ALGORITHMIC INTEGRITY: WAY FORWARD FOR A POLARIZED PRIVACY DEBATE PERTAINING ARTIFICIAL INTELLIGENCE AND MACHINE LEARNING

1. Introduction

Like many people do, I maintain four separate email accounts across three major service providers: *two* for my work as a University Lecturer (to separately communicate with colleagues, and then my students); *one* for professional interactions with clients (as an Industry Consultant), and *another* for social networking. I have been extremely cautious to compartmentalize my online communications along these four panels of my digital interactions: just to ensure that emails from my students do not end up in my social email account and get missed; and emails from my professional clients are not threated with the casual undertones that characterize the cordial relationships that I enjoy with my work colleagues. Apparently, the only place where these emails addresses converge is in the account configuration data of the four email accounts; where one or more of the other email addresses are provided as alternate accounts for retrieval purposes, in case I am locked out by 'bad guys'.

One day in February 2020, an email from my boss hit my social and work email accounts simultaneously, while I had only ever contacted him with my work email address. "How did this happen?" My investigations underscored a fairly-recent feature of modern email services that allows email servers to suggest [all] possible contact email addresses for a person whose name is being entered as recipient of an email that is about to be sent. Figure 1[1] illustrates how this operation takes place on an email client application.

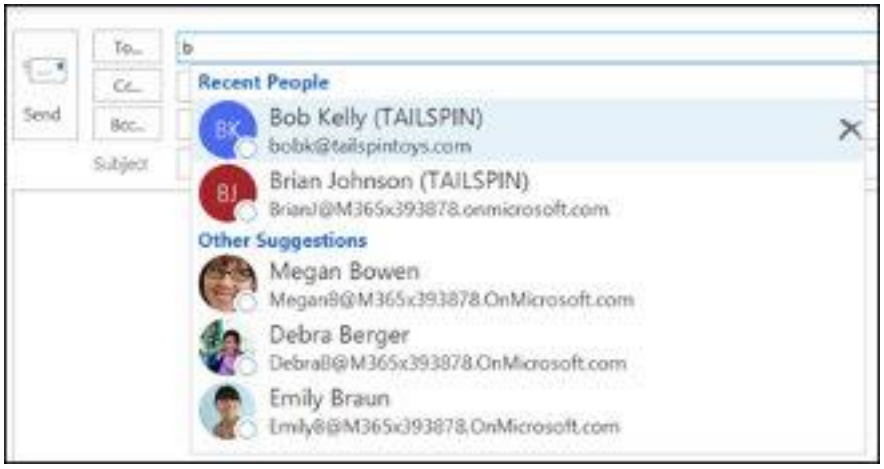

Despite the possibility that I might actually be wrong and the incident an outcome of some other random occurrence; yet, as a computer scientist, it was difficult to dismiss what I had come to know as the intricacies and capabilities of modern AI and ML algorithms. Could it be that AI had cross-recommended what it had learned from my social interactions to compromise my privacy in my work life? "*Is it possible to breach privacy without a human in the loop?*"

The rest of this Essay proceeds with a non-technical overview of BD, and the utility that it presents for AI and ML. The model General Data Protection Regulation (GDPR) of the European Union (EU) is then summarized to highlight what its goals and objectives are, and elucidate the insufficiencies of the GDPR in dealing with the modern realities of AI and ML. "The future is private"[2]; therefore, emerging concerns and contemporary arguments surrounding privacy in the era of BD and AI are presented next, in order to highlight the contending positions of what has become a polarized debate that boils down to the issue of *Trust*.

Discussions for a way forward in this debate follow after, with an attempt to categorize the various existing approaches that have been put forward as recommendations for progress, while also highlighting the limitations of these existing approaches. An additional complementary approach – *the Algorithmic Integrity Approach* (AIA) – is then proposed and discussed. The immediate benefits of this Approach are highlighted; before discussing how to chart progress for policy, innovation, and practice.

In order to clarify the context of further discussions: it is important to state that at the developmental level, AI is broadly categorized into: *Knowledge-based AI* (or Expert Systems), and *Data-driven AI*. The development of Knowledge-based AI generally involves domain experts providing their high-level knowledge and expertise to be coded into AI applications and services for particular domains, through a process known as knowledge engineering that helps machines make decisions within contexts that have almost no ambiguity. The capabilities of knowledge-based AI are often broadly restricted to the domain for which they have been developed such as self-driving cars and specialist robotic arms, to name a few. However, data-driven AI, which is the context of discussions in this Essay, involves the use to complex algorithms to process and analyse large amounts of data and arrive at logical inferences that are not only rational and reliable, but most times also remarkably accurate – even in the midst of several contextual ambiguities; through a pervasive process known as ML.

2. Big Data, Artificial Intelligence, Machine Learning, and the European Union's General Data Protection Regulation (GDPR)

AI and ML have continued to exhibit astounding and really remarkable capabilities for problem-solving in modern times; particularly in the areas of prediction, forecasting, and analysis. Thanks to Big Data. *Volume*, *Variety*, *Velocity*, and *Variability* are three terms that have been used to explain the power of big data[3]. By virtue of the *volume* of the data (which is typically accumulated from millions of sources), depth, granularity, and accuracy is greatly amplified. *Variety* multiplies this capacity by widening the pool of possibilities through internal & external aggregations and micro-aggregations that enable connections, inferences, and distinctions that serve to enhance precision and minimize error. The real-time, on-time processing and delivery of outcomes and results of analyses and operations is owed to its *Velocity*. While *variability* describes how validity, accuracy, and delivery are preserved even when there is a change in other characteristics of the data.

It is this sophisticated technology that powers the search engines, virtual reality systems, recognition & identification systems, recommender systems, profiling activities, targeted advertising, enterprise resource systems, as well as security & intelligence operations that form an integral part of daily life. They are generally classified as decision (support) systems. Underneath these sophisticated digital systems is a complex working of computational neural algorithms trying to make sense of the vast amounts of sensitive digital data that is being accumulated about individuals and organizations across several sources, including the Internet. The aim being to proffer innovative solutions (even on-the-fly) to complex problems in business, entertainment, science & technology, engineering, management, politics, commerce, and economics; with little-to-no human aid. This is ML, and it is the foundational principle of AI[4].

Machine Learning could be generally classified as: *supervised* (wherein machines are guided to correctly predict outputs by learning from input instances that are pre-labelled); *reinforced* (whereby machines get better at performing specific tasks by repeated involvement/participation); or *self-supervised/unsupervised* (where machines are able to correctly predict outcomes by enhancing the task-specific sophistication of learning models through un-labelled random observations). Extreme forms of these Learning types have commonly been referred to as *Deep Learning*.

From early on, AI was maligned based on misconceptions that attributed to machines a level of autonomy that matched the intelligences of human agents; which were further exacerbated by practices that tended to gradually mask the human agencies (a Socio-technical Blindness) in the AI process and operation chain[5]. But as AI evolved, it became obvious how this innovative technology could compromise privacy by using personal information in ways that violate widely-accepted ethical and legal principles.

This began to magnify existing concerns about ethics and legality as it applied to the technology. Concerns pertaining: the agency of control pertaining rights to seclusion and selective identification/un-identification; contexts surrounding how personal data is to be used and who holds ownership and control of such data; and definitions of what constitutes misuse/unethical use of personal data. The legal world seemed particularly unprepared to deal with the realities that began to emerge at the time; even as AI grew in complexity, sophistication, surreptitiousness, and pervasiveness – featuring capacity for behaviours and decisions that were almost entirely autonomous. It was on the heels the model GDPR[6] of the EU was passed in 2018.

It is notable that the GDPR sought to, among other things: (1) establish privacy and data protection as a fundamental right of individuals; (2) hold organizations responsible for sensitive personal data that they collect, control, and/or process within their products and services[7], and prescribe strict legal conditions to regulate how these data are collected and used – making them liable for securing and safeguarding these data, as well any acts of misuse of such data; (3) regulate the secure exchange, processing, and use of personal data among EU member states (4) outline various protocols for dealing with breaches, with varying penalties for non-compliance at different degrees of impact to data owners; (5) expand users' control over their data, to include rights like 'the right to know', 'the right to opt out and be forgotten', 'the right to access personal data collected', etc.; and (6) recommend ethical best-practices (like *pseudonymization*, and *data minimization*) for dealing with sensitive data.

But the GDPR was not without insufficiencies. It did not seem to have sufficiently addressed the complex realities that, today, characterize operations within this space. For example: to what extent can organizations and corporations be held culpable/liable for the "automated decisions" made by autonomous learning/intelligent algorithms? Especially considering the fact that a growing number of these algorithms are now able to learn, adapt, respond, and make decisions using open-source or open-access data that is not necessarily being 'collected' and 'kept' by the company itself.

3. Emerging Concerns, Contending Debates & Viewpoints

Indeed, digitalization greatly changed the way and manner in which policy issues are engaged with; especially because certain digital data can be used to uniquely identify, and accurately profile an individual for various purposes. It was the multi-jurisdictional, global reach of the Internet that expanded what used to be a pretty clear and relatively convergent perspective to these discussions; even as the advent of AI and ML further complicated the terrain.

- On the leeward side where consumers face, there is a broad consensus that "AI needs more regulation, not less"[8], as transparency, accountability, and fairness are critical to its continued growth and innovativeness. Also decrying the insufficiencies and often reactionary approach of existing regulation in dealing with contemporary issues of data protection and digital privacy[9] in the age of AI. But then, there does not seem to be clear direction on how to proceed – as a tension ensues between the call for a globally-unified technology legislation[10] that creates a level playing field, enabling everyone cope with the modern realities of globalization and the trans-national impact of digitalization; and whether such global legislation is even attainable, citing factors such as contextual disparities in technological development caused by the digital divide, vested sovereign political interests, and the ethics of legality in a tension of autonomies[11].
- However, on the windward side where AI companies and service providers face, there seems to be a fairly general consensus that stringent privacy and data protection policies are likely to stunt the development and innovative capabilities of AI and ML, or even entrap privacy legislation in complex social and political quagmires[12] that might be counter-productive in the long run.

Balancing all these legitimate concerns with minimal trade-offs has continued to pose a crucial challenge for regulators and policymakers; and has now transformed into a multi-faceted debate that has been polarized even within specific communities, but where many players have come to participate from across the converging digital (telecommunications, media, and technology) ecosystem.

- In the age of AI, the civil society and human rights community, have, on the one hand, expressed concerns about the applications of identity and recognition systems powered by sensitive personal (biometric) data in the personalization of social services, unification of citizenship records, as well as for cross-border identification and verification[13][14]. They cite increasing cases of digital authoritarianism, privacy invasion through digital surveillance, and rising statistics of harassment, brutality, and threats to journalists, communicators, and activists[15]. Arguing that the capacity for AI to be misused in the line of these operations by governments and state authorities for broad-based infractions against privacy, human rights, and social justice, is greatly amplified[16].
- On the other hand, government security agencies and lawmakers participating in this debate have also argued that dealing with matters of national security that bother on rebellious dissidence, political insurrection, violent extremism, and civil restiveness orchestrated by the fake news media, may sometimes warrant that AI is leveraged to accurately identify and box-in rebels, extremists, terrorists, and dissidents through broad scale digital surveillance and crackdowns[17]. Within this same circle, however, there exists the more popular counter-position by policy experts from the wider private and public sectors, arguing that more surveillance and data-enabled privacy infractions are scarcely the solution to these problems; but rather more international cooperation and intelligence sharing based on whatever limited data is available, and within the ambience of existing legislation[18][19].

In the multiplicity of these contending viewpoints, one key issue is brought to the fore – the concept of *Trust*, which is the framework that guarantees privacy.

> "To *Trust* is to have confidence both in the Integrity and Abilities of the Trusted [paraphrased]" – **Stephen R. Covey**: *The SPEED of Trust*

Trust is a critical concern that is usually not difficult to spot in digital privacy debates; yet crucial to building sustainable digital futures. This concern is often subtly expressed in the language of questions like: At what point, and under what circumstances would privacy invasions be justifiable? What (new or existing) policy / legal frameworks are needed to safeguard how private information is obtained and used under these circumstances? How would redress be pursued in the event of violations? How is it guaranteed that what AI learns about our work-lives or health cannot be carried over and used to malign our social reputation, or compromise our healthcare insurance, respectively?

4. Way Forward for a Polarized Debate

With the human in the loop, it seemed relatively easy to find progressive answers and logical compromises for these trust-related questions in privacy debates. But by obscured the human in the loop, even the best answers, when they do exist, are often speculative, conjectural, and incomplete; especially for untested applications and use-cases. In order words, how do we deal with issues of trust and privacy in AI and ML? A number of recommendations have already been proffered.

These recommendations can be broadly categorized into *ethical approaches*, *regulatory approaches*, and *non-regulatory approaches*.

4.1 Ethical Approaches

Ethical Approaches recommend adopting certain ethical codes by AI companies and service providers for accountability. Pledging to treat consumer data with dignity and transparency. Implementing safeguards ("ethics by design"[20][21]) like human reviews, bias evaluation, reproducibility for audit purposes, structures for handling displacement, and transparency of computational operations. In order to build trust and ensure that machines make no 'harmful' decisions[22][23].

However, '*harm*' remains a fundamentally contested concept in the digital policy space: At what point can it be said that '*harm*' is being or has been done? Could such '*harm*' be the climax of aggregated micro-'*harms*' done in the past

that went unnoticed? That is, is it possible that this '*harm*' is a ripple effect originating from lower-level digital micro-operations? What safeguards could have helped to avert/forestall such '*harm*'? These concerns stalk the operations of the *Clearview AI* facial recognition software[24], and the suicide detection algorithm that was rolled out by Facebook in 2017[25].

Whereas, it seems foolhardy to expect that AI companies and service providers would do what is right with regards to ethics in the products/services that they offer. Especially since '*right*' is a subjective reality that is often gauged along a continuum between intent that is geared towards creating utility, and outcomes with implications having magnitudes that are often impossible to entirely foresee, *ab initio* – a reality reflected in a recent incident which saw the Massachusetts Institute of Technology take down a 12-year-old dataset that used strong racial and misogynistic slurs to train AI detection and profiling systems[26].

4.2 Regulatory Approaches

Regulatory Approaches apply legislation to demand better stewardship, with the goal of building trust through disclosure, transparency, and informed consent. By establishing culpability for infractions & misuse, and recommending technical processes to safeguard, secure, and ensure the scientific integrity and quality of data that guarantees algorithmic accountability in AI. Basically, regulatory approaches help guarantee that ethical approaches are followed – balancing consumer concerns ways that do not stifle innovation and development. The GDPR was one such approach.

However, a well-known struggle of regulatory approaches is being able to keep up with the rapid evolution of modern technology; and also being at pace with new applications/use-cases, and the new realities that often emerge.

4.3 Non-Regulatory Approaches

Non-Regulatory Approaches recommend avenues for appropriately addressing risk posed by certain AI applications, when existing regulations are sufficient or the benefits of a new regulation do not immediately or eventually justify its costs.[27]

As an alternative to inaction, organizations may assess and manage risks within the boundaries of 'acceptable harm'; keeping sight of the flexibility required to cope with new domains and realities. By way of seeking: (1) direction in the recommendations and policy statements of non-regulatory authorities; (2) waivers from regulatory authorities or approval to run pilot programs for specific applications within safe boundaries; and (3) gauging risk through voluntary consensus standards by other stakeholders in the private and other sectors.

However, foreign to the philosophy that informs this approach is the moral imperative of preserving/enhancing public trust by seeking informed consent beforehand; while also applying other measures that hinge on openness and accountability to protect privacy.

To a large extent, the European Union's General Data Protection Regulation (GDPR) managed to blend all three approaches to create arguably the finest existing regulation for privacy and data protection; and which has been lauded and adopted as a model by many emerging national and trans-national regimes. However, while emphasizing the need for a globally harmonized framework for privacy and data protection aligned with the GDPR[28], *Mark Zuckerberg* underscored the presence of critical gaps in the GDPR for dealing with how data protection and privacy applies to new technologies like AI – when there is no human in the loop. Because, although the GDPR references the need for providing 'data protection by design' in new digital products and technologies; a recent report by the

European Parliamentary Research Service (EPRS) finds that the GDPR seems lacking in useful and concrete guidelines and recommendations regarding how these might be pursued at the lower levels of design[29]. Consequently, I wish to propose an additional and complementary piece to these existing approaches – the *Algorithmic Integrity Approach*.

4.4 The Algorithmic Integrity Approach (AIA)

> "Integrity... is an internal guidance system that will never guide you astray!" – **Kelley Kosow**: *The Integrity Advantage*

The **AIA** emphasizes the importance of ensuring algorithmic integrity by standardizing internal algorithmic processes and interactions, and developing algorithmic best practices for sensitive use-cases/applications of ML; as a way of cultivating trust and building confidence in AI behaviours and decisions. Figure 2 illustrates how AIA integrates with the Ethical, Regulatory, and Non-regulatory Approaches.

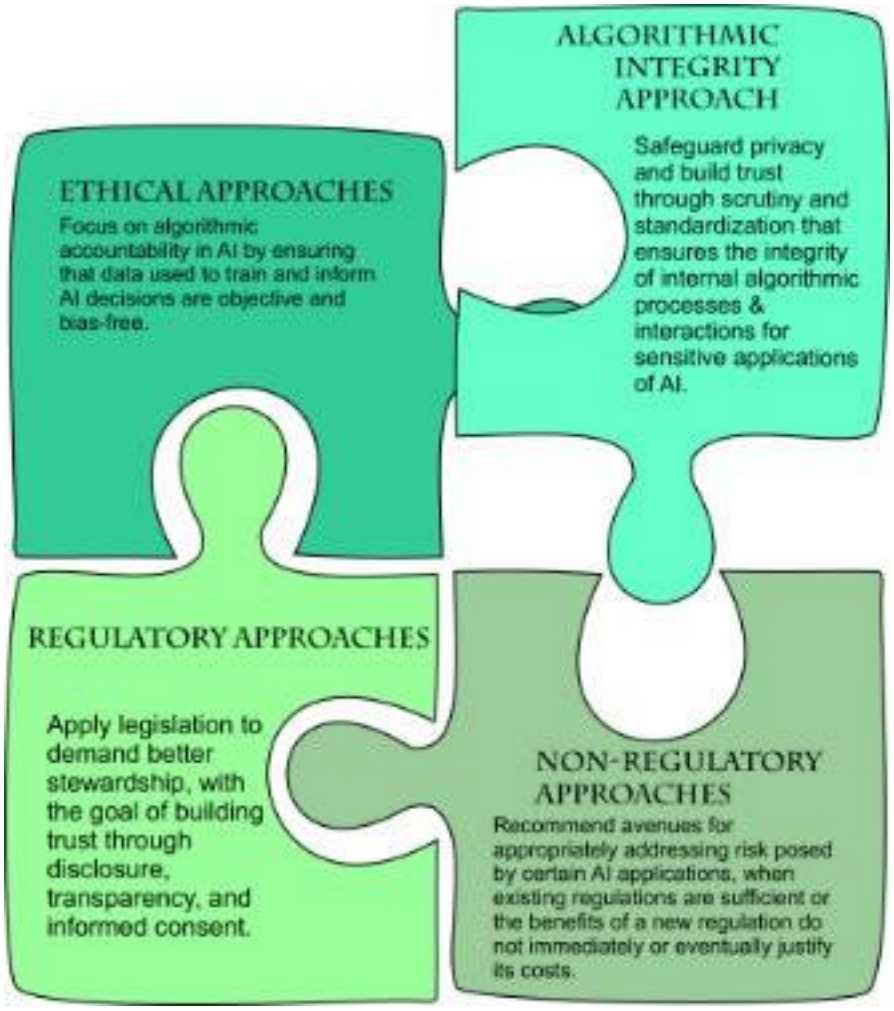

It has been decried how AI regulatory/policy efforts are largely reactive in scope, and have, at best, dealt only with superficialities pertaining the features and outcomes of AI – failing to grapple with the fundamental issues at the crux of consumer concerns regarding privacy and data protection in ML. This concern is not particularly difficult to understand; especially considering the fact that policy experts are often not also trained in the technical background that is often needed to unpack, understand, and grapple with the internal workings, as well as the intricacies and complexities of ML. In this regard, Melissa Whitney underscores the importance of bringing the judicial community up to speed with the technicalities of AI algorithmic processes and their far-reaching capabilities[30].

For example: What is the nature of the data that are parsed between the various internal learning functions/constructs of these algorithms? How are these data handled and secured during these parsing operations? How are they

indexed, called, and logged? How are the internal log files that record these meta-information about algorithmic operations handled? These are some fundamental questions that inform the need for the AIA, and are crucial for tackling the underlying issues of privacy and trust that have trailed AI and ML applications.

A lot of "black box" interactions and manoeuvrings characterize the operations of many machine learning algorithms and techniques that are featured in AI products and services. However, proprietorship is usually protected by law through patents; making it difficult to scrutinize these operations in an open, accountable manner to answer these fundamental questions. Nonetheless, engaging actively and objectively with these rudimentary questions holds tremendous progress for policy efforts in dealing with contemporary privacy and data protection concerns in AI.

Spotlighting my experience shared at the Introduction of this Article: recommender AI systems running on email servers can (separately) learn my social and work email addresses from a personal buddy who is also a colleague at work, and whom I have contacted in separate capacities using both email addresses. Then apply that knowledge in recommending my social email address to my boss without my consent. Furthermore, when *Spectre and Meltdown* was discovered in 2018, the world also saw how vulnerabilities even in computing hardware (processors, in this case) made it possible for sensitive information including passwords, emails, messages, and documents to be lifted from underneath algorithms and programs while in execution[31]. *Spectre and Meltdown* trailed a variety of computing platforms ranging from personal computers and mobile devices, into the cloud.

Thus, it becomes easier to understand how sensitive user information provided as input parameters deep within algorithmic processes could be parsed/visible to more global class constructs during execution. Also how the things that AI learns from particular aspects of our lives can be used to compromise our privacy in other areas of our lives. Hence, sensitive user information becomes susceptible to vulnerabilities like *Spectre and Meltdown*. Otherwise, they could be either learned from separate user interactions and cross-recommended without consent; or logged/indexed during algorithm compilation and execution – forming part of the mass of big (meta-) data used to inform learning decisions in AI products and services. For sensitive use-cases involving healthcare data[32], personally-identifiable information, and financial records, the danger to privacy is colossal; even without a human in the loop.

These possibilities confirm that indeed privacy can be breached, even without a human-in-the-loop. Perhaps, it is therefore necessary to begin to consider the need to, for example: standardize the extent of visibility for sensitive information within algorithmic structures; recommend/develop secure methods for parsing such information amongst different algorithmic constructs; regulate if, when, and how such information is logged; compartmentalize what is learned about individuals along the various panels of their digital interests and interactions; as well as prescribe controls to safeguard how log files are accessed and handled during/after algorithmic operations.

To this end, a multi-stakeholder community-based model that involves the active participation of voluntary standards organizations, and integrating technical experts from across the information and communications technology and engineering communities, is recommended. Because the success of the AIA would depend on an integrated multi-sectoral strategy that is bottom-up coordinated, and involves the technical information technology sector, as well as the digital governance, engineering, and communications sectors.

This is particularly crucial because, in the age of AI and BD, it is no longer enough that privacy and data protection policies wait for the manifestations and outcomes of breaches and infractions before they kick in (as decried of regulatory approaches); nor is it enough to expect that AI companies pledge to apply ethical best practices and safeguards in their products and operations (like ethical approaches might prefer). Neither has it also proven beneficial

to long-term public trust and acceptance of AI to relegate or bypass informed consent at any point (as seems convenient for non-regulatory approaches).

The immediate benefits of the Algorithmic Integrity Approach include that:

a. It highlights considerations through which policy efforts can be able to deal more proactively and progressively with the dangers to privacy and security from the foundational (algorithmic) levels of AI.
b. It poses no foreseeable challenges for Competition amongst players in the Industry.
c. It inspires AI companies to balance the innovative goals of their products and services, with the nitty-gritties of privacy and security.
d. It galvanizes greater participation from the wider multi-stakeholder digital ecosystem (information technologists, governance & policy experts, system engineers, administrators & executives, and government).
e. It provides a way for AI companies/service providers to be tangibly accountable regarding the behaviours and decisions of their featured algorithms.
f. It ensures guard rails in AI algorithms that help guarantee its integrity.
g. Ultimately, it would help to rebuild consumer trust and restore dwindling confidence towards AI products and services.

The thinking behind the AIA is consistent with the trajectory of several emerging positions regarding ethics and privacy in AI, which: (1) recognize the risks that algorithmic interactions in AI decision systems could portend for privacy, democracy, human rights[33]; (2) advocate the importance of building foundational safeguards into AI to help preserve trust[23]; (3) decry the almost-entirely profit-driven motives that underscore AI products and services, as well as the collaborations between the companies that develop and use them[32]; (4) call for more (concrete) regulation that adopts a proactive before-the-fact approach[8], and aligns with data governance approaches for 'privacy by design' in AI algorithmic processes[12]; and (5) recommend the establishment of clear standards to guide how AI algorithms are applied to personal data[29].

4.4.1 Progress for Policy, Innovation, and Practice

Many companies (over 80%[34]) continue to adopt AI on a massive organization-wide scale, with at least 90% of them indicating plans to incorporate innovative AI capabilities into their operations for process automation, and AI-enabled conversational systems in the very near future[35]. Other areas including supply-chain management, marketing and sales, service operations, and product development, among other business processes[36]. The concern being that these companies might apply algorithms that have been suspected of inappropriate[37] and illegal[38] behaviour, and would report huge returns on investment with record profit margins.

While these companies decry that excessive regulation could stifle the development and applications of AI; there has also been the very legitimate fear on the side of users, regarding what might be the consequences to privacy by virtue of the growing capabilities of AI. The magnitude of this fear is captured in a recent survey[39] where at least half of Facebook users expressed concern and discomfort over how they were being profiled by the social media Platform, with 71% not even being aware that they were being profiled at all. A similar report also discovered that 90% of Americans remain concerned about privacy and security with regards to how their data are used online, with at least 60% calling for strict(er) national privacy laws[40]. Further, a 2018 survey[41] revealed that 84% of Americans believe that AI should not be left loose in its operations and interactions.

Thus, the implications of the AIA for innovation and practice can essentially be condensed to focus on the tensions between: (1) growing organizational adoption of AI to provide utilities, products, and services; (2) scepticism towards algorithmic processes and activities that are shielded by copyright and intellectual property laws; and (3) the legitimate fear of clients and consumers regarding how AI uses and interacts with their data. Essentially, a well-founded mistrust by consumers towards the behaviour and activities of AI on the one hand; and on the other hand a seeming frenzied and almost-entirely-profit-motivated scramble by organizations to harness the innovative problem-solving capabilities of AI using algorithms that might be behaving inappropriately, but are shielded by law.

Angel Fu has called for "new rules of the road" that protect consumers' interests, without stifling innovation and digital development[10]. A way forward with the AIA that upholds this viewpoint is to begin first by exploring deployments for sensitive applications like those involving healthcare data and financial records – where privacy and trust[42] are prevailing concerns inhibiting wider acceptance of AI, despite a growing demand[43][44]; then analysing behaviour, and gauging performance & impact through broad-based multi-stakeholder participation.

It would also be important to engage the public in the process of developing and formalizing codes of practice in this dimension[45] – through neutral ombudsmanship (as has been seen in the UK's "Office for Responsible Technology"[46]); combined with more initiatives like OpenAI[47] that advocate an open-source community-involved approach to AI development. This way, multi-stakeholder deliberations can progress when AI algorithmic processes are suspected of inappropriate behaviour, without necessarily violating proprietorship.

5. Conclusion

> "Integrity is a concept of consistency of actions, values, methods, measures and principles, expectations and outcomes..." – **So-Young Kang**[48]

Privacy International has recently requested the EU[49] to block a proposed merger between Google and the fitness technology Company – Fitbit – citing critical concerns about privacy and human rights that would be swelled by Google's AI, following such a merger[50].

It has become fairly straightforward to boil down the contemporary issues and debates about privacy and trust in AI to the uncertainties surrounding how underlying big data and machine learning algorithms behave and interact. While it seems like a fairly clear problem definition, it is one that feeds back from a complex convolution of social, political, ethical, and legal issues[51]. In part, a consequence of growing mistrust towards the intent and motives of companies that develop AI products and services; but also a fallout of the realities that the world is gradually coming to terms with regarding the automated capabilities and decisions of AI algorithms. I propose that a potential solution lies in pursuing a progressive, integrated Approach to Algorithmic Integrity.

The *Nirvana Fallacy* argues about the danger of perpetuating a problematic status quo simply because a (possibly-better) alternative seems inconvenient and (equally) imperfect[52] – in a choice between an existing problematic situation and a progressive alternative that is neither perfect nor unachievable[53].

Thus, it might pose a *Nirvana Fallacy* to continue in the *status quo* with rife privacy concerns and growing mistrust towards the behaviours and decisions of AI, in hopes that someday regulations would eventually catch up with the rapid developments that characterize AI evolution; or, perhaps, wish that by some inexplicable means, the public and consumers would suddenly develop the level of trust for AI behaviours and decisions that would become the "tipping point"[54] for the technology. While, alternatively, we could begin to gradually pursue a non-blanket and achievable approach to algorithmic integrity – beginning with sensitive applications. In the long run building trust by galvanizing multi-stakeholder participation and co-accountability on a broad scale.

The **AIA** does not claim the magic wand that would summarily, singularly, or independently settle debates surrounding privacy in AI. Unfortunately, there no such approach. Rather, when used in synergy with existing approaches that have been discussed in this Article, including the ethical paradigm that retains the human-in-the-loop, AIA contributes a mutually-complementary point-of-focus to help chart a path for progress in debates surrounding digital privacy and trust in AI.

References