

Some Ethical Reflections on the EU AI Act

Marc M. Anderson¹

¹ LORIA, UMR 7503, Université de Lorraine, Inria and CNRS, Campus Scientifique, 615 Rue du Jardin-Botanique, 54506 Vandœuvre-lès-Nancy, FRANCE

Abstract

This article considers the European Commission AI Act proposal development process and the aims and conceptual vision explicitly and implicitly embedded in the Act, specifically from the point of view of ethics as a practice. I argue that the often-encountered tendency by many authors to invoke ethics and law together ‘in the same breath,’ needs to be re-examined, both with regard to ethics of AI, and here, in a more concrete context, with regard to the AI Act proposal. On the basis of sustained reflection on the Act in light of ethics, there are a number of reasons to reject the claim of the AI Act regulation as being ethically grounded. Among these reasons are: the proposal’s characterization of the objectives of the AI Act, the proposal’s vision and use of public and even democratic consultation, and finally, the arbitrary embedding of what I shall call the ‘speed paradigm’ in both the regulatory process and the view of AI as a technology.

Keywords

EU AI Act, Ethics, Artificial Intelligence, Law

1. Introduction

The 2021 European Commission *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*² [5] is a major step toward the regulation of AI, which some legal scholars, e.g. Battista [3] have called ‘pioneering.’ The Act is not yet adopted, and is still in a process of evolving [10]. It has already been criticized in depth by certain legal scholars, most prominently Veale and Borgesius [18].

My aim here is not to review the specific choices that the creators of the AI Act have made in terms of the scope, the obligations, the enforcement, or the harm categorization approach adopted (Risk). It is rather to consider the AI Act development process and the explicit and implicit aims and conceptual vision imbedded in the Act. Moreover, I will specifically consider them from the point of view of ethics as a practice.

There is a recurring tendency to invoke ethics and law together ‘in the same breath,’ both with regard to ethics of AI, and here with regard to the AI Act, e.g. Townsend [17]. That tendency is at work within the AI Act proposal as well. But I will argue that *on the basis of reflection on ethics as practice, there are a number of reasons to reject the claim of the AI Act regulation as being ethically grounded.* Among these reasons are: the proposal’s characterization of the objectives of the AI Act, the proposal’s vision and use of public consultation, and finally, the embedding of what I shall call the ‘speed paradigm’ in both the regulatory process and view of AI as a technology. The latter, briefly, is the unconsidered assumption that a process or development which appears inherently fast – here the technology of AI – necessitates a social engagement and integration into human affairs whose tempo matches the process in question, thus setting up a self-perpetuating cycle.

IAIL 2022: 1st International Workshop on Imagining the AI Landscape After the AI Act, June 13, 2022, Amsterdam, Netherlands
EMAIL: marc.anderson@inria.fr (M. M. Anderson)



© 2020 Copyright for this paper by its authors.
Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).
CEUR Workshop Proceedings (CEUR-WS.org)

² Hereafter ‘AI Act’ or ‘Act’

2. The Relation between Law and Ethics advanced by the AI Act

The link made between the AI Act as regulatory legislation and the ethical explorations of the (European Commission) is clear and unambiguous in intent. The explanatory memorandum states that the Act builds upon the preparatory work of the High-Level Expert Group on AI (HLEG), whose task, among others, was to prepare the HLEG ethics guidelines for Trustworthy AI (2019). The latter took into account “more than 500 submissions from stakeholders. The key requirements reflect a widespread and common approach, as evidenced by a plethora of ethical codes and principles developed by many private and public organizations in Europe and beyond, that AI development and use should be guided by certain essential value-oriented principles” [5]. Again, later: “the proposed minimum requirements are already state-of-the-art for many diligent operators and the result of two years of preparatory work, derived from the Ethics Guidelines of the HLEG” [5]. Thus, the proposed regulatory legislation is characterized as part of ‘a process which begins in ethics.’ Strong emphasis is placed on the ‘consultative’ nature of the process moreover. Finally, the proposed Regulation is viewed as ‘a supporting foundation for ethical efforts’: “[the Regulation] ensures the protection of ethical principles, as specifically requested by the European Parliament” [5].

The consultation aspect of the process is not specifically described as a democratic aspect, but that reading is emphasized indirectly in the careful and detailed description of the online public consultation which was carried out, with specific mention of *citizens*, and *representation* [5]. Bradford has pointed out that in a general sense, the EU tends to emphasize: “the strong democratic backing for its regulatory stance. The European Commission has described the EU’s commitment to further its social agenda as part of its trade policy as ‘forging collective preferences’” [4]. There is thus a good weight of evidence for viewing the consultation process of the EU AI Act and of the preparatory work, as subsidiary democratic processes.

A self-supporting ‘loop’ is thus envisioned, with ethics feeding into law, followed by law ‘solidifying and protecting’ ethics, all the while informed by a broad plurality of representative input from stakeholders, including citizens. Parts of this loop have been argued for elsewhere. Ferretti, for example, has argued from several angles that, as an ethical priority, private agents should help governments in creating legislation regarding AI, since governments have advantages of scale and efficiency above all [6]. But if ethics is viewed in this light, there are inconsistencies and contradictions which arise, as I shall argue below.

3. Specific Ethical Problems within the AI Act or its Development

In what follows I will reflect upon some of the ethical problems which the AI Act gives rise to. The issues noted are not exhaustive. Nor are they on the level of ethical consequences arising from the technical formulation of the articles of the Act. Rather they are at the level of the interplay between law and ethics, in terms of contradictions and tensions between intentions, explicit and implicit, developed in the Act. The legal technical tensions and contradictions within the Act itself are well summarized and considered in Veale and Borgesius [18]. Satisfactory ethical treatments of the more specific formulations of the Act, beyond summary outlines, are difficult to find, but will no doubt accumulate gradually.

3.1. The AI Act Objectives

The primary objective of the AI Act, legal consistency, is betrayed indirectly in the text of the Act. Even though the authors of the document lay out four specific objectives in the Explanatory Memorandum of the proposal, there are indications in the document that achieving “legal certainty to facilitate investment and innovation in AI” [5] acts as a *primus inter pares* which grounds the objectives. We find, for example “The primary objective of this proposal is to ensure the proper functioning of the internal market by setting harmonized rules . . .” [5]. And further, “the objectives of this proposal can be better achieved at Union level to avoid a further fragmentation of the Single Market into potentially

contradictory national frameworks . . . only common action at Union level can also protect the Union’s digital sovereignty and leverage its tools and regulatory powers to shape global rules and standards” [5]. Thus, the primary objective of the Act coalesces around internal legal consistency, which is to lead in turn to unambiguous sovereignty over AI development within the EU. The offshoot of this objective is also revealed indirectly as a wielding of regulatory might to mold development practices beyond European boundaries. This tendency is indirect here, but it is not accidental [4].

How does this square with the ostensible inclusion of ethics within the process of AI regulation development? Arguably not very well. It reduces the ethical contribution to a secondary and weak contribution. Law may not need ethics at all to achieve internal legal consistency, as legal positivists hold. In that view the latter process is purely a matter for law itself. It can be argued that law does need ethics, in the form of moral correctness, as Robert Alexy argues for example. Alexy suggests that discourse theory provides a middle way between the impossible demand of proving the connection of moral norms to law and viewing legal content as arbitrary. He also recognizes that there can only be an approximative approach toward reasonable discourse as an ideal, never a reaching of it [2]. The strongest approximations to the ideal which he presents – those that give moral clarity in terms of the impossibility or necessity of certain norms – are nonetheless only strong relative to the accepted framework of discourse theory, i.e. however strong in terms of internal consistency with the accepted framework of discourse theory, they are no more than norms for the possibility of discourse itself. The rest are reasonably possible opinions from which it is difficult to choose, since discourse theory cannot pass its own boundary of reasonableness, as Alexy admits. And even these opinions are ineffectual unless externalized into legal certainty applying to all, a process which morality demands in its weakness [2].

On a very generous reading – an acceptance of discourse theory as a solution for a moral backing to law – could we view the development process of the AI Act in light of something like the process described by Alexy? Only very superficially so. We indeed have the ‘word’ of the framers of the AI Act that the moral groundwork preceded the legal process. But as we saw above, legal certainty is strongly emphasized in the AI Act, both directly and indirectly, and in ways which should lead us to suspect that it was this urge to legal certainty which acted as the true goad to the AI Act development.

If that is the generous reading, the less generous is this: that the objective of legal certainty colored the ostensibly foundational moral work which it appealed to in terms of the process of that foundational work and in terms of the latter’s outcome, and in a number of ways. First, there is no formal, or even tacit, acceptance of discourse theory to be found in discussions of the ethical consultations around AI which are said to precede the AI Act, nor such an acceptance of any other ethical theory. Second, the development process of the AI Act appeals to law first in its urge to be democratic and majoritarian, a point I shall expand further upon below, as indeed did the ethical work of the High-Level Expert Group. Third, the solidification of the AI Act, in terms of content³ and definitions, including purported ethical content, is a passing beyond the border of a general framework of agreement for ethical discourse and a passing beyond which appeals to lawlike processes. In other words, rather than law drawing upon ethics as it finds it, the process undergone is one which explicitly tries to engage ethics in order to have it solidified into law, because ethics is too weak for what is desired. Fourth, there are whole areas of AI application – military applications being the most obvious – which are arbitrarily removed from engagement in what on a generous discourse theory reading would be the field of ‘reasonably possible opinions.’

Thus, even if we accepted the approach of Alexy, contrary to simply admitting an unashamedly positivistic legal view, nonetheless, here, in the development process of the AI Act, the appeal to ethics is not approached as an ideal of discourse which might guide ethical debate and consolidation. Rather it is approached as a mere open process of consultation among all interested stakeholders, some of which need not have any interest in a framework of reasonable discourse. Indeed, as ‘stakeholders’ accepted into the consultation, they are already accepted as passing beyond the framework of such a realm of ethical discourse.

In other words, taken from various angles – more on these angles later – it is an appeal to ethics which itself, in the notions of interests, representation, and consultation, recurs to a process common to the development of legislation, a lawlike process applied to ethics. It is thus, as I will argue further,

³ And this solidification includes the exclusion of certain areas of AI application from consideration by the Act.

plagued by tensions which an *ethics as law process approach* subject it to, e.g. the vastly superior and inferior scales of knowledge available to the different stakeholders, and the difference between stakeholders as natural individuals and as enterprises. Thus, the appeal to ethics here is the appeal to a chimaera.

The driver of this urge to internal legal consistency is not presented as a matter of fairness, which it might be if the emphasis were on the ethical good of all ‘stakeholders’ contributing to the regulatory process. On the contrary the avowed primary objective of the proposal, emphasizes that aspect which is most likely to give rise to unfairness relative to stakeholders, i.e. the ‘ease of investment and innovation’ which legal certainty will bring to the market for AI. That aspect immediately favors the cohort of enterprises as stakeholders.

3.2. Consultation and Democratization in Law and Ethics

The Act makes much of the process of consultation undergone preliminary to its development, as I noted above. This consultation process in itself has consequences both in terms of the effect upon the inclusion of ethics and in terms of inconsistencies between the use of ethics claimed by the regulatory process here and the actual use of ethics.

Some who will be affected by AI will not have had access to submit their views. Others would not have submitted views, simply because they have a limited or non-existent understanding of how AI affects them. A recent Australian poll, which can probably be taken to be indicative of the average level of knowledge of AI in European populations, indicates that 61% of Australians had a low understanding of AI issues [11].

The comparison with democratic style consultation relative to law raises larger issues. In the European civil law tradition, law is democratic insofar as derived from legislation imposed by elected assemblies. In the AI Act this tendency is ostensibly reinforced at another level through direct consultation of interested parties during the development of that law. Those parties included citizens, but by no means as a majority. They also included a large cohort of businesses within which a majority were large enterprises, including some of the largest industry players based outside of Europe, e.g. Qualcomm, Intel, Google, IBM, and Microsoft. Given that, there are contradictory tendencies at play.

On the one hand, if it be argued that in its purported links to law, ethics can be supported by democratic practices, then in the democratization of the process toward legislation of the AI Act – the consultative process – we have a situation where the weight of opinion in terms of opinions and position papers cannot have been anything but unbalanced in favor of special interests. On the other hand, if we assume that the democratized process of consultation was internally balanced by those who formulated the Act, we can still ask whether such a process is ethical in a strong enough sense to back up the desired linkage between law and ethics. In other words, in the first instance, we may have an ethical appeal to democratic practices which were not very democratic, i.e. *‘an equal say’ among participants, in which some enjoy massive advantages over others that directly contradict the supposed move toward equality*. In the second instance, we may accept that they were democratic practices but insist that such practices have little to do with ethics, i.e. *we adopt democratic practices in order to be ethical, but adopting them both plays favorites ethically and weakens ethics*.

Let us consider the first instance. The discussion of the consultation process is clear that the will of the majority was taken into account in formulating the proposal, e.g. agreement on the need for action, request for a precise definition of AI, favoring of a risk-based approach, and specific types of enforcement [5]. The majority was a stakeholder majority however, rather than a majority among equals; the large enterprises in particular commanding far greater resources and arguably greater understanding of the legal process and technical aspects of AI, than most individual citizens. Moreover, it was a majority in which many of the parties, including some of the largest parties giving opinions – the largest corporations – were engaging the process from outside of the EU jurisdiction where the democratization of the process was putatively occurring through consultation. Accordingly, outside opinions might not have the best interests of the individuals in the process as primary motives. There is thus a tension with regard to the question: who is the AI Act really for? As noted earlier, the answer, directly or indirectly, may be business interests primarily. This in itself may not be an issue if it is

clearly stated and if it remains within the legal realm. It becomes an issue however, for ethics, if the link between ethics and law is invoked in order to suggest a false sense of ethical responsibility in terms of the public good.

If we move to the second instance, we see that at the very least the democratization of the process, tends to skew it toward a utilitarian ethical outlook. Riley, for example, notes many thinkers who accepted or argued for the link between Utilitarianism and democracy [16].

John Stuart Mill himself upheld this link, but he also added caveats which serve as a locus for questioning this democratization of the development process relative to ethics. First, he disagreed that democratic participation should have equal weight [12], a viewpoint which, if countenanced, leads easily enough into justifying the unbalanced participation of special interests in the process noted above. Second, he viewed – without prejudice to the process – the results of the process of democratization as desultory. “No government by a democracy . . . either in its political acts or in the opinions, qualities, and tone of mind which it fosters, ever did or could rise above mediocrity, except in so far as the sovereign Many have let themselves be guided (which in their best times they always have done) by the counsels and influence of a more highly gifted and instructed One or Few” [13] If Mill is right here, then adopting the process of democratization has a corresponding weakening effect upon the results of consultative processes applied to legal regulation such as the AI Act. The urge toward a linkage of ethics and law, subsequently transfers this watering down effect into the results of ethical efforts derived through public consultation. Mittelstadt [14] and Hagendorff [8] have highlighted the weakness in achieving only high-level ethical principles, although most authors have not linked this explicitly to a democratized consultative process.

Moreover, if the legal regulation through consultation is a variety of democratization and if the latter is strongly associated with Utilitarianism, and influences public ethics, then public ethics itself is subject to a variety of special interest, i.e. special interest internal to ethics as theory. In other words, this trajectory from accepted practice in law to accepted practice in ethics⁴ amounts practically to favoring one type of ethical theory, utilitarianism, over the various other historic and contemporary ethical traditions. But theoretical ethical foundations are multiple and disputed. It is doubtful whether the ethics of the High-Level Guidelines for Trustworthy AI (HLEG) for example, [9] which the Act refers to, could be distilled from the one ethical theory. But if not, it is a mistake to evoke the link between ethics and law.⁵

3.3. The Need for Speed

One of the implicit developmental aims of the AI Act is to implement binding regulation upon the field of AI development in a swift and timely manner. This need is characterized indirectly in terms of market forces, in terms of choosing regulatory options which reduce costs for compliance so as not to slow compliance. The upshot is that: “the European Union will continue to develop a fast-growing AI ecosystem of innovative services and products embedding AI technology or stand-alone AI systems . . .” [5]. Moreover, the definition of AI is chosen “taking into account the fast technological and market developments related to AI” [5]. And what is the primary characteristic of that definition? It is – as sought by stakeholders – a narrow and exact definition of AI, which the Act calls “a single future-proof definition.” [5] In other words, the definition is to obviate the need for its own future evolution and development.

Thus, similarly to the self-supporting loop sought for the ethical and legal processes, here we have a self-supporting loop envisioned in terms of the tempo of the regulatory engagement of AI. AI is, de facto seemingly, *a fast evolving and growing technology*, the defining characteristic of the chosen definition of AI is that it takes this characteristic into account, in fact *makes an end run around it by being ‘future-proof’*, and finally, the definition allows *a swiftness of market compliance and adaptation*

⁴ And in fact, it is also invoked in the other direction in the AI Act, i.e. ethics is viewed as being behind the regulatory development process, and then, indirectly, law becomes a support for similar consultative processes in deriving ethical frameworks.

⁵ At another level, which is too broad to address here, democratic processes have historically instituted un-ethical actions often enough, the classic case of the sentencing of Socrates springing to mind immediately.

to the regulation which further facilitates the development of AI as a fast-growing technology, thus closing the loop.

From a legal standpoint, this emphasis on speed may or may not be acceptable. What I am interested in here is the ethical status of this interwoven tendency toward speedy results and developments. We can engage it first from the angle of the definition. Without providing a specific definition of AI, we can suggest that there are at least two ways of viewing AI with regard to tempo. The word *intelligent* has indeed been equated with speed, with notions such as quick witted, quick to comprehend, etc. But it has just as often been equated with the slow and more ponderous gathering and processing of information, e.g. in military intelligence. In fact, *wisdom*, with little or no connotation of speed, has sometimes served as a passable substitute for intelligence, and the notion of wisdom as a deeper seeing, a disclosure and understanding of hidden inner relationships, can arguably be applied to many contemporary uses of algorithms, particularly in deep learning.

When we move to the technological development of AI, the issue is even more clear. The technological development of AI as a swift process is not a de facto state of affairs. It may very well be a recent state of affairs in terms of the application of AI to various problems, but we should not confuse recent with swift. The direct groundwork for AI, as we have it, has taken between 60 and 70 years. This is on par with inventions such as the incandescent lightbulb and decidedly slower than others such as the electrical telegraph and the internet. No one could reasonably argue that that groundwork has not been part of the evolution of AI. Nor should we confuse the rapid diffusion of AI into the range of its potential applications with any particular swiftness. That diffusion, again, is comparable with a great many human inventions. Finally, we should not confuse the capacity to self-evolve of some AI, insofar as it exists, with swiftness. That capacity, in the most prominent current manifestations of AI, depends upon data, which must be gathered, even if the diffusion of the gathering process into the field from which it is gathered – the taking over of a seemingly ‘free domain’ particularly by Big Tech – gives an illusion of swiftness. The process of the data gathering cannot consistently be taken as independent to the process of evolution if the notion of intelligence is to be the characteristic which joins them. AI may self-evolve, but it does not self-evolve swiftly, all things considered.

If the imputed swiftness of AI technology is neither inherent to its character as ‘intelligent’, nor to its character as evolving, there remains the third of our self-supporting loop: swiftness of market compliance to the regulation. This may be the real source of the speed paradigm of the AI Act. Low cost regulatory options allow enterprises to quickly develop AI technologies within the low or minimal risk category, with an eye to broad diffusion in the EU market. The future proof definition meanwhile, essentially a list of AI techniques given in AI Act Annex 1 [5] is so flexible as to allow a wide range of products and services to be classified as AI. The list can simply be expanded quickly by further techniques without having to reflect – ethically – upon the nature, evolution, or evolving use of the technique, as long as prohibited uses are avoided and high risk uses follow regulation to the letter of the law. As Veale and Borgesius warn, in the name of free trade, and because it is the supreme law: “the Draft AI Act may disapply existing national digital fundamental rights protection. It may prevent future efforts to regulate AI’s carbon emissions or apply use restrictions to systems the Draft AI Act does not consider ‘high-risk’. Counterintuitively, the Draft AI Act may contribute to deregulation more than it raises the regulatory bar” [18].

If the development cycle of AI is characterized as swift, as the Act characterizes it, there is therefore no justification for invoking ethics in its relation to law. Indeed, just the opposite, law gives up its claim to partner with ethics insofar as it degrades to its mere regulatory character, as here in the AI Act. The appeal to AI as a special type of technique or product disappears if examined more closely – as above – and cannot be used to justify a link with ethics. Speed is a choice here. As a tendency it is consistent with regulation, but inconsistent with ethics.

The idea of legislation as regulation is a notion from civil law, dear to the EU law tradition. It fits well enough with a speed paradigm. But to impose regulation – particularly regulation whose main elements are ‘future-proof’ – is a narrowing of time for ethical engagement akin to saying: ‘we will get it done and over with now and forever.’

This can be a matter of ending conversation and debate. The latter is most strongly tied to approaches to ethics such as discourse ethics, with its emphasis on a slow process of continuing to question and answer publicly and a building up of ethical foundations in discourse consistent opinions, based upon

that slow process. The ending of conversation and debate, practically, if not explicitly, is a narrowing of time which works against discourse ethics.

But this narrowing can also be a matter of ending reflection and adjustment based on patiently waiting for outcomes and on listening. This aspect of narrowing arising from the speed paradigm is directly counter to an ethics of care, for example, where attentiveness and responsibility are emphasized, among others. To care is to be willing to attend to what actions have led to, while to be responsible is to move beyond the narrow sense of obligations eligible to be forgotten once fulfilled – in their legal regulatory sense – into a slower and broader time within which what I or we have done before, perhaps a long time ago, is constantly remembered as being accessory to the creation of the situation I or we now find ourselves in.

Even a hard deontological ethics, e.g. the Kantian, needs its slow time to work out the implications for action in the hard rules it imposes, and that working out, though imposed, can be imposed nowhere but internally, until such ethics strays into the externalized impositions of regulatory law.

In other words, regardless of the flavor of ethics, a speed paradigm, when partnered with ethics, opts for a negative view of ethics. On this view ethics is prohibitory, as a cursory theoretical working out of that which regulation will then prohibit by force. It stands against a positive mode of ethics, as a practice, a mode which all ethical theories take part in at their best. Ethics as a practice could be defined as: *the gradually evolving and expanding set of reflective interpretations, assessments of logical consistency, moderated generalizations, practical methods, and recommendations for action, which arise from working over examples of the ethical (or-unethical) disclosed in lived experience in order to smooth away value conflicts and bring higher value to that experience.*

In other words, ethics as practice is what ethics has been as a developing strand of human interest, a slow and gradual accumulation of wisdom regarding what to do, a building of the good, which is arrived at through learning and experiencing and ethical growth inside the human individual while acting in human society, no matter through which ethical theory.

All of this takes time, and has taken time. Ethical time, is different from legal time, even though the two overlap at their best. *Ethical time* is analogous to geological time. *Legal time* in its regulatory aspect is analogous to anthropocentric time.

In terms of the ethical effect of regulation the negative view of ethics leads to a complacency which can be described in the terms: “the bounds of our arena are now set; within that arena, as long as we follow these prohibitions to the letter, we are now free to do whatever we like without any further reflection upon the consequences.”

The effect of this thinking is to transform ethical engagement into a matter of warring camps without a middle ground. Law as regulation then comes to predominance – for some narrowed span of time. An example of negative ethics woven into regulation is the era of Prohibition in 1920s America. This era and its ethical issues seem quaint to us now, but it illustrates well enough the effect of the negative ethical route: the split into opposing camps (the ‘drys’ and the ‘wets’), the gradual hardening of those camps, and the circumvention of ethics ‘solidified’ into the narrowed time of regulation by those who will not accept regulation despite its penalties.

Ironically, Big Tech intuitively recognizes the positive mode of ethics as a force to be reckoned with. It even, seemingly, encourages the positive mode, with the assumption that its molasses slow workings pose no threat to unethical actions being carried out under the speed paradigm of technological advancement insofar as the positive ethical mode can be guided, see e.g. Abdalla and Abdalla [1]. In the case of AI Ethics, the result, arguably, has been an explosion of AI Ethics, which is largely an illusion insofar as it has failed to come down out of the clouds of vague principles divorced from context. As Yeung et al. put it: “the prevailing approach to AI ethics amounts to little more than a marketing exercise aimed at demonstrating that the tech industry ‘takes ethics seriously’ in order to stave off external regulation” [19].

This disingenuousness is gradually being unmasked however, precisely because Big Tech works with the speed paradigm, a paradigm which visibly warps ethics when applied to the latter. One can imagine what the great ethicists, e.g. Socrates or Kant, would have thought of this exponentially exploding ‘instant ethics’ of AI, created within half a decade, through the disingenuous urging of Big Tech.

Recently, Luciano Floridi has indirectly brought this divergence in ethical and legal regulatory time into relief. He has argued that the time for a certain type of ethical engagement – which I have here

called the positive ethical mode – has passed, to be replaced by regulatory engagement [7]. Viewed locally, in the sense of the ‘momentary’ and disingenuous efforts of Big Tech to encourage an ‘ethics of AI principles industry’ mimicking the positive mode of ethical engagement, his reading is right. But in a larger picture I believe Floridi misses something, namely that in *ethical time*, the engagement of ethics with AI, if it is to be beneficial, is only the barest beginning of what will be a very long and patient process.

The point being made here is that the acceptance of the speed paradigm by the creators of the AI Act, not merely in the effort toward the regulatory legislation, but in the characterization of that legislation, disallows those creators from the claim that the AI Act is well founded on a preparatory ethics. This characterization, and the legal results of the Act may well be good law, but they are not good ethics.

It might be argued that the AI Act is after all only a very early attempt to grapple with ethical issues in AI and to include ethics somehow, but on the terms of the AI Act, and thus we shouldn’t expect too much from the initial attempt.

Viewed in the light of ethics as I have characterized it, we could indeed grant this if it were so and if it were disclosed in an open and honest manner. On the other hand, if the appeal to ethics is – as I have suggested it is – one made so as to add a respectable veneer to what might otherwise be merely arbitrary regulation, then it would be better if those developing AI regulation admitted openly to a positivistic legal foundation for regulation, went about it, and left ethics out of the process to go its own direction. The result might be healthier for both sides. This is even more the case if ethics is warped to mere legal ends by a process which engages it through a process of lawlike development, as I have argued.

There may be multiple ethical approaches which can be combined to achieve ethical uses of AI. Any one of them may achieve good results, e.g. helping avoid destructive AI uses, while promoting uses that can bridge conflicts in value. The preliminary step in developing ‘good ethics’ for AI, is thus to recognize the commonality between these approaches, i.e. the manner in which ethics has hitherto proceeded relative to whatever it has achieved. If it has achieved nothing, then regulation does not need it and should avoid it. If it has achieved some good results, then at a general level it has achieved them through the way it proceeds, regardless of particular success. In this sense ethics – the practice of ethics – has been: slow, patient, reflective, retrospective, re-interpretive of past efforts and suggestions, and appeals to a mode of wisdom seeking in human experience. You come to it when you give up the urge to quick acceptance of certainties regarding value conflicts, in order to pull back and rethink things. It thus shares the temperament of philosophy generally. In these characteristics it is in opposition to the current temperament of development of AI as well as many other technologies. Except in name only, it cannot be twisted into the latter temperament and still remain ethics.

It may be countered that the ethical temperament evoked is impossible in light of the nature of AI and other digital technologies. If so, then the latter will not be developed ethically, though they may well be regulated. They will develop at their own pace and temperament, causing various destructive and self-destructive harms all the while. Then eventually, as with some other human advances they may exhaust themselves and turn for council to various ethical approaches which have grown up independently of them. That may be the unavoidable default. The more optimistic scenario is that with enough interest we might develop AI differently by an appeal to the temperament of ethics: slowing down, re-considering, reflecting on the purposes for the technology, etc. That would be ‘good ethics’ for AI development.

4. Conclusion

In the recent bloom of Artificial Intelligence ethics concern, law and ethics have often been lumped together, as if the practice of one almost necessarily complements the other. No doubt through the best of intentions, this conflation of law and ethics is also present in the EU AI Act proposal. In the forgoing I have argued that, as appealing as such a claim might be, there is little justification for it relative to the AI Act proposal, and in fact there are tendencies, both implicit and explicitly stated in the Act which should lead us to reject such a claim.

The primary objective of the Act, as a tool toward internal market consistency only requires legal development. Its appeal to ethics is demonstrably secondary, and indeed, warps ethics into the mold of a legal development process naturally open to influence by special interests.

The process of consultation meanwhile, evoking a democratic style consensus by majority in support of the role of ethics in the AI Act, rationalizes a disproportionate weighting of some stakeholders (enterprises), many of whom are acting from outside – literally – the domain of best interests of the legal jurisdiction of the EU. This appeal to a democratic participatory process rationalizes an imbalance of interests under a Utilitarian ethical outlook in particular. The appeal results in a weakening of the AI Act regulation as derived through consensus as a legal tool and in a weakening – both antecedently and subsequently – of the ‘new’ ethical practices and frameworks modeled upon them, even to the point of transferring a legal style acceptance of special interests into ethics itself by favoring a utilitarian ethical foundation over other ethical options.

Finally, the ‘speed paradigm’ embedded in the vision behind the development of the Act and its characterization of Artificial Intelligence, is arbitrary and profoundly counter ethical when examined in the light of ethical practice.

As the eminent 19th century Harvard scholar Charles Eliot Norton put it: “Moral progress must, under any circumstances, be very slow. Nor is there anything more opposed to real advance than hasty attempts to secure it . . . the progress which is permanent is made step by step, and not stride by stride. The great moral changes among men are like the great physical changes in the earth” [15] So it must be with AI ethics, if there is to be a satisfactory AI ethics, and for AI ethicists, hitching our wagon to the fast burning and glittering star of legal regulation is probably a mistake.

5. Acknowledgements

The project leading to this research has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 957391.

6. References

- [1] Abdalla, Mohamed., Abdalla, Moustafa.: The Grey Hoodie Project: Big Tobacco, Big Tech, and the Threat on Academic Integrity. In Proceedings of the 2021 AAAI/ACM Conference on AI, Ethics, and Society (2021).
- [2] Alexy, Robert.: Legal Certainty and Correctness. *Ratio Juris* 28 (4), 441-451 (2015).
- [3] Battista, L.: The European Framework Agreement on Digitalisation: a tough coexistence within the EU mosaic of actions. *Italian Labour Law E-Journal*. 14 (1), 105–121 (2021). <https://doi.org/10.6092/issn.1561-8048/13357>
- [4] Bradford, Anu.: The Brussels Effect. 107(1), *NW. U. L. REV.* 107(1), (2012).
- [5] European Commission.: 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. (2021).
- [6] Ferretti, Thomas.: An institutionalist approach to AI ethics: justifying the priority of government regulation over self-regulation. *Moral Philosophy and Politics* (2021)
- [7] Floridi, Luciano.: The end of an era: from self-regulation to hard law for the digital industry (2021). (forthcoming in *Philosophy & Technology*). Available at SSRN: <https://ssrn.com/abstract=3959766> or <http://dx.doi.org/10.2139/ssrn.3959766>
- [8] Hagedorff, T.: The Ethics of AI Ethics: An Evaluation of Guidelines. *Minds & Machines* 30, 99–120 (2020).
- [9] High-Level Expert Group on AI (HLEG), Ethics Guidelines for Trustworthy AI. (2019).
- [10] Kazim, Emre., Gucluturk, Osman Gazi., R. S. Almeida, Denise., Kerrigan, Charles., Lomas, Elizabeth., Koshiyama, Adriano., Hilliard, Airlie., Trengove, Markus.: Proposed EU AI Act - Presidency Compromise Text - Select Overview and Comment on the Changes to the Proposed Regulation (2022). Available at SSRN: <https://ssrn.com/abstract=4060220> or <http://dx.doi.org/10.2139/ssrn.4060220>

- [11] Lockey, S., Gillespie, N., Curtis, C.: Trust in Artificial Intelligence: Australian Insights. The University of Queensland and KPMG Australia (2020). doi.org/10.14264/b32f129
- [12] Mill, John Stuart. Considerations On Representative Government. Longmans, Green. London (1861).
- [13] Mill, John Stuart.: On Liberty. Longmans, Green, London (1867).
- [14] Mittelstadt, Brent.: Principles Alone Cannot Guarantee Ethical AI. Nature Machine Intelligence (2019). Available at SSRN: <https://ssrn.com/abstract=3391293> or <http://dx.doi.org/10.2139/ssrn.3391293>
- [15] Norton, Charles Eliot.: Considerations on Some Recent Social Theories. Little, Brown, and Co., Boston (1853).
- [16] Riley, Jonathan David Charles.: Utilitarian Ethics and Democratic Government. Ethics 100, 335-348 (1990).
- [17] Townsend, Bev.: Decoding the Proposed European Union Artificial Intelligence Act. The American Society of International Law. 25 (20), (2021).
- [18] Veale, Michael., Zuiderveen Borgesius, Frederik J.: Demystifying the Draft EU Artificial Intelligence Act — Analysing the good, the bad, and the unclear elements of the proposed approach. Computer Law Review International 22, 97-112, (2021).
- [19] Yeung, K., Howes, A., Pogrebna, G.: AI Governance by Human Rights–Centered Design, Deliberation, and Oversight. (2020):