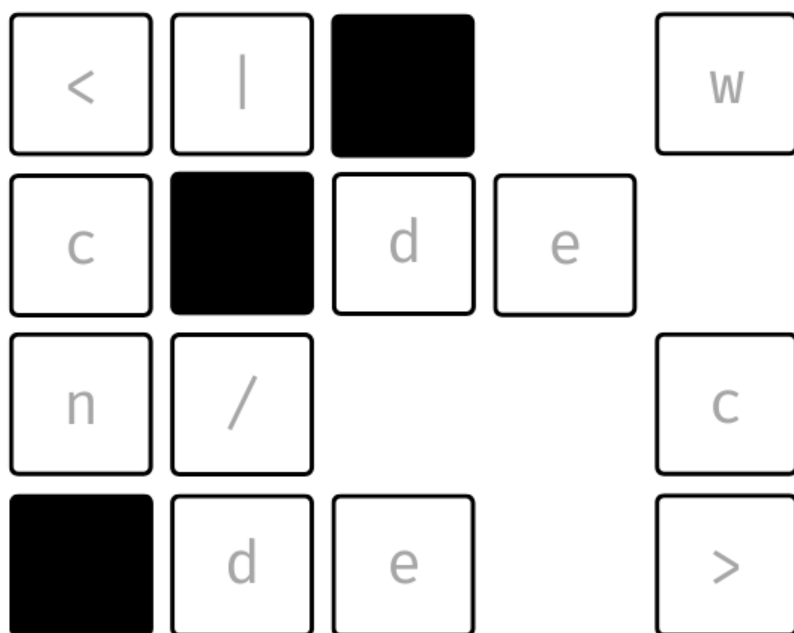


Samuela Marchiori

MIND

THE



GAP

Enhancing the ethical regulation
of low-code/no-code AI platforms

Mind the gap. Enhancing the ethical regulation of low-code/no-code AI platforms

Samuela Marchiori

*A thesis submitted in partial fulfilment of the requirements for the degree of
Master's of Science in Philosophy of Science, Technology and Society*

Supervisor: Prof. dr. P. A. E. Brey

Second reader: Dr. D. Babushkina

Word count: 23039

University of Twente

August 2022

To Cristine,

for encouraging me to take the leap.

To my family, and to Andrea,

*for trusting that I had a good reason to jump,
and supporting me all the way.*

Acknowledgements

Two years ago, I tiptoed my way into philosophy and into UT.

Today, I would like to express my gratitude to those who welcomed me into the former and make it so hard to say goodbye to the latter.

First and foremost, to Philip, for pushing me to be a better researcher and leading by example. For helping me open doors I didn't know existed and setting impossibly high standards for any future supervisors.

To Dina, for the expert guidance, the fierce criticism, and the unwavering support.

To Yashar, Maren, YJ, Peter, Pim, Annemarie, Jeroen, Peter-Paul, and Dr. Dainow, for the stimulating conversations and the generous guidance along the way.

To my friends, old and new, for the support.

Thank you.

Sam

Cover design: Fawaz Ahmed Sha

Table of contents

Summary	5
1. Introduction	6
1.1 Problem statement	8
1.2 Intended research methodology	9
1.3 Structure of the thesis	12
2. Low-code/no-code AI platforms. Technical features and applications	14
2.1 Artificial intelligence	14
2.2 Cloud computing	16
2.3 Low-code/no-code development	17
2.4 Applications	21
2.5 Conclusion	22
3. Ethical issues	24
3.1 Ethical issues raised by low-code/no-code AI platforms	24
3.2 Transparency	27
3.3 Bias and discrimination	31
3.4 Responsibility	34
3.5 Conclusion	37
4. The EU regulatory landscape	39
4.1 Charter of Fundamental Rights of the European Union	40
4.2 AI Act	41
4.3 General Data Protection Regulation	45
4.4 Ethics Guidelines for Trustworthy AI	47
4.5 Conclusion	51
5. Ethical regulation of low-code/no-code AI platforms	52
5.1 The benefits of legal regulatory tools	52
5.2 Filling the regulatory gap	56
5.3 Anticipatory regulation	60
5.4 Conclusion	62
6. Conclusion	64
Appendix 1	67
Bibliography	71

Summary

Low-code/no-code AI platforms allow virtually anyone with a computer and an internet connection to develop AI systems autonomously in a fast, easy, and inexpensive way, without the need for expert human supervision. This results in AI systems that are likely to give rise to a wide range of ethical issues but are not routinely checked for ethical shortcomings before being implemented. This is concerning in that it effectively delegates ethically charged development choices to individuals who may not have the necessary skill set to grasp their significance.

My aim in this thesis is twofold. On the one hand, I avail myself of a standard applied ethics approach to examine the extent to which the current EU regulatory landscape provides adequate tools to mitigate the ethical concerns raised by low-code/no-code AI platforms and their applications. On the other hand, I aim to propose adjustments to enhance such a framework to the extent that it does not. Specifically, I focus my attention on three categories of ethical concerns, namely, the lack of transparency, the presence of bias and discrimination, and the lack of responsibility.

I ultimately argue that the regulatory framework currently in place in the European Union is overall inadequately equipped to mitigate the ethical issues arising from the design, development, and deployment of low-code/no-code AI platforms and the resulting AI systems. On the one hand, non-legal regulatory tools are inadequate due to their voluntary nature not guaranteeing adherence to their guidelines. On the other hand, regulatory measures of legal nature lack the appropriate level of granularity to effectively mitigate the ethical concerns under investigation. I conclude the thesis by proposing two adjustments to the AI Act that would enhance the ethical regulation of low-code/no-code AI platforms in the EU.

1. Introduction

Imagine a time where children can design, develop, and deploy an AI system for object recognition as part of a short course on artificial intelligence offered by their elementary school curriculum. This may seem like a scene from a somewhat distant future. Instead, it already took place. The year was 2020. The students were sixth graders from Kontiolahti's elementary school in Finland. The AI system in question was a software able to detect berries and mushrooms, which the students were able to design, develop, and deploy after three 2.5-hour workshops (Toivonen et al., 2020). This was made possible by low-code/no-code AI platforms.

Low-code/no-code AI platforms are a subcategory of development platforms, i.e. environments that provide comprehensive tools for the development of software, which are usually geared towards professional software developers (Aarno & Engblom, 2014). Specifically, low-code/no-code AI platforms are a combination of low-code/no-code development platforms and AI platforms. On the one hand, low-code/no-code development platforms simplify software development through visual development: platform users are not expected to rely heavily (or at all) on programming languages, and are presented with pre-written code in the form of graphics and icons. On the other hand, AI platforms focus on the development of AI systems specifically, as opposed to any kind of software. These elements are combined to create platforms that offer cloud-based services that allow users with vastly different degrees of experience in computer science to create AI systems in a fast, easy, and inexpensive way, without the need for expert human supervision.

Unlike traditional development platforms, low-code/no-code AI platforms have a larger userbase, as they can also be used by amateur software developers and laypeople (the so-called *citizen developers*), and allow for the development of a smaller range of software, i.e., AI systems exclusively. This leads to two ethically significant peculiarities that distinguish low-code/no-code AI platforms from other digital technologies.

Firstly, low-code/no-code AI platforms are highly versatile, as they allow their users to create virtually any kind of AI systems, e.g. ranging from (semi-)automated business workflows to AI systems with visual object recognition capabilities (Toivonen et al., 2020), and are not restricted to a particular domain. While it would seem unreasonable to rely on the

services offered by a low-code/no-code AI platform to develop AI systems that may be used for health or defence purposes, this is currently not outside of the scope of such platforms, despite them not being specifically marketed to such domains. In the primary school children's application, the AI system was able to recognise mushrooms and berries. A similar object recognition system could be developed to recognise potentially cancerous moles, or individuals carrying weapons.

This raises ethical concerns as it is unclear whether the technical versatility of low-code/no-code AI platforms is at all counterbalanced by ethical constraints that guarantee the ethical design, development, and deployment (from here onwards, DDD) of such platforms and their applications. In fact, the lack or inadequacy of any such constraints, e.g., no provision that such platforms may not be used to develop AI systems for defence purposes, may create a gap between what is technically feasible and what is ethical, with negative repercussions on both the individuals impacted by the applications and society at large. In this thesis, I will focus on ethical constraints of regulatory nature to determine whether regulation can be a solid starting point to ensure the ethical DDD of the technologies under investigation. This choice does not mean to undermine the instrumental role of non-regulatory measures such as ethics education and training (Garrett et al., 2020), ethics-based auditing (Mökander & Floridi, 2021; Mökander & Axente, 2021; Mökander et al., 2021), and ethics-by-design methodologies (Jansen et al., 2021; Dignum et al., 2018; d'Aquin et al., 2018).

Secondly, these platforms allow users who may not have a solid grasp of software development to create sophisticated AI systems. In fact, low-code/no-code AI platform provide such a simplified development process that users with no prior knowledge of software development or programming languages do not have to undergo any kind of training before using them, nor do they need any expert supervision throughout the development process. Simply put, these platforms make the DDD of AI system within the reach of anyone who has access to a computer, tablet, or smartphone with an internet connection, regardless of their background and technical expertise. This is problematic in that it effectively delegates ethically-charged development choices to individuals who may not have the necessary skill set to grasp their significance and, without being trained to spot and mitigate possible ethical issues related to the DDD of the AI systems they are developing, may only have a rudimentary understanding of the possible negative repercussions of their applications.

Despite the relevance of the ethical concerns raised in relation to the DDD of both low-code/no-code AI platforms and their applications, no studies have been conducted on such technologies from an ethical perspective. This is where this thesis finds its main justification: as low-code/no-code AI platforms are becoming increasingly more common (Gartner, 2021) and ethical issues related to them are likely to arise in the near future, it is crucial for an appropriate regulatory framework to be in place to mitigate or eliminate the ethical issues raised by both such platforms and their applications. To this end, in this thesis, I address and fill the gap in the ethics of technology literature concerning the ethical DDD of low-code/no-code AI platforms, by shedding light onto the yet unexplored ethical implications raised by both such platforms and the AI systems they give rise to, as well as the adequacy of the current regulatory framework governing such technologies in the European Union.¹ Unless specifically mentioned otherwise, I will refer to “low-code/no-code AI platforms” to include both such platforms and the resulting AI systems. In general terms, as it lies at the intersection of applied ethics, technology, and regulation, this research is inserted in the responsible research and innovation (RRI) framework (Owen et al., 2020). Moreover, to the extent that it focuses on measures for the preemptive regulation of low-code/no-code AI platforms, this study incorporates the (vastly unexplored) notion of anticipatory regulation (Aczel et al., 2022).

1.1 Problem statement

In this thesis, I will examine the extent to which the current EU regulatory landscape provides adequate tools to mitigate the ethical issues arising from low-code/no-code AI platforms, and

¹ Given the interconnectedness of various levels of regulation (e.g., national, supranational, and international), the adequacy of the EU regulatory framework partially depends on additional factors, such as the presence of solid national and international regulation, which lie beyond the scope of this investigation. As for the former, EU regulations provide a benchmark for ethical safeguards in the DDD of AI systems vis-à-vis all Member States. My interest lies in understanding whether the level of mitigation of the EU regulatory framework is adequate regardless of the means that individual Member States may decide to adopt to supplement and refine it. As for the latter, the EU regulatory framework is significantly more fine-grained and comprehensive than its international counterpart, which is limited to soft law instruments, and thus provides a weak foundation for the mitigation of ethical issues concerning AI.

propose improvements to the extent that it does not. To this end, the following sub-questions need to be answered.

Firstly, identifying the ethical concerns raised by the technology under investigation requires gaining an adequate understanding of its distinctive features. What are low-code/no-code AI platforms and how do they work?

Subsequently, both the parameters and the regulatory tools considered to assess the adequacy of the ethical mitigation provided by the regulatory landscape need to be clarified. What are the main ethical concerns that low-code/no-code AI platforms give rise to? What are the regulatory measures currently in place to govern low-code/no-code AI platforms in the EU?

Only at this point will it be possible to determine the adequacy of the ethical regulation provided by the EU regulatory framework, by which I mean regulation capable of anticipating and mitigating ethical concerns. Does such a framework adequately mitigate the ethical concerns raised by low-code/no-code AI platforms? If that is not the case, which measures could enhance the ethical regulation of such platforms?

1.2 Intended research methodology

For the purpose of this thesis, I will avail myself of what Brey (2000) describes as a standard applied ethics approach, i.e., a methodology often used in research lying at the intersection between technology and applied ethics whose aim is to "clarify and evaluate morally controversial practices through [the] application and defense of moral principles" (Brey, 2000, p. 10), principles that are not specifically identified by this approach, which is rather intended as a research structure onto which to insert the moral theories or principles to which one may subscribe. This approach follows a tripartite structure, illustrated in Table 1, whose elements are reminiscent of the generalised methodology for ethical assessment of emerging technologies proposed by Brey et al. (2020).

	Standard applied ethics approach (Brey, 2000)	Generalised methodology for ethical assessment of emerging technologies (Brey et al., 2020)
Step 1	<ul style="list-style-type: none"> - Establishment of the focus of the investigation on the ethical regulation of low-code/no-code AI platforms. - Use of the literature to outline the core technical features of the technology and illustrate examples of ethically controversial applications that can be developed through low-code/no-code AI platforms. 	<ul style="list-style-type: none"> - Step 1 (Specification of subject, aim, and scope of ethical analysis) - Step 2 (Stratification of the subject of ethical analysis) - Step 3 (Description of the subject of ethical analysis)
Step 2	<ul style="list-style-type: none"> - Description of the ethical concerns raised by low-code/no-code AI platforms based on a literature review of the ethical issues raised by the technologies that compose such platforms, i.e., artificial intelligence, cloud computing, and low-code/no-code development. - Integration of the ethical issues raised by low-code/no-code's enabling technologies. - Analysis of three categories of ethical issues raised by the DDD of low-code/no-code AI platforms, i.e., lack of transparency, bias and discrimination, lack of responsibility. - Description of the most significant regulatory measures currently in place within the EU that may act as mitigation tools for such ethical issues. 	<ul style="list-style-type: none"> - Step 4 (Identification of potential impacts and stakeholders) - Step 5 (Identification and specification of potential ethical issues) - Step 6 (Analysis of ethical issues)
Step 3	<ul style="list-style-type: none"> - Analysis of the existing EU regulatory framework to evaluate to what extent it can adequately mitigate the ethical concerns raised by the DDD of low-code/no-code AI platforms highlighted in Step 2. - Proposal for amendments to the current regulatory framework that would enhance its effectiveness in achieving the ethical regulation of low-code/no-code AI platforms. 	<ul style="list-style-type: none"> - Step 7 (Evaluation and recommendations for ethical decision-making)

Table 1. Overview of the structure of the standard applied ethics approach as used in the thesis and its relation to the approach illustrated by Brey et al. (2020)

I argue that a standard applied ethics approach is the most suitable methodology for this research, due to its versatility and adaptability to technologies at different stages of the technology adoption lifecycle, especially in light of the peculiar status of low-code/no-code AI platforms. In fact, these platforms are currently neither completely emerging, nor fully entrenched in society. What I mean by “emerging” is a technology that is in the late ascent/early adoption phase, i.e., “in its design, research, development, or experimental stages, including beta testing” (Floridi & Strait, 2020, p. 78). On the contrary, I understand “entrenched” in terms of a technology that is influential and widely used in society (Moor, 2005). Specifically, “emerging technologies” and “entrenched technologies” correspond to the first and third stages, respectively, in Moor’s (2005) three-stage model of technology ethics:

Moor’s three-stage model of technology ethics	
Stage I	Introduction stage: technology begins to be used by early adopters
Stage II	Permeation stage: technology becomes conventional and standardised
Stage III	Power stage: technology is conventional and widely used

Table 2. Three-stage model of technology ethics (Moor, 2005).

While it is true that low-code/no-code AI platforms have not been adopted by a large user base yet, this technology nevertheless already does have a wide variety of real-world applications, as several concrete examples of AI systems that can be developed through these platforms throughout this investigation will illustrate. In this sense, it is crucial that the methodology used for this study be compatible with the current stage of the technology life cycle of low-code/no-code AI platforms.

By not being bound to a specific group of technological artefacts (e.g., digital or biomedical technologies), set of guiding moral principles, nor status of a technology (e.g., emerging, new, or entrenched), the standard applied ethics approach is highly versatile, which makes it more suitable to tackle this research topic compared to other popular applied ethics methodologies, most notably, foresight approaches for the anticipation and evaluation of new and emerging technologies, which have frequently been employed by studies on the ethical implications of such technologies from an applied ethics perspective (Floridi & Strait, 2020),

such as the ETICA approach (Stahl & Flick, 2011; Stahl et al., 2013) and the Anticipatory Technology Ethics (ATE) approach (Brey, 2012).

Indeed, unlike many of the technologies that are investigated using foresight approaches, it is, at least to a reasonable extent, arguable whether a foresight approach is a necessary first step to conducting a comprehensive investigation of the ethical implications raised by low-code/no-code AI platforms. While foresight could ultimately prove beneficial to this investigation and bring valuable inputs, as it could anticipate potential future applications and uses of this technology that may raise important ethical questions, the current state of the technology allows for this preliminary foresight analysis to be removed without the risk of falling into a merely speculative investigation.

Additionally, despite the lack of research on the ethical implications of low-code/no-code AI platforms specifically, the body of literature on the ethical implications of AI at large, cloud computing, and software development, is sufficiently solid to lay the foundations for the work to be carried out in this research, and a sensible starting point for the detection of such platforms' ethical concerns. In this sense, even in comparison with the benefits of the most comprehensive and nuanced foresight approaches for the anticipation and evaluation of emerging technologies that have become prevalent within the ethics of technology literature (Floridi & Strait, 2020), a standard applied ethics approach, which is integrated with a literature review, remains the most fitting methodology for the investigation at hand. Indeed, performing a foresight analysis of the technology at hand would not be a necessary requirement for an overall comprehensive investigation of the ethical issues raised by low-code/no-code development platforms, in that real-world applications of this technology already allow for such an analysis to be conducted at a very pragmatic level, and thus, to not be (highly) speculative in nature.

1.3 Structure of the thesis

This investigation will be structured as follows. In Chapter 2, the key features of low-code/no-code AI platforms and their applications will be illustrated, as well as concrete examples of such applications. Subsequently, Chapter 3 will provide an overview of the ethical concerns stemming from these technologies, with a particular focus on issues related to lack of transparency, bias and discrimination, and lack of responsibility. Chapter 4 will focus on the

current regulatory framework governing low-code/no-code AI platforms in the European Union, and on the question of whether it provides adequate tools to mitigate the ethical issues raised by the technology at hand. Specifically, I will analyse a selection of relevant regulatory measures, both of legal and non-legal nature, and reflect on the extent to which they can address and mitigate the specific ethical issue highlighted in the previous section, in relation to the above-mentioned application examples. In Chapter 5, I will formulate recommendations regarding which regulatory measures should be put in place to safeguard the ethical impact of low-code/no-code development platforms. In the final chapter, I will retrace my steps, summarise the results of this investigation, and present the limitations of this study.

2. Low-code/no-code AI platforms. Technical features and applications

Low-code/no-code AI platforms provide their users with tools to develop applications of various degrees of sophistication. To be able to provide such a vast array of services, the apparent simplicity of the platform as it appears to its users is not matched by its extremely complex and multi-layered structure. In this Chapter, I will explore low-code/no-code AI platforms' enabling technologies, i.e., artificial intelligence, cloud computing, and low-code/no-code development, to identify the key technical features of such platforms. I will subsequently present an overview of the main application domains of the AI systems that can be developed through such platforms.

2.1 Artificial intelligence

As the name suggests, artificial intelligence (AI)² plays a crucial role in the functioning of low-code/no-code AI platforms. As already mentioned, the focus of this investigation are platforms that allow their users to develop applications that are, at least partially, powered by AI, with a focus on machine learning. This choice is justified by a number of considerations.

Firstly, given the vast range of AI systems that can be developed through low-code/no-code AI platforms, it does not seem feasible to touch upon all the relevant ethical issues that may arise during the DDD of both such platforms and their applications. Nevertheless, it seems reasonable to expect that highly sophisticated applications will give rise to a greater number of ethically charged issues compared to lower-level applications, as well as, arguably, to more problematic ethical implications. Specifically, it seems reasonable to assume that the majority of (ethical) issues raised by low-level applications will be common to both lower-

² By artificial intelligence I mean both AI as a discipline and AI systems (HLEG, 2018). Specifically, the former encompasses numerous different approaches and techniques, including learning, reasoning, and decision making techniques, and robotics (HLEG, 2018, p. 7), while the latter refers to systems, designed by humans, with the ability to perceive the surrounding (digital or physical) environment, collect and interpret data from such an environment, and ultimately “reaso[n] on the knowledge derived from this data [to decide] the best action(s) to take to achieve [a] given goal” (HLEG, 2018, p. 7). This can also include analysing the consequences of their own actions and learning to adapt their own behaviour (HLEG, 2018, p. 7).

and higher-level applications, while additional issues may arise that are specific to higher-level applications.

At the same time, it is worth keeping in mind that, as a result of the highly simplified development process offered by low-code/no-code AI platforms, the actual degree of technical sophistication of the applications developed by users may not always be self-evident, in that the complexity of the development process may not always be an accurate reflection of the level of sophistication of the resulting application. In light of this, it seems necessary to identify a distinguishing factor that, whenever present in an application, may serve as a reliable indicator of high levels of sophistication. For the purpose of this investigation, I propose that such a factor will be the presence of AI in the form of machine learning, or lack thereof.

Briefly, machine learning (ML) is a set of techniques that rely on data and algorithms to mimic human learning processes and make predictions or classifications with gradually improving accuracy levels, which will ultimately be used to drive decision-making processes. Specifically, the learning process occurs over three steps. Firstly, the ML algorithm relies on input data to make a classification, e.g., to detect whether an email should be considered to be spam or whether a transaction is legitimate or fraudulent, or a regression, e.g., to predict stock market performance or real estate market trends. Subsequently, provided that known examples of correct predictions or classifications are available, such examples are compared to the ML algorithm's prediction of classification to assess the level of accuracy of the ML model. Lastly, the level of accuracy of the ML algorithm can be gradually improved by reducing the discrepancy between the model estimate and known examples, thus updating and optimising the learning process.

It is also worth mentioning that, just like ML is a subset of AI, ML itself contains subfields: most notably, (artificial) neural networks and deep learning (DL). As for the former, neural networks can be described as ML algorithmic systems composed of an input layer, one or more hidden layers, and an output layer. Depending on their complexity, the number of hidden layers of a neural network may vary. Whenever such a number is greater than one, the neural network is called a deep neural network, which leads us to the field of deep learning. As a subset of ML, DL shares the same fundamental features of *classical* ML, while being less reliant on human intervention to learn. The main distinctive feature of DL is the presence

of multiple hidden layers that are not visible to the designers of the algorithm, which means that nobody knows (nor can know) exactly what happens between the input layer and the output layer, and thus, how the final output has been reached. Due to their inherent opacity, hidden layers are known as the black box. In this sense, the greater the number of hidden layers, the greater the complexity of the algorithm and its opacity.

All these features being considered, I argue that any application created through low-code/no-code AI platforms that employs ML will entail a significant degree of sophistication.

2.2 Cloud computing

The second key feature of low-code/no-code AI platforms lies in their reliance on cloud computing (CC), which allows them to make highly complex and costly services accessible to their users at a fraction of the cost, such as access to ML models. By cloud computing, I mean access to computing services in an online environment for the purposes of increased speed, flexibility, and affordability.

As already mentioned in the previous section, the complex structure underlying ML models, whose creation requires deep knowledge in AI, is only part of the equation. In fact, ML models require large amounts of data and oftentimes high computational power to work, both of which are hardly accessible to the average low-code/no-code AI platform user. However, cloud computing is able to capture and meet the ever-growing need for accessible ML solutions.

In fact, with cloud computing, ML capabilities are made more accessible by allowing users to make use of ML-as-a-service, i.e. to have access to machine learning tools in a cloud environment without users being required to have the expertise or financial resources that would be necessary for them to create such tools by themselves. Specifically, ML-as-a-service can be addressed from three perspectives, which reflect the tripartite structure of AI systems (Table 3).

AI systems	ML-as-a-service
AI hardware, i.e. all the physical resources required to support AI algorithms, which includes a combination of data centres, edge computing systems, (embedded) personal computers, and Internet-of-Things devices	ML-infrastructure as a service, which provides access to the AI hardware
AI software environment, i.e. all the frameworks, platforms, and tools that are required to design, develop, and execute AI applications effectively and efficiently	ML-platform-as-a-service, which grants access to ready-to-use toolboxes, libraries, and frameworks
AI software application, i.e. the application that is running in the system	ML-solution-as-a-service, which makes ready-to-use ML solutions available to the users

Table 3. Tripartite structure of AI systems and ML-as-a-service.

As such, with ML-as-a-service, anyone with an internet connection can have access to the tools needed to design and develop AI systems powered by ML, with little to no constraints in terms of expertise and resources.

Despite its significant advantages, ML-as-a-service is not without its drawbacks, which may also capture ethical issues. In order to make use of cloud computing services, both a stable internet connection and a high bandwidth are required, which may prove challenging whenever users do not have access to high-speed internet, e.g. due to disparities in broadband access depending on the users' location (Fredline et al., 2020). Moreover, inability to access cloud computing services may occur as a result of downtime, i.e. the time during which a system is unable to fulfil its primary function, either due to preventive maintenance or updates or technical issues (King, 2004). Additional issues raised by cloud computing concern security and privacy, as it may not always be possible to process data locally, for instance when one wishes to execute ML techniques on IoT units (also known as “tiny” ML). In this sense, processing data and training ML models in the cloud may lead to security or privacy issues, e.g. whenever data sets include sensitive data.

2.3 Low-code/no-code development

The third and last key feature of low-code/no-code AI platforms concerns their reliance on low-code/no-code (LCNC) development. The expression *low-code/no-code development*

refers to a particular type of programming in which traditional linear coding is either partially or completely replaced by visual development tools. Specifically, while low-code development relies on a combination of traditional coding and visual development, no-code development relies entirely on drag-and-drop visual development.³

In their simplest form, low-code/no-code AI platforms provide their users with libraries of cloud-based tools to develop AI systems. Thanks to low-code/no-code development, the development process is simplified to allow users with either no background or a limited background in computer science to still be able to create AI systems on their own without the need for any expert human supervision, and without having to undergo any training. What this means in practice is that, visually, the low-code/no-code AI platform looks like a collection of building blocks that can be dragged and dropped, not unlike Lego pieces, to create a specific AI application. Additionally, based on the specifics of the platform they are using and the application they are developing, platform users may (in the case of low-code development) or may not (in the case of no-code development) be allowed to access the source code of such tools and edit it to further personalise their software.

What this means from a practical standpoint is that, while low-code development requires a certain, albeit moderate, level of prior programming knowledge, no-code development allows anyone to develop applications regardless of such users having any prior programming experience. Furthermore, no-code development enables platform users to create applications quickly and easily without undergoing conventional development training, by relying entirely on drag-and-drop visual development. Figure 1 illustrates the difference between traditional coding, low-code development, and no-code development.



Figure 1. Traditional coding, low-code and no-code development (OS system, 2021).

³ The expression “no-code development” indicates that the codebase lying at the core of the software development process is entirely abstracted and translated into a “What-You-See-Is-What-You-Get” visual interface that allows platform users to build applications without every seeing any code.

Overall, low-code/no-code AI platforms present significant advantages regardless of the technical background of their prospective users, or lack thereof. Specifically, one can identify three main categories of platform users. Firstly, professional software developers,⁴ who have a solid grasp of software development and can find such platforms beneficial to carry out their work more quickly and efficiently. Secondly, individuals who have a basic knowledge of coding, but are not professional software developers and specialise in developing applications through low-code/no-code AI platforms (let us call them *professional low-code/no-code developers*). Such platforms present these users with an attractive midway solution between programming their application from scratch through linear coding, which would be beyond their capabilities, and only relying on drag-and-drop visual development, which may not allow them to personalise their application to the extent that they are technically capable of doing it. Thirdly and lastly, low-code/no-code AI platforms simplify the development process to such an extent that people with no background in programming can use them to create sophisticated applications on their own; individuals belonging to this user category are also known as *citizen developers*.

It is also worth noting that users of low-code/no-code AI platforms are at no point required to fill any technical gaps they may have to acquire basic knowledge of neither coding nor software development at large. Moreover, low-code/no-code AI platforms are set up in such a way that users are both able and encouraged to go through the entire application development process without any expert supervision. While some platforms may, in some instances, provide users with the possibility to consult with a professional software developer to solve any problems they may encounter while working on their projects, the possibility of human supervision is usually not offered at any point during the development process (Doerrfeld, 2021), and users are oftentimes discouraged from reaching out to professionals in the first place, e.g. due to burdensome or lengthy procedures to get in touch with an expert via the platform. This seems particularly significant given platform users' ability to build sophisticated applications through such platforms without having a solid background in software development, and without their applications being checked by a professional before being deployed.

⁴ Software developers' responsibilities may include design tasks.

Lastly, a terminological clarification is in order. One may argue that “low-code/no-code AI platforms” is a redundant expression, as low-code development already includes and implies the possibility of no-code development. However, to preserve clarity, even to the expense of brevity, I will opt for the most extensive form throughout this investigation. The reasoning behind this choice is twofold.

On the one hand, as it will become apparent over the course of the investigation, certain issues raised by low-code development are specific to no-code development. Moreover, such issues are not only extremely problematic from an ethical perspective, but are oftentimes different in nature, complexity, and intensity compared to issues raised by low-code development at large. As such, I argue that ethical issues stemming from no-code development specifically deserve to be addressed and discussed separately. In this sense, it is important to remind the Reader of the significant distinction between low-code and no-code development, hence the extended expression *low-code/no-code* AI platforms.

On the other hand, while both low-code and no-code AI could, in theory, be used independently from each other, this is rarely the case, and most AI platforms offer a combination of both approaches. Hence, I do not consider focusing on no-code specifically to be a viable solution, as restricting the scope of this investigation to no-code AI platforms would result in an excessively narrow set of platforms and applications. In fact, removing all references to low-code development would lead to the exclusion of the overwhelming majority of AI platforms, as most of them offer a combination of low-code and no-code development features.

Moreover, while it may very well be the case that deeper or more controversial ethical issues might stem from no-code development compared to low-code development, such issues could easily be present in applications developed using a combination of low-code and no-code development. In this sense, it would be short-sighted to ignore the fact that problems specific to no-code development do not exclusively apply to applications based solely on this type of development. Indeed, they may also be present in applications which are based predominantly on low-code development, but also make residual use of no-code development.

2.4 Applications

Before moving on to the illustration of the ethical issues raised by both low-code/no-code AI platforms and the resulting applications, it seems beneficial to present some examples of the kind of applications that can be developed through such platforms. As already mentioned, low-code/no-code AI platforms allow for the creation of virtually any kind of AI system. However, depending on the category to which platform users belong, such AI systems can be traced back to a cluster of key application domains. Specifically, it is important to distinguish between users who avail themselves of low-code/no-code AI platforms to develop applications that they will be deploying themselves from those who develop applications on behalf of third parties.

To the first category belong both professional software developers and professional low-code/no-code developers. These users avail themselves of low-code/no-code AI platforms to streamline and fast-track their work. Depending on the individual developer's preferences, such platforms can be used to develop either AI systems for vastly different fields of application, or extremely niche applications within a specific domain, e.g., cybersecurity or e-commerce. As such, the resulting AI systems do not have a prevalent application domain, as low-code/no-code AI platforms do not limit the range of possible applications that can be created by these category of users.

To the second category belong citizen developers. This category can be further narrowed down depending on whether such users avail themselves of low-code/no-code AI platforms to develop AI systems for personal or professional purposes. The predominant application domains are productivity, utilities, and (cyber)security, and business and cyber(security), respectively (Gartner, 2021). Figure 2 presents an overview of the main application domains of the AI systems developed by such users.

Lastly, for a concrete illustration of a type of application that can be developed through low-code/no-code AI platforms, I refer the Reader to the scenarios illustrated in Appendix 1, which will be addressed in the following chapters.

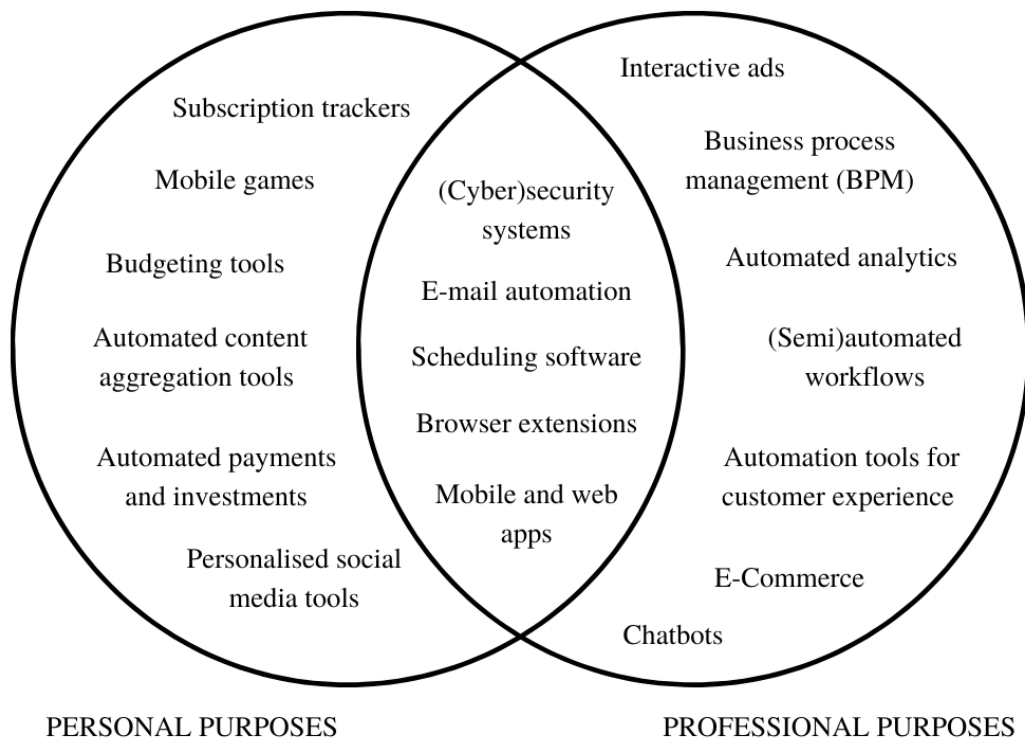


Figure 2. Main application domains of AI systems developed by citizen developers.

2.5 Conclusion

In this Chapter, I have illustrated the key technical features that can be identified as both foundational to, and distinctive of, the technology at hand, which descend from low-code/no-code AI platforms' enabling technologies.

Briefly, AI provides platform users with the ability to train ML models and use them in their applications. Cloud computing grants users access to both computational power, models, and libraries, which they would otherwise either not be able to use due to inadequate hardware and software at their disposal, or need to create from scratch. Low-code/no-code development greatly simplifies the development process and ultimately allows users with vastly different skill sets to develop applications without expert human supervision.

This means that, not only do low-code/no-code AI platforms provide their users with the tools to develop sophisticated AI applications, the development process is simplified to such an extent that is it accessible to virtually anyone with a computer and an internet connection, regardless of their knowledge or experience with programming and software development.

Lastly, the prevalent application domains of low-code/no-code AI platforms vary significantly depending on the specifics of their users.

3. Ethical issues

This Chapter will analyse the ethical concerns raised by low-code/no-code AI platforms. The importance of addressing these ethical issues stems from the combination of this technology's simplified and opaque development process and lack of expert human supervision, which results in applications that are likely to present a wide range of ethical issues, but are not routinely checked for ethical shortcomings before being deployed. In fact, low-code/no-code AI platforms make use of a one-size-fits-all approach, which provides their users with tools to develop a large variety of AI systems autonomously, but is not routinely equipped with features that allow users to detect the presence of possible ethical concerns in the applications that they are developing, issues which could be promptly identified by a human expert.

In the following sections, I will present an overview of the peculiarities of the ethical issues raised by low-code/no-code AI platforms, and focus on three key issues (bias and discrimination, lack of transparency, and lack of responsibility that emerge from low-code/no-code AI platforms' enabling technologies, but are severely exacerbated by the peculiar features of such platforms.

3.1 Ethical issues raised by low-code/no-code AI platforms

Over the past decade, rising interest in AI platforms resulted in several studies on their technical features (Di Ruscio et al., 2022; Salvaris et al., 2018; Fadaee et al., 2020). However, studies on the ethical implications of this technology (as well as development platforms at large) have been, perhaps due to a combination of its novelty and technical complexity, entirely missing from the scholarly debate. With the sole notable exception of research focusing on AI platforms from an economics and business management perspective (Mucha & Seppala, 2020; Yablonsky, 2020), even the few research projects that did address socio-technical issues related to AI platforms, did so superficially, due to the main focus of the research being the technical investigation of the technology (Fjeld et al., 2020). Overall, studies on AI platforms have either predominantly focused on the technical examination of their implications from a strictly computational perspective, or have nevertheless failed to

comprehensively address the socio-technical implications of this technology that go beyond the mere design and management of AI platforms and their products (Lee & Ha, 2018).

In their simplest form, low-code/no-code AI platforms are a combination of different technologies, namely: artificial intelligence, cloud computing, and low-code/no-code development. Within the applied ethics of technology literature to which this investigation pertains, extensive research has been carried out regarding both the ethics of artificial intelligence⁵ and cloud computing technologies⁶ and, to a lesser extent, software development at large⁷. Notwithstanding the lack of studies on the ethical implications of low-code/no-code AI platforms specifically, the body of scholarly literature on the ethical implications of their enabling technologies can lay solid foundations for the detection of low-code/no-code AI platforms' ethical concerns.

On the one hand, it seems reasonable to expect a significant overlap between the ethical issues raised by low-code/no-code AI platforms and those raised by their enabling technologies. Specifically, among the numerous issues raised by low-code/no-code AI platforms' enabling technologies, three emerge for their salience, i.e., those related to transparency, bias and discrimination, and responsibility, which will be further analysed in the following sections. Not only do such issues come into play in relation to such technologies,

⁵ Scholars identify the most significant ethical concerns raised by AI (and ML) in terms of issues relating to autonomy (Dignum, 2017; Calvo et al., 2020; Bjørlo et al., 2021; Jones et al., 2018), freedom (Ashraf, 2022; Hagendorff, 2020; Sekiguchi & Hori, 2020), privacy and data protection (Zhang et al., 2021; Stahl & Wright, 2018; Bartoletti, 2019), justice and fairness (Benjamin, 2019; Robert et al., 2020; Bennett & Keyes, 2020); safety and security (Santosh & Gaur, 2022; Leslie, 2020); transparency and explainability (Coeckelbergh, 2020; Lee, 2020; Burrel, 2016), bias and discrimination (O'Neil, 2016; Salminen et al., 2020), unreliability (Kearns & Roth, 2019), and responsibility (De Cremer & Kasparov, 2021; Constantinescu et al., 2021; Bogina et al., 2021; Cooper et al., 2022; Shah, 2018).

⁶ The key ethical issues highlighted by the literature on cloud computing concern privacy and data protection (Whitworth & de Moor, 2003; Baig, 2021; Stark & Tierney, 2014), safety and security (Firdhous et al., 2012; Ali et al., 2015), reliability of services and digital divide (Turilli & Floridi, 2009), data possession and ownership (de Bruin & Floridi, 2017), responsibility (Faragardi, 2017); the considerable power imbalances within the industry (Murphy & Rocchi, 2021), and the environmental impact of cloud computing (Moorthy et al., 2015).

⁷ The literature identifies privacy and data protection (Brey, 2000), accuracy (Poulton, 1994; Forester & Morrison, 1994; Simon, 1979), ownership (especially in regards to intellectual property and bandwidth ownership) (Mason, 1986), justice and autonomy (Brey, 2000) as the main areas of ethical concern concerning software development (Thomson & Schmoldt, 2001; Lurie & Mark, 2015).

but they are exacerbated whenever such enabling technologies are combined in low-code/no-code AI platforms.

On the other hand, the ethical issues raised by low-code/no-code AI platforms do not fully overlap with the ethical issues raised by AI, cloud computing, and low-code/no-code development. In fact, the peculiar features of low-code/no-code AI platforms lead to a wide variety of possible ethical issues, which partially transcend the issues commonly raised by their enabling technologies. To get a complete picture of all the ethical issues that such platforms may give rise to, it is not enough to merely add up the ethical problems raised by the different technologies that comprise them, but it is crucial to also consider the ethical issues arising from the intersection of such technologies within the specific context in which the technology is deployed.

Specifically, I claim that low-code/no-code AI platforms give rise to a relevant issue that transcends its enabling technologies, and can be identified in the presence of citizen developers. I argue that the degree to which ethical issues related to low-code/no-code AI platforms, such as the three categories of concerns identified above (transparency, bias and discrimination, responsibility), manifest themselves varies drastically depending on the different types of users of such platforms. More precisely, such issues are exacerbated to an unprecedented degree by the presence of citizen developers in the DDD of AI systems.

By this, I do not imply that the risk of technology created by professional software developers being unethical or ethically controversial is negligible. Indeed, one of the key issues that are regularly highlighted in the ethics of technology literature is that of technology designers not being mindful of the possible negative ethical repercussions of their work (van den Hoven et al., 2015; van de Poel & Verbeek, 2006). However, I argue that, in the case of low-code/no-code AI platforms, this issue is exacerbated by the fact that this technology allows people with no background in programming to develop sophisticated AI systems autonomously, without requiring such individuals to undergo any training on the matter. Specifically, I argue that non-professional users give rise to partially different and more problematic ethical concerns compared to their professional counterparts. I claim that, contrary to traditional software developers, who have a solid grasp of the specifics of the technologies they develop, but may lack equally strong ethical foundations, citizen developers have a double blind spot. Not only are they, as professional software developers may also be,

most likely unaware of the full spectrum of ethical implications of the technological applications they are developing, but they are also unaware of the mechanisms underlying the highly simplified software development process that allows them to create applications in the first place. Simply put, they do not know *what* they are doing, and they do not know *why* are are doing it.

I argue that an exceptionally limited understanding of the technical makeup and workings of a specific artifact translates to an inability to fully grasp the significance of the ethical concerns that such a technology may give rise to (such as transparency, bias and discrimination, and responsibility), as well as an inability to correctly and promptly spot such potential issues, and effectively act upon them as to limit their negative potential. Thus, it seems reasonable to assume that citizen developers may lead to applications that raise more serious ethical issues compared to professional software developers. Therefore, I argue that what is problematic and peculiar about low-code/no-code AI platforms is not a specific ethical concern as much as it is the way in which ethical issues come to the fore for such platforms in ways that are different compared to artifacts that employ the same enabling technologies without the presence of citizen developers.

3.2 Transparency

The first ethical issue that I will consider is the lack of transparency of low-code/no-code AI platforms. By transparency, I mean a property of a technology such that its aims, inputs, and internal operating mechanisms, are intelligible with respect to its stakeholders (Ryan et al., 2019). This leads to numerous questions, e.g., what about the internal operating mechanisms needs to be transparent? Is it the algorithm, the hardware, how the technology processes information, or a combination of all of these? Additionally, is transparency an absolute criterion or is it relative to different stakeholders, e.g., can something be transparent for a computer scientist but not for a layperson? In this sense, explainability is often identified as a component of the principle of transparency (OECD, 2019; Resseguier et al., 2021), but does not help to answer the question, as is it remains unclear in relation to whom a technology needs to be explainable.

Moreover, why is transparency (or lack thereof) something to worry about? The dominant view argues that transparency does not hold any moral significance in and of itself, but is instrumental to achieve the ethical DDD of technological artifacts in that it plays a crucial role in the realisation of other values that do hold moral significance (Heald, 2006; Ryan et al., 2019). For instance, the principle of transparency can be considered in terms of a means towards autonomy and control, to the extent that it limits peoples' ability to make informed choices (Jansen et al., 2019). Despite the moral basis of transparency being debatable, the principle of transparency is sufficiently accepted within the AI ethics field, to which this research pertains, where it is considered a key ethical concern (Stahl, 2021). In the case of low-code/no-code AI platforms and their applications, the relevance of assessing these technologies through the lenses of their adherence to this principle relates to transparency being a gateway to several ethical issues. In this sense, lack of transparency can later translate into issues related to bias (Schmidt & Biessmann, 2020; Daneshjou et al., 2021), (moral) responsibility (Hayes et al., 2022; Coeckelbergh, 2020; Ryan et al., 2019; Ryan, 2020; Santana & Wood, 2009) going unnoticed. Conversely, a focus on transparency can help spot and mitigate such concerns.

With regards to low-code/no-code AI platforms, one can distinguish different levels of transparency (or lack thereof). Firstly, a technology can be technically transparent, in the sense that it is theoretically feasible to gain access to its inner workings. However, it does not necessarily follow that people will be able to understand its inner workings readily. Despite being *explainable*, the technology might not be always easily *explained*. For instance, it might be the case that, for a person to be able to explain the functioning of a technology, this may require an exorbitant computing power to “translate” the inner workings of the technology into a format that is understandable to humans, an unreasonable amount of time, or both. Secondly, the comprehensibility of a technology is often relative to a group with certain levels of technological awareness. In this sense, a technology may be transparent with regards to professional developers, but not with respect to citizen developers (van Nuenen et al., 2020; Ehsan et al., 2021; Andrada et al., 2021; Bogina et al., 2021; Walmsley, 2020).

Overall, I argue that the issue of transparency is more pressing in relation to low-code/no-code AI platforms and their application compared to regular AI systems, due to the convergence of the technologies that compose them. Specifically, the lack of transparency

displayed by low-code/no-code AI platforms and the resulting applications can be traced back to three main factors, i.e., the technology's inherent features, its design choices, and citizen developers.

To the first category belong transparency issues raised by ML, whose inherent opacity cannot (at least, at the time of writing) be fully eliminated, despite numerous efforts in this sense (Adadi & Berrada, 2018; von Eschenbach, 2021; Roscher et al., 2020). In this regard, low-code/no-code AI platforms, which rely on ML, do not differ significantly from their enabling technologies in terms of the transparency-related ethical concerns they give rise to. Since ML systems hold an inherent degree of opacity (Cf. Section 2.1), any technology using ML will display some degree of opacity as well. As it is inextricably linked to the techniques and approaches being used, the lack of transparency is inescapable.

To the second category belong issues resulting from a deliberate choice from part of the technology designer, which contributes to exacerbating the technology's opacity. This concerns first and foremost the platform provider and, residually, platform users. Importantly, this is where the ethical issues raised by low-code/no-code AI platforms in relation to the principle of transparency deviate significantly from the concerns traditionally raised by their enabling technologies. Specifically, that such issues can be traced back to two key categories of design choices.

On the one hand, platform providers may limit low-code/no-code AI platforms' transparency to offer an accessible and user-friendly service for users with little to no experience in software development. Specifically, they may decide to trade transparency for the sake of practicality, thus opting for a design that is more functional than it is transparent. One may argue that this may not be a primary concern, as separating users from the complexities of both the platforms and the applications' inner workings is an integral part of the reason why such platforms are so attractive to users of various levels of programming skills in the first place.. However, this choice is morally problematic, as it reduces platform users' ability to self-determine freely, by not allowing them to be fully aware of the scope and significance of the choices they make during the development process. As such, whether to prioritise accessibility or transparency is a decision that should lie outside the domain of platform providers.

On the other hand, this choice might be motivated by reasons of convenience, on two main grounds. Firstly, platform providers may opt for a higher degree of opacity to prevent third parties from gaining access to their proprietary systems (e.g. algorithms and models) (Stahl, 2021), which may also help them maintain a competitive edge over different providers. Practically, this might mean that platform users may be unable to access the source code of the applications they develop, to avoid third parties using such code as a blueprint to develop similar services, thereby entering into competition with the platform providers. However, this also means that platform users may not be able to understand exactly how the platform and the resulting applications work.

Secondly, the lack of transparency of the low-code/no-code AI platform may result from a lack of affordable transparent alternatives from part of the platform providers. In fact, relying on third parties' services is a considerable part of the reason why low-code/no-code AI platforms can offer their services at an affordable rate. However, the downside of this choice is that this means retaining less control over the extent to which the services offered by the platform are transparent, in that this will be, to a great extent, dependent on the choices and preferences of the third parties who offer such services to the platform providers. In this sense, whether third parties' services are extremely opaque or moderately transparent will be almost entirely outside of platform providers' control. Again, this might mean that platform users may be unable to access the source code of one or more of the services and functionalities offered by low-code/no-code AI platforms. However, in this case, this would not result from a decision taken by platform providers, but it would be a consequence of a design choice made by the providers of a service that is offered by such platforms, e.g., ML-as-a-service.

In both instances, this leads to an increase in the degree of opacity of the technology, which is problematic and should instead be mitigated, in that it contributes to limiting the extent to which (especially) platform users can fully understand the functioning of the AI systems they are developing. The higher the degree of opacity of low-code/no-code AI platforms, the greater the chance that platform users may ultimately create applications that are not ethically sound, but cannot be checked for ethical issues due to their low level of transparency.

Lastly, to the third category belong transparency issues connected to the presence of citizen developers. These issues result from the combination of the issues mentioned above, i.e., black box technology, platform providers' design choices that make the development process opaque, with developers who have an extremely limited understanding of both the AI systems they are developing and the technology they are using to develop them.

3.3 Bias and discrimination

The second key ethical issue raised by low-code/no-code AI platforms concerns discrimination and bias, both of which are extensively covered by the AI ethics literature (Stahl, 2021; Jansen et al., 2019; O'Neil, 2016; Kearns & Roth, 2019). By discrimination, I mean the unjustified unfavourable treatment of groups of people, especially based on their gender, race, and sexual orientation (Jansen et al., 2019). This notion is closely related to that of bias, by which I mean an unjustified either positive or negative prejudice towards someone or something (Jensen et al., 2020). In fact, technological artefacts displaying biases often lead to discriminatory outcomes (Resseguier et al., 2021; John-Mathews et al., 2022; Feuerriegel et al., 2020; Yapo & Weiss, 2018; Baeza-Yates, 2022). For example, instances of discrimination may also occur as a consequence of a technology's limited usability (also known as functional bias; Jensen et al., 2020). For instance, an artifact may display bias in that it may not be universally accessible, e.g., stairs are biased against people using wheelchairs and strollers. This can result in the technology perpetrating discrimination in the form of ableism, leading to one or more categories of individuals not being able to use it. Additionally, it may also be the case that an artifact can be used by everybody, but it is going to disproportionately benefit certain groups over others, e.g., by proving more useful for individuals from a certain socio-economic background, profession, or with certain interests. For instance, while most people can use elevators, they will prove significantly more beneficial to individuals who have mobility issues and are unable to use the stairs, who would otherwise not be able to access the upper floors of a building, compared to individuals who use them out of convenience instead of necessity.

With reference to low-code/no-code AI platforms, the most significant instances of discrimination occur predominantly as a result of algorithmic bias, i.e., the presence of bias in

the output generated by an algorithm with respect to either (categories of) individuals whose data was processed by the algorithm, (groups of) individuals otherwise affected by the output of the algorithm, or both (Mehrabi et al., 2021). An example of algorithmic bias can be found in Amazon's 2014 AI-enabled recruiting engine, in relation to which male applicants scored overall higher than female applicants. This stemmed from the recruiting algorithm being trained with data that showed the company's hiring practices for the previous decade. While the algorithm was not trained to prefer male applicants over female applicants, it nevertheless derive a correlation between one's gender and one's likelihood of being hired from the data sets with which it was trained (Kordzadeh & Ghasemaghaei, 2022). The algorithm interpreted the company's trend to hire more men as a preference towards a specific gender, which resulted in a positive bias towards men and a negative bias against women (Dastin, 2022). This resulted in women being penalised during the application process, and ultimately discriminated against, as male candidates were routinely preferred and hired over female candidates even if the latter were more qualified (Tilmes, 2022).

With specific reference to low-code/no-code AI platforms, the presence of algorithmic bias is worrisome on three main grounds.

Firstly, and more generally, the presence of bias is a source of ethical concern in that it increases the chances of the technology leading to unfair and unjust outcomes. In fact, one may argue that bias and (to a lesser extent) discrimination are not ethical issues in and of themselves (Jansen et al., 2019). However, the relevance of addressing such issues from an ethical perspective is justified by them being closely linked to more fundamentally philosophical principles such as fairness and justice (Christian, 2020). Uncovering the presence of bias can both allow to avoid instances of discrimination and ultimately help to avoid scenarios in which artefacts contribute to, or perpetrate, unjust practices, as in the example of someone not being access the upper floors of a public building due to them having mobility issues and the building only having stairs, or unfair ones, as in the Amazon example. In this sense, algorithmic bias can negatively impact either (or both) the users of an AI systems or third parties that are affected by it (Christian, 2020; Aysolmaz et al., 2020), while also undermining the accuracy of an algorithm and its related applications (O'Neil, 2016; Cerrato et al., 2022). In this respect, the concerns raised by low-code AI platforms mimic those raised by traditional AI systems.

Secondly, the presence of algorithmic bias in a system is not easily detectable, as it has numerous potential sources. Specifically, algorithmic bias can be traced back to two main subcategories, depending on whether it is caused by social bias or by technical factors. In the first scenario, algorithmic bias stems from the (oftentimes systemic) perpetration of cultural biases in the form of social norms and assumptions (Ferrer et al., 2021; Rivera, 2012). In the second scenario, algorithmic bias originates from either the algorithm's input data, its processing of data, or its data output (Fazelpour & Danks, 2021).

As a way of example, this means that algorithmic bias may occur whenever a dataset being used to train a ML model does not accurately reflect the environment in which the application will be deployed (Barocas & Selbst, 2016; Brynjolfsson & Mitchell, 2017). For instance, if a ML model for facial recognition is trained with a data set containing primarily photos of adult males without glasses, this will most likely result in the model displaying significantly less accuracy whenever tasked with identifying categories of people that were not at all included in the data set used for training, or merely marginally included, e.g. individuals wearing glasses, women, and children (Cf. Appendix 1).

This issue may be exacerbated whenever the AI system developer projects their own prejudices and expectations on the sample being selected (Danks & London, 2017), e.g. a male developer may not think of including photos of people of different genders in the dataset used to train the ML model to recognise non-family members, just as a fair-skinned developer may not think to include photos of people of different complexions, which may result in the ML model being more accurate in detecting certain categories of people (Amini et al., 2019). This also points to a systemic issue of AI technologies, which oftentimes reveal a racial (Benjamin, 2019; Noble, 2018), gender (D'Ignazio & Klein, 2020), and ability bias (O'Neil, 2016). Specifically, facial recognition technologies have traditionally been trained with data sets composed by the overwhelming majority of caucasian, able-bodied males. As such, they are significantly less accurate when identifying certain categories of individuals, among which women, people of colour, and people with disabilities. Moreover, whenever a ML model is trained with data containing social bias, this may lead to the unintentional perpetration and exacerbation of social bias by the AI system (Stahl, 2021).

Thirdly and lastly, it should be noted that algorithmic bias is often unintentional (Koene, 2017). As such, what contributes to making it hard to detect is the fact that the

designers of the biased algorithm themselves may not be aware of its presence before the discrimination occurs (Johnson, 2020), as it was the case with Amazon's recruiting algorithm (Dastin, 2022). Indeed, if they were, they would correct the algorithm to remove it.

This issue is exacerbated by the fact that the implementation of mechanisms for the mitigation of algorithmic bias is significantly more challenging in regard to low-code/no-code AI platforms compared to traditional AI systems, especially in the case of applications developed by citizen developers (Yapo & Weiss, 2018). Indeed, a distinctive feature of low-code/no-code AI platforms is the lack of training (ethical or otherwise) required by their users to develop AI systems, combined with the opacity of both the platform and the applications that can be developed. In this sense, the presence of citizen developers increases the risk of biases going unnoticed, as it is highly unlikely that individuals with little to no background in computer science will be able to detect potential biases, especially when they are technical in nature. This leads to both issues of bias and discrimination being exacerbated exponentially. In this regard, the ethical issues raised by low-code/no-code AI platforms depart significantly from those raised by traditional AI systems.

3.4 Responsibility

The third main ethical concern raised by low-code/no-code AI platforms relates to the issue of responsibility. By (moral) responsibility, I mean the condition of being worthy of blame or praise as a result of one's actions or omissions (Ryan et al., 2019). Specifically, the attribution of responsibility requires the subject to whom an action or omission is imputable to be a moral agent, i.e., to be able to discern right from wrong and display autonomous intentions over one's conduct (Sullins, 2011). In this regard, a responsibility gap emerges whenever someone or something (e.g., an AI system), who is not considered a moral agent, engages in conduct that is blameworthy or praiseworthy.

Given the applied nature of this investigation, the relevance of such a gap in relation to low-code/no-code AI platforms, especially in terms of blameworthy actions and omissions, stems primarily from its substantive practical implications, which extend beyond issues of responsibility for an AI system or (opaque) algorithm, and include a more general responsibility for the impact of the technology within the broader socio-technical context in

which it operates. The importance of clarifying algorithmic responsibility in relation to the DDD of low-code/no-code AI platforms closely relates to the very possibility of effectively mitigating the ethical concerns raised by such technologies. As the choices, actions, and omissions of an opaque AI system that can negatively impact both its users and third parties, as well as society at large, e.g., by undermining autonomy, justice, and fairness (Cf. Section 3.2-3), not being able to identify who is responsible for such choices contributes to exacerbating the difficulty of mitigating such issues and safeguarding the individuals negatively impacted by them. In fact, it is often necessary to know who is responsible for an issue to effectively mitigate such problem at its roots.

With reference to low-code/no-code AI platforms, the main issue raised by such technologies relates to the responsibility of low-code/no-code AI platforms as machines. Since machines, unlike humans, are not moral agents, they cannot be held morally responsible. As a consequence, whenever performed by machines such as low-code/no-code AI platforms and the AI systems that can be created through such platforms, actions and decisions that normally include moral responsibility lead to a responsibility gap. This gap is further exacerbated by the opacity of such artefacts that goes beyond the inherent opacity of ML models and permeates the design of both the platforms and the resulting applications (Cf. Section 3.2) (Kossov et al., 2021; Blacklaws, 2018; Stahl, 2021). However, even when a system is not opaque, but fully transparent, there can be a responsibility gap as a consequence of the system acting in place of a human and making decisions that are normally made by humans.

Traditionally, one can try and fill such a gap by stipulating that humans, such as those who create or use a machine, are going to be morally responsible in its stead (Santoni de Sio & Mecacci, 2021). This raises the questions of (1) whether such a stipulation can and should be made in the case of low-code/no-code AI platforms and their applications, as well as (2) what such a stipulation may entail. I argue that such a stipulation *can* be made, in the sense that low-code/no-code AI platforms do not differ significantly enough from their enabling technologies to suggest that should not be the case (Kroll, 2020; Comandé, 2019; Cooper et al., 2022), and that is *should* be made, to fill the responsibility gap. However, the content and terms of such a stipulation raise further questions.

For instance, it is not necessarily clear who are the creators of a machine. In the case of low-code/no-code AI platform, are the creators the platform providers, or are they rather considered to be users, as they rely on algorithms created by third parties (Cf. Section 2.2), or are they both creators *and* users? Similarly, while it is easier to argue that platform users are the users of a technology (regardless of whether platform providers or third parties have created it), are they the sole creators of the applications they develop, or are such AI systems created by both users and providers (and/or third parties, depending on how one solves the previous issue)? Moreover, does discerning between professional users and citizen developers contribute to solve this question or does it make this question more nuanced?

Overall, the attribution of responsibility with respect to low-code/no-code AI platforms should be mindful of their peculiarities. For example, the responsibility gap is partially a result of the opacity of such technologies. In fact, a dominant view is that a party can only be ascribed responsibility for the behaviour of a machine if they have control over it, and if the action taken by the machine was intended by such individual (Matthias, 2004; Jansen et al, 2019). Determining the presence of the elements of control and intention becomes challenging whenever the functioning of the machine is partially or fully inaccessible, e.g., in the case of opaque AI systems, both due to the inherent opacity of ML models and as a result of the complexity of their inner workings and their multi-layered structure (Cf. Section 3.2).

Similarly, the presence of numerous individuals involved, in various capacities, in the DDD of the technology complicates the issue of the attribution of responsibility further (Hansson et al., 2021; McManus & Rutchick, 2019; Mokrian & Schuelke-Leech, 2021). It might be the case that a particular action or omission: (a) does not result from a single individual's conduct, but from a combination of several people's conduct, but it is impossible to determine if that is the case; and/or (b) results from the conduct of only some of the individuals involved in the process, but it is impossible to determine which ones.

Moreover, even being able to make such determinations may not necessarily clarify the attribution of responsibility. One may argue that, whenever a platform user develops an application using the services provided by the platform, and deploys it, they are the only person to ever work on that application for its entire life cycle. However, this leads to an

additional issue, as it is unclear whether moral responsibility for such systems should be placed on such individuals alone, on three main grounds.

Firstly, some platform users (namely, citizen developers) may lack intention, due to them not being required at any point to have a solid grasp of the technicalities of the system they want to create. Secondly, due to the simplified development process offered by low-code/no-code AI platforms, platform users are not actually creating applications from scratch. Instead, they are always relying on predefined tools, design and development choices made by third parties, e.g., platform providers and companies providing services to such platforms, thus lacking a meaningful control over the AI systems they are developing, e.g., their choices may be severely restricted by the platform's drag-and-drop visual development, and they may not be able to access the source code of their applications. Thirdly and lastly, it is worth noting that moral responsibility does not coincide with causal responsibility (Talbert, 2019; Babushkina, 2020).

Thus, it seems reasonable to hypothesise that the responsibility for the low-code/no-code AI platforms' conduct should be shared among several individuals, and that different agents may be responsible for different actions or omissions. However, this does not solve the issue of the responsibility gap in low-code/no-code AI platforms, which requires further attention. The absence of philosophy of technology literature on the topic contributes to leaving such question partially unanswered.

3.5 Conclusion

In this Chapter, I argued that the ethical concerns raised by low-code/no-code AI platforms partially transcend those raised by their enabling technologies. In this sense, the presence of citizen developers within the unsupervised development process offered by low-code/no-code AI platforms is especially problematic in that it contributes to exacerbating ethical concerns already present in these platforms' enabling technologies. This is due to such individuals' double (both technical and ethical) blindspot, compared to professional developers' technical proficiency, but (frequent) ethical shortcomings.

Subsequently, I briefly analysed three key categories of ethical concerns raised by low-code/no-code AI platforms. With regards to issues related to transparency (or lack

thereof), both low-code/no-code AI platforms and their applications are opaque in three respects. Firstly, because they make use of techniques, such as ML, which are by definition opaque. Secondly, due to (mostly) platform providers' design choices, which increase their degree of opacity. Thirdly, because of citizen developers' limited technical skill set. Moreover, the presence of citizen developers can contribute to increasing the difficulty of both identifying biases and discriminatory instances resulting from applications developed through low-code/no-code AI platforms, as well as addressing them and, ultimately, mitigating them. Lastly, low-code/no-code AI platforms echo and exacerbate responsibility concerns raised by their enabling technologies, and give rise to a severe responsibility gap.

4. The EU regulatory landscape

The aim of this Chapter is to analyse the extent to which the ethical issues illustrated in the previous Chapter can be mitigated by the regulatory framework⁸ governing AI in the European Union. Specifically, I will focus on four different levels of regulation: (a) the Charter of Fundamental Rights of the European Union (European Parliament, Council and Commission, 2000), which sets out the substantive rights and principles that are foundational to the EU, and places them at the core of EU law; (b) the Artificial Intelligence Act (European Commission, 2021), which is a proposed regulation to govern the DDD and use of AI systems in the European Union; (c) the General Data Protection Regulation (European Parliament and Council, 2016), which is an EU regulation on data security and privacy; and (d) the Ethics Guidelines for Trustworthy AI (HLEG, 2019), which is a set of guidelines for the ethical DDD and use of AI drafted by the European Union High-Level Expert Group on AI.⁹

⁸ I intend “regulatory framework” to include regulatory tools of both legal and non-legal nature, i.e., regulatory measures that are binding within a legal system, and voluntary measures that are an expression of self-regulation, respectively. In the interest of completeness, the AIA is included in this framework despite being a *proposed* regulation, as it is considered as the first comprehensive regulation of AI systems.

⁹ I argue that these are the most salient measures to illustrate the ethical regulation of AI in the EU. Additional hard law measures include the Product Liability Directive (85/374/EEC), governing civil liability issues following damage caused by defective products, including AI systems, and the European Parliament’s Resolution with recommendations to the Commission on a civil liability regime for AI (2020/2014 Legislative Initiative Procedure), differentiating between strict liability for damage caused by a high-risk AI system (Cf. Section 4.2), and a fault regime for damage caused by other AI systems. As neither of these tools have a clear connection to AI ethics, I limit my analysis to the CFR, GDPR, and AIA, which stand out for their numerous references to ethics. Conversely, among relevant soft law tools and non-legal regulatory tools emerge, respectively, the Recommendation on Artificial Intelligence (OECD, 2019), i.e., intergovernmental standards for the trustworthy use of AI, and the ISO and IEEE standards (Resseguier et al., 2021), i.e., standards for the design and development of trustworthy AI. Instead, I focus on the HLEG’s Guidelines for Trustworthy AI, which both references the benefits of standardisation, e.g., as part of its non-exhaustive Trustworthy AI Assessment List (HLEG, 2019), and provides a comprehensive and detailed overview of the ethical principles and best practices that should guide the DDD or AI systems in the EU.

4.1 Charter of Fundamental Rights of the European Union

The Charter of Fundamental Rights of the European Union (CFR; European Union, 2012) is a set of fundamental principles and substantive rights of EU law. There seems to be a consensus among scholars and professionals working in the AI field according to which an ethical framework for the regulation of AI in the EU should start from the principles and rights listed in the CFR (European Union Agency for Fundamental Rights & Council of Europe, 2018; Ulnicane, 2022; European Commission, 2018; Wagner, 2018; Koulu, 2020). However, especially when considered as the sole regulatory measure to mitigate the ethical issues raised by low-code/no-code AI platforms, the CFR is unable to provide a solid background for the ethical regulation of such platforms due its regulatory force being weakened by the combination of a low level of granularity and a lack of explicit references to AI.

In fact, while the CFR does provide a vast array of fundamental rights and principles, several of which have a direct or indirect connection to the ethics of AI, e.g., human dignity (art. 1), privacy (art. 7) and data protection (art. 8), equality (art. 20), and non-discrimination (art. 21), this connection is never rendered explicit. In fact, the formulation of such rights and principles is general and abstract, but is meant to provide foundational guidelines for the safeguard of EU citizens' rights at large, rather than address specific concerns. The interpretative effort required to link the safeguard of AI ethics principles and values to the rights and principles presented in the CFR decreases the regulatory strength of the latter, in that chains of interpretation open the way to ambiguity.

For instance, let us consider Article 21 of the CFR, which illustrates the safeguard of non-discrimination, and consider it in relation to the scenarios illustrated in Appendix 1. The first clause of Article 21 reads:

“Any discrimination based on any ground such as sex, race, colour, ethnic or social origin, genetic features, language, religion or belief, political or any other opinion, membership of a national minority, property, birth, disability, age or sexual orientation shall be prohibited” (European Union, 2012).

While it seems reasonable to state that discriminatory practices should be avoided, what constitutes discrimination? At least two problems seem to emerge when trying to answer this question.

Firstly, with regard to low-code/no-code AI platforms, instances of discrimination could be a result of algorithmic bias, e.g., whenever an individual is not identified by an intelligent security system due to issues with the dataset used to train the related ML model, which is not sufficiently reflective of the real world to provide accurate results (Christian, 2020). However, in order to understand whether an AI system is leading to the creation or perpetration of instances of discrimination, one would need to first understand how the ML model has reached its decision, i.e., how a specific input leads to a specific output. Yet, due to the opacity of ML models, uncovering the exact step-by-step process that governs ML models decision-making is not technically feasible by the very nature of such models. This makes it impossible to identify with certainty the presence of a discriminatory outcome.

Secondly, some scholars suggest that it is impossible to determine with a high degree of accuracy whether an AI system has engaged in discriminatory behaviour (European Union Agency for Fundamental Rights, 2020; Christian, 2020) without having access to sensitive personal data. However, the CFR also protects privacy (Article 7) and personal data (Article 8). This is an example of how, by being general and abstract, the principles and rights set out in the CFR can end up conflicting with each other as a result of a series of interpretative chains. This undermines the CFR's ability to effectively mitigate ethical concerns, e.g., if its safeguard of data protection is interpreted as prevailing over the need for transparency required to uncover instances of bias and discrimination, and ultimately fill a responsibility gap.

4.2 AI Act

The AI Act (AIA; European Commission, 2021) is a piece of proposed EU legislation aiming to set a comprehensive regulatory framework for the regulation of AI within the European Union. The AIA lays out a set of harmonised rules for the DDD of AI systems in the EU following a proportionate risk-based approach, which groups AI systems into four categories based on the severity of the risks that they may give rise to, which depends on whether such

AI systems pose a threat to the safety of their users, or to their fundamental rights: minimal risk, limited risk, high risk, and unacceptable risk. According to the severity and nature of the risk posed on their users, AI systems are then subjected to varying levels of regulatory interventions.

The proposal has received mixed reactions from scholars and professionals in AI ethics (Stahl et al., 2022; Johnson, 2021; Roberts et al., 2021; Whittlestone et al., 2021), a significant number of which have been highly critical of it (most notably, European Digital Rights et al., 2021; AlgorithmWatch, 2021). I argue that this also holds true when considering how the AIA governs the regulation of low-code/no-code AI platforms, which seems to be highly ambiguous and, ultimately, ineffective, for two main reasons.

The first reason relates to the levels of risk illustrated in the AIA and, more specifically, to all the examples of low-code/no-code AI platforms' applications illustrated in Appendix 1 belonging to the category of minimal-risk AI systems, regardless of the background of their users (from citizen developers to professional developers), their degree of sophistication (from modest to advanced), and their intended use (personal or professional). This is problematic as it means that the AIA does not affect the regulation of low-code/no-code AI platforms. The AIA does not specifically mention low-code/no-code AI platforms, but indicates how to determine what belongs to which risk category. I will now illustrate how I reached this conclusion.

In the scenarios outlines in Appendix 1, platform users used low-code/no-code AI platforms to develop different kinds of intelligent security systems (ISS) with biometric identification purposes. According to Article 5(1)(d) of the AIA, AI systems for biometric identification are prohibited when such identification takes place remotely, in real time, in publicly accessible spaces, and for the purpose of law enforcement. This does not apply to the scenarios at hand. As such, neither low-code/no-code AI platforms nor their applications are prohibited.

Furthermore, Annex III, 1(a) of the AIA, states that AI systems for biometric identification qualify as high-risk whenever such biometric identification, be it “real time” or “post” identification, takes place remotely, and is intended to be used for the identification of natural persons. According to Recital 8 of the AIA, remote biometric identification systems should be intended in terms of AI systems for the “identification of natural persons at a

distance through the comparison of a person’s biometric data with the biometric data contained in a reference database” (European Commission, 2021, p. 19). This aligns with the scenarios described in Appendix 1, since biometric data include facial recognition and the key function of the intelligent security systems developed by Alpha, Beta, and Gamma is to identify individuals whose biometric data is present in a database, e.g., a dataset containing images of family members or employees. Additionally, Recital 8 states that such biometric identification should occur “without prior knowledge whether the targeted person will be present and can be identified,¹⁰ irrespectively of the particular technology, processes or types of biometric data used”. Whether this is consistent with the scenarios at hand depends on what is understood by “knowledge over the presence of the targeted person”, which may vary greatly depending on the application scenario that is being considered.

In scenarios 1 and 2, Alpha and Beta are creating applications to identify family members through a smart security camera. However, it is unclear whether both individuals expect their applications to identify all family members. For instance, they may have trained the ML model to recognise family members regardless of the likelihood of them actually being physically in the proximity of their home, e.g., in the case of family members living in a different State. Therefore, it may be argued that Alpha and Beta do not know whether the individuals who can be identified by their AI system will ever be identified.

In scenarios 3 and 4, Gamma develops ISS to identify specific individuals and let them into a building, be it their home, in the case of their children (scenario 3), or their place of work, in the case of their employees (scenario 4). In these scenarios, since Gamma’s main concern is letting their children or employees into a building, biometric identification is a necessary requirement: Gamma expects that anyone who they want to be let in automatically by the ISS will be first identified. Therefore, while Gamma may not know exactly *when* the ISS will identify their children and employees, they nevertheless expect that they *will* be

¹⁰ It is also worth mentioning that, as for Alpha and Beta’s knowledge of the ability of their AI systems to identify such individuals, due to their lack of deep technical knowledge concerning the ways in which AI systems are created and how they work, it is unclear whether both individuals may grasp the possible shortcomings of their systems in terms of the systems’ inability to (correctly) identify their family members. As a result of this, while it cannot be stated that Alpha and Beta know that their intelligent security systems may not be able to identify their family members, it can also not be excluded that, although they do not know *why* their systems are not accurate, they may still be sufficiently aware of their shortcomings as to not expect them to always correctly identify whether an individual is one of their family members.

identified at some point. Additionally, Gamma is concerned with their security. In the examples, they set the accuracy threshold for the identification of the targeted individuals up to 80% (scenario 4) and 90% (scenario 3) for the door of the building to open. All these factors being considered, it seems reasonable to expect that Gamma will only include targeted individuals in their database if they want, and expect, such individuals to be identified and let into the building at some point. Hence, I argue that Gamma knows who will be identified by their ISS. As such, neither scenario would be consistent with the requirements necessary to qualify as a high-risk AI system.

Moreover, one may argue that the AI systems developed and deployed in scenarios 3 and 4 cannot fall under the high-risk category as a result of the intelligent security systems being used in combination with a smart lock, which means that they concern the biometric *authentication* of individuals, be them family members or employees, not their *identification*. The difference between identification and authentication can be found in the different aims of the procedures. While biometric identification describes the process of determining the identity of a person using the biometric data of a database, biometric authentication describes the process of *confirming* the identity of an individual through the same procedure. When Gamma's children are identified by the system before being let into their home, what the AI system is ultimately doing is proving their identity. Ultimately, only scenarios 1 and 2 are consistent with the definition of remote biometric identification system as stated in Recital 8.

However, this leads to the second main shortcoming of the AIA, which relates to the individuals to whom the AIA applies, which do not include individuals who develop and use AI systems for their own personal non-professional activity.

In fact, Article 2(a) of the AIA states that the regulation applies to *providers* and *users* of AI systems. According to Article 3(2), one should understand *providers* in terms of agents that develop AI systems or have them developed by a third party with the aim to place them on the market or “put them into service under [their] own name or trademark”. In this sense, it is debatable whether low-code/no-code AI platform providers fit this definition. While AI platforms provide their users with ML models that could be trained for biometric identification purposes, such platforms are technically not developing biometric identification systems themselves. Hence, platform providers do not offer their users biometric identification systems, but merely the tools to create such systems.

At the same time, according to the definition contained in Article 3(4), the individuals using AI systems for personal non-professional activities as *users*. As a result, despite the applications developed through low-code/no-code AI platforms being partially consistent with the scenarios described in the AIA (e.g., in the context of scenarios 1 and 2), that neither low-code/no-code AI platforms nor their applications are high-risk AI systems from the perspective of the AIA.

Moreover, these platforms do not belong to the limited-risk AI system category, as they do not belong to the types of AI systems illustrated in Article 52 and do not concern biometric categorisation, but biometric identification. In this sense, according to the AIA, low-code/no-code AI platforms belong to the category of minimal-risk AI systems, which means that they are not subjected to any ad-hoc regulatory measures contained in the proposed legislation. As such, the AIA is grossly ill-equipped to guarantee any degree of mitigation of the ethical issues illustrated in Sections 3.2-4, as it has no impact on the (ethical) regulation of low-code/no-code AI platforms and their applications.

4.3 General Data Protection Regulation

The General Data Protection Regulation (GDPR, Regulation 2016/679) is a EU Regulation that governs the use, processing, and storage of personal data in the European Union.¹¹ I argue that it GDPR is only partially able to mitigate the ethical concerns raised by low-code/no-code AI platforms, as it is only well-equipped to safeguard ethical issues of low-code/no-code AI platforms that relate to the (mis)use of personal data. However, such issues remain overall marginal when compared to the vast array of serious ethical concerns raised by such platforms and their application (Cf. Sections 3.2-4, Appendix 1).

To illustrate the type of problems that the GDPR is equipped to address and mitigate, let us consider Scenario 2' illustrated in Appendix 1, in which Beta trains a ML model to identify their family members. The GDPR is well equipped to deal with such problems, by providing tools both for their mitigation and prevention. To do so, it relies on three key features.

¹¹ It is worth highlighting that the AIA refers to the GDPR both implicitly and explicitly: most notably, when individual rights are concerned, which the AIA does not include.

Firstly, the GDPR is unlike the AIA, which (for the most part) is not compatible with the peculiar dynamics intercurrent between the parties involved with low-code/no-code AI platforms which, I argue, results in the under-regulation of these technologies. In fact, the somewhat complex dynamics between AI platform user, AI platform provider, and (if present) ML-as-a-service provider, is well-aligned with the GDPR. In fact, the GDPR considers (among others) three relevant categories of agents:

Data subject	The subject to which the personal data refers
Data controller	The subject who determines the purposes and means for the processing of personal data
Data processor	The subject who processes personal data on behalf of the data controller, which includes operations such as data storage

Table 3. Categories of agents in the GDPR.

In the example at hand, Gamma is the data subject, Beta is the data controller, and the data processor is the party who provides ML-as-a-service to Beta, i.e., either the AI platform provider or a third party.

Secondly, the GDPR requires the data controller to comply with several obligations, among which two fundamental ones, laid down in Article 25. First, as stated in the first provision of the article, appropriate technical and organisational measures must be put in place at the design stage of both information systems and personal data processing mechanisms (privacy by design). Second, as illustrated in the second provision of the article, the data controller must implement adequate technical and organisational measures to ensure that personal data will not be made accessible to third parties without the intervention of the data subject (privacy by default). In prescribing the adherence to both a privacy by design approach and a privacy by default approach, the GDPR not only lays down a strong foundation for the safeguard of any ethical principles that may be endangered as a result of actions or omissions that such approaches seek to restrict and eradicate, but it also ensures that these requirements are considered at any and all stages of the technology development process, from its conception to its implementation. In this sense, the GDPR is ultimately

effective in contributing to reducing the range of ethical problems that may arise from low-code/no-code AI platforms, and not just mitigating such ethical concerns.

Thirdly, the GDPR provides a clear set of guidelines to illustrate what counts as appropriate measures for the use, processing, and storage of personal data, which can be summarised into seven sets of overarching principles, presented in Article 5, which are illustrated in Table 4.

GDPR Article	Measure
Article 5(2)	Principle of accountability in the use, process, and storage personal data
Article 5(1)(a)	Principles of lawfulness, fairness, and transparency through the lawful, fair, and transparent use with respect to the data subject
Article 5(1)(b)	Principle of purpose limitation, in regards to the collection and processing of personal data for specific, explicit, and legitimate purposes
Article 5(1)(c)	Principle of data minimisation to ensure that personal data is adequate, relevant, and limited to what is required to achieve the intended purposes
Article 5(1)(d)	Principle of accuracy, to ensure that data is accurate and, if needed, updated
Article 5(1)(e)	Principle of storage limitation, which ensures that data can only be stored in a way that allows the identification of the data subject only for the time required to achieve the purpose for which personal data is processed
Article 5(1)(f)	Principles of integrity and confidentiality, to ensure that personal data is processed in a way that ensures adequate data security

Table 4. Principles for the use, processing, and storage of personal data according to the GDPR.

Overall, due to its narrow scope, the GDPR does not specifically address a wide variety of the issues raised by low-code/no-code AI platforms. To the extent that it does, it does not seem to present any significant regulatory gaps. However, it does not contribute meaningfully to the mitigation of the issues illustrated in Sections 3.2-4.

4.4 Ethics Guidelines for Trustworthy AI

The Ethics Guidelines for Trustworthy AI (from now on, the Guidelines; HLEG, 2019) are a set of guidelines drafted by the EU High-Level Expert Group on AI, first published in 2018 and ultimately finalised in 2019, with the aim of promoting trustworthy AI through the ethical

DDD of AI systems in the European Union. These guidelines seem to follow the ELSA (ethical, legal, and societal aspects) approach to AI ethics, as they identify trustworthy AI in terms of AI systems that are aligned with EU ethical principles and values (ethical aspects), while also adhering to existing regulation (legal aspects), and being mindful of the wider technical and social implications of the technology (societal aspects). By considering these guidelines in parallel with the legal instruments analysed in the previous sections, the Guidelines seem to be, in theory, the most comprehensive tool for the ethical regulation of low-code/no-code AI platforms, in that they are well-structured, clear, and nuanced. That is especially true on two grounds.

Firstly, they clarify the strong and direct link between fundamental (legal) rights and ethical principles and values. In doing so, the Guidelines also provide an interpretative key in the light of which to infer principles for the ethical DDD of AI systems directly from fundamental rights recognised within the European Union, such as the already mentioned legal rights and principles contained in the CFR. By offering a unified interpretation of how such principles and values could be considered through the lenses of AI ethics, the Guidelines make the content of the CFR more accessible and, by reducing their margin of ambiguity through clear and precise interpretation, ultimately increase its regulatory force. For example, the Guidelines draw from the CFR to identify four key ethical principles which have to be accounted for to design, develop, and deploy trustworthy AI, namely:

Respect of human autonomy	This principle derives from an interpretation of Articles 1 and 6 of the CFR, on human dignity and the right to liberty and security, respectively (HLEG, 2019)
Prevention of harm	This principle relates to Article 3, on the right to the integrity of the person
Fairness	This principle refers in particular to Articles 21 of the CFR, on non-discrimination, and to Titles IV and VI, on solidarity and justice, respectively
Explicability	This principle derives from an interpretation of the content of Article 47 of the CFR, on the right to an effective remedy and to a fair trial

Table 5. Key ethical principles (CFR).

As a result of this, the Ethics Guidelines help to uncover how the rights and principles contained in the CFR can be traced back to ethical principles and values, even (rather, especially) when the connection between the two is not self-evident, e.g., in regards to the principle of explicability which, in its formulation in Article 47 of the CFR, would seem to hold an overwhelmingly legal connotation, and to be extremely hardly linkable to AI ethics. Additionally, it helps clarify how such principles are ultimately not conflicting with each other, and may be harmonised through a nuanced interpretative effort (HLEG, 2019).

Secondly, the Guidelines provide a degree of elasticity that is lacking in legal regulatory tools, such as the AIA. For instance, these guidelines avoid stringent definitions of the stakeholders involved in the DDD of AI systems, which helps to avoid situations like the one described in Section 4.3, which can lead to exhaustive (yet not comprehensive) lists of stakeholders, ultimately leading to significant interpretative issues when they are applied to real-case scenarios. In this sense, the Guidelines provide a sensible and actionable middle ground solution between fundamental principles that may be too vague and not easily implementable, and very stringent regulations that may be too specific and leave low-code/no-code AI platforms unregulated.

In fact, not only do the Ethics Guidelines outline the theoretical basis for understanding the rationale behind the guidelines, but they also provide a list of key ethical requirements for achieving trustworthy AI¹² and actionable steps-by-step instructions¹³ to illustrate how the ethical DDD of AI systems may be achieved. These include guidelines addressing the safeguard and implementation of ethical practices concerning the issues of transparency (requirement 4), non-discrimination (requirement 5), and responsibility (requirement 7), which were previously illustrates as key ethical concerns raised by low-code/no-code AI platforms (Cf. Sections 3.2-4). Importantly, the Guidelines clarify very explicitly that any measures they propose, including the Trustworthy AI Assessment List, are *not* to be considered sufficient to guarantee ethical and trustworthy AI, and stress that the requirements

¹² The main requirements are seven: (1) human agency and oversight; (2) robustness and safety; (3) privacy and data governance; (4) transparency; (5) diversity, non-discrimination and fairness; (6) societal and environmental well-being; and (7) accountability (HLEG, 2019).

¹³ The Trustworthy AI Assessment List echoes the Anticipatory Technology Ethics Checklist from Brey (2012). Moreover, it seems to be a direct application of Brey's suggestion to employ ad hoc lists, e.g., reflecting specific sets of values in relation to specific kinds of technologies.

for trustworthy AI will never be exhaustive, in that guidelines will never be able to fully encapsulate and govern issues that are ultimately linked to technologies, such as AI, that are constantly evolving (HLEG, 2019).

However, the Guidelines present a crucial shortcoming insofar as they are at once extremely burdensome and not legally binding, as they can only find application through voluntary instances of self-regulation from part of the stakeholders involved in the DDD of AI systems (HLEG, 2019). I argue that, in the absence of legal requirements and/or market incentives, self-regulation alone does not hold enough regulatory strength to lead to the mitigation of the ethical concerns raised by low-code/no-code AI platforms. Simply put, regulatory measures are only effective if they are applied. Based on the current state of affairs, complying with the Guidelines would be so onerous on stakeholders that it would arguably not be worthwhile to do so, unless specifically required by binding regulatory measures.¹⁴ While market incentives may prove powerful enough to justify a significant expenditure in terms of both the time and resources necessary to comply with the Guidelines, a significant number of AI system developers and providers are not based in the EU. Therefore, whether

¹⁴ Let me illustrate this point through an example, by considering the issue of transparency. The Ethics Guidelines ultimately place the onus on AI application developers to make sure that the AI system they develop comply with the regulations in force in the European Union, as well as with the guidelines and best practices laid down in the document. As such, in relation to transparency, AI developers should be able to show: (a) how the dataset that was used to train the ML model was tested for bias; (b) how that bias, if present, was removed; (c) the reasons for choosing the particular methodology employed to this end.

This means that low-code/no-code AI platform providers, who rely on third parties' services (e.g., ML-as-a-service), would have to account for both the choices that have been made during the design and development of their platforms and those related to any third parties' AI-based systems to which the low-code/no-code AI platform grants access. This might work if third parties could provide the appropriate documentation to show that they have met all the necessary requirements, at which point the platform providers could proceed to provide a similar set of documentation for their own work.

However, I argue that platform providers would hardly be able to satisfy such requirements, which are, to a significant extent, beyond their control, e.g., as this would require third parties' compliance, in addition to being (at times prohibitively) onerous both in terms of time and resources. For instance, AI systems located in the United States do not have to comply with similar requirements. This means that there is a high likelihood that AI-based services could be built without paying an adequate degree of attention to such considerations. If that were the case, such AI systems would have to retrofit their design and development choices. While this could be done in theory, this would also require a significant investment both in terms of time and resources. As such, this would likely require significant market motivation and it is not necessarily the case that Europe will be a big enough market for providers of AI platforms and AI systems-as-a-service to consider this effort worthwhile.

Europe is a big enough market to justify a similar choice is still to be determined. Similarly, whether the Guidelines can deliver on their promises still remains to be seen. For the time being, they seem to lack the regulatory force necessary to ensure the ethical regulation of low-code/no-code AI platforms.

4.5 Conclusion

In this Chapter, I tested the adequacy of the EU regulatory framework governing low-code/no-code AI platforms for the mitigation of the ethical concerns illustrated in Sections 3.2-4 and concluded that this framework is ill-equipped to ensure the ethical regulation of low-code/no-code AI platforms. More specifically, it is fundamentally inadequate in that it is not able to fully capture the peculiarities of low-code/no-code AI platforms and the AI systems that can be created through such platforms, which are oftentimes not accommodated by the four regulatory measures taken into account, which leads to problematic practices being overlooked, and to an overall inadequate mitigation of the issues related to transparency, bias and discrimination, and responsibility. This is especially true in the case of the AIA which, despite being portrayed as a comprehensive tool for the regulation of AI systems, seems to have no visible impact on the regulation of low-code/no-code AI platforms and their application.

5. Ethical regulation of low-code/no-code AI platforms

In this Chapter, I will argue that, notwithstanding its shortcomings, the current EU regulatory framework can still provide a good starting point for the ethical regulation of such platforms. In particular, I will start by illustrating the benefits of employing legal regulatory tools to achieve the mitigation of the ethical concerns raised by the technologies under investigation. Subsequently, I will move on to consider the AIA specifically and propose two adjustments that would significantly improve both its effectiveness and adequacy of the EU regulatory framework at large. Lastly, I will conclude the Chapter by justifying the proportionality of such adjustments.

5.1 The benefits of legal regulatory tools

Overall, due to the considerable burden of complying with the regulatory measures' requirements in regard to the DDD of AI systems in the EU, regardless of whether they are legally binding or not, hard law measures would seem to be the most appropriate tools to guarantee such compliance (Wilms, 2014). This is especially true whenever such regulations impact on companies that are based in an extra-EU Country. In fact, such companies would have to drastically change their practices to enter the EU market or remain in it, while reasonably also providing their products or services to different markets with much lower standards in terms of AI ethics. Additionally, it is not necessarily the case that the EU market will be a big enough incentive to justify substantial investment in terms of time and resources to adapt to stringent ethics requirements, especially whenever they are not required by the law. In this sense, I argue that, to ensure the ethical regulation of low-code/no-code AI platforms, such requirements should be first and foremost rooted in solid legal regulatory tools.

On the one hand, despite lacking the appropriate level of granularity to effectively mitigate the ethical concerns under investigation in this thesis, the CFR provides an account of the rights and principles that are foundational to both the European Union and the ethics of AI. As such, it sets the legal and ethical bases for a regulation of low-code/no-code AI platforms that is mindful of the ethical concerns that this technology may give rise to, not

only by laying down principles such as of non-discrimination, which has a direct link to the ethics of such technologies, but also by highlighting overarching principles that are cornerstones of the AI ethics field as a whole, such as dignity, which is often interpreted in terms of a human-centred approach to the DDD of AI. Most notably, an example of this can be found in the HLEG’s Guidelines for Trustworthy AI, which refer to the CFR to provide a comprehensive and accessible illustration of both good practices that can help to achieve the ethical DDD of AI systems, as well as the reasoning behind such guidelines, thus providing a nuanced overview of the landscape of AI ethics regulation.

Additionally, while being very limited in scope, the GDPR provides a solid framework for the regulation of AI systems within the realm of data protection, which can help mitigate the following ethical issues:

Ethical issue	How the GDPR mitigates it
Lack of transparency	The GDPR requires data subjects to be informed about the instances in which their personal data is processed as well as the purposes for which it is processed, and asked to express their explicit consent to such practices
Bias and discrimination	The GDPR puts measures in place that can effectively limit the extent to which AI systems can negatively impact on individual rights, such as the role of human oversight in avoiding discriminatory practices resulting from biased data sets (Castets-Renard, 2019)
Lack of responsibility	The GDPR holds individuals and entities who process personal data accountable regardless of whether they process such data for professional or personal non-professional reasons, thus significantly reducing the responsibility gap.

Table 6. Ethical issues mitigated by the GDPR.

On the other hand, while the AIA seems to be currently unable to mitigate the ethical concerns raised by both low-code/no-code AI platforms, by considering this technology in terms of a minimal-risk AI system, I argue that adjustments could be made to the AIA that would allow it to significantly mitigate the ethical concerns that may arise from the DDD of low-code/no-code AI platforms and their applications. In the following section, I will expand on such adjustments.

As a last remark, it is worth clarifying that, while I argue that legal instruments are the most effective measure to achieve the ethical regulation of low-code/no-code AI platforms, I do not imply that non-legal regulatory tools or non-regulatory measures may not be as suitable for the mitigation of the issues raised by such technologies. Indeed, it would be short-sighted to consider legal measures as the only way in which such a mitigation can be achieved. Rather, the measures proposed in this thesis should be complemented by non-legal and non-regulatory instruments to achieve a holistic framework for the mitigation of the ethical concerns emerging from low-code/no-code AI platforms, which can effectively permeate and govern every stage of the technology’s life cycle (Cf. Section 1).

As a way of example, let us briefly consider how technical standards may mitigate the ethical issues raised by low-code/no-code AI platforms in comparison to legal regulatory tools. In short, standards define (stringent) technical requirements that must be followed in relation to the specific product, process, or service to which they apply (Resseguier et al., 2021; Allen & Sriram, 2000). Such standards can be internal to companies or have a broader impact, e.g., in the case of supranational (e.g., CEN standards) or international (e.g., ISO and IEEE standards) standards. Table 7 illustrates the core differences between standards and regulations.

Standards	Regulations/guidelines* [*note that, by “guidelines”, the Authors intend “soft law measures”, which are legal tools]
Based on recommendations	Based on legislation
Adoption is usually voluntary	Adoption is mandatory for regulations and potentially so for guidelines (soft law)
Established by consensus of all parties concerned, including relevant industry sectors	Developed by a regulatory authority, usually involving consultation
Based on consolidated results of science, technology and experience	Guidelines provide technical specifications either directly or by reference, e.g. to standards
Approved and published by recognized standardization body	Adopted by a legal authority
Oversight by independent third party certification	Oversight by formal government-appointed regulatory bodies

Table 7. Differences between standards and regulations/guidelines (Tait & Banda, 2017, p. 3).

With relation to low-code/no-code AI platforms, the relevance of their impact can be identified in their ability to “codify” ethics within their technical requirements (Resseguier et al., 2021), as it is the case for IEC SEG 10 (a standard on ethics in autonomous and AI applications) and ISO/IEC JTC1 SC 27 (a standard on information security, cybersecurity, and privacy protection). Table 8 provides an overview of the relative advantages of standards in comparison to regulations.

Advantages of standards	Advantages of regulations/guidelines* [*note that, by “guidelines”, the Authors intend “soft law measures”, which are legal tools]
Standards can act as infrastructures for coordination; a common language for interoperability and compatibility	Regulations have the force of law, and compliance is compulsory and enforceable
Standards as routines (usually internal standards) can govern behaviour required for certain activities/routines	Easier to diffuse through inter-country, regional or international treaties and conventions
Standards as technology can reduce variety and enhance economies of scale thereby reducing transaction costs	Regulations are prescriptive, and sometimes are linked to specific guidelines and/or standards which, if adhered to, constitute compliance
Standards can be an innovation to achieve market dominance	

Table 8. Relative advantages of standards and regulations/guidelines (Tait & Banda, 2017, p. 3).

As the table illustrates, standards may be effective both as an alternative to regulation and in combination with regulatory measures. Specifically, standards may succeed where (non-legal) regulatory tools prove ineffective, e.g., whenever ethical choices concerning the DDD of low-code/no-code AI platforms are not legally binding and too onerous to be followed, by creating market incentives in this sense. Additionally, standards could integrate the legal measures considered in this section to strengthen their mitigation capabilities, e.g., concerning the sound use of the risk-based approach proposed by the AIA (ISO/IEC JTC1/SC42, on AI risk management).

5.2 Filling the regulatory gap

In this section, I will illustrate the adjustments to the AIA that would allow it to strengthen the regulatory landscape surrounding low-code/no-code AI platforms so as to effectively mitigate the ethical concerns that these technologies and their applications may give rise to. Specifically, I argue that two amendments should be introduced to the AIA.

Firstly, I argue that low-code/no-code AI platforms should be considered in terms of high-risk AI systems, as opposed to minimal-risk AI systems. This is because AI platforms provide their users with the means to develop a wide variety of AI applications for potentially vastly different purposes, e.g., ranging from automated workflows to intelligent security systems, without the platform providers exercising any form of supervision over the resulting applications at any point during the development process.

This is further aggravated by the fact that these platforms are targeted at an extremely broad audience, which includes both individuals with a solid background in AI and software development, and individuals with no prior programming experience nor a solid grasp of computer science. In this sense, the ethical concerns that are inherent to the DDD of AI systems at large are exacerbated by the presence of unskilled amateur developers whose applications do not routinely undergo any form of ethical inspection before being implemented. However, the AIA does not fill such a regulatory gap, by considering low-code/no-code AI platforms and their applications in terms of AI systems that give rise to a minimal level of risk, thus not requiring platform providers to take any specific measures to ensure the ethical DDD of such platforms, beyond merely voluntary ones.

By including low-code/no-code AI platforms in the category of high-risk AI systems, platform providers would need to comply with the provisions set out in the AIA, which would diminish the potential for ethical concerns being raised by low-code/no-code AI platforms. Specifically, the obligations placed on providers, distributors, and users of high-risk AI systems are divided into two macro categories, depending on whether such obligations arise before or after the AI system enters the EU market.

Prior to placing a product on the market, high-risk systems must undergo a conformity assessment, which includes the following elements, each of which help to address either one or a combination of the issues of transparency, bias and discrimination, and responsibility:

Element	Description	Article	Issues addressed
Risk management system	The AI system provider must implement a risk management system to: (a) identify and analyse risks associated with the use of the AI system; (b) estimate them; and (c) take appropriate (protective) measures. Such a system must be maintained throughout the entire life cycle of the high-risk AI system	Art. 9	Transparency; bias and discrimination; responsibility
Data governance	In order to minimise the potential for harmful bias and discrimination, the datasets used by AI system providers for training, validation, and testing must be managed appropriately, and must also satisfy high quality requirements in terms of relevance, representativeness, correctness and completeness	Art.10	Transparency; bias and discrimination
Technical documentation	Providers must produce and maintain complete and up-to-date technical documentation that demonstrates compliance with the requirements laid down in the AIA	Art. 11	Transparency; responsibility
Record-keeping	In order to mitigate the black box effect and the opacity associated with the operation of high-risk AI systems, providers must provide automatic logging mechanisms to ensure the verifiability and traceability of decision and process implementations by high-risk systems	Art. 12	Transparency; responsibility
Transparency and information to users	The transparency of high-risk AI systems vis-à-vis users must be ensured through the provision of instructions for use that must accompany any high-risk AI systems. As for their content, such instructions must be relevant, complete, accurate, and in line with the elements specifically identified in Article 13(3)(a-e). As for their form, such instructions should be clear, concise, accessible, and understandable to users	Art. 13	Transparency
Human oversight	AI system providers must put in place appropriate measures to allow effective human supervision of the operation of high-risk AI systems	Art. 14	Responsibility
Accuracy, robustness, and cybersecurity	AI systems that pose a high level of risk must be designed and developed in a way that ensures a high level of accuracy, robustness, and cybersecurity throughout the entire life cycle of the AI system	Art. 15	Transparency; responsibility

Table 9. Elements of the conformity assessment for high-risk AI systems.

Additionally, providers are required to verify the *ex ante* compliance of the high-risk AI system with said requirements. Depending on the type of AI system, this verification may be carried out by third parties or through an internal verification process. If the evaluation procedure is successful, providers must register an EU declaration of conformity, affix the conformity mark to their high-risk AI system, and register the high-risk AI system in a central public EU database.

However, compliance is considered as a continuous, never-ending process. As such, it is expected that regular updates and revisions will be necessary. To this end, monitoring systems should be implemented by the AI system providers so that usage data generated by the system can be collected and analysed. Based on such data, if necessary, high-risk AI system providers must apply appropriate corrective measures to ensure the continuous compliance of the AI system with the regulation. Moreover, users themselves must comply with the instructions for use that are provided to them by the AI system provider, and limit the use of their high-risk AI system to its intended purposes only. Additionally, the AIA stipulates that, were users not to follow these instructions, they would be treated as providers, with all the obligations that the role entails.

Therefore, including low-code/no-code AI platforms within the category of high-risk AI systems would sensibly reduce the scope and degree of the ethical concerns that may result from the use and misuse of such platforms. In particular, having to provide their users with a set of instructions as to how the low-code/no-code AI platform works and how platform users should make use of it, combined with the requirement of human oversight, would lead to platform providers ultimately setting limits on platform users' freedom with regard to the range of AI applications that could be created through such platforms and the ways in which such applications could be developed. This would ultimately help mitigate or even eradicate (at least some of) the ethical concerns raised by such technologies, while also sensibly reducing the number of AI applications that can be (lawfully) created through such platforms and which can still give rise to severe ethical concerns.

However, since the AIA only considered users in terms of individuals who deploy AI systems for professional purposes, this alone would not be sufficient to ensure the compliance of all the AI systems that can be developed through low-code/no-code AI platforms with such requirements. In this sense, an additional measure should be taken to ensure full adherence to

both the specific obligations set out in the proposed regulation and the core objectives that underlie the AIA as a whole, among which the safe and lawful use of AI systems with respect to EU fundamental rights and values, as illustrated in the explanatory memorandum of the AIA (European Commission, 2021).

This leads the second amendment. I argue that the AIA should broaden its scope in terms of the subjects to which it applies, by employing an approach à la GDPR and not restricting the range of such subjects *a priori* based on rigid categories (e.g., individuals who use it for personal vs professional purposes), which would not leave any room to consider the ethical impact of the AI systems designed, developed, and deployed by platform users on a case-by-case basis. Instead, while defining the scope of the proposed regulation, the focus should be on the purposes and possible ethical concerns raised by the AI systems under consideration by considering their specific application domain. As such, the regulations contained in the AIA should apply to low-code/no-code AI platform users to the extent that the applications they develop and deploy may give rise to severe ethical concerns. This should also include instances in which platform users develop and deploy AI systems for personal non-professional purposes, which are currently beyond the scope of the AIA.

Specifically, this would provide a second barrier to limit the ethical issues raised by low-code/no-code AI platforms. In fact, this would mean that, going back to the illustration of the obligations of users of high-risk AI systems as illustrated in the AIA, not only would platform users be required to follow the instructions for use supplied by platform providers but, were they not to comply with such instructions, they would be considered for all intents and purposes as high-risk AI system providers, and would thus become legally responsible for the compliance of such AI system to the regulations set out in the AIA.

Let me illustrate through an example how the inclusion of low-code/no-code AI platforms in the category of high-risk AI systems and the expansion of the definition of users to include individuals who develop and deploy AI systems for personal non-professional purposes can enhance the current ethical regulatory framework of such platforms. Specifically, let me consider scenarios 1 and 2 (Cf. Appendix 1), in which the ISS developed by Alpha and Beta are less effective with individuals with dark skin tones, which is an example of discrimination stemming from functional bias (Cf. Section 2.3).

The current formulation of the AIA does not help to mitigate this issue, due to both applications being developed for non-professional purposes, which excludes them from the category of high-risk AI systems. However, the adjustments to the AIA allow it to be more effective in mitigating this concern. For the reasons illustrated in Section 4.2, both Alpha and Beta's systems belong to the category of high-risk AI systems, despite being developed for personal purposes. This means that both ISS have to comply with the requirements envisaged by the conformity assessment for such AI systems (Cf. 5.2). As a result, Alpha and Beta need to gain a better understanding of the functioning of low-code/no-code AI platforms and, since such platforms provide their users with the possibility of using ML-as-a-service, this will include learning the basics of ML as well as how to use this service according to best practices. This is likely to contribute to Alpha and Beta developing AI systems that are less prone to instances of bias and discrimination.

5.3 Anticipatory regulation

In this section, I will justify the proportionality of the measures proposed above, by briefly focusing on one of the thorniest challenges associated with the regulation of new and emerging technologies, which concerns the problem of under-regulation and over-regulation. Specifically, while I argue that the current regulatory landscape is an example of the former, the regulatory measures proposed in the previous section may seem to be an example of the latter.

One may argue that the measures proposed in Section 5.2 might be too burdensome on the part of both platform providers and platform users. As for the former, due to low-code/no-code AI platforms not *creating* AI systems giving rise to ethical issues, but merely *facilitating* their development. As for the latter, due to the knowledge imbalance between platform providers and platform users in terms of the functioning of AI systems. That is to say, if platform users develop AI systems through simplified visual development, it might seem unreasonable to expect them to comply with the norms laid down in the AIA, which presume the ability to grasp the technicalities of the AI systems involved, which is not only beyond the abilities of such individuals, but incompatible with the very reason why such individuals are using low-code/no-code AI platform to begin with.

However, I argue that both measures are proportionate to the risks involved with the DDD of low-code/no-code AI platforms, and should not be considered as an instance of over-regulation. The main reasons as to why this is the case are two.

Firstly, addressing low-code/no-code AI platforms in terms of high-risk AI systems can fill the current regulatory gap concerning the governance of such technologies without jeopardising their existence. When regulating new and emerging technologies, such as low-code/no-code AI platforms, one has to try and foresee the possible future problems arising from these technologies and take them into account. This is especially true in the case of legal regulatory tools, which need to provide regulatory measures to govern future problems before such issues arise.

In the case of the technology at hand, the individual applications might be considered more problematic from an ethical point of view compared to the AI platform, due to the presence of citizen developers. However, the root of the ethical concerns that materialise in the AI applications can be found in the AI platform itself. Specifically, this results from low-code/no-code AI platforms allowing users with little to no knowledge of software development to create and deploy AI systems without them needing to undergo any technical nor ethical training and without such applications ever being checked before their implementation. As such, it seems reasonable to regulate the low-code/no-code AI platform upstream (i.e., *ex ante*) rather than all its applications downstream. In this sense, it does not seem disproportionate to require platform providers to follow the rules illustrated in the AIA with regards to high-risk AI systems.

Secondly, I argue that the burden placed on platform users is proportionate and ultimately justified by a trade-off between users' ability to design, develop, and deploy AI systems without being restricted by any constraints, and the safeguard of the ethical principles endangers by such AI systems from the perspective of society at large.

On the one hand, it is true that not all platform users are, from the outset, fully aware of the ethical implications connected to the applications they intend to develop and deploy, and that this ignorance is currently fostered by the way in which low-code/no-code AI platforms are marketed and structured. However, once platform providers supply platform users with clear, comprehensive, and accessible instructions for the correct use of the AI platforms and the services that it offers to its users, their technical blindspot will be filled.

Furthermore, it is true that putting the onus on individual users to adapt to the rules imposed by the AIA in relation to the DDD of high-risk AI systems could defeat the very reason why these users turn to low-code/no-code AI platforms in the first place. However, It is also true that, when platform users create high-risk AI systems, the right of the individual to develop an application without having heavy regulatory burdens placed upon them should not be taken as a priority over the need to protect society at large from the ethical problems arising from the DDD of such AI systems.

As a last remark, it is worth stressing that the AIA does account for the progress of technology and future needs, and envisages the possibility of expanding the category of high-risk AI systems, if the need were to arise to do so. At the same time, no studies have been conducted on the risks and ethical concerns raised by low-code/no-code AI platforms, which, as a result, do not currently satisfy the criteria to be included in the list of high-risk AI systems. In this sense, the regulatory measures proposed in this thesis anticipate ethical concerns that, while already present to some extent, are still in their infancy. As such, the proposed measures can be considered as an example of anticipatory regulation (Aczel et al., 2022), i.e., an approach that aims to narrow the gap between the fast pace of technological advancements and the slow pace of legislative measures to govern new and emerging technologies, through “future-facing and proactive” regulation (Armstrong & Rae, 2017, p. 1).

Specifically, this thesis incorporates the notion of anticipatory regulation in that it focuses on preemptive regulatory measures that aim to mitigate and (in certain instances) eliminate the ethical concerns raised by low-code/no-code AI platforms before they can cause significant harm. In this sense, promoting stringent rules for the ethical DDD of such platforms and the related AI systems before they become entrenched in society may foster the responsible innovation of such technologies (van den Hoven et al., 2014), and successfully limit their disruptive potential (Hopster, 2021).

5.4 Conclusion

In this Chapter, I argued that legal regulatory measures are ultimately more effective to guarantee the ethical regulation of low-code/no-code AI platforms compared to their non-

legal counterparts. I then proposed two amendments to the current content of the AIA. That is, on the one hand, the inclusion of low-code/no-code AI platforms within the category of high-risk AI systems; on the other hand, the extension of the category of users to which the obligations laid down in the AIA apply, in order to include individuals (such as citizen developers) who design, develop, or deploy AI systems for private, non-professional purposes. I argued that such adjustments can sensibly mitigate the ethical issues illustrated in Sections 3.2-4, and ultimately enhance the ethical regulation of low-code/no-code AI platforms. Lastly, I concluded that such measures, despite anticipating future needs, should not be interpreted as an instance of over-regulation of this technology.

6. Conclusion

In this thesis, I have investigated the question of whether the existing regulatory framework governing low-code/no-code AI platforms and the resulting applications in the EU is adequately equipped to mitigate the ethical concerns related to the design, development, and deployment of such technologies. To this end, I have first illustrated the key features of low-code/no-code AI platforms and provided examples of the types of AI systems that can be created through such platforms.

I have then moved on to considered the ethical concerns raised by the technology under investigation, arguing that such issues do not fully overlap with those raised by low-code/no-code AI platforms' enabling technologies, but partially transcend them, especially as a result of the presence of citizen developers. Subsequently, I have briefly analysed three categories of key ethical concerns raised by such platforms, namely, the lack of transparency, the issue of bias and discrimination, and the lack of responsibility. With regards to the former category, I concluded that low-code/no-code AI platforms and the resulting AI systems are opaque in three respects: (1) as they use techniques that are inherently opaque; (2) as a result of platform providers' design choices; (3) due to the technical limitations of citizen developers. As for the second category, I once again attributed the difficulties of identification, addressing, and mitigating instances of bias and discrimination to the presence of citizen developers. With reference to the third category, I concluded that the technology under investigation echoes the responsibility issues raised by its enabling technologies, and exacerbates them significantly.

I have subsequently analysed the EU regulatory framework and concluded that is not adequately equipped to fully mitigate the issues that may arise from low-code/no-code AI platforms. Specifically, I have argued that non-legal regulatory measures are not particularly effective to ensure the ethical regulation of low-code/no-code AI platforms, despite their high level of granularity, due to their lack of binding force. Instead, I argued that legal regulation is the most effective way to ensure compliance with ethics requirements regarding the design, development, and deployment of low-code/no-code AI platforms. However, I concluded that the such tools may also present sensible shortcomings, particularly with regards to their level of granularity. On the one hand, the Charter of Fundamental Rights of the European Union

displays too low of a degree of granularity to be effectively operationalised. On the opposite end of the spectrum, the GDPR presents a high level of granularity, but has too narrow a scope to effectively mitigate the above-mentioned categories of ethical issues raised by low-code/no-code AI platforms. As for the AI Act, the proposed regulation is promising in theory, both for its focus on AI systems specifically and for its risk-based approach, according to which higher the level of risk, the more stringent the regulations governing the AI system. However, it currently classifies low-code/no-code AI systems in terms of AI technologies presenting a moderate amount of risk, for which no specific requirements are in place. While it stipulates that obligations whose satisfaction is required in relation to high-risk AI systems can also be applied to AI systems belonging to lower categories of risk, by being a voluntary requirement, this encounters the same issues already highlighted in relation to non-legal tools.

I have subsequently proposed two measures to enhance the ethical regulation of low-code/no-code AI platforms, by suggesting two amendments to the AI Act. Firstly, I argued that low-code/no-code AI platforms should be included in the category of high-risk AI systems, as they give rise to a wide range of severe ethical concerns and allow users with neither technical nor ethical training to create applications that do not undergo any kind of supervision before being implemented. Secondly, I argued that scope of the AI Act should expand to include AI systems designed, developed, and deployed for non-professional purposes. Implementing these changes would lead to platform providers having to comply with the requirements for the ethical design, development, and deployment laid down in the AI Act, which include providing platform users with clear and comprehensive instructions on how to use the platform in conformity with best practices and ethical guidelines. Additionally, if their applications were to be considered high-risk AI systems, users would also need to comply with the obligations laid down in the AI Act regardless of personal or professional purposes of such applications.

I argued that implementing these changes would lead to a much stronger framework for the ethical regulation of low-code/no-code AI platforms in the EU. I also concluded that these amendments are examples of anticipatory regulation, in that they aim to bridge the gap between the rapid progress of technology and the slow regulatory process, by setting norms capable of ensuring the ethical design, development, and deployment of low-code/no-code AI

platforms in light of ethical concerns raised by such technologies that, while significant, remain still unexplored.

Lastly, this study contains three main limitations. Firstly, due to the partially emerging nature of low-code/no-code AI platforms, the premises of this research retains a degree of speculation, especially in terms of the ethical concerns raised by this technology. Whether they will sustain the test of time remains to be seen. Until then, its conclusions should be considered tentative. Secondly, this research would have benefitted from a more in-depth investigation of the ethical issues raised by low-code/no-code AI platforms. The main research question helps to (partially) fill this gap. The aim of this thesis was to determine whether the EU regulatory landscape is equipped to mitigate certain ethical concerns. I argue that the unrefined way in which most of the regulatory measures (with the sole exception of the Guidelines for Trustworthy AI) deal with such concerns is indicative of their inability to fully grasp the nuances of the ethical issue sketched in Chapter 3. As such, it is inadequate. However, the extent to which (if at all) my proposed adjustments can enhance the regulatory landscape in this regard ultimately depends upon a more thorough analysis of such issues, which deserve further attention. Thirdly, while outside of the scope of this research, the analysis of the mitigation of the ethical issues raised by low-code/no-code AI platforms would have greatly benefitted from the inclusion of non-regulatory measures alongside regulatory tools.

Appendix 1

The following scenarios illustrate realistic examples of applications that can be developed through low-code/no-code AI platforms. To reflect the peculiarity of low-code/no-code AI platforms in terms of their large user base, the scenarios depict the same application based on increasingly complex development processes.¹⁵ The AI system in question is an application with object recognition purposes, which can be applied to both the security and the business domains. For the sake of completeness, they include a brief description of the circumstances leading to development of the AI systems, and realistic issues that might arise in relation to the application developed.

Scenario 1

An individual (let us call them Alpha) is interested in developing an application which can identify whether a photo or video features one or more of their family members. They do not have a big budget and decide to use a low-code/no-code AI platform to develop such an application themselves in a fast, easy, and inexpensive way. Due to such low-code/no-code AI platforms hiding the codebase powering the development process from their users, Alpha is presented with a finite number of steps in order to create their application. Additionally, due to Alpha lacking programming experience, they are able to rely entirely on the platform's drag-and-drop visual development tools. During the development process, while training the ML model, they are asked to upload an unspecified number of photos of both their family members and other people. Being unfamiliar with the technical features of the application development process, they are unaware of the principles of machine learning. As such, without giving it much thought, they upload one photo for each family member they want the application to be able to identify, which, for reasons of convenience, they select based on each person's social media profile picture, as they are all close-ups of the people in question. As for non-family members, Alpha selects the first ten photos they found after typing "person" in the

¹⁵ As already mentioned in Section 2.1, due to the distinctive features of low-code/no-code AI platforms, the level of sophistication of the applications that can be created through such platforms is not always indicative of the complexity of the underlying development process. As such, sophisticated applications do not necessarily require a complicated development process.

search engine of their laptop and going to the “images” section. Incidentally,¹⁶ all ten photos depict well-lit, fair-skinned individuals. After completing the development process, they deploy their application, with mixed results. In fact, the application works moderately well whenever it is presented with photos and videos of family members where the surrounding environment is well-lit, but does not work just as well whenever photos and videos are taken in the early mornings and evenings. Additionally, the application sometimes does not work whenever a family member is wearing glasses or other noticeable accessories (such as beanies and scarves) that they did not wear in their profile picture, or when they are not wearing accessories that they wore in the photo. Moreover, the application does not work well with dark-skinned family members, and is significantly more likely to misidentify a non-family member for a family member whenever both individuals have a dark skin tone.

Scenario 2

An individual (let us call them Beta) is interested in getting an intelligent security system (ISS) able to identify whenever the people captured by their smart security camera on their front door are family members. They have some basic knowledge of software development, and decide to develop an ISS on their own through a low-code/no-code AI platform. They train a ML model to recognise their family members by uploading photos of both family members and non-family members. Unlike Alpha, Beta provides five photos for every person they want the system to identify. Like Alpha, they proceed to upload ten photos of strangers taken from the “images” section of their browser. All the photos uploaded by Beta show people in a well-lit environment. After completing the development process, they deploy their application with their smart security camera, with mixed results. The application works moderately well whenever family members stand in front of the camera and the surrounding environment is well-lit, but does not work just as well in the early mornings and evenings. Additionally, just like Alpha’s application, it sometimes does not work when family members are wearing noticeable accessories that they consistently did not wear in the pictures uploaded by Beta, as well as when they do not wear accessories that they wore in all the pictures. Once again, the application performs particularly poorly with dark-skinned family members, and is

¹⁶ Arguably, these results are not incidental (Diaz, 2008; Bourg, 2018; Introna & Nissenbaum, 2000).

significantly more likely to misidentify a non-family member for a family member whenever both individuals have a dark skin tone.

Scenario 2'

Among the photos used to train the ML model are personal photos taken during private family gatherings, which have never been published on social media and are not publicly accessible through any other means. One of these photos depicts one of Beta's family members (let us call them Delta) wearing a t-shirt with a racist slogan. After Beta has developed their ISS application, the low-code/no-code AI platform suffers from a data breach during which the photos included in the datasets used to train ML models are made public. As a result of the data breach, Delta's photo becomes publicly available, which leads to them facing workplace discrimination.

Scenario 3

An individual (let us call them Gamma) is a professional developer. Gamma decides to develop an ISS that they will use it in combination with both a smart security camera and a smart lock to not only identify their family members, but let them into the door automatically. Specifically, as their children go to school close to home, Gamma decides to use this application to let their children in automatically when they cannot open the door, e.g. as they are busy or running errands, without the children needing a key. Once again, they upload photos of the family members they would like the ISS to let in automatically, as well as of strangers, using the methodology illustrated in the previous example. However, knowing that the ML model is not infallible, they decide to set stringent benchmarks for the application. Specifically, to lower the risk of false positives, they decide to assign a 90% confidence rate in order for the system to let someone who has been identified as a family member into their home. Once they deploy their application, they are met with mixed results. In addition to the issues displayed by the application as depicted in the previous example, the 90% threshold results in a moderate amount of false negatives whenever family members are either wearing

or not wearing accessories or the environment is poorly lit, and a moderate amount of false positives whenever dark-skinned family members are concerned.¹⁷

Scenario 4

Gamma owns a small business and decides to develop an ISS that will identify employees and let them into their place of work without them needing a key. Additionally, this would avoid distracting fellow employees by having them answer the intercom whenever their colleagues need to enter the building, and would overall improve their focus and efficiency. Gamma decides to avail themselves of a low-code/no-code AI platform to simplify and speed up their work. For every employee, Gamma trains the ML model by uploading hundreds of stills taken from a short video featuring each employee in front of the entrance of the building. Depending on when such videos were shot, some of them are very well-lit, while others are not. They then make use of a pre-existing image dataset that they were able to gain access to online to train the ML model to discern between employees and non-employees, and set the threshold to open the door at 80%. Unlike Beta, they deploy the application in combination with a smart security camera, a smart lock, and a smart visual assistant, to allow for the automation of tasks such as “whenever an employee is at the front door during their shift or within a 30-minute time frame before or after their shift, let them in automatically” and “if an employee is standing at the entrance for more than 10 minutes during their shift or within a 30-minute time frame before or after their shift, notify me”. In this instance, the application seems to be working reasonably well, despite some false negatives, whenever employees are either wearing or not wearing accessories or the environment is poorly lit, and false positives, most often featuring dark-skinned non-employees mistakenly identified as employees.

¹⁷ It should be noted that, as the core elements of the development process illustrated in this scenario are closely linked with the use of ML-as-a-service, the only procedural difference between Alpha and Beta’s approaches relates to the way in which the application is implemented, i.e. with or without a smart security camera and a smart lock. In this sense, it is worth noting that, had Alpha decided to create an ISS, they would have been able to share the ML they trained with a software developer, thus saving both time and money compared to devolving the creation of the entire application to a professional.

Bibliography

- Aarno, D., & Engblom, J. (2014). *Software and system development using virtual platforms*. Waltham, MA: Morgan Kaufmann.
- Aczel, M., Heap, R., Workman, M., Hall, S., Armstrong, H., & Makuch, K. (2022). Anticipatory Regulation: Lessons from fracking and insights for Greenhouse Gas Removal innovation and governance. *Energy Research & Social Science*, 90, 102683. <https://doi.org/10.1016/j.erss.2022.102683>
- Adadi, A., & Berrada, M. (2018). Peeking inside the black-box: A survey on explainable artificial intelligence (XAI). *IEEE Access*, 6, 52138–52160. <https://doi.org/10.1109/ACCESS.2018.2870052>
- AlgorithmWatch. (2021). *Draft AI Act: EU needs to live up to its own ambitions in terms of governance and enforcement*. Available at <https://algorithmwatch.org/en/eu-ai-act-consultation-submission-2021/> (last accessed August 1, 2022)
- Ali, M., Khan, S., Khan, S. U., & Vasilakos, A. V. (2015). Security in cloud computing: Opportunities and challenges. *Information Sciences*, 305, 357–383. <https://doi.org/10.1016/j.ins.2015.01.025>
- Allen, R. H., & Sriram, R. D. (2000). The role of standards in innovation. *Technological Forecasting and Social Change*, 64(2–3), 171–181. [https://doi.org/10.1016/S0040-1625\(99\)00104-3](https://doi.org/10.1016/S0040-1625(99)00104-3)
- Amini, A., Soleimany, A. P., Schwarting, W., Rus, D., & Bhatia, S. N. (2019). Uncovering and mitigating algorithmic bias through learned latent structure. *Proceedings of the 2019 AAAI/ACM Conference on Artificial Intelligence, Ethics, and Society (AIES), 27-28 January, 2019 Honolulu, Hawaii, US*, 19, 289–295. <https://doi.org/10.1145/3306618.3314243>
- Andrada, G., Clowes, R. W., & Smart, P. R. (2022). Varieties of transparency: Exploring agency within AI systems. *AI & Society*. <https://doi.org/10.1007/s00146-021-01326-6>
- Armstrong, H., & Rae, J. (2017). A working model for anticipatory regulation. A working paper. *Nesta*. Available at https://media.nesta.org.uk/documents/working_model_for_anticipatory_regulation_0.pdf (last accessed August 1, 2022)
- Ashraf, C. (2021). Exploring the impacts of artificial intelligence on freedom of religion or belief online. *The International Journal of Human rights*, 26(5), 757–791. <https://doi.org/10.1080/13642987.2021.1968376>

- Aysolmaz, B., Dau, N., & Iren, D. (2020). Preventing algorithmic bias in the development of algorithmic decision-making systems: A Delphi study. *Proceedings of the 53rd Hawaii International Conference on System Sciences*, 5267–5276.
- Babushkina, D. (2020). Robots to Blame? In M. Nørskov, J. Seibt, & O. S. Quick (Eds.), *Culturally Sustainable Social Robotics: Proceedings of Robophilosophy*, 305–315. <https://doi.org/10.3233/FAIA200927>
- Baeza-Yates, R. (2022). Ethical Challenges in AI. *Proceedings of the Fifteenth ACM International Conference on Web Search and Data Mining*, 1–2. <https://doi.org/10.1145/3488560.3498370>
- Baig, M. M. A. (2021). Cloud computing ethical issues: A review paper to investigate and provide suggestions for solving data privacy issues of cloud computing. *Turkish Journal of Computer and Mathematics Education*, 12(7), 1433–1438.
- Barocas, S., & Selbst, A. D. (2016). Big data's disparate impact. *California Law Review*, 104, 671. <https://doi.org/10.15779/Z38BG31>
- Bartoletti, I. (2019). AI in healthcare: Ethical and privacy challenges. *Lecture Notes in Computer Science*, 11526 LNAI, 7–10. https://doi.org/10.1007/978-3-030-21642-9_2
- Benjamin, R. (2019). *Race after technology: Abolitionist tools for the new Jim Code*. Cambridge: Polity Press.
- Bennett, C., & Keyes, O. (2020). What is the point of fairness? *Interactions*, 27(3), 35–39. <https://doi.org/10.1145/3386383>
- Bertino, E., Kundu, A., & Sura, Z. (2019). Data transparency with blockchain and AI ethics. *Journal of Data and Information Quality*, 11(4). <https://doi.org/10.1145/3312750>
- Bjørlo, L., Moen, Ø., & Pasquine, M. (2021). The role of consumer autonomy in developing sustainable AI: A conceptual framework. *Sustainability*, 13, 2332. <https://doi.org/10.3390/su13042332>
- Blacklaws, C. (2018). Algorithms: Transparency and accountability. *Philosophical Transactions Royal Society A*, 376, 21170451. <https://doi.org/10.1098/rsta.2017.0351>
- Bogina, V., Hartman, A., Kuflik, T., & Shulner-Tal, A. (2021). Educating software and AI stakeholders about algorithmic fairness, accountability, transparency and ethics. *International Journal of Artificial Intelligence in Education*. <https://doi.org/10.1007/s40593-021-00248-0>
- Brey, P. (2000). Disclosive computer ethics. *ACM Sigcas Computers and Society*, 30(4), 10–16.

- Brey, P. A. (2012). Anticipatory ethics for emerging technologies. *NanoEthics*, 6(1), 1–13.
- Brey, P., King, O., Jansen, P., Dainow, B., Erden, Y. J., Rodrigues, R., Resseguier, A., Diez Rituerto, M., Hatzakis, T., & Matar, A. (2020). Generalised methodology for ethical assessment of emerging technologies. *Deliverable D6.1 of the SIENNA project*. 112 pp.
- Brey, P., Lundgren, B., Macnish, K., Ryan, M., Andreou, A., Brooks, L., Tilimbe Jiya, Klar, R., Lanzareth, D., Maas, J., Oluoch, I., & Stahl, B. (2021). D3.2 Guidelines for the development and the use of SIS. *SHERPA Project*. <https://doi.org/10.21253/DMU.11316833>
- de Bruin, B., & Floridi, L. (2017). The ethics of cloud computing. *Science and Engineering Ethics*, 23(1), 21–39. <https://doi.org/10.1007/S11948-016-9759-0>
- Brynjolfsson, E., & Mitchell, T. (2017). What can machine learning do? Workforce implications: Profound change is coming, but roles for humans remain. *Science*, 358(6370), 1530–1534. <https://doi.org/10.1126/SCIENCE.AAP8062>
- Burrell, J. (2016). How the machine ‘thinks’: Understanding opacity in machine learning algorithms. *Big Data & Society*, 3(1), 1–12. <https://doi.org/10.1177/2053951715622512>
- Calvo, R. A., Peters, D., Vold, K., & Ryan, R. M. (2020). Supporting human autonomy in AI systems: A framework for ethical enquiry. *Philosophical Studies Series*, 140, 31–54. https://doi.org/10.1007/978-3-030-50585-1_2
- Castets-Renard, C. (2019). Accountability of algorithms in the GDPR and beyond: A European legal framework on automated decision-making. *Fordham Intellectual Property, Media and Entertainment Journal*, 30(1), 91–137. Available at <https://ir.lawnet.fordham.edu/iplj/vol30/iss1/3> (last accessed August 1, 2022)
- Cerrato, P., Halamka, J., & Pencina, M. (2022). A proposal for developing a platform that evaluates algorithmic equity and accuracy. *BMJ Health & Care Informatics*, 29(1), e100423. <https://doi.org/10.1136/bmjhci-2021-100423>
- Christian, B. (2020). *The alignment problem: Machine learning and human values*. New York, NY, US: W. W. Norton & Company
- Coeckelbergh, M. (2020). Artificial intelligence, responsibility attribution, and a relational justification of explainability. *Science and Engineering Ethics*, 26(4), 2051–2068. <https://doi.org/10.1007/s11948-019-00146-8>
- Comandé, G. (2019). Multilayered (accountable) liability for artificial intelligence. *Liability for Artificial Intelligence and the Internet of Things*, 165–184. <https://doi.org/10.5771/9783845294797-165>

- Constantinescu, M., Vică, C., Uszkai, R., & Voinea, C. (2022). Blame it on the AI? On the moral responsibility of artificial moral advisors. *Philosophy & Technology*, 35(2), 35. <https://doi.org/10.1007/s13347-022-00529-z>
- Constantinescu, M., Voinea, C., Uszkai, R., & Vică, C. (2021). Understanding responsibility in Responsible AI. Dianoetic virtues and the hard problem of context. *Ethics and Information Technology*, 23, 803–814. <https://doi.org/10.1007/s10676-021-09616-9>
- Cooper, A. F., Moss, E., Laufer, B., & Nissenbaum, H. (2022). Accountability in an algorithmic society: Relationality, responsibility, and robustness in machine learning. *2022 ACM Conference on Fairness, Accountability, and Transparency, June 21-24, 2022, Seoul, Republic of Korea*. <https://doi.org/10.1145/3531146.3533150>
- Cowgill, B., Dell'Acqua, F., Deng, S., Hsu, D., Verma, N., & Chaintreau, A. (2020). Biased programmers? Or biased data? A field experiment in operationalizing AI ethics. *Proceedings of the 21st ACM Conference on Economics and Computation*, 679–681. <https://doi.org/10.1145/3391403.3399545>
- D'Aquin, M., Troullinou, P., O'Connor, N. E., Cullen, A., Faller, G., & Holden, L. (2018). Towards an 'Ethics by Design' Methodology for AI Research Projects. *Proceedings of the 2018 AAAI/ACM Conference on AI, Ethics, and Society*, 54–59. <https://doi.org/10.1145/3278721.3278765>
- D'Ignazio, C., & Klein, L. F. (2020). *Data feminism*. Cambridge: MIT Press.
- Daneshjou, R., Smith, M. P., Sun, M. D., Rotemberg, V., & Zou, J. (2021). Lack of transparency and potential bias in artificial intelligence data sets and algorithms: A scoping review. *JAMA Dermatology*, 157(11), 1362–1369.
- Danks, D., & London, A. J. (2017). Algorithmic bias in autonomous systems. *Proceedings of the Twenty-Sixth International Joint Conference on Artificial Intelligence*, 4691–4697. <https://doi.org/10.24963/ijcai.2017/654>
- Dastin, J. (2022). Amazon scraps secret AI recruiting tool that showed bias against women. In K. Martin (Ed.), *Ethics of Data and Analytics* (pp. 296–299). Boca Raton: Auerbach Publications. <https://doi.org/10.1201/9781003278290-44>
- Dignum, V. (2017). Responsible autonomy. *Proceedings of the Twenty-Sixth Joint Conference on Artificial Intelligence*, 4698–4704. <https://doi.org/10.24963/ijcai.2017/655>
- Dignum, V., Baldoni, M., Baroglio, C., Caon, M., Chatila, R., Dennis, L., Génova, G., Haim, G., Kließ, M. S., Lopez-Sanchez, M., Micalizio, R., Pavón, J., Slavkovik, M., Smakman, M., van Steenbergen, M., Tedeschi, S., van der Toree, L., Villata, S., & de Wildt, T. (2018). Ethics by Design: Necessity or curse? *Proceedings of the 2018*

- AAAI/ACM Conference on AI, Ethics, and Society*, 60–66. <https://doi.org/10.1145/3278721.3278745>
- Di Ruscio, D., Kolovos, D., de Lara, J., Pierantonio, A., Tisi, M., & Wimmer, M. (2022). Low-code development and model-driven engineering: Two sides of the same coin? *Software and Systems Modeling*, 21(2), 437–446. <https://doi.org/10.1007/s10270-021-00970-2>
- Doerrfeld, B. (2021, June 10). Is low-code/no-code a threat to existing developers? *Acceleration Economy*. Available at <https://accelerationeconomy.com/low-code-no-code/is-low-code-no-code-a-threat-to-existing-developers/> (last accessed August 1, 2022)
- Douglas, D. M., Howard, D., & Lacey, J. (2021). Moral responsibility for computationally designed products. *AI and Ethics*, 1(3), 273–281. <https://doi.org/10.1007/s43681-020-00034-z>
- Ehsan, U., Liao, Q. V., Muller, M., Riedl, M. O., & Weisz, J. D. (2021). Expanding explainability: Towards social transparency in AI systems. *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*, 1–19. <https://doi.org/10.1145/3411764.3445188>
- von Eschenbach, W. J. (2021). Transparency and the black box problem: Why we do not trust AI. *Philosophy & Technology*, 34(4), 1607–1622. <https://doi.org/10.1007/s13347-021-00477-0>
- European Commission. (2018). *A European approach on artificial intelligence*. Available at http://europa.eu/rapid/press-release_MEMO-18-3363_en.htm (last accessed August 1, 2022)
- 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 (Com/2021/206)*.
- European Group on Ethics in Science and New Technologies, European Commission, & Directorate-General for Research and Innovation (2018). Statement on artificial intelligence, robotics and 'autonomous' systems. *European Commission*. Available at <https://op.europa.eu/it/publication-detail/-/publication/dfebe62e-4ce9-11e8-be1d-01aa75ed71a1#> (last accessed August 1, 2022)
- European Union. (2012). *Charter of Fundamental Rights of the European Union*. 2012/C 326/02. Official Journal of the European Union.
- European Union. (2016). *Regulation (EU) 2016/679* (General Data Protection Regulation)
- European Union Agency for Fundamental Rights & Council of Europe. (2018). *Handbook on European non-discrimination law*. Luxembourg: European Union

- Fadaee, M., Gureenkova, O., Barrera, F. R., Schnober, C., Weerkamp, W., & Zavrel, J. (2020). A new neural search and insights platform for navigating and organizing AI research. *Proceedings of the First Workshop on Scholarly Document Processing, Online, November 19, 2020*, 207–213. <https://doi.org/10.18653/v1/P17>
- Faragardi, H. R. (2017). Ethical considerations in cloud computing systems. *Proceedings of the IS4SI 2017 Summit 'Digitalisation for a Sustainable Society', Gothenburg, Sweden, 12–16 June 2017*, 12–16. <https://doi.org/10.3390/IS4SI-2017-04016>
- Fazelpour, S., & Danks, D. (2021). Algorithmic bias: Senses, sources, solutions. *Philosophy Compass*, 16(8). <https://doi.org/10.1111/phc3.12760>
- Ferrer, X., Nuenen, T. van, Such, J. M., Cote, M., & Criado, N. (2021). Bias and discrimination in AI: A cross-disciplinary perspective. *IEEE Technology and Society Magazine*, 40(2), 72–80. <https://doi.org/10.1109/MTS.2021.3056293>
- Feuerriegel, S., Dolata, M., & Schwabe, G. (2020). Fair AI: Challenges and opportunities. *Business & Information Systems Engineering*, 62(4), 379–384. <https://doi.org/10.1007/s12599-020-00650-3>
- Firdhous, M., Ghazali, O., & Hassan, S. (2012). Trust management in cloud computing: A critical review. *International Journal on Advances in ICT for Emerging Regions (ICTer)*, 4(2), 24. <https://doi.org/10.4038/ictcr.v4i2.4674>
- Fjeld, J., Achten, N., Hilligoss, H., Nagy, A. C., & Srikumar, M. (2020). Principled artificial intelligence: Mapping consensus in ethical and rights-based approaches to principles for AI. *Berkman Klein Center for Internet & Society*. Available at https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3518482 (last accessed August 1, 2022)
- Floridi, L., Strait, A. (2020). Ethical foresight analysis. What it is and why it is needed? *Minds & Machines*, 30, 77–97. <https://doi.org/10.1007/s11023-020-09521-y>
- Forester, T., & Morrison, P. (1994). *Computer ethics: Cautionary tales and ethical dilemmas in computing*. Cambridge, London: MIT Press.
- Friedline, T., Naraharisetti, S., & Weaver, A. (2020). Digital redlining: Poor rural communities' access to fintech and implications for financial inclusion. *Journal of Poverty*, 24(5–6), 517–541. <https://doi.org/10.1080/10875549.2019.1695162>
- Garrett, N., Beard, N., & Fiesler, C. (2020). More than 'if time allows': The role of ethics in AI education. *Proceedings of the AAAI/ACM Conference on AI, Ethics, and Society*, 272–278. <https://doi.org/10.1145/3375627.3375868>
- Gartner, (2021, January 22). Forecast analysis. Low-code development technologies. *Gartner Research*. [Not available for distribution]

- Hagendorff, T. (2020). The ethics of AI ethics: An evaluation of guidelines. *Minds and Machines*, 30(1), 99–120. <https://doi.org/10.1007/S11023-020-09517-8>
- Hansson, S. O., Belin, M.-Å., & Lundgren, B. (2021). Self-driving vehicles. An ethical overview. *Philosophy & Technology*, 34(4), 1383–1408. <https://doi.org/10.1007/s13347-021-00464-5>
- Hayes, P., van de Poel, I., & Steen, M. (2022). Moral transparency of and concerning algorithmic tools. *AI and Ethics*. <https://doi.org/10.1007/s43681-022-00190-4>
- Heald, D. (2006). Transparency: The key to better governance? *Proceedings of the British Academy*, 135, 59–73.
- Hopster, J. (2021). What are socially disruptive technologies? *Technology in Society*, 67, 101750. <https://doi.org/10.1016/j.techsoc.2021.101750>
- van den Hoven, J., Doorn, N., Swierstra, T., Koops, B. J., & Romijn, H. (2014). *Responsible innovation 1. Innovative solutions for global issues*. Dordrecht: Springer.
- van den Hoven, J., Vermaas, P. E., & Van de Poel, I. (Eds.). (2015). *Handbook of ethics, values, and technological design: Sources, theory, values and application domains*. Dordrecht: Springer.
- Independent High-Level Expert Group on Artificial Intelligence (HLEG). (2019). *Ethics guidelines for trustworthy AI*. Brussels: European Commission.
- Jansen, P., Brey, P., Fox, A., Maas, J., Hillas, B., Wagner, N., Smith, P., Oluoch, I., Lamers, L., Van Gein, H., Resseguier, A., Rodrigues, R., Wright, D. and Douglas, D. (2019). Ethical analysis of AI and robotics technologies. *Deliverable D4.4 of the SIENNA project*. 226 pp. <https://doi.org/10.5281/zenodo.4068083>
- Jansen, P., Henschke, A., Erden, Y. J., Marchiori, S., Brey, P., and Hoefsloot, M. (2021). Ethics by design and research ethics for AI. *Deliverable D5.7 of the SHERPA project*. 48 pp. <https://doi.org/10.5281/zenodo.4068071>
- Jensen, S. R., Nagel, S., Brey, P., Kudlek, K., Ditzel, T., Oluoch, I., Zuiderduin, A.-C., Wagner, N.-F., Wiseman, H., and Miah, A. (2020). Ethical analysis of human enhancement technologies. *Deliverable D3.4 of the SIENNA project*. 158 pp. <https://doi.org/10.5281/zenodo.4068071>
- John-Mathews, J. M., Cardon, D., & Balagué, C. (2022). From reality to world. A critical perspective on AI fairness. *Journal of Business Ethics*, 178, 945–959. <https://doi.org/10.1007/S10551-022-05055-8>
- Johnson, G. M. (2021). Algorithmic bias: On the implicit biases of social technology. *Synthese*, 198(10), 9941–9961. <https://doi.org/10.1007/s11229-020-02696-y>

- Johnson, K. (2021). The fight to define when AI is ‘high risk’. *Wired*. Available at <https://www.wired.com/story/fight-to-define-when-ai-is-high-risk> (last accessed August 1, 2022)
- Jones, M. L., Kaufman, E., & Edenberg, E. (2018). AI and the ethics of automating consent. *IEEE Security & Privacy*, 16(3), 64–72. Available at <https://ieeexplore.ieee.org/abstract/document/8395123/> (last accessed August 1, 2022)
- Kearns, M., & Roth, A. (2019). *The ethical algorithm: The science of socially aware algorithm design*. New York: Oxford University Press.
- King, B. (2004). Performance assurance for IT systems. *Performance Assurance for IT Systems*, 1–350. <https://doi.org/10.4324/9780203334577>
- Koene, A. (2017). Algorithmic bias: Addressing growing concerns. *IEEE Technology and Society Magazine*, 36(2), 31–32. <https://doi.org/10.1109/MTS.2017.2697080>
- Kordzadeh, N., & Ghasemaghaei, M. (2022). Algorithmic bias: Review, synthesis, and future research directions. *European Journal of Information Systems*, 31(3), 388–409. <https://doi.org/10.1080/0960085X.2021.1927212>
- Kossow, N., Windwehr, S., & Jenkins, M. (2021). Algorithmic transparency and accountability. *Transparency International Anti-Corruption*. Available at https://knowledgehub.transparency.org/assets/uploads/kproducts/Algorithmic-Transparency_2021.pdf (last accessed August 1, 2022)
- Koulu, R. (2020). Human control over automation: EU Policy and AI Ethics. *European Journal of Legal Studies*, 12(1), 9–46. <https://doi.org/10.2924/EJLS.2019.019>
- Kroll, J. A. (2020). Accountability in computer systems. In M. D. Dubber, F. Pasquale, & S. Das (Eds.), *The Oxford Handbook of the Ethics of AI* (pp. 181-196). New York: Oxford University Press
- Lee, R. S. T. (2020). AI ethics, security and privacy. *Artificial Intelligence in Daily Life*, 369–384. https://doi.org/10.1007/978-981-15-7695-9_14
- Leslie, D. (2019). *Understanding artificial intelligence ethics and safety. A guide for the responsible design and implementation of AI systems in the public sector*. The Alan Turing Institute. <https://doi.org/10.5281/zenodo.3240529>
- Lurie, Y., & Mark, S. (2015). Professional ethics of software engineers: An ethical framework. *Science and Engineering Ethics*, 22(2), 417–434. <https://doi.org/10.1007/S11948-015-9665-X>
- Mason, R. O. (1986). Four ethical issues of the information age. *Management Information Systems Quarterly*, 10(1), 5–12. <https://doi.org/10.2307/248873>

- Matthias, A. (2004). The responsibility gap: Ascribing responsibility for the actions of learning automata. *Ethics and Information Technology*, 6(3), 175–183. <https://doi.org/10.1007/S10676-004-3422-1>
- McManus, R. M., & Rutchick, A. M. (2019). Autonomous vehicles and the attribution of moral responsibility. *Social Psychological and Personality Science*, 10(3), 345–352. <https://doi.org/10.1177/1948550618755875>
- Mehrabi, N., Morstatter, F., Saxena, N., Lerman, K., & Galstyan, A. (2021). A survey on bias and fairness in machine learning. *ACM Computing Surveys*, 54(6), 1–35. <https://doi.org/10.1145/3457607>
- Mökander, J., & Axente, M. (2021). Ethics-based auditing of automated decision-making systems: Intervention points and policy implications. *AI & Society*. <https://doi.org/10.1007/s00146-021-01286-x>
- Mökander, J., & Floridi, L. (2021). Ethics-based auditing to develop trustworthy AI. *Minds and Machines*, 31(2), 323–327. <https://doi.org/10.1007/s11023-021-09557-8>
- Mökander, J., Morley, J., Taddeo, M., & Floridi, L. (2021). Ethics-based auditing of automated decision-making systems: Nature, scope, and limitations. *Science and Engineering Ethics*, 27(4), 44. <https://doi.org/10.1007/s11948-021-00319-4>
- Mokrian, K., & Schuelke-Leech, B.-A. (2021). Ethical decision-making responsibility in Canadian autonomous vehicle policies. *2021 IEEE International Symposium on Technology and Society (ISTAS)*, 1–7. <https://doi.org/10.1109/ISTAS52410.2021.9629123>
- Moor, J. H. (2005). Why we need better ethics for emerging technologies. *Ethics and Information Technology*, 7(3), 111-119. <https://doi.org/10.1007/s10676-006-0008-0>
- Moorthy, J., Lahiri, R., Biswas, N., Sanyal, D., Ranjan, J., Nanath, K., & Ghosh, P. (2015). Big data: Prospects and challenges. *Vikalpa: The Journal for Decision Makers* 40(1), 74–96. <https://doi.org/10.1177/0256090915575450>
- Mucha, T., & Seppala, T. (2020). Artificial intelligence platforms. A new research agenda for digital platform economy. *ETLA Working Papers* 76, 1-14.
- Murphy, B., & Rocchi, M. (2021). Ethics and cloud computing. In T. Lynn, J. G. Mooney, L. van der Werff, & G. Fox (Eds.), *Data Privacy and Trust in Cloud Computing. Palgrave Studies in Digital Business & Enabling Technologies* (pp. 105-128). Cham: Palgrave Macmillan. https://doi.org/10.1007/978-3-030-54660-1_6
- Noble, S. U. (2018). *Algorithms of oppression. How search engines reinforce racism*. New York: New York University Press.

- van Nuenen, T., Ferrer, X., Such, J. M., & Cote, M. (2020). Transparency for whom? Assessing discriminatory artificial intelligence. *Computer*, 53(11), 36–44. <https://doi.org/10.1109/MC.2020.3002181>
- OECD. (2019). *Recommendation of the Council on Artificial Intelligence*, OECD/LEGAL/0449
- O’Neil, C. (2016). *Weapons of math destruction: How big data increases inequality and threatens democracy*. Portland: Broadway Books.
- Owen, R., Macnaghten, P., & Stilgoe, J. (2020). Responsible Research and Innovation: From Science in Society to Science for Society, with Society. In *Emerging Technologies: Ethics, Law and Governance* (pp. 117-126). Routledge.
- van de Poel, I., & Verbeek, P.-P. (2006). Editorial: Ethics and engineering design. *Science, Technology, & Human Values*, 31(3), 223–236. <https://doi.org/10.1177/0162243905285838>
- Poulton, E. (1994). *Behavioral decision theory: A new approach*. New York: Cambridge University Press
- Resseguier, A., Brey, P., Dainow, B., & Drozdowska, A. (2021). Multi-stakeholder strategy for ethical AI and robotics. *Annex 1 to Deliverable D5.4 of the SIENNA project*. 20 pp.
- Rivera, L. A. (2012). Hiring as cultural matching: The case of elite professional service firms. *American Sociological Review*, 77(6), 999–1022. <https://doi.org/10.1177/0003122412463213>
- Robert Jr., L. P., Bansal, G., Melville, N., & Stafford, T. (2020). Introduction to the special issue on AI fairness, trust, and ethics. *AIS Transactions on Human-Computer Interaction*, 12(4), 172–178. <https://doi.org/10.17705/1thci.00134>
- Roberts, H., Cowls, J., Hine, E., Mazzi, F., Tsamados, A., Taddeo, M., & Floridi, L. (2021). Achieving a ‘good AI society’: Comparing the aims and progress of the EU and the US. *Science and Engineering Ethics*, 27, 68. <https://doi.org/10.1007/s11948-021-00340-7>
- Roscher, R., Bohn, B., Duarte, M. F., & Garcke, J. (2020). Explainable machine learning for scientific insights and discoveries. *IEEE Access*, 8, 42200–42216. <https://doi.org/10.1109/ACCESS.2020.2976199>
- Ryan, M. (2020). In AI we trust. Ethics, artificial intelligence, and reliability. *Science and Engineering Ethics*, 26(5), 2749-2767. <https://doi.org/10.1007/s11948-020-00228-y>
- Ryan, M., Brey, P., Macnish, K., Hatzakis, T., King, O., Maas, J., Haasjes, R., Fernandez, A., Martorana, S., Oluoch, I., Eren, S., and Van Der Puil, R. (2019). Ethical tensions

- and social impacts [of smart information systems]. *Deliverable D1.4 of the SHERPA project*. 165 pp. <https://doi.org/10.21253/DMU.8397134>
- Salminen, J., Froneman, W., Jung, S. G., Chowdhury, S., & Jansen, B. J. (2020). The ethics of data-driven personas. *Proceedings of the Conference on Human Factors in Computing Systems*, 1–9. <https://doi.org/10.1145/3334480.3382790>
- Salvaris, M., Dean, D., & Tok, W. H. (2018). *Deep learning with Azure. Building and deploying artificial intelligence solutions on the Microsoft AI Platform*. New York: APress
- Santoni de Sio, F., & Mecacci, G. (2021). Four responsibility gaps with artificial intelligence: Why they matter and how to address them. *Philosophy & Technology*, 34(4), 1057–1084. <https://doi.org/10.1007/s13347-021-00450-x>
- Santosh, K., & Gaur, L. (2021). Privacy, security, and ethical issues. In *Artificial Intelligence and Machine Learning in Public Healthcare. Opportunities and Societal Impact* (pp. 65–74). https://doi.org/10.1007/978-981-16-6768-8_8
- Schmidt, P., & Biessmann, F. (2020). Calibrating human-AI collaboration: Impact of risk, ambiguity and transparency on algorithmic bias. In A. Holzinger, P. Kieseberg, A. M. Tjoa, & E. Weippl (Eds.), *Machine Learning and Knowledge Extraction*, 12279, 431–449. https://doi.org/10.1007/978-3-030-57321-8_24
- Sekiguchi, K., & Hori, K. (2018). Organic and dynamic tool for use with knowledge base of AI ethics for promoting engineers' practice of ethical AI design. *AI & Society*, 35, 51-71. <https://doi.org/10.1007/s00146-018-0867-z>
- Shah, H. (2018) Algorithmic accountability. *Philosophical Transactions of the Royal Society A*, 376, 20170362. <https://doi.org/10.1098/rsta.2017.0362>
- Simon, H. A. (1979). *Models of Thought*. New Haven: Yale University Press.
- Stahl, B. C. (2021a). Addressing ethical issues in AI. In *Artificial intelligence for a better future. An ecosystem perspective on the ethics of AI and emerging digital technologies* (pp. 55–79). https://doi.org/10.1007/978-3-030-69978-9_5
- Stahl, B. C. (2021b). Ethical Issues of AI. In *Artificial intelligence for a better future. An ecosystem perspective on the ethics of AI and emerging digital technologies* (pp. 35–53). https://doi.org/10.1007/978-3-030-69978-9_4
- Stahl, B. C., Antoniou, J., Ryan, M., Macnish, K., & Jiya, T. (2022). Organisational responses to the ethical issues of artificial intelligence. *AI and Society*, 37(1), 23–37. <https://doi.org/10.1007/S00146-021-01148-6>
- Stahl, B. C., & Flick, C. (2011). ETICA workshop on computer ethics. Exploring normative issues. In S. Fischer-Hübner, P. Duquenoy, M. Hansen, R. Leenes, & G. Zhang

- (Eds.), *Privacy and identity management for life, Volume 352* (pp. 64–77). Berlin: Springer. https://doi.org/10.1007/978-3-642-20769-3_6.
- Stahl, B., Jirotko, M., & Eden, G. (2013). Responsible research and innovation in information and communication technology. Identifying and engaging with the ethical implications of ICTs. In R. Owen, J. Bessant, & M. Heintz (Eds.), *Responsible Innovation: Managing the Responsible Emergence of Science and Innovation in Society* (pp. 199-218). <https://doi.org/10.1002/9781118551424.ch11>
- Stark, L., & Tierney, M. (2014). Lockbox: Mobility, privacy and values in cloud storage. *Ethics and Information Technology, 16*, 1–13. <https://doi.org/10.1080/01972243.2015.1107167>
- Stahl & Wright, 2018
- Sullins, J. P. (2011). When is a robot a moral agent? In M. Anderson & S. L. Anderson (Eds.), *Machine Ethics* (pp. 151–161). Cambridge University Press. <https://doi.org/10.1017/CBO9780511978036.013>
- Tait, J., & Banda, G. (2014). *Proportionate and adaptive governance of innovative technologies*. BSI Group, Department for Business Innovation & Skills, Innogen.
- Talbert, M. (2019). Attributionist theories of moral responsibility. In D. K. Nelkin & D. Pereboom (Eds.), *The Oxford Handbook of Moral Responsibility* (pp. 53-70). Masquis: Oxford University Press
- Thomson, A. J., & Schmoldt, D. L. (2001). Ethics in computer software design and development. *Computers and Electronics in Agriculture, 30*(1–3), 85–102. [https://doi.org/10.1016/S0168-1699\(00\)00158-7](https://doi.org/10.1016/S0168-1699(00)00158-7)
- Tigard, D. W. (2021). Responsible AI and moral responsibility: A common appreciation. *AI and Ethics, 1*(2), 113–117. <https://doi.org/10.1007/s43681-020-00009-0>
- Tilmes, N. (2022). Disability, fairness, and algorithmic bias in AI recruitment. *Ethics and Information Technology, 24*(2), 21. <https://doi.org/10.1007/s10676-022-09633-2>
- Toivonen, T., Jormanainen, I., Kahila, J., Tedre, M., Valtonen, T., & Vartiainen, H. (2020). Co-designing machine learning apps in K–12 with primary school children. *2020 IEEE 20th International Conference on Advanced Learning Technologies (ICALT)*, 308–310. <https://doi.org/10.1109/ICALT49669.2020.00099>
- Turilli, M., & Floridi, L. (2009). Cloud computing and its ethical challenges. *SSRN Electronic Journal*. Available at https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3850031 (last accessed August 1, 2022)

- Ulnicane, I. (2022). Artificial intelligence in the European Union: Policy, ethics and regulation. In T. Hoerber, G. Weber, & I. Ulnicane (Eds.), *The Routledge Handbook of European Integrations* (pp. 254–269). <https://doi.org/10.4324/9780429262081-19>
- Wagner, B. (2018). Ethics as an escape from regulation. From “ethics-washing” to ethics-shopping? In E. Bayamlioglu, I. Baraliuc, L. Janssens (Eds.), *Being Profiled. Cogitas Ergo Sum. 10 Years of ‘Profiling the European Citizen’* (pp. 84–88). Amsterdam: Amsterdam University Press. <https://doi.org/10.25969/mediarep/13281>
- Walmsley, J. (2021). Artificial intelligence and the value of transparency. *AI & Society*, 36(2), 585–595. <https://doi.org/10.1007/s00146-020-01066-z>
- Whittlestone, J., Belfield, H., Ó hÉigeartaigh, S., Maas, M., Hagerty, A., Burden, J., & Avin, S. (2021, April 28). Comment on the EU’s world-first AI regulation: ‘An historic opportunity’. *Centre for the Study of Existential Risk, University of Cambridge*. Available at <https://www.cser.ac.uk/news/eus-world-first-ai-regulation/> (last accessed August 1, 2022)
- Whitworth, B., & de Moor, A. (2003). Legitimate by design: Towards trusted socio-technical systems. *Behavior and Information Technology*, 22(1), 31–51.
- Wilms, H. C. (2014). The assumption of scientific responsibility by ethical codes. An European dilemma of fundamental rights. In J. van den Hoven, N. Doorn, T. Swierstra, B.-J. Koops, & H. Romijn (Eds.), *Responsible Innovation 1* (pp. 89–96). Dordrecht: Springer. https://doi.org/10.1007/978-94-017-8956-1_6
- Yapo, A., & Weiss, J. (2018). Ethical implications of bias in machine learning. *Hawaii International Conference on System Sciences 2018 (HICSS-51)*. Available at https://aisel.aisnet.org/hicss-51/os/topics_in_os/6 (last accessed August 1, 2022)
- Yablonsky, S. A. (2020). AI-driven digital platform innovation. *Technology Innovation Management Review*, 10(10), 4-15. <http://doi.org/10.22215/timreview/1392>
- Zhang, Y., Wu, M., Tian, G. Y., Zhang, G., Lu, J. (2021). Ethics and privacy of artificial intelligence: Understandings from bibliometrics. *Knowledge-based Systems*. <https://doi.org/10.1016/j.knosys.2021.106994>