’. Upon “reasoned request” (AIA, Annex VII), the notified body must also be granted access to the source code of the AI system and may require that the provider supply further evidence or carry out further tests to enable a proper assessment of conformity. If and when an AI system is deemed to be in conformity, an EU technical documentation assessment certificate is issued by the designated body, who is also tasked with conducting periodic audits to ensure the provider maintains and continues to apply the QMS. Any change or modification to the system that could affect its compliance or its intended purpose must be approved by the designated body which issued the assessment certificate [24]. Systems that continue to learn after being placed on the market or put into service “shall not constitute a substantial modification” (AIA, Article 43). It is notable that the ‘provider’ of an AI system is not necessarily the developer or manufacturer of an AI system. The AIA thus *distributes responsibility* for compliance across the supply chain.

“Any distributor, importer, user or other third-party shall be considered a *provider* for the purposes of this regulation and shall be subject to the obligations of the provider ... if any of the following circumstances apply: a) they *place on the market* or put into service a high-risk AI system under *their name or trademark*; b) they *modify the intended purpose* of a high-risk system already placed on the market or put into service; or c) they make a *substantial modification* to a high-risk system” (AIA, Article 28, our emphasis).

The distribution of responsibility imposes legal obligations on importers to ensure a conformity assessment has been carried out by the provider (AIA, Article 26). If importers cannot be identified, providers based outside the EU must appoint an authorised representative in the EU prior to making their systems available on the EU market (AIA, Article 25). Distributors must verify that high-risk AI systems comply with the obligations set out in the regulation (AIA, Article 27). Users must inform the provider or distributor if they identify any serious incidents or malfunctioning and interrupt the use of the system in such circumstances, and they must suspend the use of the system if they have reason to consider it presents a risk to human health, safety and the protection of fundamental rights (AIA, Article 29).

These measures are cumulatively designed to create an “ecosystem of trust in AI in Europe” by putting in place a legal framework “centred on a well-defined risk-based regulatory approach” that “imposes regulatory burdens only when an AI system is likely to pose high risks to fundamental rights and safety” (Explanatory Memorandum, AIA). This is seen as “a better option than blanket regulation of all AI systems” (ibid.). Whilst the regulation covers providers and users located in the EU, like GDPR, it will have global reach, and includes mechanisms to expand its scope beyond the EU. For example, providers who put AI systems on the market or into service in the EU, whether they are established in the EU or third countries, will be subject to these rules. Also, where there are users or providers of AI who are based in a third country but the outputs of the AI system are utilised in the EU, they will again be covered by the rules (AIA, Article 2). Furthermore, like GDPR, failure to comply with the regulation will be subject to significant administrative fines, this time up to 30 million euros or 6% of total worldwide annual turnover, whichever is higher. These measures, including the introduction

of a risk-based regulatory approach, third party conformity assessment of quality management systems, the distribution of responsibility across the supply chain, and global reach are seen as key to mechanisms for promoting “trustworthy AI.”⁵

3 THE LEGAL CONCEPT OF TRUSTWORTHINESS

Before we move on to consider the specific design and development requirements mandated by the AIA for HRAIS, we wish to briefly consider the legal notion of trustworthiness that frames the proposed regulation. The legal framing of trustworthiness is related to and informed by but not the same as or reducible to ethical conceptions of trustworthiness, which have informed technical agendas in AI. An ethical framing of trustworthiness is encapsulated in the High-Level Expert Group (HLEG) on Artificial Intelligence report *Ethics Guidelines for Trustworthy AI* [36]. The HLEG brought together a self-nominated group of experts from industry and academia to consider the principal ethical challenges confronting the widespread uptake of AI in society [75] and subsequently they defined trustworthy AI as having 3 essential components. That AI is *lawful* and complies with fundamental rights including the freedom of the individual; respect for human dignity; citizens’ rights; equality, non-discrimination and solidarity; and respect for democracy, justice and the rule of law. That AI is *ethical* and respects 4 fundamental principles including respect for human autonomy; the prevention of harm to human beings; fairness and ensuring equality of opportunity and the right to redress; and explicability or the need to be transparent about AI decision-making. That AI is *robust* from a social and environmental as well as technical perspective, not only being well engineered but also being fit for purpose within the context in which a system will be deployed and operate [36]. These components create a complex framing of trustworthiness where multiple and even competing values are at play and create distinct challenges for design [29]. For example, from the point of view of robustness there is strong alignment with *trustworthy computing* [31] and that systems are therefore secure, respect privacy, are reliable, have business integrity, are thus used in responsible ways [51], and align with the earlier notion of *dependable computing*. As Randell [55] puts it,

“... the trustworthiness of a computer system [is] such that reliance can justifiably be placed on the service it delivers. Dependability thus includes such properties as reliability, integrity, privacy, safety, security, etc., and provides a convenient means of subsuming these various concerns within a single conceptual framework.”

These concerns, which are far from new or novel and in no way confined to AI, are writ large in the HLEG report and are closely linked to the ethical principle of preventing harm. Robustness, and especially technical robustness is closely aligned then with human safety and leads to consideration of further AI-specific issues including the resilience of AI systems to attack and the need to implement fall back plans to minimise unintended consequences and errors. Trustworthiness thus comes to have a particular framing from the point of view of robustness that emphasises the engineering of “safe, secure and reliable” systems that “prevent any unintended adverse impacts” [36].

It is notable that Mundie, author of Microsoft’s *White Paper on Trustworthy Computing*, reminds us that “making something trustworthy requires a social infrastructure as well as solid engineering” [51]. It is perhaps in this context that the ethical framing of trustworthiness is best understood. This certainly applies to the law, to which ethics in AI is closely bound (but not isomorphic). As the HLEG report puts it, its ethical principles are “rooted in fundamental rights”, which are “to a large extent already reflected in existing legal requirements” but go “beyond formal compliance with existing

⁵ One immediate consequence, and indeed some may see it as positive benefit, is that many emerging applications of AI in society are likely to be seen and treated as high-risk. However, there is concern about whether or not the net should be more widely cast. For example, commercial uses of emotion sensing systems are currently low risk [46], yet regulatory bodies such as the EU Data Protection Board and EU Data Protection Supervisor have called for these systems to be banned [27].

laws” [36]. Indeed, they do, bleeding quite recognisably into technical agendas in AI. Thus, in building on the principles outlined in the HLEG report, it is not only recommended that trustworthy AI incorporate *human oversight* to ensure that AI systems do not undermine human autonomy, but also that robust *data governance* be implemented to control access to personal data, ensure the quality and integrity of data, and protect data subjects’ privacy. It is recommended too that the process and outcome of automated decision-making be *transparent*, which includes explaining AI processes and outcomes, tracing and auditing decisions, and clearly communicating AI capabilities and limitations. AI systems should respect diversity and be non-discriminatory and thus *fair*, which involves tackling algorithmic bias, and they should be *accountable* both in their development and use, particularly to “public enforcers”. Trustworthiness thus comes to have a particular framing from the point of view of ethics that emphasises the engineering of a principled social infrastructure that intersects with and informs AI design and development. This intersection is often characterised by the ‘FAccT’ (Fairness, Accountability and Transparency) agenda, which sees the technical community trying to develop ways and means of tackling principled ethical concerns [e.g., 43, 60, 61, 62]. However, Birhane et al [6] analysis of the ACM FAccT Conference series papers finds ethics is often treated in an abstract manner e.g., fairness as a mathematical or philosophical problem, whilst neglecting the situated historical and societal power structures, actual harms, and the lived experiences of marginalised groups impacted by AI systems.

Fairness, accountability, and transparency are not only ethical and technical concepts⁶; they are key to *good governance* which requires governments and industry treat people equally without discrimination, ensure ongoing answerability for decision-making, and provide openness around decisions [13]. Thus, whilst there is a need to discuss ethical and technical dimensions of building trustworthy AI, *policy and law* are key mechanisms for enabling governance [11] and the AIA makes it clear that **trustworthiness in the eyes of EU governance is fundamentally about risk management**. The AIA configures trustworthiness as the management through legislation of risks to human health, safety, and fundamental rights occasioned by algorithmic discrimination in the use of digital services and the use of AI in safety components of all manner of products, devices, machinery and machines including large-scale IT systems (AIA, Explanatory Memorandum). The idea that trustworthiness is a matter of risk management will be familiar to many in HCI, where risk has previously been considered as a *personal* disposition, attitude, or belief grounded in rational choice and the calculus of risk [30, 12]. However, the legal concept has more in common with *societal* arrangements of trust, which locate trustworthiness in such things as the social contract, social capital, and social institutions [68, 54] amongst other mechanisms of social order [48]. Thus, while informed by ethical conceptions and related technical developments, the legal notion of trustworthiness is distinct from these framings. The AIA creates a formal legal basis mandating AI providers take certain actions to manage the risks posed by AI systems backed by the authority of law. What constitutes trustworthiness is no longer a matter of personal or professional opinion but ineluctably defined in law. The AIA cannot be dismissed or ignored, then, and will impact practices of ‘ethics washing’ [80] where AI providers have sought to avoid regulation by ascribing to high level principles and adopting ethical codes of practice without meaningfully tackling actual or potential social harms.

4 LEGAL PROVOCATIONS: MANDATORY DESIGN AND DEVELOPMENT REQUIREMENTS

While the AIA implements a legally-binding version of what constitutes and provides for trustworthiness it is not, as we have said, insensitive to ethical or technical framings. Indeed, the AIA is remarkable in the degree of granularity it adopts in setting out mandatory ‘design and development requirements’ for high-risk AI systems (Chapter 2), which have clearly been informed by ethical concerns. For example, algorithmic bias and its discriminatory consequences have been one of

⁶ These concepts are legal principles and frameworks too e.g., the GDPR Art 5(2) Accountability Principle.

the primary concerns amongst technical and ethical communities around emergence of AI systems to date. Mechanisms to address these risks are incorporated in the AIA too, where legal standards of **data quality criteria** are defined. Designers are legally obliged to develop training, validation, and testing data sets where suitable data sets are available and in sufficient quantity to meet the intended purpose of the system. They need to ensure data sets are representative and have appropriate statistical properties that take specific geographical, behavioural, or functional characteristics of the setting where the HRAIS is to be deployed (e.g., training data reflects the population that will be subject to the system). They also need to make sure data sets are free of errors and that any shortcomings, including gaps in the data, are identified along with how to address them. Lastly, data sets should be examined for possible biases, and where training data sets are concerned, all of this should be done prior to model development (AIA, Article 10).⁷ Designers must subsequently measure levels of accuracy and publish accuracy metrics in the instructions of use, and develop feedback loop based mitigation measures for systems that continue to learn after being put into service, to circumvent emergent bias (AIA, Article 15).⁸

The focus on improving quality of training datasets and preventing discrimination by mandating common data governance practices is “complemented with obligations for **testing, risk management, documentation and human oversight** throughout [an] AI systems’ lifecycle” (AIA, Explanatory Memorandum). In addition to model development, HRAIS must be tested prior to being put on the market to ensure they perform consistently against “preliminarily defined metrics and probabilistic thresholds that are appropriate to the intended purpose” (AIA, Article 9). More radically, HRAIS must also be tested “for the purposes of identifying the most appropriate risk management measures” and this includes identifying foreseeable risks when a system is used according to its intended purpose, and risks that may arise under conditions of *reasonably foreseeable misuse* (AIA, Article 9). The latter includes testing that systems are resilient to attempts by unauthorised third parties to alter their use or performance, including cyberattacks that try to manipulate training data (‘data poisoning’) and inputs designed to cause the model to make a mistake (‘adversarial examples’) (AIA, Article 15). The reasonably foreseeable misuse clause puts the systematic use of adversarial testing centre stage in the development of HRAIS. As the UK’s National Cyber Security Centre reminds us, this is not just a matter of understanding and analysing the risks to data models but risks to the system as a whole occasioned by its entire array of external dependencies [53].⁹ The reasonably foreseeable misuse clause represents a significant provocation to the design and development of trustworthy AI, to think beyond intended applications, to anticipate what might trigger these, when, and to design appropriate mitigations and management measures.

Significant too is the legal obligation that the designers of HRAIS put a **risk management system (RMS)** in place prior to a high-risk system being put on the market and that the RMS be updated and maintained throughout the system’s lifetime (AIA, Article 9). The RMS requires designers to “take account of the state of the art” and implement technical measures to eliminate risks to human health, safety, or fundamental rights. Alternatively, if risks cannot be wholly eliminated, then risks should be reduced through the implementation of “mitigation and control measures” and any “residual risks” should be clearly communicated to users. Furthermore, the RMS should evaluate risks emerging over the course of a high-risk system’s lifetime by putting a **post-market monitoring system (PMS)** in place (AIA, Article 9). The PMS must “systematically collect, document and analyse data” that is provided by users or collected through other sources “on the performance of a system throughout its lifetime” (AIA, Article 61). Complementing this, the AI Act requires that HRAIS

⁷ It has been noted that these mandatory data governance practices are not sufficient and need to be further bolstered through regulation to take into account “techniques such as unsupervised or reinforcement learning [which] do not rely on validation and testing data sets” [67].

⁸ It has been noted too that accuracy is not the only measure of model performance, e.g., “One of the performance metrics used in reinforcement learning is its reliability. The requirements in Article 15 should accommodate such performance metrics” (ibid.).

⁹ It is worth noting, as found in a recent study, that popular off-the-shelf machine learning frameworks, including Caffe, TensorFlow and Torch “contain heavy dependencies on numerous open-source packages.” Indeed, “Caffe is based on more than 130 depending libraries ... and Tensorflow and Torch depend on 97 Python modules and 48 Lua modules respectively [84].

“be designed and developed with capabilities enabling the automatic recording of events (‘logs’) while the system is operating” (AIA, Article 12). Logging “should facilitate post-market monitoring” (ibid.) and thus enable the provider to monitor the continuous compliance of the system and situations that may result in the system posing risks to health, safety or fundamental rights or that occasion the need for substantial modification.

The RMS and PMS are part and parcel of the QMS, described above, and must be described in accompanying **technical documentation**, which should again be drawn up before a high-risk system is put on the market or into service (AIA, Article 11 and Annex IV). Technical documentation should minimally provide a *general description* of the system e.g., its intended purpose, alongside a *detailed description* of the system and the process for its development. This includes details such as *design specifications* e.g., key design choices; *assumptions* about the data and data subjects; the system *architecture*; *data requirements* e.g., around data characteristics and provenance; *training, testing and validation procedures* e.g., tests logs and accuracy metrics; *assessment of human oversight measures*; *monitoring, functioning and control* e.g., foreseeable unintended outcomes and sources of risk to health and safety (AIA, Article 11 & Annex IV).

Monitoring, functioning, and control requirements are not only matters of documentation. Designers will both have to incorporate elements of the RMS into the development process (e.g., data governance and testing) and develop and implement elements of the RMS and PMS as technical components within HRAIS (e.g., state of the art risk mitigation measures and automatic logging). The proposed regulation will also oblige designers to build-in measures that enable people to exercise control over HRAIS (AIA, Article 13). This requires **human oversight measures** should be identified and built into HRAIS by the provider before a system is placed on the market (AIA, Article 14)¹⁰. This includes the development of contextually or situationally appropriate *human-machine interface tools* that enable effective oversight by natural persons during the period in which a system is in use, to prevent or minimise risks to health, safety or fundamental rights.¹¹ This includes detecting and addressing *anomalies, dysfunctions and unexpected performance*, being able to *disregard, override or reverse* the system output to even interrupting and *stopping* a system from operating.¹² Generally, it is expected that HRAIS should be responsive to human operators, be designed and developed so that people can oversee their function and, as appropriate to context and risk, operational constraints should be built-in to ensure HRAIS “cannot be overridden by the systems itself” (Recital 48). This expectation is seen as key to creating trustworthy AI, which “can only be achieved by ensuring an appropriate involvement by human beings in relation to high-risk AI applications” [22].

To summarise, common mandatory design and development requirements in the AIA frame trustworthiness of HRAIS as risk management, built around the design and development of key requirements, as shown in Table 1.

Table 1

Key Requirement Description	Summary
Quality criteria and data governance	Ensure acquisition, training, validation and testing of data <i>sets meet strict standards</i> and ensure development of models from data minimise the risk of algorithmic harms such as discrimination.
Testing	Ensure <i>testing</i> for accuracy, consistent performance with respect to intended use of the HRAIS, and identification of risks, including those occasioned by reasonably foreseeable misuse of the data, the model, or the system as a whole.
Risk management	Create <i>state of the art risk management measures</i> that eliminate risks to human health, safety or fundamental rights wherever possible, or put mitigation and control measures in place if risks cannot be entirely eliminated. This can include automated logging across the life cycle of the system.

¹⁰ Or, that they be identified by the provider and implemented by the user if it is appropriate to do so.

¹¹ Whether a system is used for its intended purpose or under conditions of foreseeable misuse.

¹² Article 14 also states that human oversight of biometric systems used for remote identification should include verification and confirmation by *at least two natural persons*.

Human Oversight	Create <i>human-machine interface tools</i> enabling natural persons to exercise oversight and control systems during use, including supporting actions such as pausing or stopping the system.
Technical documentation and quality management system	Provide detailed documentation around a system at different layers of granularity from high level system overviews to accounting for specific design decisions.

5 IMPLICATIONS FOR HUMAN COMPUTER INTERACTION

In its framing of trustworthy AI, the AIA presents a number of significant provocations for design and development which require to be addressed. These focus on data governance, testing, risk management, human oversight and technical documentation, but how to address each of these remains a matter of debate and recalibration of existing design practice. Each of these areas presents implications for HCI, some perhaps more obviously than others. As brief background, there has been exploration over the years of the relationship between HCI and technology governance. Lazar and Hochheiser [37] have highlighted that HCI practitioners can play a significant role by engaging in and responding to public policy issues e.g., around accessibility. Interaction has been particularly pronounced between law and HCI around privacy and security [4; 41]. For example, improving citizen consumer and privacy interests has been a key tenet of the usable privacy and security research community. This has included addressing shortcomings in design of privacy policies, improving usability of password management, and developing privacy engineering methodologies to build privacy safeguards directly into systems [38; 69]. Legal calls to implement privacy by design have also prompted HCI researchers to consider socio-technical dimensions of building for privacy and creating design tools to do so, like ideation cards [45] or speculative design processes [83].

In responding to **data quality requirements**, the longstanding interest in HCI in information management [52] could be key, alongside growing research in Human Data Interaction (HDI) which seeks to put people at the centre of the flow of data in society [49] through designing for the concepts of agency, legibility and negotiability. The role of HDI in human-centred data governance is, as yet, under-explored particularly in the context of federated machine learning systems [34, 50], but could support goals of the AIA. Similarly, ‘research through design’ artefacts and provocations could help creatively explore both value tensions between data quality requirements and current AI data pipelines, to help consider what more human-centred, legally compliant AI systems might look like.

In turning to testing, while research has investigated how adversarial **testing** can be used to disrupt human-computer action [32], the role of HCI in adversarial testing is less pronounced though the possibility clearly exists. As Shneiderman [66] points out, a key part of the testing process, particularly the verification process, is to “develop test scenarios to detect adversarial attacks, which would prevent malicious use.” Given this, it would seem evident too that HCI may have a significant role to play in **risk management**, both prospectively and retrospectively. Prospectively, there is long tradition of risk assessment in the human factors literature spanning a broad range of domains, which includes methods for identifying and assessing risks as part of the process of usability evaluation prior to systems deployment [e.g., 10, 40, 59, 81]. The prospective identification and assessment of risk in AI systems would appear to be a strong area for continued HCI research and methodological development. Retrospectively, there is also a long tradition of studying systems deployed in the wild through fieldwork and this approach may prove useful in post-market monitoring of HRAIS. For example, a recent study of the deployment of AI to detect diabetic retinopathy [3] revealed the impact of workflow, infrastructure, and culture on the performance of the system and its ability to actually detect diabetic retinopathy. HCI may, then, play a

formative role in managing the risks that foreseeably accompany the development of AI systems and those that otherwise emerge after their introduction into actual settings of use.¹³

There may also be a role for HCI to play in the **technical documentation** of AI systems. This may not seem particularly obvious or attractive. However, the issue here concerns the role played by technical documentation in regulation, which is to demonstrate the compliance of a system with the AIA. The demonstration particularly concerns Articles 8 to 15 which include data governance, risk management, human oversight, and cybersecurity and is a legal obligation which, similar to the *accountability* principle enshrined in GDPR (Article 5(2)), is key to showing that a provider is in compliance with the law. Declaration is not sufficient; a provider must be able to demonstrate compliance on demand, even if they self-certify. Previous work at the interface between law, technology and HCI has explored how accountability might take different forms, being demonstrated, for example, through a human-centred architecture that communicates features of system design and how they map onto different GDPR legal requirements [73]. Accountability in GDPR covers various cohorts including, at least the data subject and supervisory authorities but could be broader e.g., between data processors and controllers. The AIA focuses on accountability to supervisory authorities (notified bodies and regulators) who, like the provider, have a technical interest in determining the compliance of HRAIS. There may, then, be a significant role for HCI that extends beyond literal documentation to shaping the architectures of human-centred AI systems in collaboration with developers and technology law scholars to enable novel demonstrations of compliance. There will be no one-size fits all human-centred AI architecture but many and this would be a of substantial strand research in its own right, cutting to the heart of what it is to *make* AI human-centred across different classes of autonomous system embedded within diverse application domains subject to various regulations.

We save **human oversight** until last as it is the most obvious point of connection between AI, its regulation, and HCI. The requirement that HRAIS “be designed and developed to ensure that their operation is sufficiently transparent to enable users to interpret the system’s output and use it appropriately” (AIA, Article 13) aligns strongly, on the face of it, with ethical concerns with explicability and the perceived need for transparency in AI decision-making to engender trust. This bleeds into two closely related technical agendas: interpretability in machine learning (IML) [44] and explainable AI (XAI) [72]. If there is a difference between the two it might be that IML has a broader focus than XAI. IML is of longstanding concern, reaching back at least 30 years and spans various “objectives” including understanding the causal nature of automated decision-making, the transferability and generalisation of models and outputs, their ability to inform human decision-making and assessing whether or not automated decisions are fair as well as trustworthy [44]. XAI, on the other hand, emerged more recently and does not seem to cover quite the same ground by focusing, for example, on “the machine’s current inability to explain their decisions and actions to human users” [72]. Transparency, like interpretability and explainability, is multi-faceted and implicates models, their individual components, and algorithms [44]. As Abdul et al. [1] note in their extensive literature review of papers on explainability,

“This has produced a myriad of algorithmic and mathematical methods to explain the inner workings of machine learning models ... However, despite their mathematical rigor, these works suffer from a lack of usability, practical interpretability and efficacy on real users.”

The challenge of ensuring that the operation of AI systems is sufficiently transparent to enable users to interpret their outputs and use them appropriately is underscored by other researchers too, who found in a recent study of machine learning deployments that explainability methods are, for the most part, only used by engineers to debug machine learning models,

¹³ Revisions to the EU Machinery Directive 2006 [25] also extend HCI’s role in managing risks to include those created by collaborative robots and driverless mobile machinery, which chimes strongly with HCI’s ongoing interest in human-robot interaction [58].

improve their performance, and to sanity check their work [5]. There is then, as Abdul et al. [1] point out, a pressing need for the HCI community to take the lead and ensure AI systems do not ignore the “human side” of AI and thus ensure that systems are explainable to their users as well as those who develop and maintain them, as will soon be required by law. Before turning to further points around the human side of AI, table 2 summarises some of the opportunities for HCI from the AIA design and development requirements.

Table 2

HRAIS Requirement	Opportunities for HCI
Designing and developing data sets in compliance with quality criteria and data governance practices.	Understand how to increase user agency, mediation and control over flows of data in HRAIS through more usable interface design. This could include exploring the role of human-data interaction [49] in complex federated machine learning systems. Reimagining more legally compliant AI data supply chains through design research.
Testing AI systems for accuracy , consistent performance with respect to intended use , and the identification of risks .	Explore the socio-technical risks, harms, and wider implications users face from interacting with AI systems to ensure laws reflect lived experiences of different groups. This can help build evidence for evaluating if assumptions in the law, and the design responses they require, are indeed the best solutions to address user needs. If not, then this could generate insights to enable alternative solutions and approaches to be iterated. Develop scenarios for adversarial testing, the prospective identification and assessment of risk through usability testing, and retrospective monitoring of risk through post-deployment field studies.
Novel forms of technical documentation supporting accountability and the demonstration of compliance.	Consider interactive and creative alternatives to purely textual documentation of the system operation. Research could explore how to demonstrate capabilities and compliance of AI systems with the AIA in more contextually situated ways e.g. exploring when and how to present notifications about system behaviour or reflecting on how AI system affordances could mediate new forms of human-AI communication. HCI could also provide insights into divergences between how these systems are formally documented, and how they are operated and the practices that emerge around their use ‘in the wild’.
Human Machine Interface Tools	Explore how to design and develop human-machine interface tools that go beyond focusing on interpretation and explanation of intelligent interfaces. Instead, consider the role of collaborative decision-making solutions developed through contextually sensitive, human-centred design methods and processes that enable users to oversee systems in use and take action to prevent or minimise risks to health, safety or fundamental rights. There are significant gaps in qualitative understanding of how AI systems operate in the organisations and contexts covered by the AIA, hence ethnographic studies would be invaluable here.

5.1 The Human Side of AI

If HCI is to take a lead and help AI meet its legal design and development obligations, it is important that attention is paid to the fine detail of the proposed regulation. The explainable AI agenda stemmed in part, for example, from the technical community’s (mis)reading of GDPR, which was seen to mandate a “right to an explanation” of automated decision-making using personal data [33]. This led to debates about how to make automated decision-making explainable to citizens and spurred the growth of new techniques, including non-machine based techniques such as counterfactuals [79], in a bid to help people understand how AI systems arrive at decisions. Despite an active technical agenda in XAI, legal debate has continued to discuss *if* a right to an explanation exists or not in GDPR, its breadth, on what legal basis it might exist, and if it is even the best remedy to protect user interests and build trust in autonomous systems [78, 63, 19, 64, 16, 15]. As these discussions show, the law is interpretative and rarely provides clean rules that tell us what the ‘correct’ response from

designers should be. The shift to ‘by design’ regulation is an attempt to bridge this gap between legal rules and design practice, though we still need to be sensitive to the fine print. It thus becomes apparent that the AIA situates transparency and interpretation *in* human oversight:

“High-risk AI systems shall be designed and developed in such a way, including with appropriate human-machine interface tools, that they can be effectively overseen by natural persons during the period in which the AI system is in use.” (AIA, Article 14)

It is important to recognise the situated nature of transparency and interpretation. It contextualises the meaning of the Article 13 transparency and interpretation requirement and what the raft of measures it specifies should enable. Seen from within the context of human oversight, explanation is not framed as a guarantor of trust then, but is part of a package of human-machine interface measures that are to be designed and developed to support *action*. It is not action (or interaction) with the machine per se, but action in the context of machine use. Thus, what the AIA requires by way of transparency and interpretation is that human interface tools support human oversight and thus enable human operators to take action to prevent or minimise risks to human health, safety or fundamental rights occasioned either by intentional use or reasonably foreseeable misuse of a high-risk AI system. **Action not explanation is imperative**; transparency and interpretation are subservient to this goal.

The transparency and interpretability requirement of the AIA should not be taken at face value then and understood to segue neatly with ethical principles, IML, or XAI methods, or the need to extend them beyond developers to actual users. On the contrary, the mandate is not an invitation to dive down the rabbit hole of interpretability, but to design and develop human-machine interface tools that *enable users to take action to prevent or minimise actual or potential harms*. Explanation has a role to play in this achievement, but there is much more to it than explanation can provide. While there is, as Abdul et al. [1] point out, both scope and need to help AI develop methods of interpretation and explanation that are usable by non-developers in the course of their practical activities, the legal requirement for human oversight opens up the design space. The emphasis on explanation might be complemented by the design of *intelligent user interfaces* [77] and human-machine interface tools that support *collaborative decision-making* [85]. These approaches have been leveraged to bridge between AI and HCI in the effort to support interpretation and explanation and build the human into the decision-making loop. Both approaches also face fundamental challenges which must be addressed if they are to get to grips with the human oversight requirement. New design frameworks are required to strike a balance between excessive automation and excessive human control [65], for example, and there is pressing need to address the fundamental challenge of aligning collaborative solutions with situated human reasoning [9]. However, both the legal and practical need to build interfaces usable by human operators into AI systems speaks to further HCI expertise.

Of particular relevance is HCI expertise in understanding the context of use in which systems are deployed and shaping computing systems (and AI systems *are* just computing systems, despite the hype) around users to support situations of use.

“High-risk AI systems should be designed and developed in such a way ... [as to be] responsive to the human operator ... the natural persons to whom human oversight has been assigned [should] have the necessary competence, training and authority to carry out that role.” (AIA, Recital 48)

This recital, which provides interpretative guidance, may be read literally as placing certain obligations on human operators or it may be understood to reflect the fundamental requirement that human-machine interface tools should be designed and developed to be responsive to user competence and training. Not competence in using an AI system per se, though this is

clearly part of the mix (AIA, Article 29), but competence in using the system for some *purpose*. This recognises that AI systems have an indexical relationship to and will thus be embedded in organised settings of use. They therefore need to be designed and developed to support human oversight with respect to those organised settings of use and what persons therein will *do* with them. As Harper puts it in describing the role of HCI in the age of AI, “*doings with AI* are more important than the *mechanics of AI*” [35].

Designing human-machine interface tools to support situated “doings with AI” (including doing oversight) gives rise to the need to develop what Harper calls “new grammars of action” to mediate human-machine interaction (*ibid.*). It is an idea that has previously been used to characterise interdisciplinary design approaches that leverage Wittgensteinian thinking and orients HCI researchers to ordinary language concepts as resources for the study of organised settings and the design of technologies to support the everyday activities which articulate those concepts [e.g., 71, 14, 20]. In more prosaic terms, understanding and designing for grammars of action involves ethnographic study and cooperative prototyping, both of which are commonplace in HCI along with an arsenal of other methods for understanding users, the contexts in which they operate, their organisational needs and for tailoring technological support accordingly through the active participation of users in the design process. The need to take account of users, the contexts in which they operate and their organisational needs in design and development is also writ large in the AIA:

“In eliminating or reducing risks related to the use of the high-risk AI system, due consideration shall be given to the technical knowledge, experience, education, training to be expected by the user and the environment in which the system is intended to be used.” (AIA, Article 9)

The AIA also seeks to “encourage and facilitate” the uptake of voluntary codes of conduct to foster “stakeholders participation in the design and development of the AI systems” along with sustainability and accessibility requirements (AIA, Article 61). Whether mandated by name or not, the human oversight requirement reaches beyond a technical concern with interpretation and explanation to situate design and development in broader **human-centered processes and methods** that are the hallmark of HCI. It is not sufficient to design and develop interfaces that instruct users how to use AI systems, those interfaces must fit in with and support the organised ‘doings’ the user and the system are embedded in and enable users to take action to prevent or minimise harm to human health, safety, or fundamental rights in that context. The requirement to contextualise AI design and development – to give due consideration to the user and their environment – opens up AI to HCI and its wealth of experience to ensure AI systems are both practically and legally fit for purpose.

At the current moment in time, it appears as though the users of HRAIS are envisaged by the AIA to be persons acting in a professional capacity only.

“For the purpose of this Regulation, the following definitions apply ‘user’ means any natural or legal person, public authority, agency or other body using an AI system under its authority, except where the AI system is used in the course of a personal non-professional activity.” (AIA, Article 3)

While this is not an impediment to HCI, indeed it very much plays to its strengths in studying work, organisations and computer-supported cooperative work, it has led technology law scholars to suggest the term ‘deployers’ should be used instead of users [18]. More importantly, the lack of reference to non-professional users of HRAIS leaves a gap in the scope of the AIA requiring oversight by consumers or citizens [8]. Indeed, it has been argued this group have no rights in the AIA and lack means to engage with approval processes or to influence whether or not products should be allowed onto the market in the first place [18]. It is worth noting that unlike GDPR, which confers fundamental rights on citizens, the AIA mainly focuses on product safety requirements across the supply chain and mediating market access. It is not a matter of

argument that the AIA does not confer rights on citizens, it simply does not. Whilst this is a gap, many HRAIS systems will likely not only be subject to the AIA but also other regulatory regimes which do provide more direct citizen protections such as equality, human rights, data protection, and consumer protection laws. Furthermore, there is a strong and enduring argument to be made for including citizens in *responsible innovation* [70] processes to ensure the design and development of AI systems (high-risk or not) that are socially desirable as well as legally trustworthy, and this too plays to HCI expertise in co-design.

6 CONCLUSIONS

The law is often aligned with the encoding of rules in design [42], which has given rise to the idea that “code is law” or that it is possible to effect regulation through programming [47] and the “architecture of the internet” [57]. This is not the only possible relationship, however. The law is not simply a body of rules, but much more importantly a lively interpretive framework in which rules are invoked and their applicability contested. There is more than one possible relationship between law and design and we have sought to unpack here how proposed regulation of AI opens up possibilities for HCI that go beyond established concerns with the ethics of AI and explanation and what it means to design for the ‘human in the loop’ from a legal perspective. We have thus presented a number of legal provocations from the law, based on common mandatory design and development requirements by the proposed European AI Act (AIA) for high-risk AI systems (HRAIS) that present a risk to human health, safety, and fundamental rights (such as privacy).

We are aware of efforts to align HCI and AI [e.g., 82, 35, 65], that the relationship between the two fields is being explored through dedicated conferences including, for example, the International Conference on AI in HCI and the Annual Conference on Intelligent User Interfaces, and that HCI & AI courses proliferate. However, in articulating the various possibilities for HCI in the design of AI systems, this paper should not be read as sitting at the intersection between HCI and AI. We are not treading familiar ground even though we have hopefully touched upon it in places. Distinctively, our interests lie in a different direction at the *intersection between HCI and the regulation of AI*. We have thus sought to consider how new rules put HCI at the front and centre of regulation, particularly with respect to human oversight and the need to situate the design and development of human-machine interface tools in context, to enable *design for human action*.

HCI has a lead role to play in the design and development of trustworthy AI systems that prevent or minimise risks to human health, safety, and fundamental rights and it is in this respect that two further provocations are occasioned by the legal imposition of design and development requirements. First, while it is clear that regulation anticipates a key role for human-centred design, AI systems development is currently dominated by ethical and technical agendas that shape the design process. This presents a fundamental challenge as to how legally mandated design and development requirements are to be *communicated within the design process* to shape trustworthy AI. This might be addressed through interdisciplinary collaboration between HCI researchers and legal scholars [e.g., 45]. Second, and perhaps more challenging, a further provocation lies in questioning the assumption that compliance with design and development requirements will actually produce trustworthy autonomous systems, as the functional steps specified in the AIA do not necessarily make significant contact with fundamental rights [28]. While health and safety requirements are covered by legislation relevant to a specific AI product, the possibility exists that a provider might put extensive resources into building a system that complies with design and development requirements but contravenes wider fundamental rights and is thus, by definition, untrustworthy. The design and development requirements of the AIA mandate what are basically good engineering practices around high-risk AI systems. They do not translate what designers need to do to comply with the largely non-functional requirements of fundamental rights and what steps developers need to take to protect them. This

opens up the question of the kind of tools needed to ground the design and development of trustworthy autonomous systems in fundamental rights and the potential role of HCI in *assessing the impact of AI systems on fundamental rights*.

Seen from a regulatory perspective, and as elaborated in our consideration of legal provocations furnished by the AIA, HCI has a central role to play in ensuring that AI can comply with mandatory design and development requirements. These legal requirements make it clear that there is a lot more to the “human side” of AI than making methods of interpretation and explanation usable by non-developers. The legal provocations presented here not only project a potential role for HCI in the technical heartlands of AI and novel explorations of the accountability requirements mandated by technical documentation, but a critical role in developing human-machine interface tools to support human oversight. This is not only a matter of enabling users to interpret and understand how the machine works but also, and distinctively, of enabling users to take the situationally-relevant and contextually-appropriate *actions* needed to accomplish the organised ‘doings’ they are tasked with and to prevent or minimise risks to human health, safety, and fundamental rights in the process. The requirement that due consideration be given to the user and their environment creates new opportunities for HCI in AI systems development and mandates a fundamental turn to human-centred design methods and processes. The turn to HCI also projects further key roles for the community including developing means to communicate mandatory design and development requirements within the design process and assessing the impact of AI systems on fundamental human rights. Ultimately, the introduction of ‘by design’ regulation presents significant challenges to the providers of AI systems to ensure they are engineered around users and the environments in which they operate and that human-machine interface tools allow them to oversee the risks occasioned by the use of AI systems therein. In doing so, the regulation of AI opens up its design and development to interdisciplinary challenges that turn upon HCI’s unique expertise. Thus, and to borrow once more from Harper [35], we might say “*the future is not AI; it can only be an AI enabled through HCI.*”

ACKNOWLEDGMENTS

This work was supported by the UK Engineering and Physical Sciences Research Council [grant numbers EP/V026607/1, EP/T022493/1] and UK Economic and Social Research Council [grant number ES/T00696X/1].

Data Access Statement. No research data is associated with this paper.

REFERENCES

- [1] Ashraf Abdul, Jo Vermeulen, Danding Wang, Brian Lim, Mohan Kankanhalli. 2018. Trends and Trajectories for Explainable, Accountable and Intelligent Systems. In Proceedings of the ACM Conference on Human Factors in Computing Systems (CHI), ACM, Montreal, Canada, 1-18. <https://doi.org/10.1145/3173574.3174156>
- [2] Shaowen Bardzell, Jeffrey Bardzell, Jodi Forlizzi, John Zimmerman, and John Antanitis. 2012. Critical design and critical theory: the challenge of designing for provocation. In Proceedings of the Designing Interactive Systems Conference (DIS '12). Association for Computing Machinery, New York, NY, USA, 288–297. <https://doi.org/10.1145/2317956.2318001>
- [3] Emma Beede, Elizabeth Baylor, Fred Hersch, Anna Iurchenko, Lauren Wilcox, Paisan Ruamviboonsuk, and Laura Vardoukakis. 2020. A human-centred evaluation of a deep learning system deployed in clinics for the detection of diabetic retinopathy. In Proceedings of the ACM Conference on Human Factors in Computing Systems (CHI), ACM, Honolulu, HI, USA. <https://doi.org/10.1145/3313831.3376718>
- [4] Victoria Bellotti, and Abigail Sellen. 1993. Design for Privacy in Ubiquitous Computing Environments. In: de Michelis, G., Simone, C., Schmidt, K. (eds) Proceedings of the Third European Conference on Computer-Supported Cooperative Work 13–17 September 1993, Milan, Italy ECSCW '93. Springer, Dordrecht. https://doi.org/10.1007/978-94-011-2094-4_6
- [5] Umang Bhatt, Alice Xiang, Shubham Sharma, Adrian Weller, Ankur Taly, Yunhan Jia, Joydeep Ghosh, Ruchir Puri, José Moura, and Peter Eckersley. 2020. Explainable machine learning in deployment. In Proceedings of the Conference on Fairness, Accountability, and Transparency (FAT), ACM, Barcelona, Spain, 27–30. <https://doi.org/10.1145/3351095.3375624>
- [6] Abeba Birhane, Elayne Ruane, Thomas Laurent, Matthew S. Brown, Johnathan Flowers, Anthony Ventresque, and Christopher L. Dancy. 2022. The Forgotten Margins of AI Ethics. In 2022 ACM Conference on Fairness, Accountability, and Transparency (FAccT '22). Association for Computing Machinery, New York, NY, USA, 948–958. <https://doi.org/10.1145/3531146.3533157>
- [7] Barry Brown, Alexandra Weilenmann, Donald McMillan, and Airi Lampinen. 2016. Five Provocations for Ethical HCI Research. In Proceedings of

- the 2016 CHI Conference on Human Factors in Computing Systems (CHI '16). Association for Computing Machinery, New York, NY, USA, 852–863. <https://doi.org/10.1145/2858036.2858313>
- [8] Bureau Européen des Unions de Consommateurs. 2021. Protecting European Consumer in the World of Connected Devices. Retrieved March 15, 2022 from https://www.beuc.eu/publications/beuc-x-2021-091_protecting_european_consumers_in_the_world_of_connected_devices.pdf
 - [9] Carrie Cai, Samantha Winter, David Steiner, Lauren Wilcox and Michael Terry. 2019. “Hello AI” – Uncovering the onboarding needs of medical practitioners for human-AI collaborative decision-making. *Human-Computer Interaction* 3, Issue CSCW (November 2019) Article No. 104. <https://doi.org/10.1145/3359206>
 - [10] Stuart Card, Allen Newell, and Thomas Moran. 1983. *The Psychology of Human-Computer Interaction*. Lawrence Erlbaum Associates Inc., Mahwah, NJ, USA.
 - [11] Aurelia Ciachhi and Dietmar von der Pforden. 2011. Exploring the Relationship between Law and Governance: A Proposal. Retrieved March 15, 2022 from <https://www.rug.nl/rechten/congressen/archief/2011/governancemeetslaw/workingpapers/papercolombiacchipfordten.pdf>
 - [12] James Coleman. 1990. *Foundations of Social Theory*. Belknap Press, Cambridge, MA, USA.
 - [13] Council of Europe. 2022. 12 Principles of Good Governance. Retrieved 15 March, 2022 from <https://www.coe.int/en/web/good-governance/12-principles>
 - [14] Andy Crabtree. 2000. Talking work: language games, organisations and computer supported cooperative work. *Computer Supported Cooperative work: The Journal of Collaborative Computing* 9, 2 (May 2000) 215-237.
 - [15] Andy Crabtree, Lachlan Urquhart and Jiahong Chen. 2019. Right to an explanation considered harmful. Social Science Research Network. Retrieved 15 March, 2022 from <https://ssrn.com/abstract=3384790>
 - [16] Finale Doshi-Velez, Mason Kortz, Ryan Budish, Chris Bavitz, Sam Gershman, David O’Brien, Kate Scott, Stuart Shieber, James Waldo, David Weinberger, Adrian Weller, Alexandra Wood. 2019. Accountability of AI under the law: the role of explanation. arXiv. Retrieved 15 March, 2022 from <https://arxiv.org/abs/1711.01134>.
 - [17] Antony Dunne and Fiona Raby. 2001. *Design noir: the secret life of electronic objects*. Birkhäuser, Berlin
 - [18] Lilian Edwards. 2022. Regulating AI in Europe: Four Problems and Four Solutions. Retrieved March 15, 2022 from <https://www.adalovelaceinstitute.org/wp-content/uploads/2022/03/Expert-opinion-Lilian-Edwards-Regulating-AI-in-Europe.pdf>
 - [19] Lilian Edwards and Michael Veale. 2017. Slave to the algorithm? Why a ‘right to an explanation’ is probably not the remedy you are looking for. *Duke Law & Technology Review* 16 1 (March 2017) 18–84. <https://scholarship.law.duke.edu/cgi/viewcontent.cgi?article=1315&context=dltr>
 - [20] Per Ehn. 1988. Language-games: a Wittgensteinian alternative. In *Work-Oriented Design of Computer Artefacts*, 103-122. Arbetslivscentrum, Stockholm, Sweden.
 - [21] European Commission. 2016. Regulation (EU) 2016/679 of the European Parliament and of the Council of 27 April 2016 on the Protection of Natural Persons with Regard to the Processing of Personal Data and on the Free Movement of such Data, and Repealing Directive 95/46/EC (General Data Protection Regulation). *Official Journal of the European Union* 119 (May 2016), 1-88. <https://eur-lex.europa.eu/eli/reg/2016/679/oj>
 - [22] European Commission. 2020. White Paper on Artificial Intelligence: A European Approach to Excellence and Trust. Retrieved 15 March, from https://ec.europa.eu/info/sites/default/files/commission-white-paper-artificial-intelligence-feb2020_en.pdf
 - [23] European Commission. 2021. Proposal for a Regulation of the European Parliament and of the Council Laying Down Harmonised Rules on Artificial Intelligence (Artificial Intelligence Act) and Amending Certain Union Legislative Acts. Retrieved 15 March, 2022 from <https://eur-lex.europa.eu/legal-content/EN/TXT/?qid=1623335154975&uri=CELEX%3A52021PC0206>
 - [24] European Commission. 2021. Annexes to the Proposal for a Regulation of the European Parliament and of the Council Laying Down Harmonised Rules on Artificial Intelligence (Artificial Intelligence Act) and Amending Certain Union Legislative Acts. Retrieved 15 March, 2022 from <https://eur-lex.europa.eu/legal-content/EN/TXT/?qid=1623335154975&uri=CELEX%3A52021PC0206>
 - [25] European Commission. 2021. Proposal for a Regulation of the European Parliament and of the Council on Machinery Products. Retrieved March 15, 2022 from <https://ec.europa.eu/docsroom/documents/45508/attachments/2/translations/en/renditions/native>
 - [26] European Commission. Nando (New Approach Notified and Designated Organisations) Information System. Retrieved March 15, 2022 from <https://ec.europa.eu/growth/tools-databases/nando>
 - [27] European Data Protection Supervisor. 2021. EDPB & EDPS call for ban on use of AI for automated recognition of human features in publicly accessible spaces, and some other uses of AI that can lead to unfair discrimination. Retrieved March 15, 2022 from https://edps.europa.eu/press-publications/press-news/press-releases/2021/edpb-edps-call-ban-use-ai-automated-recognition_en
 - [28] European Union Agency for Fundamental Human Rights. 2020. Getting the Future Right: Artificial Intelligence and Fundamental Rights. Retrieved 15 March, 2022 from https://fra.europa.eu/sites/default/files/fra_uploads/fra-2020-artificial-intelligence_en.pdf
 - [29] Batya Friedman and David Hendry. 2019. *Value Sensitive Design; Shaping Technology with Moral Imagination*. MIT Press.
 - [30] Diego Gambetta. 1988. Can We Trust Trust? In *Trust: Making and Breaking Cooperative Relations*, 213–237. Basil Blackwell, Oxford, UK.
 - [31] Bill Gates. 2002. Trustworthy Computing. Retrieved 15 March, 2022 from <https://www.wired.com/2002/01/bill-gates-trustworthy-computing>
 - [32] Yuan Gong and Christian Poellabauer. 2018. An Overview of Vulnerabilities of Voice Controlled Systems. arXiv. Retrieved March 15, 2022 from <https://arxiv.org/abs/1803.09156>
 - [33] Bryce Goodman and Seth Flaxman. 2016. EU regulations on algorithmic decision-making and “a right to an explanation”. In *Proceedings of the ICML Workshop on Human Interpretability in Machine Learning (WHI)*, ICML, New York, USA, 26–30. <https://arxiv.org/pdf/1606.08813.pdf>
 - [34] Google AI Blog. 2017. Federated Learning: Collaborative Machine Learning without Centralized Training Data. Retrieved 15 March, 2022 from <https://ai.googleblog.com/2017/04/federated-learning-collaborative.html>

- [35] Richard Harper. 2019. The Role of HCI in the Age of HCI. *International Journal of Human-Computer Interaction* 35, 15 (June 2019) 1331-1344. <https://doi.org/10.1080/10447318.2019.1631527>
- [36] High-Level Expert Group on Artificial Intelligence. 2019. Ethics Guidelines for Trustworthy AI. Retrieved March 15, 2022 from https://ec.europa.eu/newsroom/dae/document.cfm?doc_id=60419
- [37] Harry Hochheiser & Jonathan Lazar. 2010. HCI and Societal Issues: A Framework for Engagement, *International Journal of Human Computer Interaction* 23(3), 339-374
- [38] Giovanni Iachello, and Jason Hong, 2007. End User Privacy in Human Computer Interaction, *Foundations and Trends in Human Computer Interaction* 1(1), 1-137
- [39] ISO 9001. Quality Management. Retrieved 15 March, 2002 from <https://www.iso.org/iso-9001-quality-management.html>
- [40] Barry Kirwan and L K Ainsworth. 1992. *A Guide to Task Analysis*. Taylor & Francis, London, UK.
- [41] Marc Langheinrich. 2001. Privacy by Design — Principles of Privacy-Aware Ubiquitous Systems. In: Abowd, G.D., Brumitt, B., Shafer, S. (eds) *UbiComp 2001: Ubiquitous Computing*. UbiComp 2001. Lecture Notes in Computer Science, vol 2201. Springer, Berlin, Heidelberg. https://doi.org/10.1007/3-540-45427-6_23
- [42] Ronald Leenes and Federica Lucivero. 2014. Laws on robots, laws by robots, laws in robots: regulating robot behaviour by design. *Law, Innovation and Technology* 6, 2, 193–220. <https://doi.org/10.5235/17579961.6.2.193>
- [43] David Leslie. 2019. Understanding Artificial Intelligence Ethics and Safety: A Guide for the Responsible Implementation of AI Systems in the Public Sector. Alan Turing Institute. Retrieved March 15, 2022 from https://www.turing.ac.uk/sites/default/files/2019-08/understanding_artificial_intelligence_ethics_and_safety.pdf
- [44] Zachary Lipton. 2016. The mythos of model interpretability. In *Proceedings of the ICML Workshop on Human Interpretability in Machine Learning (WHI)*, ICML, New York, USA, 96–100, 2016. <https://arxiv.org/pdf/1606.03490.pdf>
- [45] Ewa Luger, Lachlan Urquhart, Tom Rodden and Michael Golembewski. 2015. Playing the legal card: using ideation cards to raise data protection issues with the design process. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI)*, ACM, Seoul, South Korea, 457-466. <https://doi.org/10.1145/2702123.2702142>
- [46] Andrew McStay and Lachlan Urquhart. 2019. ‘This time with feeling?’ Assessing EU data governance implications of out of home appraisal based emotional AI. *First Monday* 24, 10 (October 2019). <https://doi.org/10.5210/fm.v24i10.9457>
- [47] Dens Merigoux, Nicolas Chataing and Jonathan Protzenko. 2021. Catala: A Programming Language for the Law. arXiv. Retrieved 15 March, 2022 from <https://arxiv.org/pdf/2103.03198.pdf>
- [48] Barbara Mitzal. 1992. The Notion of Trust in Social Theory. *Policy and Society* 5, 1 (December 1992) 6-15. <https://doi.org/10.1080/10349952.1992.11876774>
- [49] Richard Mortier, Hamed Haddadi, Tristan Henderson, Derek McAuley, Jon Crowcroft and Andy Crabtree. 2016. Human-Data Interaction. *Interaction Design Encyclopaedia*, Chapter 41. <https://www.interaction-design.org/literature/book/the-encyclopedia-of-human-computer-interaction-2nd-ed/human-data-interaction>
- [50] Richard Mortier, Hamed Haddadi, Sandria Servia and Liang Wang. 2021. Distributed data analytics. In *Privacy by Design for the Internet of Things: Building Accountability and Security* (eds. Crabtree, A., Mortier, R., and Haddadi, H.), Chapter 8. Institution of Engineering and Technology, London, UK. https://doi.org/10.1049/PBSE014E_ch8
- [51] Craig Mundie. 2002. Trustworthy Computing. Microsoft White Paper. Retrieved 15 March, 2022 from http://download.microsoft.com/documents/australia/about/trustworthy_comp.doc
- [52] National Coordination Office for IT Research and Development. 2003. Human-Computer Interaction and Information Management Research Needs. Retrieved 15 March, 2022 from https://www.nitrd.gov/pubs/hci-im_research_needs_final.pdf
- [53] NCSC. 2019. Secure Design Principles. Retrieved March 15, 2002 from <https://www.ncsc.gov.uk/collection/cyber-security-design-principles>
- [54] Bart Nooteboom. 2007. Social capital, institutions and trust. *Review of Social Economy* 65, 1 (March 2007) 25-53. <https://doi.org/10.1080/0034676061132154>
- [55] Brian Randell. 1992. System dependability. *Lecture Notes in Computer Science*, vol 653. Springer, Heidelberg. https://doi.org/10.1007/3-540-56320-2_51
- [56] Dimitrios Raptis, Rikke Hagensby Jensen, Jesper Kjeldskov, and Mikael B. Skov. 2017. Aesthetic, Functional and Conceptual Provocation in Research Through Design. In *Proceedings of the 2017 Conference on Designing Interactive Systems (DIS '17)*. Association for Computing Machinery, New York, NY, USA, 29–41. <https://doi.org/10.1145/3064663.3064739>
- [57] Regulation by Code. Retrieved March 15, 2022 from <https://cs.stanford.edu/people/eroberts/cs181/projects/2010-11/CodeAndRegulation/sec2.html>
- [58] Fuji Ren and Yanwei Bao. 2020. A review of human-computer interaction and intelligent robots. *International Journal of Information Technology and Decision Making* 19, 1 (February 2020) 5-47. <https://doi.org/10.1142/S0219622019300052>
- [59] Mary Beth Rosson and John Carroll. 2002. *Usability Engineering: Scenario-Based Development of Human Computer Interaction*. Morgan Kaufmann Publishers, San Francisco, CA, USA.
- [60] Royal Society. 2017. Data Management and Use: Governance in the 21st Century. Retrieved 15 March, 2002 from <https://royalsociety.org/-/media/policy/projects/data-governance/data-management-governance.pdf>
- [61] Royal Society. 2019. Explainable AI: The Basics. Retrieved March 15, 2022 from <https://royalsociety.org/-/media/policy/projects/explainable-ai/AI-and-interpretability-policy-briefing.pdf>
- [62] Royal Society & National Academy of Sciences. 2019. AI: An International Dialogue. Retrieved 15 March, 2022 from <https://royalsociety.org/>

/media/policy/topics/open-science-data/ai-international-dialogue.pdf?la=en-GB&hash=A888136F57440CA9833F3DA79D683320

- [63] Andrew Selbst and Julia Powles. 2017. Meaningful information and the right to an explanation. *International Data Privacy Law* 7, 4 (November 2017) 233–242.
- [64] Andrew Selbst and Solon Barocas. 2018. The intuitive appeal of explainable machines. *Fordham Law Review* 87, 3 (Autumn 2018) 1085–1139. <https://ir.lawnet.fordham.edu/cgi/viewcontent.cgi?article=5569&context=flr>
- [65] Ben Shneiderman. 2020. Human-centred artificial intelligence: reliable, safe and trustworthy. *International Journal of Human-Computer Interaction* 36, 6 (March 2020) 495-504. <https://doi.org/10.1080/10447318.2020.1741118>
- [66] Ben Shneiderman. 2020. Bridging the gap between ethics and practice: guidelines for reliable, safe and trustworthy human-centred AI systems. *ACM Transactions on Interactive Intelligent Systems* 10, 40 (October 2020) article 26. <https://doi.org/10.1145/3419764>
- [67] Kris Shrishak. 2022. Technical Errors in the AI Act. Irish Council for Civil Liberties. Retrieved 15 March, 2022 from <https://www.iccl.ie/digital-data/technical-errors-in-the-ai-act>
- [68] Brian Skyrms. 2008. Trust, risk and the social contract. *Synthese* 160, 1 (January 2008) 21-25. <https://doi.org/10.1007/s11229-006-9075-3>
- [69] Sarah Spiekerman, & Lorrie Faith Cranor. 2009. Engineering Privacy. *IEEE Transactions on Software Engineering* 35 (1), 67-82
- [70] Jack Stilgoe, Richard Owen and Phil Macnaghten. 2013. Developing a framework for responsible innovation. *Research Policy* 42, 9 (November 2013) 1568-1580. <http://dx.doi.org/10.1016/j.respol.2013.05.008>
- [71] Peter Tolmie. 2003. “Everyday Intimacy” – The Social Organisation of an Ascriptive Device. PhD Thesis, Sociology Department, Lancaster University.
- [72] Matt Turek. 2018. Explainable Artificial Intelligence (XAI). DARPA. Retrieved 15 March, 2022 from <https://www.darpa.mil/program/explainable-artificial-intelligence>
- [73] Lachlan Urquhart, Tom Lodge and Andy Crabtree. 2019. Demonstrably doing accountability in the Internet of Things. *International Journal of Law and Information Technology* 27, 1 (Spring 2019) 1-27. <https://doi.org/10.1093/ijlit/eay015>
- [74] Lachlan Urquhart and Diana Miranda (2021) Policing Faces: The Present and Future of Intelligent Facial Surveillance. *Information and Communications Technology Law*. Published online 28 October 2021. <https://doi.org/10.1080/13600834.2021.1994220>
- [75] Michael Veale. 2020. A critical take on the policy recommendations of the EU High-Level Expert Group on Artificial Intelligence. *European Journal of Risk Regulation* 11, 1 (March 2020) E1. <https://doi.org/10.1017/err.2019.65>
- [76] John Vines, 2018. Playing with Provocations. In: Blythe, M., Monk, A. (eds) *Funology 2. Human-Computer Interaction Series*. Springer, Cham. https://doi.org/10.1007/978-3-319-68213-6_8
- [77] Sarah Theres Völkel, Christina Schneegass, Malin Eiband, and Daniel Busche. 2020. What is “intelligent” in intelligent user interfaces? A meta-analysis of 25 years of UIU. In *Proceedings of the 25th International Conference on Intelligent User Interfaces (IUI)*, ACM, Caliri, Italy, 477-487. <https://doi.org/10.1145/3377325.3377500>
- [78] Sandra Wachter, Brent Mittelstadt and Luciano Floridi. 2017. Why a Right to Explanation of Automated Decision-making Does Not Exist in the General Data Protection Regulation. *International Data Privacy Law* 7, 2 (May 2017), 76–99.
- [79] Sandra Wachter, Brent Mittelstadt and Chris Russell. 2018. Counterfactual explanations without opening the black box: automated decisions and the GDPR. *Harvard Journal of Law and Technology* 31, 2 (Spring 2018) 842-887. <https://jolt.law.harvard.edu/assets/articlePDFs/v31/Counterfactual-Explanations-without-Opening-the-Black-Box-Sandra-Wachter-et-al.pdf>
- [80] Ben Wagner. 2018. Ethics as an escape from regulation: from ethics-washing to ethics-shopping. In *Being Profiled: Cogitas Ergo Sum* (ed. Hilderbrandt, M.), 84-89. Amsterdam University Press, The Netherlands.
- [81] Colin Watson and Ahmed Kharrufa. 2021. HCI – H is also for hazard: using HAZOP to identify undesirable consequences in socio-technical systems. In *Proceedings of the ACM SIGCAS Conference on Computing and Sustainable Societies (COMPASS)*, ACM, Virtual Event, 230-242. <https://doi.org/10.1145/3460112.3471959>
- [82] Terry Winograd. 2006. Shifting viewpoints: artificial intelligence and human-computer interaction. *Artificial Intelligence* 170 (November 2006) 1256-1258. <https://doi.org/10.1016/j.artint.2006.10.011>
- [83] Richmond Wong and Deirdre Mulligan. 2019. CHI '19: Bringing Design to the Privacy Table: Broadening “Design” in “Privacy by Design” Through the Lens of HCI. *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems* 1–17 <https://doi.org/10.1145/3290605.3300492>
- [84] Qixue Xiao, Kang Li, Deyue Zhang, and Weilin Xu. 2018. Security Risks in Deep Learning Implementations. In *Proceedings of the 2018 Symposium on Security and Privacy Workshops (SPW)*, IEEE, San Francisco, CA, 123-128. <https://doi.org/10.1109/SPW.2018.00027>
- [85] Pascale Zaraté. 2013. *Tools for Collaborative Decision-Making*. Wiley, London.