



UNIVERSITY OF THE
AEGEAN

SCHOOL OF ENGINEERING
DEPARTMENT OF INFORMATION
AND COMMUNICATION SYSTEMS ENGINEERING

UNIVERSITY OF THE
AEGEAN
SCHOOL OF
ENGINEERING

DEPARTMENT OF
INFORMATION &
COMMUNICATION
SYSTEMS
ENGINEERING

UNDERGRADUATE
PROGRAM

INFORMATION
SYSTEMS LABORATORY

DIGITAL GOVERNANCE
RESEARCH CENTRE



Artificial Intelligence: Dangers to Privacy and Democracy

Sigalas Markos

321/2010159, icsd10159@icsd.aegean.gr

Supervisor: Professor, Charalabidis Yannis

July 2021

www.aegean.icsd.gr/is-lab

Submitted in partial fulfillment of the requirements for the degree of the
bachelor's degree in the Department of Information & Communication Systems
Engineering, University of the Aegean

Examination Committee:

Committee chair:
CHARALABIDIS YANNIS, Supervisor

Professor

Department of Information & Communication
Systems Engineering

Committee Members:

LOUKIS EURIPIDES, Member

Professor

Department of Information & Communication
Systems Engineering

KOKOLAKIS SPYROS, Member

Professor

Department of Information & Communication
Systems Engineering

© 2021

Sigalas Markos

University of the Aegean

Abstract

With the ever-widening use of the Internet, an increasing amount of social and personal interactions between people is taking place on public, online fora, such as forums and social media. This has resulted in the movement of much of once personal socialising into the public sphere. Much of the information that was once regarded as private or personal has now become easily and publicly accessible, if not to the end-user then to the operators of social networks who often simply charge a price for access to advertising companies and other entities. The subject under study is how such practices can be abused through Artificial Intelligence systems and the repercussions that Artificial Intelligence has for a democratic society.

In order to properly evaluate these concerns, we will need to examine what constitutes a “democratic society”, the social norms and political institutions that entails. We need to examine the concept of privacy, the way it has evolved as telecommunications systems have, and how closely tied it is to our understanding of democracy. If we are to ascertain whether Artificial Intelligence poses a danger to democracy, it is necessary to be able to describe conditions in which democracy has been compromised.

Technologically, we need to address the ways in which Artificial Intelligence and Big Data intersect with personal data and public life, and the extent to which that has an effect on democracy and individual privacy. Of interest in this is the capabilities of the technology fundamentally, which is to say, how much and which data is collected, how Artificial Intelligence systems can make use of it, how accountability can be attributed in those systems, if it can at all, and ways in which ethical standards can be set and enforced in a viable way.

To ground the theoretical framework that the above analysis will establish to empirical evidence, we will look at major social networks’ data collection and use policies, examples of Artificial Intelligence being used to exploit such data, such as the Facebook-Cambridge Analytica scandal, and efforts to establish regulatory frameworks that ensure an ethical application of these technologies, such as the European Union’s General Data Protection Regulation and ethical guidelines for trustworthy Artificial Intelligence. The ultimate conclusion we arrive at is a generally weak response to protect democratic values where they are threatened by AI, and an unwillingness to confront some of the fundamental economic motivators.

Keywords: Artificial Intelligence, Internet, predictive algorithms, democracy, individual rights

Εισαγωγή

Με την συνεχώς διευρυνόμενη χρήση του Διαδικτύου, ένα αυξανόμενο ποσό κοινωνικών και προσωπικών διαδράσεων λαμβάνει χώρα σε δημόσια, Διαδικτυακά φόρα όπως οι σελίδες forum και οι σελίδες κοινωνικής δικτύωσης. Αυτό είχε ως αποτέλεσμα την μεταφορά πολλής μέχρι πρότινος προσωπικής κοινωνικοποίησης στην δημόσια σφαίρα. Πολλές πληροφορίες που θεωρούσαμε ιδιωτικές και προσωπικές πλέον γίνονται εύκολα και δημόσια προσβάσιμες, εάν όχι στους χρήστες τότε στους διαχειριστές κοινωνικών δικτύων οι οποίοι αρκετά συχνά απλά χρεώνουν την πρόσβαση σε διαφημιστικές εταιρείες και άλλες οντότητες. Το θέμα υπό μελέτη είναι πως μπορεί να γίνει εκμετάλλευση τέτοιων τεχνικών μέσω συστημάτων Τεχνητής Νοημοσύνης και τις επιδράσεις που αυτή η χρήση Τεχνητής Νοημοσύνης έχει για την δημοκρατική κοινωνία.

Για να κάνουμε σωστό έλεγχο πάνω σε αυτά τα ζητούμενα, πρέπει να μελετήσουμε το πως ορίζεται μια «δημοκρατική κοινωνία», τις κοινωνικές νόρμες και τους πολιτικούς θεσμούς που συμπεριλαμβάνονται σε αυτή. Χρειάζεται να μελετήσουμε την ιδέα της ιδιωτικότητας και το πως εξελίχθηκε μαζί με τα μέσα τηλεπικοινωνιών, και το πόσο συνδεδεμένη είναι με την δημοκρατία όπως την εννοούμε. Για να αποφασίσουμε αν η Τεχνητή Νοημοσύνη αποτελεί απειλή για την δημοκρατία, πρέπει να μπορούμε να περιγράψουμε συνθήκες υπό τις οποίες η δημοκρατία έχει θιγεί.

Τεχνολογικά, πρέπει να απαντήσουμε στο πως η Τεχνητή Νοημοσύνη και τα Μεγάλα Δεδομένα αγγίζουν τα προσωπικά δεδομένα και τον δημόσιο βίο, και το μέτρο στο οποίο αυτό επηρεάζει την δημοκρατία και την προσωπική ιδιωτικότητα. Προς αυτόν το σκοπό μας ενδιαφέρουν οι δυνατότητες που κατέχει η τεχνολογία, δηλαδή πόσα και ποιά δεδομένα συλλέγονται, πως χρησιμοποιούνται από συστήματα Τεχνητής Νοημοσύνης, πως και αν μπορεί να αποδοθεί λογοδοσία σε τέτοια συστήματα, και τους τρόπους με τους οποίους μπορούν να τεθούν και να τηρηθούν πρότυπα δεοντολογίας.

Για να βασιστεί το θεωρητικό πλαίσιο της άνω ανάλυσης σε εμπειρικά δεδομένα, θα μελετήσουμε πολιτικές δεδομένων μεγάλων κοινωνικών δικτύων, παραδείγματα χρήσης Τεχνητής Νοημοσύνης για εκμετάλλευση αυτών των δεδομένων όπως το σκάνδαλο Facebook-Cambridge Analytica, και προσπάθειες να κατασκευαστούν κανονιστικά πλαίσια που εγγυώνται την δεοντολογική εφαρμογή σχετικών τεχνολογιών όπως ο Γενικός Κανονισμός για την Προστασία Δεδομένων της Ευρωπαϊκής Ένωσης και προτεινόμενο δεοντολογικό πλαίσιο για αξιόπιστη Τεχνητή Νοημοσύνη. Τελικό συμπέρασμα στο οποίο φτάνουμε είναι μια γενικά αδύναμη απόκριση για την προστασία των δημοκρατικών αξιών και μια έλλειψη θέλησης ως προς την αντιμετώπιση ουσιαστικών οικονομικών κινήτρων.

Λέξεις-κλειδιά: Τεχνητή νοημοσύνη, Διαδίκτυο, αλγόριθμοι πρόβλεψης, δημοκρατία, ατομικά δικαιώματα

Table of Contents

1. Introduction	6
2. Methodology	8
3. Democracy	10
3.1. Defining Democracy	10
3.2. Conflicting Definitions	10
3.3. Liberal Democracy and Privacy	13
3.4. Facing the Contradictions	14
4. Artificial Intelligence and Relevant Technology	18
4.1. Basic Definitions	18
4.2. General Ethical Concerns	21
4.3. Big Data in Practice & Privacy	24
4.4. The Technological Landscape of Everyday Life	25
5. Cases of malpractice	27
5.1. Electoral Manipulation	27
5.2. Commercial misuse	28
5.3. Misuse by Authorities	28
5.4. Filter bubbles and Information Control	29
5.5. Artificial Intelligence and the Democratic Society	30
5.6. Threat Assessment	31
6. Relevant Regulatory Approaches	34
6.1. Data Privacy in the European Union	34
6.2. European Commission - High-Level Expert Group on AI	40
6.3. United States of America & Platform Self-Regulation	43
6.4. Democratic Vigilance in Public Policy	45
6.5. Policy assessment	46
7. Conclusion	49
References	51

1. Introduction

Human societies are shaped and characterised by a multitude of things. The material conditions and relations of a society have, undoubtedly, a major impact on how that society functions. The technological landscape of a society is a major determinant of material conditions in that society (consider the importance of irrigation and farming techniques in early agricultural societies for example). A social study must, by necessity, account for it in some fashion. Furthermore, the political system of a society being a major facet of it, it is expected that it would interface with technology in some fashion. The various aspects of society inevitably experience feedback loops: in our case, a society's needs drive technological progress, they determine what fields of research are prioritised and what sort of infrastructure is invested into. In turn, the technological innovations that arise, and the scope of their application, alter the material conditions of society and open new possibilities.

Currently, the state of the broader field of information technology is the result of dramatic progress that has occurred over the past 80 years. The computational power of electronic systems, the easy of communication enabled by the Internet, the ability to collect previously unfathomed amounts of data, are examples of technological advances that form a defining feature of today's global society. Consider that, since mobile phones first became a common piece of hardware in the 1990s, they have become ubiquitous and expanded their capabilities to perform as general-purpose Internet terminals¹, GPS devices, and cameras.

At the same time, traditional democratic governance in much of the world predates these developments, thus raising pertinent questions as to the compatibility of these governing systems and the rapidly evolving technological landscape they find themselves in: is democratic governance, requiring a degree of trust between government and governed, viable at a time when the Internet can be weaponised for the mass spread of misinformation? Are our civil rights irreparably compromised by such technologies as Artificial Intelligence (AI) which can easily gather and deduce information about us with dubious consent?

These questions are central to this study. Its primary object - artificial intelligence in the context of posing a danger to privacy and democracy - is a rephrasing of this question of the compatibility of our established forms of democracy and civil rights with the evolving technologies that contribute to the field of artificial intelligence and the capabilities of those technologies. In other words, we seek to study and understand the feedback loop generated between our development of artificial intelligence and the political system that it exists in.

To do this it is necessary to study all parts of the loop; the capabilities of artificial intelligence exist in the context of their application. To get a complete picture we must understand the motivations behind the development and deployment of AI so that we may determine its social role, which is the determinant of the nature of its impact on social structures and institutions such as democracy and civil rights.

We begin this study by giving particular emphasis on the political science involved. This is done for two reasons: firstly, because democracy is a broad category of ideologies that require disambiguation. Secondly, because politics is an arena of moral arguments, and thus disagreement or non-alignment in political matters means that any political subject is potentially contentious. It is thus the opinion of the author that for an honest examination of a subject that has an inherent political aspect, a detailed political context must be provided to avoid confusion and reach an understanding. It is hoped that even in case of a fundamental divide in opinion, the context provided forms a foundation for the sum of the study to be a coherent text.

The second part of the study involves examining the technological background and its impact on the democratic political values established. This involves examination of ethical concerns that are resulting

¹ A function of particular importance in countries with emerging economies, where computer penetration is lower than smartphone penetration - people in those countries are dependent on mobile networking for Internet access. (Silver et al., 2019)

from AI technologies or implied by future advances in the field, of privacy and other rights violations that are incidental to AI technologies, and of documented instances of bad faith and unethical use of AI and relevant technologies (a point to which considerable length is devoted). Recognising the role of the Internet as the medium used for much of this, an ethical examination of online speech, communication and control is included.

Finally, there are specifics of the legal landscape that must be acknowledged and examined critically. Specific attention is given to European legislation and policy initiatives that regard data protection and ethical AI, as well as law regulating the Internet, and in particular United States law with regards to freedom of speech and the obligations of platform owners. The texts examined and the feedback collected on them over the years constitute a window into the vision that state authorities are willing to pursue as regards AI, democracy and the Internet. They are a means to determine and judge state attitudes towards the issues discussed, and judge whether or not a well-intended and sufficient intent towards them is being exercised by the state.

In summary, the present study can be seen broadly in four parts:

- An exploration of the political terminology, vital for explicating and framing the political concepts involved as “democracy” can be seen and understood in a multitude of often contradictory ways
- A study on how Artificial Intelligence and relevant technologies interface with the political aspects of the study that raise ethical concerns
- Specific examination of real instances involving the use of AI and relevant technologies for unethical purposes or using unethical means
- A study of important legislative actions and officially sanctioned work for the purposes of addressing the issues discussed above or which forms a foundation for the development of the relevant technological landscape

The study as a whole is made to, aside from provide an overview of the technical aspects involved, serve as a rejection of historical determinism. The matter of AI posing a “danger to democracy” is ultimately approached as an open-ended question: the design, rather than to provide a fundamental affirmation or rejection of the premise, is to understand the social and technical mechanisms that pull Artificial Intelligence in a hazardous direction. Respecting individual and social self-determination, we seek to examine how these mechanisms are being or can be addressed.

2. Methodology

The primary concern of this study is to determine whether artificial intelligence poses a threat to democracy and individual privacy. To accomplish this, we need to clarify ambiguities in terminology. Artificial Intelligence and Big Data are technical/technological terms that are clear-cut and refer to specific concepts. The “right to privacy” as well is a fairly specific term, with a relatively recent and well-documented history. “Democracy”, however, has referred to multiple and sometimes contradictory concepts, and therefore demands contextualisation. Thus, the first question we will address is one of defining democracy.

This necessitates looking beyond a simple etymological definition: democracy, as a concept, has been a part of human history for thousands of years. As such, as foundational as its dictionary definition is, an analysis of the historical development of the concept of democracy is equally, if not more, important. This means a historical study of democracy from antiquity to today. Of relative import to this analysis are the Roman Republic, the English Civil War, the French Revolution and the transition from feudalism to capitalism in Europe.

The reason for the focus on these particular subjects is that in the historical examples they provide, a study can be made of how democracy in practice attempts to live up to democratic ideals, and the associated successes and failures in accomplishing that (for example, the transfer of power from the monarchy to Parliament in the English Civil War, and the stamping out of attempts to expand the electoral franchise that elected said Parliament).

At the end of this historical analysis is the present day, and the predominant form of democratic governance found in liberal democracy. Therefore, liberal democracy, its ideological underpinnings and historical development, are expanded upon at length, along with relevant political critiques such as those of Marxist ideologues. The concept of privacy is also examined in the context of liberal democracy.

Greater emphasis is perhaps given on the political aspects than similar studies. Consider Manheim & Kaplan’s (2019) study on the same subject: The authors are largely concerned with the state of American democracy following the 2016 Presidential election and AI-powered interference in it, often missing a broader picture of historical socio-political development. Statements such as “[t]he institutional press has, over the 20th century, developed journalistic norms of objectivity and balance” (p. 150) are presented uncritically, not accounting for the historical trajectory that has created and maintained that institutional status and potentially undeserved trust. For this reason in the present study “democracy” is not treated only as an established status quo, but as the ideal that democratic societies strive for, and which in their diversity they realise in different and varying degrees.

Completing the political analysis is one prerequisite for examining how Artificial Intelligence influences democracy. Another is understanding the technological aspects involved themselves. There are three technological categories of relevance to this. One is Artificial Intelligence itself, but as the study concerns its socio-political implications, we also find the Internet and Big Data to be equally important fields, as they represent a medium and a set of techniques through which AI’s intrusiveness into everyday life becomes realised.

Of particular concern is the observed historical development of the Internet and the concentration of online socialising on a small number of very large and often interconnected social media platforms. Those platforms’ command of a great portion of the online user base market share, and the resulting control of the generated data is thus problematised as a particular enabler of abusive use of AI.

The analysis of the more academic and theoretical aspects of the technological study is then followed by an examination of practical cases in which AI has been used to subvert the democratic process. The high-profile and well-documented case of Cambridge Analytica’s role in the campaigns for the 2016 Brexit referendum and the US Presidential election is given due consideration. On an individual level attention is given to cases where privacy or individual rights have been or can be disrespected through AI systems, followed by an examination of filter bubbles and the ways in which they are harmful to public discourse, a vital component of an open democratic society.

The feedback from society and the democratic state on these issues is also an object of this study: Artificial Intelligence is a rapidly developing field, and the historical evolution of political systems is a

perpetually ongoing process. As such, how they react to each other is of direct interest. Of particular interest is the European Union's (EU) General Data Protection Regulation (GDPR), which represents an effort to comprehensively address matters of data privacy and enshrine into law individuals' data rights. A thorough, article-by-article analysis of the GDPR is given, followed by an examination of its first two years of practical application and official evaluations, highlighting its successful application and valid criticisms levied towards it.

While the GDPR is a hallmark piece of data privacy legislation, it does not address Artificial Intelligence specifically (in fact, the legislators went to great length to ensure technologically-neutral wording in order to future-proof it); though the European Union has no adopted legislation regulating AI, it did form an advisory body, the High-Level Expert Group for AI (AI HLEG), to produce documents relating to the ethical application of AI and policy recommendations on AI. The group's work, particularly with regards to the former, is the subject of analysis after the GDPR.

Finally, we end our examination of public policy with an examination of US law's Section 230, part of legislation from 1996 which established the legal framework for online platform self-regulation is thus fundamental to online free speech. With online socialising and online political speech having become more prominent in the 25 years since the law was passed, Section 230 has proven fairly controversial, and so the discourse around it is likewise an important item for examination.

3. Democracy

The question of whether something poses a “threat to democracy” is highly loaded with meaning and subtext due to the ambiguity and contextual interpretation of terms. “Democracy” will be understood differently across geography, time, and ideology. In order to honestly engage with the research question, it is then necessary to first clear ambiguities regarding what we mean when we say “democracy”.

As such, we open this chapter with an etymological, fundamental look at the meaning of the term “democracy”. Following that is an exploration of the historical development of democracy², tracing the development of democratic institutions and the extent to which they fulfilled the democratic ideal promised by the etymology. Finally, we acknowledge that in colloquial terms “democracy” is used near-exclusively to refer to liberal democracy and as such offer a closer examination to liberal-democratic values and principles, especially that of privacy.

3.1. Defining Democracy

Etymologically, the term has its roots in Greece of antiquity, where Athenian democracy was established through Cleisthenes’ reforms following the overthrow of the tyrant Hippias (Lallement, 2000/2004). The word itself means “rule by the people”, which was intended to distinguish the regular citizenry from the political elite (Heywood, 1992). This distinction is upheld by Hobbes (1651), who in *Leviathan* distinguishes between three forms of government:

The difference of Commonwealths consisteth in the difference of the sovereign, or the person representative of all and every one of the multitude. And because the sovereignty is either in one man, or in an assembly of more than one; and into that assembly either every man hath right to enter, or not every one, but certain men distinguished from the rest; it is manifest there can be but three kinds of Commonwealth. For the representative must needs be one man, or more; and if more, then it is the assembly of all, or but of a part. When the representative is one man, then is the Commonwealth a monarchy; when an assembly of all that will come together, then it is a democracy, or popular Commonwealth; when an assembly of a part only, then it is called an aristocracy. Other kind of Commonwealth there can be none: for either one, or more, or all, must have the sovereign power (which I have shown to be indivisible) entire. (p. 114)

In this theoretical framework, “democracy” is a state in which class tensions between an “aristocracy” and the non-aristocratic population have been resolved in favour of the latter.

The French Revolution and the French First Republic were such attempts at democracy, by taking political power from the French crown and the landed aristocracy and transferring it to the general population through representation in the National Assembly. Though short-lived, much of their democratic ideals would be passed on through the introduction of the Napoleonic Code to much of Europe and form the basis for the abolition of feudalism and transition to parliamentary democracy in many states over the century following the end of Napoleon’s reign. The American revolution likewise was rooted in this thinking, with the declaration of independence asserting universal equality.

3.2. Conflicting Definitions

The distinction made above between a political elite and the citizenry is an abstracted one and highly contested, both throughout history and even today. The concept of “citizenship” itself, while with regards to the aristocracy is describing a someone who does not possess their position of privilege, can be quite limiting indeed, and even within it further restrictions on democratic rights can be imposed.

If Athenian democracy serves as a prototypical form of democracy, its franchise was very limited, with only adult male citizens having the right to vote, making up only a fraction of the total population of the city-state: slaves and metics (resident aliens) were not considered citizens, and society heavily discriminated against women, excluding them from political participation (Thorley, 1996). It is clear that

² The focus here is entirely on the development of western democratic institutions and philosophy, its exportation from Europe to a global reach acknowledged as a result of the bloody legacy of European colonialism and imperialism. Democratic ideas and institutions that trace their lineage through non-European beginnings, such as Islamic democracy and the shura councils (Ansary, 2009), are hereby acknowledged, but sadly beyond the scope of this study.

while political power was transferred away from an aristocracy, it was transferred to only a portion of the general population and not the whole.³

The Roman Republic (and later Empire), also a prototypical democratic state of antiquity, shaped much of European history. Through its republican period, unlike the Greek *polis*, the former aristocracy survived and maintained power as the Senatorial class, with democratic institutions with a wider franchise complementing rather than replacing it (Anderson, 2013). When the republic transitioned into the empire, those institutions and the Senate gradually lost their power, reduced largely to a ceremonial role and completely sidelined by the reign of Diocletian, and democratic institutions in western Europe were left to develop through those of the various Germanic kingdoms that succeeded the western empire.

As antiquity transitioned to the middle ages and rulership over the land from imperial administration to feudalism, society was generally structured into the three estates: the nobility, the clergy, and the commoners, who were usually represented by urban, bourgeois classes, rather than the peasantry⁴. These estates often yielded power over the monarchy and provided a structure that would, in many countries, evolve into the modern parliamentary system. For example, in 17th century Britain conflict between King Charles I and Parliament led to the English Civil War, in which the Parliamentary forces won, Charles I was executed, and the monarchy even abolished from 1649 to 1660. Though the electoral franchise at the time was very restrictive and society highly unequal, this period of anti-monarchical sentiment inspired some radical groups who argued for general equality, notably the Levellers and the Diggers. (Hill, 1980)

Social stratifications that run contrary to democratic ideals can be observed throughout history in democratic states even as late as the last century. A closer look at the more immediate origins of modern democracy is therefore warranted. The American (1776) and French (1789) revolutions were mentioned earlier as hallmarks of the development of modern democracy. The US declaration of independence asserts that “all men are created equal”, and the French *Declaration of the Rights of Man* likewise asserted that all men are free and equal in rights.

Neither state immediately abolished slavery. The French First Republic formally abolished slavery in 1794, after the Haitian revolution, the largest slave uprising in modern history, threatened French control of the colony (Popkin, 2011). Napoleon would later re-instate slavery and Haiti would win its independence from France and be declared a free republic in 1804. The re-establishment of slavery in France would last to 1848. The United States of America also did not abolish slavery following independence from the United Kingdom. Though many states independently did so, the rural southern states with their plantation-based economy maintained slavery to the point of secession from the Union and civil war, over 80 years after the war of independence. The victory of the north in the civil war finally ended the practice and granted citizenship to the freed slaves.

With regards to gender equality, democratic progress was even slower, with the movement for women’s suffrage only achieving its goals in the 20th century around most of the world. In the above cases of the USA and France, women got the right to vote in 1920 and 1944 respectively. Prior to that, in France, the Napoleonic Code had explicitly established a patriarchal family relationship with the husband as head of the household.

If we were to follow a strict Hobbesian definition, we would have to classify most of the traditionally democratic states as aristocracies. Yet, even lacking the principle of true universal suffrage we couldn’t claim that ancient Athens, or the United States prior to the civil war, or the French First, Second, and Third Republics, were not democracies in any way the term is normally understood. They were democracies because they defined their political structure contrary to what came before them, the state and the political elite that they overthrew, with power transferred from them to the citizenry, however limited or internally hierarchical the concept of the citizenry might have been at the time.

³ For a detailed material analysis in the role played by slavery in the development of ancient Greek society, including the development of ancient city-state democracy, see Anderson (2013).

⁴ This specific breakdown is mostly associated with the Kingdom of France, though it has parallels elsewhere. The English equivalent socio-political structure for instance was King, Lords, and Commons, which integrated the crown into the estates, and with the clergy and nobility represented as a single class, though in the context of the English parliament the French model also applies, even as the Lords Temporal and Lords Spiritual sit in the same assembly.

And if democracy is then treated as a process, one of transferring power from an elite to commoners, rather than a state of being, it is possible to identify an opposite process, the creation and establishment of a new elite, as undemocratic. Furthermore, so long as hierarchies within the commoners exist, the ideals of the democratic process cannot truly be fully realised, as the more privileged classes have greater freedom to practice their democratic rights. One example of this is offered by the United States, where even after the end of slavery, many of the former slave states adopted literacy tests as a prerequisite for voting rights, designed with the purpose of discriminating against black Americans (Kates, 2006).




Figure 1. The transfer of power from an elite class to the general population can be seen as the practical application of democracy. The concept of a “threat to democracy” could then be seen as the reverse process, the consolidation of power in a society into the hands of an elite.

Marxist political analysis attempts to address these contradictions on the basis of economic class, through a breakdown of society along economic class, separating between capitalists and the proletariat. In it, the bourgeois democratic state is regarded as a “dictatorship of the bourgeoisie”, as the capitalists are the de facto ruling class, to be contrasted with a hypothetical proletarian democracy - a “dictatorship of the proletariat” - which would serve as a stepping stone to the abolition of the state and classes. Writing in 1918 in defence of the Russian Revolution in his polemic *The Proletarian Revolution and the Renegade Kautsky*, Lenin spells out this criticism of bourgeois democracy:

Bourgeois democracy, although a great historical advance in comparison with medievalism, always remains, and under capitalism is bound to remain, restricted, truncated, false and hypocritical, a paradise for the rich and a snare and deception for the exploited, for the poor. It is this truth, which forms a most essential part of Marx’s teaching, that Kautsky the “Marxist” has failed to understand. On this—the fundamental issue—Kautsky offers “delights” for the bourgeoisie instead of a scientific criticism of those conditions which make every bourgeois democracy a democracy for the rich.

Defining democracy then becomes highly subjective: from the point of view of a traditional liberal, we already live in a democratic society, whereas a Marxist would say that bourgeois democracy is inherently undemocratic. An ethno-nationalist might try to assert a belief in the idea of the nation-state and see the exclusion of foreigners or minorities from citizenship as an expression of their conception of “true” democracy.

Even in countries where one would assume that liberal democracy is taken for granted, suffrage is not always truly universal. The right to vote being conditional on age can be handwaved in the sense that the restriction is equally applied to everyone. However, many countries’ laws contain conditions for more selective disenfranchisement, such as in the case of prisoners (a common situation in the United States, where many states’ laws disenfranchise people convicted of felonies even after they have fully served their sentences, and where these laws disproportionately affect people based on race (Uggen et al.,

2016)). Additionally, other issues such as gerrymandering (the drawing of electoral district borders to produce an artificial advantage for one political party over its competition) complicate the practical application of democracy.

These are the abstract terms in which we must consider the research question. Whichever one way we choose to understand the question of a “threat to democracy”, it is possible for our interpretation of what democracy even is to not be universal.

3.3. Liberal Democracy and Privacy

Having considered the above, we will make an arbitrary decision to focus on liberal democracy, as it is the democratic form most associated with the generic term in casual, non-academic conversation, and whose culture of individualism is most relevant to the concept of privacy (Kymlicka, 2002/2006).

With the decline and end of feudalism, and secularisation of the state, the estate system is gone; the commons are free to shape their own destiny. This vague concept of “liberty” is at the heart of liberal democracy, and we need to examine how the liberal democratic state attempts to manifest it in practice. Perhaps one of the most fundamental and influential concepts in this is that of the separation of powers, which French enlightenment philosopher Montesquieu expanded upon in his 1748 work *The Spirit of Laws*. In *The Spirit of Laws* (1748/1899) Montesquieu defines political liberty as “tranquillity of mind arising from the opinion each person has of his safety” (p.151). Defining the three powers of the government as the legislative, the executive, and the judiciary, Montesquieu argues that these powers must be separate for that tranquillity of mind: the legislative and executive being held by the same person or body engenders the possibility of tyranny; the judiciary and the other two, that of arbitrary judgements and the judge being emboldened to act oppressively. Lamenting that the various Italian republics at the time, such as Venice, did not separate these powers of state, he wrote:

In what a situation must the poor subject be in those republics! The same body of magistrates are possessed, as executors of the laws, of the whole power they have given themselves in quality of legislators. They may plunder the state by their general determinations; and as they have likewise the judiciary power in their hands, every private citizen may be ruined by their particular decisions. (p.152)

A relevant concept of particular interest to this study is that of the press and the media as the “fourth estate”, referring to their role as a mediator between the government and the governed. Though an informal ideal, this function elevates the press and places it in a position of responsibility. And because it is an informal ideal, in practice the media often fall short of it. Hampton (2010) notes the tendency for commercialisation and oligopoly, as well as interference from and collusion with the state, as contradictory forces that increasingly hamper the press’s ability to truly realise its fourth estate function, with the more historic examples of the press holding the government accountable forming exceptions rather than the rule. Because of this study’s subject matter, it is of particular interest to question this role of the media as an increasing number of people receive their news online, often through search results and social media feeds.

The transition from feudalistic monarchy into bourgeois democracy is also marked by economic development, as the mercantile classes of Europe started amassing greater fortunes and political influence (as we discussed in the case of parliament in the English Civil War). This process was sped up by the French revolution and the industrial revolution; by the end of the 19th century, between the industrial pull towards urbanisation and growing liberal political demands in the wake of the French revolution, serfdom had been abolished throughout Europe, and any remnants of the feudalistic mode of production had given way to the capitalist mode of production. Thus we have the rise of the private economy over that of the land-owning nobility, and a first encounter with one of the definitions of “privacy” as freedom from the state, though we usually conceive the “right to privacy” in social rather than economic terms.

We mentioned earlier that the third estate was usually represented by the bourgeoisie, even though it referred to commoners as a whole. Indeed, at the time of the English Civil War, the electoral franchise was limited to landowners. One of the demands of the Levellers, one of the factions that arose in the Parliamentary faction and had widespread influence in Parliament’s New Model Army, was the

expansion of the franchise⁵, but they were seen as a radical threat and were forcibly dissolved. When the Levellers' draft constitution was debated by the Army Council, the counterargument posed by Commissary-General Ireton against expanding the franchise was that the right to vote ought to be limited to those with "a permanent fixed interest in this kingdom". (Hill, 1980)

The franchise did eventually expand, and universal or near-universal suffrage became the norm in countries that adopted liberal democratic forms of government, posing a contradiction: the liberal state which is built on ideas of equality among those it represents, and the reality of inequality among them. An attempt to address that contradiction is through redistributive welfare systems: though achieving the true ideal of liberal equality is not truly possible with post-facto corrective measures, they make an attempt to permit everyone equal participation in society where poverty might be an obstacle thereof. (Kymlicka, 2002/2006)

Privacy as freedom from the state was mentioned earlier. Freedom from state surveillance and control, and the implicit or realised intimidation and violence that arise from that, are important civic rights, but the right to privacy in these terms can be expressed both individually and communally. There is then a certain independence of what is called the public sphere from state authority that is related to, but not the same, as the individual's freedom from the state. Similarly, an individual's right to privacy is not only that with relation to the state, but also to that public sphere; it is an individual's right to privacy from society. The increasing ease of communication and access to media following the industrial revolution heightened awareness of this: one of the earliest modern legal essays arguing for a right to privacy was written in 1890, in reaction to gossip columns in newspapers (Glancy, 1974).

The distinction is important: In American law, one of the hallmark Supreme Court decisions with regards to privacy was *Griswold v Connecticut* (1965), which protected the liberty of married couples to use contraception. Kymlicka (2002/2006) notes that this interpretation of privacy as that of the married household towards the state is restrictive: it fails to protect personal privacy and autonomy within the household, and even potentially restricts the state's ability to interfere in family life to do so. The individual's right to privacy from society then needs to be argued for independently.

One last political concept that is important to the research question is that of freedom of speech, generally regarded as one of the more fundamental democratic liberties. This is not an absolute freedom: democratic societies tend to impose restrictions on speech so that cases like threats of bodily harm or incitement to crime are not typically protected. Freedom of expression is, nevertheless, important for an open and democratic society, though even that can be a contentious topic. Philosopher Karl Popper, in describing the "paradox of tolerance" in *The Open Society And Its Enemies* (1945), makes allowance for the right to suppress intolerant speech in order to prevent its speakers from abrogating tolerance. This argument rests on the idea being that such bigoted speech poses a fundamental threat to open, democratic societies:

But we should claim the *right* to suppress them if necessary even by force ; for it may easily turn out that they are not prepared to meet us on the level of rational argument, but begin by denouncing all argument ; they may forbid their followers to listen to rational argument, because it is deceptive, and teach them to answer arguments by the use of their fists or pistols. We should therefore claim, in the name of tolerance, the right not to tolerate the intolerant. (p.265)

That kind of allowance is not universally accepted, and indeed pre-supposes ideological support for liberal, bourgeois democracy. Writing in 1938 for a Mexican magazine on occasion of the leadership of the Confederation of Mexican Workers supporting a campaign to suppress the reactionary press, Trotsky argued that a bourgeois state, given the authority to police and censor ideological speech, even if initially conceived as a means to fight against hateful speech, would then use its role as protector of political discourse to silence the workers' movement (Trotsky, 1938). While Trotsky was himself a communist revolutionary and as such an opponent of liberal democracy, the potential for censorship laws to set a precedent and be abused demands consideration.

3.4. Facing the Contradictions

As an idea that's been conceived, evolved, and practiced for more than two thousand years, democracy demonstrates a great diversity in form. We've already noted that the decision to focus on liberal

⁵ Not universal suffrage. Leveller demands excluded those without economic independence, such as wage labourers and paupers, and gender equality wasn't even a consideration.

democracy in particular is arbitrary, based on its present status as the defining form of democracy and on its close association with an individualistic worldview, which relates it deeply to the concept of privacy, a fundamental aspect of the present study. There is a diversity of form, however, even within liberal democracy as demonstrated by the (relatively recent, in historical terms) question of privacy as the right of a family or the right of the individual. Across time and geography, states characterised as liberal-democratic vary, and the term itself becomes a vague approximation of a set of values and ideals⁶.

The decision to *focus* on liberal democracy also does not mean that, as regards the rest of the study, it will be treated as the only form of democracy, or as uniquely closer from other forms to the democratic ideal. As demonstrated in its historical variety, liberal democracy itself faces shortcomings when it comes to living up to that. We used earlier the idea of the general population and an elite, and the transfer of power from one to another, as a founding block for coming up with a generic definition of what a democracy is, as a state and as a process. Indeed, the end of feudalism and the loss of Europe's traditional elite class, the nobility, marked a process of democratisation. But even before the end of feudalism the Third Estate was not uniform in power; the transition from a feudal to a capitalist economy has elevated the bourgeoisie into the elite of the restructured society as it developed over the past two centuries. Democratic tensions thus inevitably arise along economic class, between the bourgeoisie and the proletariat.

That Marxist critique of bourgeois democracy finds the expression of its ideological opposite in neoliberalism. A political as well as economic ideology, according to neoliberal thought the markets and private economy need to be protected from government intervention. National governments need to be bound by rules, overseen by supranational institutions whose role is to ensure the functioning of the capitalist economy without "political intervention" (a role realised in modern times by such institutions as the International Monetary Fund (IMF), the World Bank, or the European Union) (Mullan, 2020). That separation of the economy from politics is, of course, political in itself, and demonstrates a fundamental distrust in democracy. Where Lenin would decry "the hypocrisy of bourgeois democracy" (Lenin, 1918) with regards to political rights when those threatened the capitalist economy, a neoliberal thinker would applaud at the restrictions of such liberties for the sake of economic ones.

Such contradictions are not, in and of themselves, fatal to a democratic governing system. They are a historical reality and the social antagonisms that arise from it are something that every democracy has had to contend with. But the conceptualisation of an "elite" by itself necessitates an existing power disparity; thus, the preservation of the status quo without an ongoing democratisation process forms by necessity a resolution of the contradiction in favour of the current elite class, as the power disparity between them and the rest is likewise preserved. We've seen this in how the Roman Republic resolved the democratic contradictions in its political system in favour of the senatorial class, for example, and it could easily be argued that the Marxist critique of liberal democracy as a "dictatorship of the bourgeoisie" holds true, with the class antagonisms between capitalists and workers being resolved in favour of the former within an otherwise democratic system.

In formulating the hypothesis that a reverse transfer of power, that is to say, from the general population to the elite, would constitute a threat to democracy, we are also begging the question as to whether there is historical precedent we can point to. Indeed, the 20th century provides a wealth of examples of democratic systems dramatically overthrown in the name of anti-communism, from the Greek military junta to Augusto Pinochet's dictatorship in Chile. Moreover, Ronald Reagan's presidency of the United States of America and Thatcher's terms as Prime Minister of the United Kingdom proved foundational to neoliberalism's rise to prominence, which, with the collapse of the Soviet Union in 1991 which had up to that point served as the standard-bearer of an opposing dogma, became the guiding ideology for the global economic order. The impact of this ideology in prioritising the interests of the capitalist class can be seen in the ways wealth disparity has increased over the past 40 years in the first world⁷: according


⁶ Consider as a thought experiment the United States of America: did it qualify as a liberal democracy from the moment of its founding, or did slavery disqualify it? Should it be considered a liberal democracy following the end of slavery, but prior to the end of segregation? The same country practiced "liberal democracy" in radically different ways throughout its history.

⁷ The term "first world" here is used in its original Cold War-era meaning, when it referred to the developed, pro-American capitalist countries. By contrast "second world" referred to the Soviet sphere of influence, and "third world" to those less-industrially developed (and often freshly independent

to research by Horowitz et al. (2020) for the Pew Research Center, US upper-income households have experienced greater income growth compared to medium and lower income households since 1970. Household wealth has also consolidated in the hands of upper income households, from 60% of the national aggregate wealth in 1983 to 79% in 2016. Furthermore, in the decade following the recession of 2007-2009 only the top quintile of families saw an increase in household wealth - the lower quintiles saw a decline of up to 39%.

The gaps in wealth between upper-income and middle- and lower-income families are rising, and the share held by middle-income families is falling

Median family wealth, in 2018 dollars, and share of U.S. aggregate family wealth, by income tier



Note: Families are assigned to income tiers based on their size-adjusted income.

Source: Pew Research Center analysis of the Survey of Consumer Finances.

"Most Americans Say There Is Too Much Economic Inequality in the U.S., but Fewer Than Half Call It a Top Priority"

PEW RESEARCH CENTER

Figure 2. Chart of US aggregate wealth shares by households categorised by income tier, from 1983 to 2016, by Horowitz et al. (2020). The effect of neoliberal policies adopted from Reagan's presidency onwards on wealth consolidation at the hands of the economic elite are clearly visible.

Even in spite of this reverse process then we see states subject to it that we would, without qualification, identify as democracies. Additionally, one could rightly point out that in the cases mentioned where democracy was dissolved were such that the dissolution of democracy was done in order to affect that transfer of power rather than as a result of it, and that democracy has usually been restored in those countries since, in spite of the affirmation of pro-business policies. There are two issues with this. Firstly, the resulting democratic form is one that has to contend with the pro-business reforms made by the state's former neoliberal stewards, which might have since become institutionally entrenched. It is no coincidence that, after Reagan and Thatcher and the fall of the USSR, the '90s saw many parties with social-democratic platforms embrace market logic, and the economic restrictions that came with that. For an illustration of this consider the 2011 Eurozone crisis and the policy restrictions that states such as Greece and Portugal were forced to accept: even after the election of an anti-austerity governing coalition in 2015, the Greek government found itself unable to fulfill its democratic mandate to fundamentally renegotiate its position vis-à-vis the wider economic structure (Lowen, 2015). Too tightly integrated into and dependent upon supranational organisations such as the EU and the IMF, to break away from them would be an economic shock greater than the continuing implementation of neoliberal policy.

Second, the institutional decay implied by the transfer of power need not be a short-term affair. The transition from the Roman Republic to the Roman Empire, for example, is easily marked by historians

following decolonisation) countries which formed a sort of ideological battleground for the USA and USSR to compete in.

as starting with the ascension to power by Emperor Octavian. But that rise to power followed the weakening of (elected) Senatorial control by the dictatorships of Sulla and Julius Caesar; additionally, the title adopted by Octavian was *Princeps*, styling himself as more of a “first citizen” of the Republic rather than an absolute ruler, and for the the first few centuries of the Roman Empire (referred to by historians as the “Principate”), the Roman Senate maintained significant institutional authority, and was presented with multiple opportunities in which it could have potentially wrested power back. (Duncan, 2012) To the Romans alive at the time, the question of whether the Roman Republic had ended would be a much more complicated affair than it is to historical hindsight. And so today it should be difficult to recognise whether we are part of a historical current which in time will be recognised as a development towards the end of democracy.

4. Artificial Intelligence and Relevant Technology

In this chapter, we will provide an analysis of the socio-technological landscape that is relevant to the study subject. A brief introduction into the terminology of Artificial Intelligence and related concepts is made. Following that is a look at ethical concerns raised by these technologies in modern society. Finally, greater focus is given on the ways in which Big Data often involves invasive data collection and violations of privacy.

4.1. Basic Definitions

The missile knows where it is at all times. It knows this because it knows where it isn't. By subtracting where it is from it isn't, or where it isn't from where it is (whichever is greater), it attains a difference or deviation. The guidance system uses deviations to generate corrective commands to drive the missile from a position where it is to a position where it isn't, arriving at a position where it wasn't, but is now. Consequently, the position where it is, is the position where it wasn't. So it follows that the position where it was, is the position where it isn't. In the event that the position where it is now is not the position where it wasn't, the system has acquired a variation. The variation being the difference between where the missile is and where the missile wasn't.

(Colonel (Ret) George Grill, on the GLCM missile guidance system)

Artificial Intelligence

The term “artificial intelligence” is applicable to a wide field of practices. Though popular imagination might evoke images of robotic assistants and all-powerful control systems, AI principles commonly find application in equally mundane tasks, such as data mining and analysis. While advances in the field of robotics are impressive and demand examination, the more subtle ever-presence of the science in data analytics is just as important to matters of privacy and democracy, if not more so.

At the same time as AI today doesn't refer to sentient robots, some popular usage of the term is also too simplistic to satisfy some academic criteria; in video games for instance, “AI” is often used to describe any non-player entity behaviour, regardless of how intelligent it actually is. While we lack a singular, conclusive academic definition as to what AI is, Russell & Norvig (2005/2003) break down older definitions along two axes: how an artificial intelligence *thinks* vs how it *behaves*, and whether it does so emulating *human thought/behaviour* or following *rational thought/behaviour* (pp. 31,32).

In the above quote we are given a description about the guidance system of a missile, through which we can glean certain characteristics of intelligent systems: a goal that the system is working towards (where the missile isn't), the conditions in which the system is working towards that goal (where the missile is), and the ability to make corrections and adjustments as these conditions change on the way towards accomplishing its goal (accounting for variation encountered during guidance). One element missing from that description is that of learning, of the ability of the system to adjust its own logic through iteration and improvement.

The High-Level Expert Group on AI, set up by the European Commission as an advisory body, gives the following definition for AI:

Artificial intelligence (AI) systems are software (and possibly also hardware) systems designed by humans that, given a complex goal, act in the physical or digital dimension by perceiving their environment through data acquisition, interpreting the collected structured or unstructured data, reasoning on the knowledge, or processing the information, derived from this data and deciding the best action(s) to take to achieve the given goal. AI systems can either use symbolic rules or learn a numeric model, and they can also adapt their behaviour by analysing how the environment is affected by their previous actions.

As a scientific discipline, AI includes several approaches and techniques, such as machine learning (of which deep learning and reinforcement learning are specific examples), machine reasoning (which includes planning, scheduling, knowledge representation and reasoning,

search, and optimization), and robotics (which includes control, perception, sensors and actuators, as well as the integration of all other techniques into cyber-physical systems). (High-Level Expert Group on Artificial Intelligence, 2019a, p. 6)

Big Data

AI fundamentally consists of data processing, making it intertwined with the developing field of “Big Data”: the growing volume of data being generated in today’s networked world. The amount of data makes it unwieldy for human processing and thus requires the use of mechanical data analysis and AI.

The traditional definition of Big Data involves the so-called “three V’s”, volume, velocity, and variety, referring to the massive amount of data that needs to be analysed, the high speed in which new data comes in, and the variety of formats and structures (or lack thereof) of the incoming data. Software companies that develop and sell data analysis tools, like Oracle and SAS, add more “V’s” such as “veracity” on their business websites, but the basic description remains unchanged: the data is such that computing systems need to algorithmically sort through it.

Online Fora/Internet of Things (IoT)

With the invention and proliferation of the Internet, a transformation occurred. The public sphere has changed in two ways: First, it has become universal. Public posts on the Internet are visible by anyone who has Internet access. Secondly, they are permanent. That is to say, unless you opt into deleting your personal online history (which might well not be an available option), or the administrators of whichever platforms do so, they will remain accessible by all, potentially in perpetuity.

Compared to previous forms of public communication, be it speech or writing, the ease of communication and accessibility thereof creates a novel situation with regards to the public sphere and blurs the line between public and private life. Social media especially blurs that line further by promoting the publication of private life (the satirical news website The Onion ran a story in 2012 jokingly suggesting that every potential candidate for future US presidential questions has been rendered unelectable due to the sharing of inappropriate or self-humiliating posts on Facebook (The Onion, 2012)).




Figure 3. Big Data forms a link between the Internet and Artificial Intelligence: the Internet produces vast amounts of data through its everyday use. Service providers use Artificial Intelligence to utilise that data, whether to produce predictive algorithms as part of their service's user experience or even just to sell it to business partners.

In discussing these technical terms, especially in the context of their intersection with society and politics, it is vital that we do not treat them simply mechanistically. Technological development doesn't occur in a vacuum: it is the result of its own historic context, and it occurs to fulfil material and intellectual needs, it lives and dies by them. Artificial intelligence is no exception. The promise of perfecting the processes of the planned economy was a major drive for research into cybernetics in the Soviet Union in the 1960s. Parallels were drawn to the individual automated control systems of capitalist firms, but the Soviet plan's scale was vastly greater (Levien & Maron, 1964). Threatening established institutional and bureaucratic interests, the project to create the network required for that computerisation of the planned economy (called "OGAS"), stagnated in the decade afterwards (Peters, 2016). In modern times one of the biggest drivers for the development of improved AI techniques is commercial interests. Artificial intelligence has wide applicability and its problem solving capacity provides opportunities that businesses have had good reason to exploit: the Soviet Union failed, and in the global capitalist economic model the same motives that drove the development of private enterprises' "separate systems of control" (Berg, 1960, as cited by Levien & Maron, 1964, p. 26) were left to drive the development of modern AI.

4.2. General Ethical Concerns

"You have zero privacy anyway. Get over it."

Scott McNealy, CEO of Sun Microsystems (Sprenger, 1999)

Artificial Intelligence gives rise to ethical issues that are new and unique to the field, or put a new spin into older ones. Russell & Norvig (2005/2003) give a rundown (pp. 1064-1067):

- *Automation coming at the cost of people's jobs.* Russell & Norvig refute this by claiming that the development of AI has created plenty of jobs in the field and that AI's role as an intelligent agent tends to be to assist human work rather than replace it. It is curious to wonder, however: mechanical automation and industrial outsourcing have, over the past half-century, largely shifted employment in "first world" countries from agriculture and industry to services. AI reaching the point where automation of the service economy becomes possible does not seem that distant a possibility, so the concern retains some validity.
- *Humans might end up with too much or too little free time.* A different effect of automation as a result of AI could theoretically be a reduction of working hours. Here Russell & Norvig cite both optimistic and dystopian predictions of this from the late '60s but note that the reverse trend has taken place, with the systematisation of work and the profits from additional work in an information economy driving longer workhours.⁸
- *Dehumanisation.* The idea being discussed here is that research into AI could lead to humans being themselves perceived as automata. The authors dismiss this by pointing out that other scientific advances, such as the heliocentric model and the theory of evolution have moved our understanding of the world away from more anthropocentric models without costing humanity its nature.
- *Loss of privacy.* AI systems become enablers of mass surveillance as technological fields like voice recognition advance, especially as the public reaction to the September 11 terrorist attacks created a willingness to accept civil rights violations. This will be explored at length later.
- *Loss of accountability.* The question posed here is as to who is accountable for error on the part of an AI system. Who is responsible in cases of AI malpractice? Responsibility for medical AI support ultimately lands on the doctor who signs off to it, but other cases are less clear cut. This will also be discussed later in the study.
- *AI could cause the end of mankind.* The idea of a runaway AI adopting destructive behaviour is a frequent subject in science fiction. Though fantastical, the idea rests on questions of responsible design and maintenance, and the above question of accountability.

The High-Level Expert Group on AI also provided in its work *Ethical Guidelines for Trustworthy AI* (High-Level Expert Group on Artificial Intelligence, 2019b) a list of examples of critical concerns raised by AI:

- *Identifying and tracking individuals with AI.* AI provides the capability, through the use of technologies such as facial recognition and identification through biometric data, to automatically identify individuals.

⁸ Studies actually show a reverse effect of longer working hours on productivity (Pencavel, J. (2014), Collewet M. & Sauermann J., (2017)). The drive for longer working hours not being easily explained by market rules seems to put into question the logic that AI could create a financial incentive to employers to reduce working hours.

Such automatic identification can have a chilling effect and shape the way individuals react, and by extension society, in a bad way.

- *Covert AI systems.* The ability of AI systems to imitate human behaviour is a core element of the field of artificial intelligence, and systems that do not clearly communicate that they are AI and can be mistaken for human pose a fundamental ethical challenge.
- *AI enabled citizen scoring in violation of fundamental rights.* Citizen scoring engenders a threat to personal autonomy and can provide a basis for discrimination, going against egalitarian democratic principles.⁹
- *Lethal autonomous weapon systems (LAWS).* The ongoing development of such systems risks provoking an arms race and raises major issues of accountability.
- *Potential longer-term concerns.* The potential for scientific breakthroughs to bring unrealistic scenarios into existence, such as artificial consciousness, should not be entirely dismissed out of hand.

A closer examination of the High-Level Expert Group's work on this subject is made later in the study.

Giving satisfying answers to such ethical issues is difficult, and frequently impossible, and the above lists are by no means exhaustive. Giving these issues due consideration is a necessity for resolving concerns of laws and regulations governing the use of artificial intelligence. For a look at such considerations in the field of self-driving cars (an AI technology that is presently transitioning from an experimental to a production stage, lending ethical questions greater gravitas), see Holstein et al. (2018).

Following is a table of the concerns mentioned above, grouped together where appropriate, with consequences they could have towards individuals and towards the broader society (and consequently to democratic institutions and liberal-democratic value systems).

Ethical concern	Risk to individual	Risk to society
Automation/Loss of jobs	Being made jobless, job skills made obsolete leading to long-term joblessness	Large-scale unemployment social unrest that established institutions are unprepared to address
Dehumanisation/Loss of free time	Material degradation of living conditions, mental health issues	Apathy towards human rights and civil liberties
Identification tracking/Loss of privacy	Infringement on personal autonomy and loss of civil rights, inducement of paranoia	Direct degradation of human rights and civil liberties
Covert AI systems	Loss of privacy and civil rights if combined with identification recognition, inducement of paranoia	Apathy towards human rights and civil liberties, widespread social paranoia
Citizen scoring systems/Algorithmic bias	Becoming subject to discrimination, inducement of paranoia	Direct degradation of human rights and civil liberties, discrimination as an institution
Loss of accountability	Inability to seek redress for being wronged	Risk of treating errors and wrongdoing as a fact of life, apathy towards victims' demands for redress
Lethal Autonomous Weapons	Being on the receiving end	"Risk-free" warfare can result in

⁹ Discrimination on the basis of AI is also entangled with ethical concerns about Big Data. A focus study by the European Union Agency for Fundamental Rights (2018) notes the dangers of biased or poor data sets for algorithmic processing causing them to reinforce existing forms of discrimination.

Systems		apathy towards and dehumanisation of the enemy, might provoke jingoist sentiments
---------	--	---

Table 1. A breakdown of the risks posed by the problems raised by the development of AI.

Social media and surveillance on the Internet of Things

We made reference earlier to the mixing of private and public life that has become normalised in online socialising. Mitrou et al. (2013) raise several issues arising from this blurring. The anonymity afforded by the Internet makes online speech qualitatively different from speech in more traditional public fora¹⁰. The regulation of speech on online platforms is an open question as freedom of speech and the terms of use of those platforms, which are privately owned, create a contradiction. With regards to privacy, it is hard to say whether users are aware of the degree to which they make their personal information public on social media, where they are simultaneously given the task of networking with people of their choice and self-promoting on what is otherwise an open platform.

Facebook forms an archetypal example of this dual role, according to Piskopani in Mitrou et al. (2013), having started as a networking site for college students and over time expanding into a gigantic platform hosting public speech and political campaigning, playing a role in voter engagement as early as the 2008 US presidential election, and multiple other campaigns or movements since. Social media provides an easy and cheap technical framework for organising and building an audience. (ibid)

The infrastructure that supports this kind of mass accessibility is in large part the result of advances in miniaturisation and mobility of computing systems: mobile phones have quickly become ubiquitous since the '90s and over the past decade developed into capable computing machines with remote Internet access. This remote accessibility and reliance on this technology also involve a greater variety of data than just online communications, as the user's activity is tracked in a variety of different manners. Geolocation data is being tracked through GPS services, metadata about phone and application usage is immediately available, and microphones have been known to be used to collect data from users without consent (Green, 2017).

Indeed, the nature of voice assistants necessitates a microphone-always-on configuration for that technology to function. The user is perhaps supposed to be aware of this, either through explicit means (such as labelling of applications that gain access to a phone's microphone for instance) or implicitly (for example, if a user buys an Amazon Alexa unit, a voice-activated home assistant, an always-on microphone is the essence of what they are buying)(ibid). Even in the explicit cases, however, that expectation is not necessarily warranted: it is natural for an application like Facebook's messenger, which includes VOIP functionality, to demand microphone access, but this does not distinguish between its use for VOIP services and passive data collection.

Surveillance technology also includes more traditional systems like closed-circuit cameras (CCTVs), which have also proliferated in the past decades. With more widespread Internet connectivity, the potential grows for such systems to balloon beyond proportion. In November 2020, the city of Jackson, Mississippi launched a pilot programme through which willing citizens could give access to police surveillance of their private security cameras, including doorbell cameras (Guariglia, 2020). Privacy concerns naturally arise: even if a citizen does not give the police permission to use their cameras, a neighbour's camera that might happen to be pointed at their doorstep functions just as well.

User License Agreements and consent

The traditional way in which user consent has been derived with software and services is End-User License Agreements (EULAs). These are long documents delineating permissions and restrictions on the user. The purpose of EULAs has traditionally been the assert rightsholders' copyright on the software attached, but they frequently deviate from that role. Documents of that sort, however, are not an effective mechanism through which to derive consent (Neisse et al. (2016) outline a multitude of ways in which EULAs fall short of that intended purpose). Legal enforceability of such user agreements is also questionable as the language used and the rights claimed might be interpreted differently by each jurisdiction's legal system (Corbett, 2019).

¹⁰ Writing in 2013, Mitrou was more concerned with the protections afforded by this anonymity to hateful speech. The more recent trend of botnets being used to manipulate social media interactions puts this issue in a new light.

Similar documents describing services' privacy policies suffer from the same shortcomings. Even if they are easily accessible, the user is unlikely to read several pages of Facebook's data policy, for instance. This is not so much a question of enforceability, as data and privacy policies do not involve feedback from the service back to the user¹¹, but of personal awareness as to the service one is signing up for. For their part, Facebook's data policy page and Twitter's privacy policy page include vague statements that could easily be interpreted liberally to include eavesdropping as described above.

4.3. Big Data in Practice & Privacy

AI is used to sort through and categorise the data collected from the mechanisms mentioned previously - Big Data, by definition, is too unwieldy for human processing. Many of the relevant data-collecting companies, such as Google and Facebook, use their access to tremendous amounts of user information and screen time to provide a customised advertisement feed to their users. So for instance, a company seeking to advertise through Google's services pays Google to make sure the company's adverts are served to target demographics, rather than compete for advert space aimed at potentially irrelevant demographics. While this may not broadcast user information to third parties, and the data used for advertising purposes might even be depersonalised so as to make it unable to link it to the person it came from (as for instance Google's privacy policy page states), this operational framework provides strong business incentives to collect the data in the first place: these companies' monopolistic ownership of their user data enables them to offer and improve on many of their monetised services.¹²

While data collection for the purpose of serving better ads and improving on targeted services might sound benign, the data collection involved engenders the possibility for mass surveillance, either by the services themselves or by third parties that gain access to them, such as intelligence agencies. Such concerns are not unfounded: the practice of tech firms cooperating with American authorities by supplying them with user data was revealed in dramatic fashion in 2013's Snowden leaks, where the American NSA was shown to have access to the data of several major US tech firms (Landau, 2013).

The legal status of that surveillance is questionable, with the NSA possibly involving itself in illegal activities. It has, however, faced little repercussion if so. Even if such troubling abuse is not taking place, however, the possibility for it exists, and that alone should be cause for alarm. Because of the opacity of secret programmes like those, there is little if any effective democratic oversight on the actions of such intelligence agencies. The protection of civil rights, in a democratic society, being dependent on the willingness of officials to voluntarily impose on themselves the regulations that are meant to bind them, seems contradictory - and that is to say nothing of blind spots where such rights might not exist. For example, it is legal for the NSA to spy on non-US citizens without warrant; that one's private communications might be eavesdropped by a foreign country is little consolation (ibid).¹³

China's Social Credit System (SCS) is an example of how a more interventionist approach by the state has the potential to make use of mass surveillance. Lacking a liberal, individual-minded socio-political framework, the Chinese state under the leadership of the Chinese Communist Party is at liberty to use the modern capabilities of the surveillance state to implement programmes with the aim of shaping society (Creemers, 2018). Though originally intended as a system for shoring up the state's quality control over the economy, even its earliest iterations in the 2000s sometimes involved measures of social control.

Creemers (2018) and Liang et al. (2018) note that at present the SCS is not a panopticon, that it's not an implementation of total Orwellian social control, though foreign press has frequently remarked on it as such. The design specifications for current implementations (which there are many - the SCS is not a single, unified system) and the necessary infrastructure were laid down in 2014, and in many respects the system retains the commercial

¹¹ Attempts to make feedback from the service to the user a legal responsibility of the service in some cases have been made, as we will see for example in our examination of the General Data Protection Regulation in a later chapter.

¹² Advertisements have, historically, provided a vital backbone of monetisation to social networks which, being dependent on the mass membership that is only viable by providing a service free of charge to the end-user, would otherwise have been unable to meet investor expectations. (Falch et al., 2009)

¹³ This mode of surveillance also serves to highlight the opacity of the systems used for online socialising. One has to wonder whether users downloading Muslim Pro, a Muslim prayer app, ever even had an opportunity to consider that the US military was purchasing their location data for purposes of "counterterrorism". (Cox, 2020)

focus of previous attempts. As Creemers (2018) notes about the 2014 blueprint, the system, through its use of blacklists, doesn't make affordances for incremental scoring, "neither does it contain reference to the sort of correlative big data analytics that foreign observers have ascribed to the SCS" (p. 13). However, as noted earlier, such predictive AI offers a commercial incentive to gather that data in the first place. In the example of the NSA's PRISM programme, we also see that these private databases are well within the state's reach. This state-private sector collaboration and its potentially central role in mass surveillance is made real through the use of AI.

4.4. The Technological Landscape of Everyday Life

1999: there are millions of websites all hyperlinked together
2019: there are four websites, each filled with screenshots of the other three.

(Massad, 2019)

Technological progress in computing has undoubtedly changed the world we live in. The first electronic computers were massive, purpose-built machines that took up whole rooms. Advances in circuitry resulted in smaller, general purpose personal computers becoming available, and the Internet enabled every household to connect to each other through these machines. Thus a new online layer of society was created. The further advances and cost reductions that have resulted in smartphones, smartwatches, and so on have further enmeshed this layer into everyday life. The Internet is no longer a network that has a barrier of entry in the form of a sizable, expensive terminal, that you need to put time aside for. It instead forms a permanent background element of daily life, and is increasingly integrated into it such as to be almost inescapable.

And as this development has taken place, so has the Internet, in its relatively short existence, changed with it. Microsoft's dominance in the Operating System market and Google becoming synonymous with search engines serve as preludes for the centralisation into the oligopolistic structure of the broader tech industry that we see in 2021. Though the process was gradual, there is a sharp divide between the Internet of message boards that was prevalent in the 1990s and 2000s, and the Internet of social media to which online socialising has gradually centralised over the past decade and a half. The highly public nature of social media (a necessity for discoverability, a selling point of these platforms that was not as important for the old message boards and forums), and the resulting degree to which personal information becomes both public and centralised, are demanding of reflection.

Adding to that, online interactions have also become more than our deliberate Internet communications. Mention has already been made on data collection/surveillance through geolocation and voice assistant services offered by smartphones. Different web services also often offer integrated services; for instance, the ability to login through an existing Google or Facebook account instead of making a new account for the service, or the ability to link your accounts from different services to provide automated feedback between them (an example: Spotify, a music and podcast streaming service, and Discord, a chat service, allow you to link your accounts so that your friends on Discord can see what you are listening to). This sort of interconnectedness and integration might be a pleasant quality-of-life improvement, but it contributes to the same basic concern about surveillance and centralisation of data.

The benefits granted by cloud services have likewise resulted in a push for greater Internet connectivity and integration in software. Consider office tools, such as text and spreadsheet editors. Documents that once sat on a specific machine but now make use of an online storage service are more easily accessible and avoid the risk of data loss on account of power failure. System requirements are negated: an Internet connection and a browser is all you need to access software running on a cloud service. But this also means being connected and logged on to an account. The wall separating the private from the public has changed. It is not a technical fact but a voluntary, or at best legally-mandated, policy of the platforms hosting the software. Rather than requiring an intervention (like hacking or backdoors) for digital privacy to be infringed, the default state is that of no privacy, and intervention (like policies and software restrictions) needs to occur to ensure it instead. This can turn a breach of privacy from a malicious act into the result of negligence - questions of accountability naturally thus arise and make it harder to assign responsibility.

Another vector that must be considered in this analysis of data collection is the Internet of Things. While smartphones are part of the IoT, had online connectivity been restricted to personal computers and smartphones then the access terminals to the Internet would be controllable. But the interconnectedness seen between services is replicated in hardware, with the introduction of computing systems in other household electronics, resulting in

smart TVs or smart refrigerators. Such devices similarly deliver quality-of-life improvements: the ideal promised by the “smart home” concept is to assist with and automate domestic tasks. To achieve this ideal necessitates the integration of Internet-capable devices in everyday life with all the risks thus entailed.

Finally, how the technology available is used is an inextricable part of understanding the technological landscape. We briefly mentioned earlier the typical reasoning for the data collection discussed here, serving adverts. The Internet existing as a medium for business in such a systematised way demands that we perceive it from another perspective, that of the vendor in search for an audience. The ultimate realisation here is that, to the business side of the Internet, the audience functions as a service provider, access terminals functioning both ways, giving us access to the Internet, and the Internet access to us. Accessibility to us drives monetisation, which drives development and deployment. The economic structure of the Internet dictates to tech companies to obtain as much access to us as possible, and this should be considered when assessing where and why AI gets misused.

5. Cases of malpractice

The theoretical understanding of the dangers posed by AI and Big Data that we have established can be grounded on reality through examination of real-life scenarios in which ethical concerns such as the ones discussed in the prior chapter were made manifest. In this chapter, we will examine examples such as election interference and the reinforcement and reproduction of biases by AI systems used commercially and by public authorities. Finally, we will give special attention to the concept of filter bubbles: how AI negatively impacts the Internet by transforming it from an impartial medium of information into a collection of echo chambers, and place that analysis in the context of a democratic public sphere.

5.1. Electoral Manipulation

The year 2016 provided two cases in which the use of AI played a prominent role in electoral politics: the United Kingdom's Brexit referendum and that year's United States presidential election. In these elections, data analytics firm Cambridge Analytica used techniques we have described earlier in the context of social media while working with the Leave campaign in the first case, and Donald Trump's presidential bid in the second.

This happened through another firm, Global Science Research (GSR), who in 2014 used survey tools to gain access to users' Facebook profiles. Facebook's then-lax API restrictions also allowed access to each user's friends' data, and GSR was able to build a massive data set that could be used to correlate interactions with Facebook (such as page likes) to behavioural and personality traits (Isaak & Hanna, 2018). This research could then be utilised to "microtarget" political advertisements with content tailored to elicit a specific emotional response in the viewer.

GSR's massive data collection, which is described as a "breach" by Cadwalladr & Graham-Harrison in their exposé on Cambridge Analytica on *The Guardian* (2018), was collected through, according to Facebook, legitimate means, but its sharing to third parties constituted a violation of terms. At the same time, both Facebook's UK spokesman and Cambridge Analytica's CEO denied that the data in question was Facebook user data. Regardless of the technical legality of the methodology used, however, that data analysis of user data was leveraged to influence election results remains undisputed.¹⁴

Another case of microtargeting social media political adverts is the 2017 UK general election, though in a more peculiar way, involving internal Labour party politics. At that time, Jeremy Corbyn was the leader of the Labour party and represented a left-wing agenda that alienated many of the more centrist members of the party and the party bureaucracy (the so-called "Blairite wing", due to their association with Tony Blair's rebranding of the Labour party and its shift towards the political centre during his leadership). The dispute ran deep enough that, in the 2017 election, Labour campaign officials were willing to lie to Corbyn about the party's campaigning: instead of running the advertisements they were directed to, they microtargeted the party leadership with them so as to create the illusion that they were doing so, when in fact few others saw them. (Bubsy, 2018)

Electoral manipulation on social media isn't limited to microtargeting of political advertisements. The fundamental anonymity afforded by the Internet enables its use as a medium for propaganda disguised as authentic expression of speech. The best-documented case of this is the 2016 US Presidential election, in which fictitious social media accounts operated in support of Donald Trump's ultimately successful campaign. Of particular interest is Special Counsel Robert Mueller's 2019 *Report On The Investigation Into Russian Interference In The 2016 Presidential Election*¹⁵. The investigation discovered Russian social media operations coordinated largely by the St. Petersburg-based Internet Research Agency (IRA) aimed at influencing the election (Mueller, 2019).

This electoral interference was a long-term project: the IRA began building its American-themed social media infrastructure as early as 2014, and was setting up public groups and pages on Facebook and Twitter related to US politics by 2015. These groups obtained wide reach, the largest of the Facebook groups ("United Muslims of America") surpassing 300,000 followers before being shut down in 2017 (ibid, p.26).

¹⁴ Commenting on the absence of a legal framework regulating online advertising on occasion of the Cambridge Analytica scandal, US Senator Mark R. Warner stated that "the online political advertising market is essentially the Wild West".

¹⁵ It should be noted that the publicly-available Mueller report is heavily redacted and that the observations made on it were made with no regard to its unavailable content.

On Twitter, IRA personnel-run fake accounts were successful at presenting as American citizens: their tweets were promoted by Trump campaign staff and affiliated persons, and their content frequently reached more traditional media who reported on them as genuine partisan opinions. IRA also employed bot networks - automated Twitter accounts that operate together to cross-promote and build an audience, again under the pretence of being real persons. In 2018 Twitter was able to identify 3,814 IRA-linked accounts, and a total of 50,258 Russia-linked bot accounts (Twitter, 2018; Mueller, 2019).

Due to the nature of such a covert operation, it is hard to say whether those findings represent ultimate reality. Supposing that some botnets were less conspicuous than others, we could assume that some of them avoided detection by Twitter's attempts to clean them out. Regardless, claims to the use of AI in such a scale have not been made in later election cycles, and it would be reasonable to assume that the awareness raised by the events discussed and subsequent action has successfully mitigated the problem to a considerable extent.

Nevertheless, the risk of the same tactics being used should not be discounted, especially in cases that receive less publicity. Caution should be raised by cases such as that of West Papua, where as recently as November 2020 botnets have been employed to influence and distort public perception of the debate surrounding the West Papuan independence movement by making online posts in support of Indonesian sovereignty over West Papua. These posts, made in English as well as German and Dutch, seemingly with the intent of winning over foreign influence, generally spoke in favour of special autonomy within Indonesia and against the independence movement. (Strick, 2020)

5.2. Commercial misuse

Passive data collection for commercial purposes, even assuming benign intent, poses an ongoing moral concern. Green (2017) mentions a story in which a teenager's pregnancy was outed to her father through targeted advertisements by a retail company's predictive algorithm, which, correctly determining her pregnancy, started serving the family maternity advertisements. In addition to this violation of personal, intimate privacy, devices that are used for data collection are used in all environments by, or near, people in positions of responsibility who deal with more sensitive information. A smartphone used by a doctor might inadvertently violate their patients' medical privacy, for instance. It is not difficult to imagine ways in which this breach of privacy could be damaging or ways in which bad faith actors could take advantage of it if they gain access to it.

Aside from privacy concerns about data collection, and accidental or hypothetically malicious use of that information, the standard deployment of such predictive AI also needs to be addressed critically. Studying Facebook's job advertisement delivery, Ali et al. (2019) discovered its mechanism to produce biased results; for instance, supermarket cashier job adverts were overwhelmingly marketed to women, and janitor and taxi driver adverts broke along racial lines to primarily target non-white users. Similar algorithmic bias is often claimed against hiring algorithms used by employers. While there are techniques used to reduce bias in these systems in order to satisfy legal requirements, the standards set by them cannot be treated as absolute, and the opacity of the systems makes it impossible to critically examine them in detail. (Raghavan et al., 2020)

5.3. Misuse by Authorities

The introduction of biases in machine-learning algorithms becomes an even greater concern when those are being used by the state security apparatus. Use by the police can and does lead to false arrests, as in the January 2020 case of Robert Julian-Borchak Williams, a black Detroit man, who was falsely arrested on the basis of identification by facial recognition software in a theft case (Hill, 2020). A twofold problem arises: whether the technology in itself contains biases and whether it is used appropriately. With regards to the former, a study by the US Department of Commerce (Grother et al., 2019) found consistently poorer results for people of American Indian, African, and East Asian descent compared to other backgrounds (sometimes up to a hundred times worse), and to a lesser degree poorer results for women compared to men. Facial recognition technology was also employed by police in England and Wales where courts found it to violate human rights and data protection laws (though the ruling was specifically in regards to a lacking legal framework in its use by South Wales Police, rather than the technology itself fundamentally) (Croft & Venkataramakrishnan, 2020).

As artificial intelligence then becomes employed by the state security apparatus, questions of accountability naturally re-emerge, as a wrong decision by an AI can lead to violations of civil rights. In the case of Julian-Borchak Williams mentioned above it was ultimately judged that it was the investigators' choice to trust the algorithm's results without further investigation that was at fault. But other cases might not be as clear cut, especially as advances in robotics enable the deployment of automata to physically intervene. In 2017, online

magazine Salon described San Francisco as a dystopia when a non-profit hired a security robot to patrol its parking space, with the (suspected as intentional) result of driving away the homeless people that tried to make use of the parking as a living space - the municipal government had to step in and have them stop using the robot (Lyons, 2017). Robotic technologies such as drones also find military applications.

Even as legal responsibility might eventually land somewhere, such as whoever authorised the robots' use, we start encountering greater degrees of separation between the human actor and the act. The case of the self-driving car highlights some finer details of the accountability question. Autonomous vehicles might reduce accidents overall, but who is responsible when an accident does occur? If the vehicle is faced with a trolley problem-style situation, what should its course of action be? Bonnefon et al. (2016) raise the question of there being potentially different competing algorithms on the market: if a buyer chooses a car that will prioritise the safety of the driver, are they the ones liable if, to protect them, the car deprioritised the lives of others in an accident? Other algorithms might not even be marketable: who would buy a car knowing that in case of an accident, it wouldn't prioritise their safety?

These kinds of questions were not quite applicable in the previous examples because those still involved much greater degrees of human oversight. But the point of AI is to reduce the need for human involvement in a given process and automate it, and the moral issues that arise in the case of autonomous vehicles are ones that derive from AI removing the need for human oversight. To the extent that advances in artificial intelligence also result in a reduction or elimination of human oversight in other fields, similar complications will need to be addressed. Just as self-driving cars might reduce accidents but complicate the question of accountability, hiring algorithms could well show less bias than a human - bias mitigation mechanisms form a frequent selling point of hiring algorithms to potential customers (Raghavan et al, 2020) - leaving us with a dilemma between a potentially racist human who can be held accountable, and a potentially racist, but less so, machine, that cannot be held accountable.

5.4. Filter bubbles and Information Control

Rose : Everyone withdraws into their own small gated community, afraid of a larger forum. They stay inside their little ponds, leaking whatever "truth" suits them into the growing cesspool of society at large.

Colonel : The different cardinal truths neither clash nor mesh. No one is invalidated, but nobody is right.

(Metal Gear Solid 2: Sons of Liberty, 2001)

The case of the 2017 UK general election highlights a more general danger inherent in the personalisation of the Internet. This danger is the creation of online echo chambers that might limit a user's exposure to information, such as opposing viewpoints. The term most widely used to describe this is "filter bubbles", coined by Pariser (2011), in reference to search engine results and social media feeds being algorithmically personalised to show the user content they might be interested in - at the expense of other content.

This is an inadvertent effect of the use of AI to serve more relevant content to the user. The "algorithm" builds a profile for the user based on the information available on them and produces relevant recommendations. Let's use YouTube as an example. Your YouTube search history, watch history, and so on form a basis around which your profile gets constructed. YouTube being owned by Google, your profile might already exist before you ever even use YouTube - or indeed, any Google services, if you have interacted with any of Google's partners that share information with them. The result is a recommendation feed custom-made for you, with content that the algorithm believes to already be of interest to you. This fundamentally means that information is presented to you in a biased manner, one that one-sidedly affirms your existing interests or which encourages a selective branching-out, never asking for consent.

A breakdown of the ways in which filter bubbles fundamentally undermine democratic ideals is given by Bozdag & van der Hoven (2015). In their work, they describe multiple viewpoints - and priorities - of abstracted democratic value systems and the manner in which filter bubbles antagonise them. These are:

- The “liberal”(p.250) viewpoint, in which personal freedom and individual autonomy are the fundamental values. Filter bubbles endanger this autonomy by artificially reducing choice available to individuals. The lack of agency towards, or even knowledge of, the filter, is a major violation thereof.
- The “deliberative democracy”(p.251) viewpoint, in which a public deliberative process on contentious issues is seen as important for a society to address opposing viewpoints and reach a consensus. Filter bubbles present this kind of deliberation with an obstacle, by making these opposing viewpoints less accessible.
- The “republican” viewpoint, or “contestatory democracy”(p.252), which emphasises the ability of people to contest authority. For that to be effective, access to information must be free and effective - filter bubbles, being a de facto form of information control, are in direct contradiction of that.
- The “agonist”(p.253) viewpoint, similarly to the deliberative one, demands an open discussion of contentious issues. Unlike the deliberative model, however, the goal is not to achieve consensus, but for an inclusive discussion to take place, for people to be able to air their disagreements and oppose consensus. Filter bubbles then pose a problem not only because they stifle equal access to information, but also because they promote ignorance of minority viewpoints, which from the agonistic perspective ought to be elevated and given access to the public forum that they might have been denied due to their status otherwise, in order to challenge majority opinion.

These conceptions all might seek different solutions. For instance, a mechanism that would eliminate bias in terms of content delivery might satisfy the demands arising from the first two perspectives, but fail to deliver on the contestatory demand to provide a means of challenging consensus or the agonist demand for minority views to be brought into the wider public discourse. The source of the problem is, however, shared, and it poses a challenge to the conception of the Internet as a true public sphere.

Control of the flow of information also comes in the form of platform moderation. Terms of service on a platform are, by necessity, more restrictive than laws on free and protected speech, as they apply in addition to them. Can a platform that restricts speech function as a true public forum? If the Internet is, instead of a public forum, a mere collection of private platforms, each with its own rules and terms of service, their own opaque methods of content delivery, all dependent on the rules of the capitalist market for viability, the result is a controlled medium and a controlled public discourse.¹⁶

But at the same time, mere assurances of freedom of speech are clearly not enough. We noted in the case of the 2016 US Presidential election a rampant phenomenon of inauthentic activity being used to manipulate the users of social media platforms. In fact, the use of artificial intelligence to create seemingly authentic, yet ultimately fake information, is not only limited to social media accounts. Recently techniques have developed that are able to create facial constructions that - at a glance - appear real, and can even create similarly convincing video. Given a large enough vocal sample, and an audio track can be created as well. This “deepfake” technology, while not perfect, is dangerously close to being able to create facsimiles of famous, or fake-but-real-looking, people, putting out any potential message through them. (Metz, 2019)

How do scams like that mesh with online freedom of speech? The 2016 electoral manipulation was helped by the ease involved in creating and maintaining accounts. The tools required for platform owners to combat this inauthentic behaviour by necessity empower platform owners to control information on their platform. This is an environment in which the platform owner and sections of the user base both vie for control of information. In this struggle neither side truly has an incentive to move in good faith.

5.5. Artificial Intelligence and the Democratic Society

The ultimate realisation we reach is that Artificial Intelligence can and does serve as a tool for infringing upon civil rights and interfering with democratic society. From the more flagrant examples of conscious attempts to manipulate public and electoral discourse to the less grandiose ones of AI making - truthful or false - determinations about individuals without their consent to the process, ethical concerns about AI turn from an academic subject into a factual reality. Even something as passive and innocent-sounding as serving better search results forms a credible threat to the well-functioning of democratic society in the context of filter bubbles. The question is not whether the lack of privacy realised by modern technology should be alarming, but how to mitigate the damage already being caused by it.

¹⁶ Similar criticisms could, of course, be levied against traditional media. Recognising the shared limitations of modern and traditional media to perform an idealised role is not an endorsement of the model being followed. See also the earlier discussion on the concept of the “fourth estate” and Hampton’s (2010) detailed analysis.

On an individual level there are obvious concerns with regards to personal autonomy, the ability to make decisions for one's self. Take the example of the teenage girl whose pregnancy was revealed to her father by predictive marketing algorithms. As Green (2017) rightly points out, the results of something like that in an abusive, or even simply less-supportive family, could have been disastrous. Individual privacy is necessary to be able to exercise personal self-determination and exist as equals, and the implied freedom from coercion is a cornerstone of liberal democratic values. A society lacking this spirit of egalitarianism should find it challenging to convince its members of its democratic credentials. Poor use of AI leads precisely to that; when an AI system is used to pass an automated judgement on someone without further due diligence, leading to unfair treatment on the basis of ethnicity, gender, or other such characteristics, social hierarchies and stratifications are reinforced and the inclusive, egalitarian character of liberal democracy becomes questionable.

Questions of trust in democratic values and institutions also arise on a systemic level in two forms. First, AI manifests challenges to democracy directly through hostile action, such as Cambridge Analytica's attempts to manipulate electoral politics using subterfuge and duplicity. Second is the unintended consequences of artificial intelligence's use on a mass scale in ways that have social implications - filter bubbles' effect on the development of public discourse is one example of this, but the normalisation of the breaches of individual rights mentioned in the previous paragraph can also have knock-on effects on a social scale. Biases reinforce themselves by creating a feedback loop, and algorithms that replicate them can help perpetuate them, as well as give them a veneer of cold, machine objectivity. There is furthermore the long-term danger of instilling more general apathy towards those rights, as people who grow up or live long enough in a world in which those rights are de facto lost become accustomed to it, with lack of privacy or discrimination becoming expected norms.

As we are preparing to discuss state policy towards AI and the Internet we need to give due consideration to doubts naturally arising as to the ability to effectively restrict those breaches of democratic values and norms. This is essentially a question of *technological determinism*, defined in Wikipedia as "a reductionist theory that assumes that a society's technology determines the development of its social structure and cultural values". Given AI's ability to rapidly process and cross-reference vast amounts of data, readily and often voluntarily available on the Internet, is it possible to re-assert the right to privacy and avoid the use of AI in ways that compromise the democratic character of society? In other words, is AI destined to act in such a way as to undermine democratic values as surely as the development of agriculture transformed hunter-gatherer nomadic tribes into sedentary societies, or the industrial revolution resulted in a drive towards urbanisation?

It is the opinion of the author that treating such a social transformation as inevitable is ill-advised. This is for two reasons: on a moral level, it can serve to free us, as individuals and as a society, from the responsibility towards the malapplication of technology, in this case artificial intelligence. Alternatively, it could be seen as a luddite response towards artificial intelligence, treating it as an inherently dangerous technology that cannot be used ethically. The former is tantamount to an abdication of free will and our ability to dictate our lives. Adopting that point of view would be antithetical to the premise of the current study: that something can pose a danger to democracy implies a conflict, an ongoing situation in which democracy is at stake. To then posit that artificial intelligence inevitably means the end of democracy is the same as proclaiming democracy already dead. The latter is a potential outcome of the subject matter analysis but no more than that. It has the potential to be the ultimate conclusion, but does not offer itself any insights, nor is it the only possibility. Challenging it, for instance, is the European Union's General Data Protection Regulation which, as we will see in its analysis in the following chapter, attempts to provide a technologically-neutral legal framework that guarantees data privacy rights.

The second reason is that there is a leap of logic involved in treating the social transformations facilitated by AI as deleterious to democracy and to civic rights. We made brief mention earlier of OGAS, the plan to establish a computer network across the Soviet Union in the 1960s for the purpose of perfecting the Soviet planned economy. ARPANET, the precursor to the Internet, was an American military programme. In neither of those original designs is there a glimpse of the highly commercialised vortex of personal data that is the Internet in its present form. The relatively unregulated development of the commercial Internet did not happen absent of human decision-making. While technology such as the Internet is transformative, it is our own intent and use that guides and shapes that transformation. The assumption that it must by necessity come at the expense of democracy then feels unwarranted.

5.6. Threat Assessment

Having gone through the ethical questions surrounding Artificial Intelligence and the way we see them manifest at present, we can attempt to take stock of commonalities and differences between them – in short, do

categorisation. Listing those threats posed by AI comes with certain challenges: for instance, “electoral manipulation”, while describing a straightforward concept, is something that we’ve seen in the examples studied making use of data commodification and botnets – concepts that themselves could qualify as threats. Similarly, the commodification of personal data is a prerequisite for, but not the same thing as the unregulated and unscrupulous selling of that data, such as the sale of customer location data to military clients. Even so, the use of AI for government surveillance specifically can be seen as a separate threat, as it does not always need to be the result of such a commercial exchange. It is therefore recognised that the threats identified here do not exist in a vacuum, but rather overlap, sometimes significantly so, or feed into each other.

For this analysis we have chosen the following “threats” to break down:

- *Data commodification*, referring to the general trend of treating user data as a service for advertisers or other customers
- The *Unregulated data market*, referring to the anarchic nature of data commodification in terms of regulation¹⁷
- *Botnets & Deepfakes* as tools that try to trick human users into thinking they are legitimate users through mimicking the appearance of a normal human user profile
- *Direct manipulation of public discourse*, referring to attempts to utilise AI to influence voter behaviour or alter public perception around something¹⁸
- *Government surveillance* as enabled through AI systems
- *Citizen scoring systems*, seen as an invasive system that also acts as an enabler of hierarchies and discrimination based on them
- *Filter bubbles*, seen as a threat to open public discourse in the manners outlined earlier
- *Algorithmic bias*, which may be the subconscious or unintended result of an AI system

The categories chosen for this are as follows:

- *Intent*, describing the primary motivation behind the threat – the term “Malicious” is here used for those that directly seek to accomplish their goals through the undermining of democratic principles
- *Threatened democratic values*, describing the particular facet of liberal democracy that being threatened
- *Threat source*, distinguishing whether the threat posed is a result of the involved technologies in and of themselves or of external factors (such as state policy or economic motivators)
- *Territorial scope*, distinguishing between “Local”, “National” and “International”
- *Social scope*, describing whether the threat is targeting individuals or broader society – this category is perhaps problematic, as individual-scoped systems (such as filter bubbles) will always achieve society-scale effects when scaled up to wide social groups
- *Opacity*, describing not the degree but moreso the intent, whether opacity is by design or incidental
- *Timeframe of concern*, a dual category describing whether the threat in question is of immediate concern, happening at present, and whether it has a high developmental potential as technology improves (for instance, deepfakes and botnets already exist and can trick unaware observers, but the potential for improvements that make them harder to detect is notable)

The table thus formed is as follows:

	Intent	Threatened Democratic Values	Threat source	Territorial scope	Social scope	Opacity	Timeframe of concern
Data commodification	Profit-seeking	Right to privacy	External factors	International	Individual, Society	Incidental	Happening at present
Unregulated data market	Profit-seeking	Right to privacy	External factors	National, International	Society	Incidental	Happening at present
Botnets & Deepfakes	Malicious	Open & democratic public discourse	Inherent to the technology	Case-by-case	Society	Opaque by design	Happening at present, high future potential

¹⁷ As will be seen later in the study, legislation that aims to address this such as the GDPR exists, but has not been without trouble in its implementation.

¹⁸ Commonly referred to as “psyops”

Artificial Intelligence: Dangers to Privacy and Democracy

Direct manipulation of public discourse	Malicious	Open & democratic public discourse	External factors	Case-by-case	Society	Opaque by design	Happening at present
Government surveillance	Malicious	Right to privacy, Other civil liberties	External factors	National, International	Individual, Society	Opaque by design	Happening at present
Citizen scoring systems	Malicious	Right to privacy, Freedom from discrimination	External factors	National	Society	N/A (case-by-case?)	High future potential
Filter bubbles	Benign	Open & democratic public discourse	Both	International	Individual	Incidental	Happening at present
Algorithmic bias	Case-by-case	Freedom from discrimination	Both	Local, National	Individual	Case-by-case	Happening at present

Table 2. A breakdown of various characteristics identifiable in threats posed by AI systems against democratic values.

6. Relevant Regulatory Approaches

This chapter will examine legislation and official recommendations that are applicable or relevant to the study subject. Though relevant legislation that established rules for electronic commerce such as liability for Internet content is not new, Papadopoulos & Charalabidis (2020) point out that attempts to legislate a regulatory framework for AI is a more recent phenomenon, with countries catching up to decades of discussion and responding to more recent advances, noting that “[i]n a brief period of three years, between 2017 and 2019, over 30 countries have developed national AI strategies or action plans” (p. 110). Of this wealth of legislation and related official documents, we will specifically study:

- The European Union’s General Data Protection Regulation, which is a landmark piece of legislation on privacy and personal data rights
- The work of the High-Level Expert Group on Artificial Intelligence, which was formed by the European Commission to provide policy recommendations on matters of AI; specifically we will study the first document the group delivered, *Ethical Guidelines for Trustworthy AI*
- Section 230 of US telecommunications law, which has in large part shaped the Internet in regards to responsibilities of platform operators and users, and thus forms a fundamental aspect of the legal background of the study’s subject matter

6.1. Data Privacy in the European Union

The most comprehensive legislation to date regulating artificial intelligence, albeit indirectly, is the European Union’s General Data Protection Regulation, drafted in 2016 and effective from 2018. As per its name, the GDPR is an attempt to regulate the handling of personal data, particularly in the context of information and communication technologies. This relates to AI because, as we’ve seen, many AI techniques that we’ve discussed depend on the mass collection of data. Therefore, data ethics and AI ethics are intimately intertwined and largely overlapping. We shall take a close look at GDPR, the rights granted by it, the demands placed upon service providers, and the means of enforcement it establishes. Ultimately this is intended as a critical analysis, so we will question the regulation’s efficacy and enforceability where appropriate, as well as any loopholes provided by it.

Before we proceed with that analysis, we should familiarise ourselves with the Regulation’s terminology. Article 4 of the GDPR provides a list of definitions, of which we will note some for purposes of using terminology consistent with the regulation:

- A **data subject** is a natural person, identified or identifiable by their **personal data**. The definition explicitly makes allowance for that identification to be indirect, such as through location data.¹⁹
- **Processor** refers to those who process personal data on behalf of a **controller**, which refers to those that determine “the purposes and means of the processing of personal data” (Article 4(7)).
- The term **pseudonymisation** is used for the processing of personal data that is done such that the data subject cannot be identified.
- **Profiling** in the Regulation’s context is given the definition of automated processing for the purpose of making an evaluation about a natural person.
- **Consent** with regards to processing is defined as unambiguous and informed, and resulting from a clear affirmation by the data subject. Article 7 goes into more detail about consent specifically: the data subject’s consent has to be demonstrable by the controller, and the data subject has the right to withdraw consent at any time. Additionally, paragraph 2 demands that consent in written form be “presented in a manner which is clearly distinguishable from the other matters, in an intelligible and easily accessible form, using clear and plain language”.²⁰

The material and territorial scopes of the GDPR are defined in articles 2 and 3. Territorially, the Regulation applies to controllers and processors that are conducting their activities through an establishment in the EU (even if the processing takes place outside of it), and to those outside of the EU where their processing involves data subjects within it. Article 2 provides a few exemptions, such as activities of a personal nature or those by public authorities, as the Regulation primarily aims to regulate commercial activity.

Rights granted by the GDPR

¹⁹ Though numerous examples of such identifying data are given, the definition ultimately leaves itself open to interpretation. This is purposefully done, among other things, to ensure the legislation remains future-proof and technology-proof (Mitrou, 2018).

²⁰ Article 8 provides some additional protections for children up to 16 years old.

Chapter III of the GDPR grants the data subject six fundamental rights: right of access (Article 15), right to rectification (Article 16), right to erasure (also ‘right to be forgotten’) (Article 17), right to restriction of processing (Article 18), right to data portability (Article 20), and the right to object (Article 21).

The right of access allows the data subject to obtain from the controller a copy of their personal data, confirmation as to whether they are being processed, and information relevant to that processing, including information regarding any profiling that might be related to the processing. The right to rectification allows the data subject to have their personal data corrected by the relevant controller if inaccurate or complete them if incomplete. The right to be forgotten grants the data subject the right to erasure of their personal data, and the right to restriction of processing, along with the right to object, the right to exempt their personal data from processing. The right to data portability states that controllers should, on request by the data subject, be able to transfer the data subject’s personal data “in a structured, commonly used and machine-readable format”, and if possible do so directly to another controller if the data subject wishes.²¹ Finally, under Chapter VIII (“Remedies, liability and penalties”), the GDPR grants in Articles 77 and 78 to the data subject the right to lodge a complaint to a supervisory authority if their personal data isn’t being handled per regulation, and, being that the supervisory authority’s role is one defined by the GDPR as an enforcement mechanism, the right to appeal to the justice system and contest the supervisory authority’s decisions where it concerns them in court.

At this point we should make a note on the language used to refer to the relationship between the data subject and personal data. Though in this study we refer to “the data subject’s personal data” and use possessive forms such as “their”, the Regulation avoids this in favour of “personal data concerning him or her”. That is perhaps more accurate to the extent that data and information is not something that can exactly be “owned”. However, it is the opinion of the author that the rights conferred by articles 15, 16, 17, 18, and 21 can be said to define an abstract form of “ownership” of the personal data concerning a person by that person.

For the purposes of respecting those rights, articles 12, 13 and 14 require the controller to make certain information known to the data subject. Article 12 establishes transparency requirements for providing the data subject with information related to the rights granted by Chapter 3. Article 13 and 14 require the controller to notify the data subject when their personal data has been collected, with article 14 making special note of the data having not been obtained by the data subject themselves.

Articles 9 and 10 elaborate on certain categories of data and data subjects which are placed under special protection. Article 10 restricts the control of processing of personal data relating to criminal convictions and offences to official authorities²² and asserts the illegality of private registers of criminal convictions. Article 9 lists the following categories of data for which processing is forbidden unless consent is expressly given, or if it has to be carried out for legal, health, or scientific purposes:

- Racial or ethnic origin
- Political opinions
- Religious or philosophical beliefs
- Trade union membership
- Genetic data
- Biometric data for the purpose of uniquely identifying a natural person
- Data concerning health
- Data concerning a natural person’s sex life or sexual orientation

It is perhaps curious that gender and gender identity are not explicitly covered by the wording of the legislation. It is possible that the legislators did not consider the social differences between men and women to be large enough to warrant this special protection, and did not take into account individuals whose gender identity places them outside of social norms. Alternatively, the protection of genetic and biometric data and data regarding sex life and sexual orientation could cover those categories, if interpreted broadly.

Demands of the GDPR on the controller and processor

Chapter 4 of the GDPR describes technical and policy responsibilities of the controller and processor for the purpose of respecting the Regulation’s protection of data privacy.

²¹ Allowance is made for restrictions of those rights in specific circumstances, such as to avoid obstructing legal processes or purposes of public interest, historical research and so on.

²² Making allowance for member state laws to permit otherwise so long as data protections are in place.

Article 25, “Data protection by design and by default”, calls for the controller to implement data protection measures (like pseudonymisation) as early as the design phase of processing, and to ensure that only the necessary data gets processed at every step of the processing. Article 28 describes a number of requirements for a processor to be authorised to process personal data per the regulation, such as only being allowed to subcontract the processing with explicit permission from the controller, ensuring a paper trail of processing requests, or cooperating with the controller to ensure compliance with the Regulation. Recordkeeping responsibilities of both controller and processor are expanded upon in article 30.

Article 32 establishes controller and processor responsibilities with regards to data security in processing, demanding that they ensure data undergoing processing is pseudonymised and encrypted, that they are able to demonstrate the well-functioning of the processing systems, able to restore data access in case something goes wrong, and that these systems are regularly re-evaluated to ensure data security. Articles 33 and 34 govern the notifications the controller has to make in case of a personal data breach, the former article regarding notifications to the supervisory authority (unless the breach is deemed unlikely to result in “risk to the rights and freedoms of natural persons”) and the latter to the data subject themselves (if the breach is “likely to result in a high risk to the rights and freedoms of natural persons”).²³

Article 35 demands that the controller evaluate whether a processing operation is likely to result in such a high risk and carry out a data protection impact assessment prior to processing in these cases. The article makes reference to some specific circumstances in which that impact assessment is required²⁴, and otherwise leaves it up to members states’ supervisory authorities to establish lists of other types of processing operations that require (or don’t require) an impact assessment. If such an assessment determines that such processing would result in a high risk, article 36 mandates that the controller consult the supervisory authority prior to processing, who, should they find the controller’s efforts to offset the risks insufficient, are to advise the controller and processor, and are empowered to formally investigate the situation.

Articles 37 through 39 describe the role of the data protection officer (DPO): the controller and processor must designate someone as a DPO where processing happens on a regular basis or where it involves the protected characteristics described in articles 9 and 10 (public authorities acting as processors must always have a designated DPO). The DPO is charged with the task of ensuring that data processing is compliant with the law and relevant responsibilities, and empowered to carry out that task without interference from the controller, who is forbidden from instructing the DPO on how to do so.

Articles 40 through 43 govern “codes of conduct”, which may be drawn by associations of controllers and processors in order to ensure the operations of these controllers and processors are compliant with the Regulation. Article 41 establishes criteria for a body to be accredited by a supervisory authority to monitor compliance with codes of conduct. Articles 42 and 43 provide for the establishment of further certification bodies that controllers and processors may voluntarily apply to for certification in their compliance with the Regulation.

Chapter 5 of the GDPR describes the regulatory framework of the GDPR concerning transfers of EU personal data to third countries or international organisations. Article 45 empowers the European Commission to, taking into account a set of criteria relating to continuous data protection, declare countries or international organisations safe such that data transfers to them do not require “specific authorisation”. Otherwise, the controller or processor has to prove that the transfer is happening in a manner compliant with the Regulation (described in articles 46, 47 and 49).

GDPR enforcement mechanisms

Chapter 6 of the GDPR establishes the framework governing the supervisory authorities that we’ve occasionally mentioned up to this point. Article 51 mandates that each member state establish one or more such authorities which are tasked with “monitoring the application of this Regulation”. Articles 52 through 54 establish the member state’s obligations to the supervisory authorities, establishing their independent status and transparency of appointment to them (either by parliament, the government, the head of state, or an independent body charged with the task, per each member state’s choosing).

Article 57 goes into detail about the tasks of supervisory authorities. Beyond enforcing the GDPR, they are to, among others, act in an educational capacity to the public and to controllers and processors, as advisors to other

²³The distinction between “risk” and “high risk” is unclear.

²⁴ Notably, this also includes “a systematic monitoring of a publicly accessible area on a large scale”.

state institutions, and to fulfil the obligations assigned to them earlier in the Regulation (such as drawing up the lists relating to operations that require data protection impact assessments).

Article 58 describes the powers conferred to supervisory authorities by the Regulation. Notably, supervisory authorities are given the power to carry out audits with regards to data protection and to gain access to the physical premises and equipment of controllers and processors as part of their investigative duties. They are further empowered to take corrective measures, such as (among others) ordering a ban on processing or imposing fines.

The member states come together on data protection through the European Data Protection Board (established in article 68), which consists primarily of the heads or representatives of one supervisory authority from each member state. The Board's tasks, outlined in article 70, broadly consist of monitoring the GDPR's application, particularly on cross-member state cases where disputes between supervisory authorities might arise (the duties of the Board in these cases detailed in articles 64 and 65), issuing guidelines and recommendations, and generally act as a review mechanism for the Regulation.

Finally, in adopting the Regulation and bringing it into effect, the member states are asked through article 85 to account for freedom of speech and provide exemptions where appropriate.

Exemptions granted by the GDPR & limitations

It is difficult for a piece of legislation to have such a wide scope as to cover all the ways in which its spirit might be violated. Mitrou (2018) notes the challenge for legislation to keep up with technological progress, though she concludes that, by virtue of the legislation being technology-neutral (a conscious choice by the legislators to prevent circumvention), and by requiring data protection by design, the GDPR will “contour the way AI and machine learning will be developed and applied” (p. 74).

Indeed, the relevant language of the GDPR does appear to cover for technological advances by demanding protections for privacy rather than restrictions in technology. When specific techniques are mentioned, such as pseudonymisation, they are provided as examples of the requirements for ensuring data privacy, and it so follows that their listing should not be treated as exhaustive. We should not, however, forget that the GDPR's scope is also limited by design: it is a regulation that aims to primarily regulate commercial activity. Furthermore, it is also naturally restricted geographically to such data processing as concerns the European Union in some way.²⁵

Article 2 makes an exemption from the Regulation's obligations that seems especially important in this context, stating that the Regulation does not apply to processing “by competent authorities for the purposes of the prevention, investigation, detection or prosecution of criminal offences or the execution of criminal penalties, including the safeguarding against and the prevention of threats to public security”. We have already examined examples of police and national security agencies engaging in activities that violate data privacy and which use it for profiling (as defined by the GDPR). In advocating for data protection and privacy, the GDPR's refusal to regulate data processing activities by such authorities is a glaring omission.

Article 23 re-affirms this omission by granting the right to other legislative initiatives, including by individual member states, to restrict the Chapter 3 rights conferred to the data subjects for purposes of, among others, “national security”, “defence”, and “public security”. While the language of the article does make mention of respecting “fundamental rights and freedoms”, that can also serve as an acknowledgement of the opposing tendencies at work between democratic liberties and the systemic needs of the state security apparatus. The legislators' resolution to this contradiction appears to ultimately be on the side of the latter.

A fundamental limitation of localised legislation is also its limited territorial scope. The GDPR naturally can only apply in circumstances where the EU is involved in some way; be it that the controller and processor are engaging in their activities through shops in the Union, or that their activities implicate EU citizens, as we've mentioned in the article 3 analysis. To the extent that the GDPR's authors intended to influence the development of AI and relevant technologies, this is an unavoidable shortcoming. Technological advances and techniques do not discriminate between countries when they are developed. The hope here might be that, desiring to conduct business in the EU, and not wanting to incur the costs associated with diversifying data analysis techniques on a regional basis, firms simply use the European regulation as a baseline for all their operations.

²⁵ The argument could be made that, by restricting the commercial viability of AI that fails to respect the GDPR, the overall development of such technology becomes hamstrung. Nevertheless, even if the GDPR does create such a domino effect, its reach cannot simply be assumed to be universal.

While that expectation might be grounded in the present reality, considering the European Union's relative wealth and the resulting market influence, it is also not future-proof. Any such implied influence depends on the European Union continuing to represent a market which firms choose to invest in following cost-benefit analyses that incorporate the GDPR. Additionally, while the EU remains a desirable market, for the GDPR to be effective it needs to be actively enforced.

The enforceability, then, of the GDPR, also needs to be seen in a critical light. As we've seen, the legislation mandates from the member states the establishment of supervisory authorities tasked with this. While the Regulation spends great length to establish rules for joint operations and co-operation between different supervisory authorities, and the Board's existence can provide oversight, the lack of a strong central authority is noted. The degree and fervour with which these authorities ultimately pursue their duties is also something that should not be taken for granted.

Evaluating the GDPR

The Regulation does include review processes to monitor the regulation's implementation: article 59 requires supervisory authorities to publish annual reports on their activities, article 71 mandates that the European Data Protection Board publish annual reports on data protection, and article 97 that the European Commission to draw up a report every four years on the progress of the GDPR. So far the examination of the GDPR's effects we've made has been based on speculation. These reports are a valuable resource for grounding this critical analysis on reality and precedent.

The reports of the European Data Protection Board on 2018 and 2019 list some notable cases pursued by member state supervisory authorities, giving insight into the extent to which the supervisory authorities pursue their duties. The largest of these is the fining of Google for 50 million Euros by the French supervisory authority in 2019, citing "a violation of the obligations of transparency and information" and "a violation of the obligation to have a legal basis for data processing for ad personalisation" (pages 32 and 33 of the EDPB's 2019 Annual Report).

While that is one case among many, it stands out as the only one against a major tech firm. The violations cited against Google can be seen as fairly standard practice across most online services²⁶. It thus appears that the ability of the supervisory authorities to pursue their responsibilities to their full extent is facing challenges - indeed, according to the same report, a survey of supervisory authorities found them in want for more resources (p. 37).

Target of fine	Supervisory authority	Amount fined	Reason
Google	French SA	€50 million	Failure to provide transparency and unlawful processing ²⁷
H&M	Hamburg SA	€35 million	Unlawful recording of employee return-to-work meetings
TIM	Italian SA	€27.8 million	Multiple infractions, part of an aggressive marketing campaign
British Airways	British SA	€22 million	Insufficient security

²⁶ See for example the rules for compliance with cookie rules set by the Information Commissioner's Office, the British supervisory authority. They establish that websites that pre-determine cookie settings and simply inform the user that by continuing to use their service they offer consent to their use are non-compliant. They additionally state that interfaces that include an option to reject but de-emphasise it in favour of accepting are non-compliant, describing it as "nudge behaviour".

²⁷ This is the same fine as described in the 2019 case. However, it was pushed back into 2020 following an appeal by Google that the courts ultimately dismissed.

			measures at the time of a 2018 data breach
Marriott	British SA	€20.4 million	Breach exposing the data of millions of EU residents as a result of a long-existing vulnerability
Wind	Italian SA	€17 million	Unlawful marketing (not giving customers the option to unsubscribe)
Google	Swedish SA	€7 million	Failure to remove listings subject to right-to-be-forgotten order
Allgemeine Ortskrankenkasse	Baden-Württemberg SA	€1.24 million	Failure to deliver direct advertising only to customers who gave consent

Table 3. Fines larger than 1 million Euros levied by Supervisory Authorities in 2020, up to November 15 (Tessian, 2020)

The 2020 European Commission report, while praising it as exemplary legislation that has inspired similar initiatives, is ultimately of the opinion that it is “premature at this stage to draw definite conclusions regarding the application of the GDPR” (p. 4). The report identifies cross-border co-operation as an area for improvement and notes “inconsistencies between the national guidance and the Board guidelines” (p. 6), as well as inconsistencies in the ways individual member states approach GDPR implementation²⁸. With regards to the GDPR’s application on new and emerging technologies, the report restates that the GDPR is written in a technologically neutral way, but also asserts “strong and effective enforcement of the GDPR vis-à-vis large digital platforms” to be “an essential element for protecting individuals” (p. 10).

The Commission’s report also includes a review on cross-border aspects of the GDPR. With regards to the Commission’s “adequacy decisions” (a decision on whether a third country can be considered “safe” with regards to data transfers, which the GDPR empowers the Commission to make), among others, the EU-Japan mutual adequacy decisions and the progress for similar arrangements with South Korea are touted. Adequacy with regards to the United Kingdom post-Brexit is mentioned as a concern. As regards reviewing these adequacy decisions, the report makes mention of “intense dialogue” (p. 11) with the third countries that decisions have been reached with. Additionally, the report noted a then-outstanding decision by the European Court of Justice (*Data Protection Commissioner v Facebook Ireland Limited and Maximilian Schrems*, or *Schrems II*): that decision has since been made, striking down the EU-US Privacy Shield, which formed the framework for an adequacy decision regarding the United States and rendering it invalid, on account of invasive US government surveillance programmes.

The report also remarks on foreign operators active in the EU market and their representatives in the EU, concluding that “this approach should be pursued more vigorously in order to send a clear message that the lack of an establishment in the EU does not relieve foreign operators of their responsibilities under the GDPR” (p. 12).

In its closing remarks about the future of the GDPR, the Commission report states that, with the regulatory framework in the member states still in the process of being revised, and the time since its adoption as short as it is, it is difficult to “draw definitive conclusions on the existing level of fragmentation” (p.14), the expectation being that as cases move through the judiciary a consistent interpretation will be established. With regards to

²⁸ This is at least partially a design issue. The report provides the case of differing ages of children’s consent with regards to data rights, which the GDPR allows individual member states to set freely between 13 and 16, as a complicating factor.

international co-operation, the Commission states its intent to continue its work in pursuing adequacy dialogues with third countries (furthermore, in responding to the aforementioned *Schrems II* case, Commission Vice President Jourová stated the judgement's value as a guiding tool for future developments).

6.2. European Commission - High-Level Expert Group on AI

In 2018 the European Commission appointed a group of 52 experts as the High-Level Expert Group on Artificial Intelligence to advise it on its AI strategy. This group has since produced four documents (referred to as “deliverables”):

- Ethics guidelines for trustworthy AI (High-Level Expert Group on Artificial Intelligence, 2019b)
- Policy and investment recommendations for trustworthy Artificial Intelligence (High-Level Expert Group on Artificial Intelligence, 2019c)
- Assessment List for Trustworthy Artificial Intelligence (ALTAI) for self-assessment (High-Level Expert Group on Artificial Intelligence, 2020a)
- Sectoral Considerations on Policy and Investment Recommendations for Trustworthy AI (High-Level Expert Group on Artificial Intelligence, 2020b)

Of particular interest to us are the first and third deliverables: *Ethics Guidelines* provides an examination of the abstract concepts of interest to this study, and the *Assessment List* provides a number of questions and techniques for realising them. However, a critical examination of the recommendations provided in the general body of work is also in order, similar to the issues examined with regards to the GDPR earlier.

Ethics Guidelines for Trustworthy AI

In *Ethics guidelines for trustworthy AI*, the AI HLEG establishes three components to “Trustworthy AI”: that it should be lawful, ethical, and robust. The guidelines document is then focused on the second and third components.

As already stated, the AI HLEG's work is part of the European Commission's AI strategy, and it cannot be seen independently from it. Per the AI HLEG itself in the introductory section of the ethics guidelines document the Commission's vision it is meant to support rests on three pillars: increasing public and private investment in AI, preparing for socio-economic changes, and “ensuring an appropriate ethical and legal framework to strengthen European values” (p. 4). They assert the fundamental principle that AI is a means for human needs, rather than an end in itself. It is for this reason that they stress the importance of “trustworthiness”. To pursue “trustworthy AI” in this human-centric context, the AI HLEG asserts that the three components mentioned earlier must be met by AI systems²⁹:

- AI systems must be **lawful**. While the document doesn't strictly concern itself with this component, the AI HLEG considers it “the duty of any natural or legal person to comply with laws” (p. 6) and proceeds with the assumption that the law is being upheld.
- AI systems must be **ethical**. The group recognises that the law might not be in sync with ethical norms, and therefore the first component by itself is insufficient.
- AI systems must be **robust**. Reliability and the avoidance of unintended behaviour is as important as the intent itself.

Ethical Principles

In order to provide an ethical framework to guide the second component, the AI HLEG used the EU Treaties, the EU Charter, and international human rights law. Of the human rights guaranteed by those, they consider especially pertinent to trustworthy AI those that regard five wider concepts: human dignity, personal autonomy, democracy and the rule of law, egalitarianism, and citizens' rights³⁰. Drawing from those, the AI HLEG established four ethical principles, presented in non-hierarchical order, that AI systems must follow to ensure “trustworthiness”:

- The principle of **respect for human autonomy**: “Humans interacting with AI systems must be able to keep full and effective selfdetermination over themselves, and be able to partake in the democratic process. AI systems should not unjustifiably subordinate, coerce, deceive, manipulate, condition or herd

²⁹ A note is made here that these components must refer to the entire system and not just components of it; one cannot separate an AI from the framework and context in which it operates.

³⁰ Referring to rights such as the right to vote or the right to access public documents. The AI HLEG notes here that this is not meant to detract from or exclude rights of non-citizens in EU countries.

humans. Instead, they should be designed to augment, complement and empower human cognitive, social and cultural skills.” (p. 12)

- The principle of **prevention of harm**: AI systems should protect human dignity, and physical and mental wellbeing. The component of robustness is particularly pertinent here. AI systems must be inclusive of vulnerable people. “Particular attention must also be paid to situations where AI systems can cause or exacerbate adverse impacts due to asymmetries of power or information, such as between employers and employees, businesses and consumers or governments and citizens. Preventing harm also entails consideration of the natural environment and all living beings. ” (p. 12)
- The principle of **fairness**: “While we acknowledge that there are many different interpretations of fairness, we believe that fairness has both a substantive and a procedural dimension.” (p. 12) The substantive dimension of fairness means avoiding bias and discrimination, avoiding deceiving people, and balancing means and ends. The procedural dimension means the ability to seek redress against AI systems and being able to trace accountability (similar to the values of contestatory democracy we encountered in the analysis by Bozdag & van der Hoven (2015) earlier).
- The principle of **explicability**: The understandable and transparent operation of AI systems is stressed as vital for trust, and for applying the above principle of fairness - “Without such information, a decision cannot be duly contested.” (p. 13) In the cases where that’s not entirely possible (referred to by the group as “‘black box’ algorithms”) other ways and processes to achieve explicability must be in place.

As with the three components of trustworthy AI, the possibility of the four principles being in contradiction is recognised (predictive policing is provided as an example that pits the *principle of respect for human autonomy* and the *principle of prevention of harm* against each other). The group concludes that these contradictions should be resolved through evidence-based reasoning and only if the benefits substantially outweigh the risks, with a caveat made that some principles and rights such as human dignity cannot ever possibly be seen as conditional.

Key Requirements

Based on these four ethical principles the AI HLEG determined a list of seven requirements that AI systems should meet to comply with those principles. These requirements interplay with everyone involved in the practical application of AI systems, categorised by the group into developers, deployers, and the end-users, as well as the broader society.

These seven requirements, also presented non-hierarchically and interrelated, are:

- Human agency and oversight
- Technical robustness and safety
- Privacy and data governance
- Transparency
- Diversity, non-discrimination and fairness
- Societal and environmental wellbeing
- Accountability

We will examine these requirements in more detail, as well as the technical and non-technical methods the AI HLEG suggests for their implementation.

Human agency and oversight

With regards to *human agency and oversight* the group relates the concept to the principle of human autonomy and fundamental rights, noting:

AI systems should support human autonomy and decision-making, as prescribed by the principle of respect for human autonomy. This requires that AI systems should both act as enablers to a democratic, flourishing and equitable society by supporting the user’s agency and foster fundamental rights, and allow for human oversight. (p. 15)

Cautioning that AI can be both a benefit and a risk to fundamental rights on account of their reach and capabilities, they recommend fundamental rights impact assessments prior to the system’s development if such risks are applicable, as well as opening up to external feedback on such projects.

Concern as regards human agency rests on AI’s ability to manipulate human behaviour through the use of subconscious biases. To respect human agency AI systems should empower their users to make their own independent, informed decisions, and to question the system’s judgement. The group concludes this by singling out the importance of the system not making automated decisions for users that have legal or otherwise significant consequences for them.

To achieve human oversight the AI HLEG suggests the use of oversight mechanisms. Three methods are discussed: human-in-the-loop (HITL), which involves constant human participation, human-on-the-loop (HOTL), which involves human participation in the design phase and monitoring afterwards, and human-in-command (HIC), which refers to an administrative role over when, how, and why the system is used. For systems for which less oversight is possible, the group asks that they be subject to more testing and stricter governance.

Technical robustness and safety

Technical robustness and safety are linked to the principle of prevention of harm. The design of AI systems should be such as to account for risks and unexpected situations in which harm to humans could be caused. This includes protections against both sudden environmental changes and the presence of agents hostile to the system, thus system security is paramount and consideration should be given into the unintended ways in which the system could be made to behave. Fallback plans and risk assessments should be established to ensure a proper response in case of an attack or otherwise failure and assist with minimising the impact thereof.

Other features needed for a robust system are system accuracy, reliability, and reproducibility. System accuracy refers to the system's "ability to make correct judgements" (p. 17), such as proper categorisation or correct decisions and predictions, and if errors and inaccuracies are expected, their likelihood should be communicated. A reliable system is one that works consistently under proper conditions. Reliability, along with reproducibility of results, is important for assessing and testing an AI system.

Privacy and data governance

As regards *privacy and data governance*, linked to the principle of prevention of harm, the AI HLEG asserts the need for proper data protection protocols that respect personal privacy, warning that sensitive information can often be discerned through data about personal behaviour. Also warning about self-learning systems' susceptibility to bad data, they urge for proper testing and documentation of for the purposes of avoiding the use of biased or incorrect data, or the insertion thereof by hostile agents. Lastly, assurances need to be in place to prevent the access of individuals' data by unauthorised, unqualified personnel.

Transparency

Transparency is vital to the principle of explicability. It concerns not only technical transparency of the computational systems involved, but also includes data, the system as a whole, and the business model being followed.

An important element of the technical aspects of transparency is traceability: proper documentation of the data sets and algorithms that are being used is necessary to understand why an incorrect judgement was made and to correct for it in future. Traceability also contributes to explainability, another feature required for the transparency requirements to be met. Explainability means being able to make understood to a human person "both the technical processes of an AI system and the related human decisions" (p. 18), such as the amount of influence the AI's judgement has over decision-making, the reasoning behind the system's design choices, and the reason for its use. The AI HLEG finally asserts that AI systems should avoid pretending to be human, that they should clearly communicate to users their nature as automated systems (and if appropriate, their specifications), and that the option to opt out of using an automated system and switching to a human agent in its place might be necessary in some cases in order to respect fundamental rights.

Diversity, non-discrimination and fairness

The fifth requirement, *diversity, non-discrimination and fairness*, relates to the principle of fairness. It involves respect for accessibility, equal treatment, and inclusive design practices.

A critical component of accomplishing this is being aware of and avoiding biases. The AI HLEG acknowledges that data sets "may suffer from the inclusion of inadvertent historic bias, incompleteness and bad governance models" (p. 18). Recognising and removing such biases from collected data is an important part of correcting for this and preventing their perpetuation. The development of an AI system is also a process susceptible to bias: this should be corrected by oversight and transparency. Additionally, diverse hiring practices should be utilised for the same purpose.

Accessible design is also necessary for fulfilling the requirement and avoiding unfair discrimination. Among other characteristics, special attention should be given towards making AI systems accessible to people with disabilities and enable everyone to participate. Likewise, those affected by an AI system should be invited to participate in dialogue and consultation to provide feedback, prior to and following the system's deployment.

Societal and environmental well-being

Societal and environmental well-being is a concern that relates to the principles of fairness and prevention of harm. It involves a broader perspective of AI systems as part of society and the environment and demands that it respects that wider ecosystem by being sustainable and working for the benefit of all humanity. Environmental sustainability should be assessed with regards to resource and energy usage throughout the system's use, including the full length of its supply chain.

The social impact of AI should also be considered: the AI HLEG warns that "ubiquitous exposure to social AI systems in all areas of our lives [...] may alter our conception of social agency, or impact our social relationships and attachment" (p.19). The potential of AI to compromise democratic institutions' democratic character is also recognised.

Accountability

Accountability is linked to the principle of fairness and, having been mentioned through the text so far, it is inescapably necessary for all the other requirements to be met as well. The need to be able to assess and audit AI systems and their components, and to take the necessary steps for risk minimisation, has been stressed every time it's been applicable - this group here reiterates this need, and calls for protections for whistleblowers, NGOs, trade unions or others that might raise concerns regarding AI systems. Cases where deployment of a system involves trade-offs between meeting these requirements, those need to be acknowledged, documented, and resolved in a rational way if such an option exists. Finally, mechanisms for redress should be planned for in case things go wrong.

For achieving these requirements the AI HLEG suggests several technical and non-technical methods that could be used. Described briefly, the technical methods are:

- *Architectures for Trustworthy AI*, translating the seven requirements into procedures and restrictions that can be more easily followed on a technical level
- *Ethics and rule of law by design*, integrating ethical considerations in the design phase of a system (similar to the "data protection by design" requirement of the GDPR)
- *Explanation methods*, making use of research into Explainable AI to facilitate explicability
- *Testing and validating*, which the AI HLEG recognises is a more complicated process for AI systems, whose output can be difficult to predict for a given input
- *Quality of Service indicators*, "to ensure that there is a baseline understanding as to whether they [AI systems] have been tested and developed with security and safety considerations in mind" (p. 22)

For non-technical methods, they list the following:

- *Regulation*, following on the model of similar precedent such as product safety laws
- *Codes of conduct*, so that organisations dealing with AI integrate the guidelines in their internal processes
- *Standardisation*, as standards can help establish a quality baseline, including in the field of ethical AI
- *Certification*, the issuing of which can serve as a complement to broader review processes
- *Accountability via governance frameworks*, to provide complementary accountability mechanisms alongside legal requirements
- *Education and awareness to foster an ethical mind-set*, necessary for informed participation in AI systems and their evaluation
- *Stakeholder participation and social dialogue*, such as panels including experts, consumers and workers
- *Diversity and inclusive design teams*, which "contributes to objectivity and consideration of different perspectives, needs and objectives" (p. 23)

6.3. United States of America & Platform Self-Regulation

Much of what we've discussed to this point, particularly involving online platforms and social media depends upon or makes use of online legal architecture. Most of these platforms being US-based, American law plays a vital role in mediating the duties and responsibilities of platforms owners and users. The most relevant piece of legislation as regards to that is Section 230, part of the Telecommunications Act of 1996, which provides the legal basis for platform self-regulation: it places the legal responsibility for content uploaded on a platform on the user

rather than the owner, and it gives platform owners the right to moderate content on their platform “in good faith”.³¹ Specifically, Section 230(c) states the following:

“(1) Treatment of publisher or speaker.--No provider or user of an interactive computer service shall be treated as the publisher or speaker of any information provided by another information content provider.

“(2) Civil liability.--No provider or user of an interactive computer service shall be held liable on account of--

“(A) any action voluntarily taken in good faith to restrict access to or availability of material that the provider or user considers to be obscene, lewd, lascivious, filthy, excessively violent, harassing, or otherwise objectionable, whether or not such material is constitutionally protected; or

“(B) any action taken to enable or make available to information content providers or others the technical means to restrict access to material described in paragraph (1).

With few exceptions, such as copyright and trafficking laws, Section 230 has enabled online platforms to pursue their own policies with regards to content moderation without state regulation or fear thereof. This relates to the discussion on the Internet as a public forum vs a collection of private platforms briefly touched upon in the earlier chapter on filter bubbles: at the same time as online speech becomes exempt from state regulation and platforms have no obligation or need to censor their users as they are not liable for their posts, they are nevertheless at liberty to do so. Free speech becomes a privilege for platforms to grant and restrict through their terms of service, rather than a right for users to enjoy.

We touched on a similar blurred line when we touched on the press’s role as the “fourth estate”, where we saw the press’s informal institutional role being at a tension with the private interests of those with executive power over the papers’ editorial line, where the duty to inform and educate is mediated through the needs, wants, fears, and allegiances of the papers’ ownership and editorial staff. In the current day, as social media feeds are more and more serving as an alternative to traditional news media, we identify a similar tension between online spaces’ informal³² dual role as public fora and an information source and their owners’ and administrators’ biases - with the introduction of an additional source of tension in the form of algorithmic curation.

For their part, social media such as Facebook and Twitter, seem aware of their role as a modern fourth estate. Twitter has refused to enforce its terms of service on tweets deemed “newsworthy”, such as Donald Trump’s tweets that violated the platform’s terms of service, as for instance a tweet that could be seen as threatening war against North Korea (Wagner, 2017)³³. Furthermore, these platforms have attempted to present relevant information to their users in cases of crisis or major events in order to refer users to more informative resources. This has particularly been the case since 2016 and criticism of social media platforms as platforms for dissemination of fake news and disinformation. For instance, in 2020 Twitter updated its platform to enable the possibility of labelling posts that contain misleading information (Yeol & Achuthan, 2020), and a few months later as the COVID-19 pandemic spread, adopted more specific policies with regards to posts specifically containing wrong or misleading information about the outbreak and the virus (Yeol & Pickles, 2020). Similar labels were used by Twitter and Facebook during the 2020 US Presidential election for posts containing misinformation in regards to election results, or otherwise posted information that was in conflict with or attempted to pre-empt official sources. (Reuters Staff, 2020)

³¹ Other legislative initiatives, such as the EU’s currently proposed Digital Services Act (European Commission, 2020d) operate on the opposite principle, establishing obligations for large online platforms to meet and which they can be held accountable for. In the United States as well, proposed legislation such as the Stop Internet Sexual Exploitation Act attempts to formalise obligations of platforms towards certain kinds of content (Wille, 2020).

³² Legal precedent for treating social media as public fora exists. Courts have found, for instance, that in his governmental capacity as President, Donald Trump was constitutionally forbidden from blocking users on Twitter. (Savage, 2019) It is therefore acknowledged that whether the role of online spaces as public fora is truly informal is up to debate.

³³ On the 6th of January 2021 Twitter took a firmer stance following Trump’s support of a violent takeover of the US capitol by far-right protesters, leading to his suspension two days later. (Twitter Safety, 2021a; Twitter Safety, 2021b)




Table 4. Examples of Facebook and Twitter tagging posts containing misinformation by US President Donald Trump during the US 2020 presidential election.

The tension between the Internet’s role as a public sphere, social media’s growing influence in their capacity as news delivery platforms and the social responsibilities that are implied by that role, and the legal framework regulating online spaces has been the source of debate and controversy. Twitter’s moderation policies have been controversial with regards to what is seen as a soft stance towards neo-nazis and white nationalists; commenting on that, a Twitter employee claimed that following through with an algorithmic solution would result in a disproportionate amount of content by US Republican politicians also getting flagged. (Newton, 2019)

Regardless, conservatives and right-wingers often feel as though their viewpoints are the subject of greater scrutiny and censorship online. Notably, following a labelling of his tweets similar to what was described above with regards to mail-in ballots (Mangan & Breuninger, 2020), President Trump issued an Executive Order (EO) directing relevant agencies to provide clarification as to the scope of the protections granted by Section 230(c). The EO asserts that online platforms whose moderation of content goes beyond a strict interpretation of Section 230(c)(2)(A) are engaged in editorial conduct, and thus forfeit the protections granted by paragraph (c)(1). (Exec. Order No. 13925, 2020)

Ethical considerations of freedom of speech on social media are furthermore complicated by their ability to facilitate human rights abuses: Facebook has been notoriously used in Myanmar as an instrument for inciting ethnic violence and genocide against the Rohingya people. While against its terms of service, Facebook was ultimately massively underprepared, without adequate Burmese-speaking staff for its 18 million Myanmar users in 2018. In its report, the United Nations stressed Facebook’s role in this atrocity (BBC, 2018; Miles, 2018).

6.4. Democratic Vigilance in Public Policy

The political identity of a democratic state is a complex matter; while “democratic” here is used as a descriptor, there is no binary switch that determines whether a state is actually democratic or not, and contradictions abound. Public policy, being the way in which the state transforms its identity into action, is likewise complicated. In many respects the state is perfectly willing to sacrifice democratic ideals in pursuit of another goal. The growth of the American national security apparatus following the September 11 2001 terrorist attacks immediately jumps to mind, but in analysing the GDPR, the AI HLEG’s work, and Section 230 and the public debate around it, we see similar compromises being made or suggested.

The GDPR’s exemption of authorities to its regulations, and the refusal to lay down separate regulations for them, stands out as the most blatant example of this, being part of a legislation whose entire purpose is to establish individuals’ fundamental rights to data privacy. The application of the Regulation as active policy, with an apparent unwillingness to commit the resources necessary to fully realise its scope, is likewise reflective of a competing tendency to want to protect business interests, even if they find themselves in violation of the GDPR. Much as the legislation might be new and its institutions still growing into their responsibilities, it is hard to treat it as the revolutionary, highly ambitious piece of legislation that it is often portrayed as when the investment in enacting its provisions is so lacklustre.

The AI HLEG’s work can be criticised along similar lines. While the group has indeed laid out a comprehensive set of issues that should be addressed by AI technologies in *Ethical Guidelines for Trustworthy AI*, they tend to couch their concerns by invoking proportionality, asserting the usefulness of Artificial Intelligence, or by giving

consideration to commercial interests. For instance, in assessing the *accountability* requirements as regards auditability, they assert that it “does not necessarily imply that information about business models and intellectual property related to the AI system must always be openly available” (High-Level Expert Group on Artificial Intelligence, 2019b, p.19-20) - they instead posit that evaluation by “internal and external auditors” (p. 20) is generally satisfactory for meeting the requirement. This kind of reasonable approach taken by the AI HLEG means that in spite of establishing a solid list of requirements, their actual suggestions for meeting them falls short of the tremendous ambition of establishing ethical rules for a highly controversial scientific field.

In considering the discourse around online free speech and Section 230 in American politics we encounter a likewise worrying conflict of interest: while great arguments can and have been made about amending the legal status quo, they are often drowned in bad-faith arguments (such as President Trump conflating algorithmic bias with editorialising (DiResta, 2018)) in an attempt to claim victimhood, the conversation “being co-opted and twisted by politicians and pundits howling about censorship and miscasting content moderation as the demise of free speech online” (ibid). It is difficult to not be wary of any attempt to resolve questions about censorship and online content moderation when the political establishment blatantly treats the discourse around it as a means for political gain, online censorship invoked to grab attention rather than in an effort to implement serious reform.

The Marxist critique of the bourgeois state becomes pertinent here by describing the ideological contradiction between the liberal democratic state’s egalitarian character and its bourgeois character. There are frequent cases of civic and business interests in opposition, making that contradiction manifest, and the choice of policy response made, or lack thereof, is a position taken to provide a resolution to that tension. This is not a question of maintaining a delicate balance: we should be wary of the possibility and willingness of the democratic state to sacrifice its democratic character.

6.5. Policy assessment

As with the threats of the previous chapter, having completed our review of policies and policy recommendations seeking to regulate AI and AI-adjacent fields, be it to remedy ethical issues or to address other questions regarding their subject of interest, we can create an overview of those policies and proposals.

The (non-exhaustive) list of policies, techniques and proposals that will be examined is:

- *Consent forms*, referring to widely applied practice of presenting a user with all relevant legal information and asking for their consent (such as in EULAs or cookie consent forms)
- *Pseudonymisation/Anonymisation of data*, referring to the technique of removing unique identifiers from data sets that could be used to explicitly identify the data subjects they correspond to
- *Data ownership rights*, referring to the data rights conferred by the GDPR to individuals that give data subjects agency and control over their personal data
- *Data protection by design/by default*, referring to design processes for AI systems that take consideration of data protection
- *Protected data categories*, referring to special protections for sensitive data categories (such as race, sexual orientation, etc.)
- *GDPR quality control requirements*, referring to the miscellaneous other requirements the GDPR places upon the controller and processor such as elements of designated accountability, controls for data transfers to third parties and so on
- *Human oversight*, referring to the policy of including human elements in automated systems
- *Platform self-regulation*, referring to the policy of giving online platforms free rein in content moderation

The categories examined are:

- *Classification*, describing whether the item a policy, a technique, legal rights etc.
- *Corresponding legislation*, describing whether there is legislation that demands the item or which the item is making an effort towards addressing
- *Field of interest*, describing the more general field that is being addressed or otherwise impacted (such as “data privacy” or “AI ethics”)
- *Subject concerned*, describing who the item concerns
- *Corresponding AI HLEG guideline*, describing which of the seven guidelines presented by the AI HLEG the item is interacting with

Artificial Intelligence: Dangers to Privacy and Democracy

	Classification	Corresponding legislation	Field of interest	Subject concerned	Corresponding AI HLEG Guideline
Consent forms	Legal technique for establishing consent	Consent required by GDPR	Establishing informed consent	Service providers	Transparency
Pseudonymisation/Anonymisation of data	Technique for data privacy	Explicitly required by GDPR	Data privacy	Service providers	Technical robustness and safety
Data ownership rights	Individual rights	Protection guaranteed by GDPR	Data privacy	Service providers, Data subjects	Privacy and data governance
Data protection by design/by default	Design process element	Explicitly required by GDPR	Data privacy	Service providers	Technical robustness and safety
Protected data categories	Individual rights	Explicitly required by GDPR	AI ethics	Service providers, Data subjects	Diversity, non-discrimination and fairness
GDPR quality control requirements	Data handling process requirements	Explicitly required by GDPR	Data privacy, AI ethics, Market regulation	Service providers	Accountability, Transparency, Technical robustness and safety
Human oversight	Voluntary policy	N/A	AI ethics	Service providers	Human agency and oversight
Platform self-regulation	Online content moderation policy	Section 230 protection	Online free speech	Online platform owners	N/A

Table 5. An overview of policies, techniques and proposals discussed so far.

Cross-referencing the two tables we can take a wider view on the degree to which these policies address the identified threats. For the resulting table:

- ✓ (colour-coded blue) marks that a policy comprehensively addresses or is part of comprehensively addressing the identified threat
- ~ (colour-coded yellow) marks that a policy nominally addresses the identified threat, but that its effectiveness towards accomplishing this is questionable at best
- X (colour-coded red) marks that the policy in question contributes to the threat in some fashion

	Consent forms	Pseudonymisation/Anonymisation of data	Data ownership rights	Data protection by design/by default	Protected data categories	GDPR quality control requirements	Human oversight	Platform self-regulation
Data commodification	~		~					
Unregulated data market		✓	✓	✓	✓	✓		
Botnets & Deepfakes								
Direct manipulation of public discourse								

Artificial Intelligence: Dangers to Privacy and Democracy

Government surveillance		~						
Citizen scoring systems					~			
Filter bubbles								X
Algorithmic bias					✓		~	

Table 6. Table cross-referencing the rows of tables 2 and 5.

The lack of an effective response to multiple threats by the discussed policies is visible on this table. Notably, there is no real paradigm to operate as a guideline for dealing with malicious actors and tools available at their disposal such as botnets and deepfakes. Data commodification, a fundamental issue, is essentially unaddressed, with no effective mechanism for accomplishing informed consent, itself a necessity for asserting data ownership rights. Filter bubbles meanwhile are the result of a purposefully unregulated online ecosystem – as owners of online spaces are essentially given free speech protections by Section 230, they are also freed of any obligation to democratic neutrality, something of considerable concern for very large platforms such as search engines and social media sites. The only effective correction being made is against the unregulated data market, which the GDPR by design is primarily aimed towards regulating, the provisions for protected data categories also potentially serving to address matters of algorithmic bias, though as already stated, enforcement remains a major concern.

7. Conclusion

Writing his essay *The End of History?* in the closing days of the Cold War, Francis Fukuyama argued that the defeat of fascism and the fall of communism left liberal democracy as the dominant global ideology, and furthermore, one that has the proven ability to resolve social contradictions to the exclusion of its then-dead competitors. The decades following the fall of communism have not been kind to that sentiment of liberal hubris. Following the Trade World Center terrorist attack of September 11 2001, the expansion of the state security apparatus at the expense of civil liberties in the USA became a fact. The financial crisis of 2008 and the resulting Eurozone crisis shook confidence in the EU's ability to mediate the material conflicts between its member states. At the same time, liberal democracy's success "abroad" has been questioned with the backsliding of Russia into de-facto autocracy following a turbulent decade in the '90s, the economic success of China without any accompanied democratic reforms, or the overall failure of pro-democracy movements such as the Arab Spring to maintain staying power in the face of reaction.

We have repeatedly seen through this study the use, and even weaponisation, of the Internet and related technologies such as AI, in these socio-political developments and struggles. The loosening grasp of liberal democratic ideology on the world and on society has resulted in an equivalent weakening of the ideology's accompanying ethical framework. This also presents difficulties in attempts to translate that ethical framework into a legal framework.

Let's consider the GDPR: Overall, it remains a very ambitious piece of legislation. It should not be forgotten that it is an effort at establishing a highly comprehensive set of standards common to 27 countries with different legal approaches and cultures. Nevertheless, while that statement carries explanatory power, it doesn't change the fact of the difficulty in translating the ideal of data protection into a reality of data protection. Furthermore, it is a regional solution to a global problem. It is also a sectoral solution, as it exempts law enforcement from the privacy standards it establishes.

In practical terms, we saw that there was only one major case brought against a large tech firm as a result of the GDPR, that of France's supervisory authority hitting Google with a 50 million Euro fine. The law might be relatively new, but this early performance is lacklustre, though not entirely unexpected, considering the supervisory authorities generally claimed a lack of funding. The Irish supervisory authority in particular, which is the designated lead supervisor in cases against some of the largest tech firms, has had to deal with major funding issues (receiving less than a third of its requested budget size in 2019) and extremely drawn-out procedures (Vinocur, 2019; Collins, 2020).³⁴ A law that is poorly enforced will not be able to truly achieve its goals vis-a-vis society. The documented disregard towards the law by security agencies such as the American NSA needs also be addressed for civil rights to be truly respected - anything less rings hollow.

The enforceability of legislation such as the GDPR as concerns major tech firms also runs into practical issues given the monopolistic and oligopolistic tendencies of online media. What is the recourse if Google refuses to meet its obligations under the GDPR? Beyond its search engine, Google is also services such as YouTube, Google Docs, Google Maps, or the Android OS. Would the answer in that case be to cut access to the company, with all its services, to the European market? The ideal of an open Internet and the reality of the monopolistic grasp on it that a number of companies have (and the power that comes from that) are here at odds. For regulations like the GDPR to be effective, antitrust laws need to also be enforced. There is reason to hope - an antitrust case against Facebook is currently being examined, with the US government seeking to spin Instagram and WhatsApp into separate entities. (Iyengar, 2020)

The work by the High-Level Expert Group on AI also, while invaluable for its assessment of the requirements for ethical applications of AI, proves somewhat disappointing in some regards.³⁵ In the introduction of *Ethics Guidelines for Trustworthy AI* (High-Level Expert Group on Artificial Intelligence, 2019b) the group recognises that

Just as the use of AI systems does not stop at national borders, neither does their impact. Global solutions are therefore required for the global opportunities and challenges that AI systems bring forth. We therefore encourage all stakeholders to work towards a global framework for Trustworthy AI, building international consensus while promoting and upholding our fundamental rights-based approach. (p. 5)

³⁴ These cases have slowly progressed through the pipeline, with Twitter receiving a €450,000 fine in December 2020. (Schechner, 2020)

³⁵ For a detailed critique on the AI HLEG's policy recommendations and overall attitude towards achieving ethical AI, also see Veale (2019).

However, in the later deliverables only minimal consideration is given to foreign policy. In *Policy and Investment Recommendations for Trustworthy AI* (High-Level Expert Group on Artificial Intelligence, 2019c) the shaping of the global AI landscape is considered almost entirely through the lens of capitalist competition and investment, of turning the European Union into a world leader in AI research and development (indeed, they assert that “it is essential that the EU remains a champion of free trade and investment in the world” (p. 46)). Speaking in an interview in 2018, in the earlier stages of the AI HLEG’s work, the AI HLEG chair, former Nokia president Pekka Ala-Pietilä, stated that “ethics and competitiveness are intertwined, they’re dovetailed” and suggested that preemptive legislation on AI should be avoided to prevent placing the EU in a competitive disadvantage. (Delcker, 2018)

In examining the shortcomings and failures of democratic institutions to keep up with this rapidly evolving technological environment it is tempting to ask the question: is democracy a lost cause? Doing so brings us back to the political analysis of democracy at the beginning of the study and asks us, in turn, to look at democracy itself with a critical eye.

In our initial look at historical democracies we came up with an abstract definition of “democracy as a process” of transferring power from an oligarchy to a wider population. Simultaneously we proposed that this implies the reverse process - that of transferring power from a wider population to an oligarchy - is what we could define as a “threat” to democracy. The global dominance of the capitalist mode of production and the growing economic disparity between the wealthier members of society and the working class forces us to face the reality of a capitalist elite as a socio-political force, both in a global as well as a regional/national scope. It also follows that it is this elite that is best positioned to invest in, and reap the rewards of, artificial intelligence. For all the good intent of legislation such as the GDPR and recommendations such as those of the AI HLEG, this is a fundamental fact that remains unchanged without a wider re-negotiation of socio-economic relations.

Such a re-negotiation is, as for example noted earlier regarding the AI HLEG’s policy recommendations, often beyond the ambitions of those in position of authority. In short, we propose that the research, development, and deployment of AI at present, and at a global level, contains a capitalist class character. It was the demands of capitalist rationality that led to the development of advertisement algorithms focused on generating profits, rather than embodying liberal democratic ideals such as the right to privacy, that in turn permitted Cambridge Analytica’s successful propaganda campaign. Attempts to correct for things such as that after the fact, without fixing the power structures that cause them, seem then like an attempt to patch over a problem without addressing the systemic issues that cause it.

Ultimately we have to recognise the futility and hubris of any end-of-history type of thinking. History has seen empires, such as Rome or China, span millenia, and finally fall. Systems of government such as monarchies, tracing their legitimacy to the divine right of kings, have also come and gone. Our liberal democratic institutions, values and traditions, are much younger by comparison - and certainly younger than democracy itself as a broader concept. Even assuming that the specific arrangement of liberal-democratic society transforms and civil liberties are eroded, that does not mean that its replacement will not represent some other, different form of democracy. This is not so much to make a pessimistic, fatalistic statement, as it is to serve as a reminder that these values are vulnerable to the march of history, and that maintaining them requires action and praxis.

Part of that march of history is technological progress, including the further development of AI. We assert that that technological development is not ideologically neutral, but shaped by the dominant ideology. The work of the AI HLEG shows that a theoretical ethical framework for AI that respects liberal-democratic values is conceivable. The only thing needed is the will to vigorously pursue it, perhaps without the willingness to compromise those values to appease the markets those same experts suggest.

References

1. Ali, M., Sapiezynski, P., Bogen, M., Korolova, A., Mislove, A., Rieke, A. (2019), *Discrimination through optimization: How Facebook's ad delivery can lead to skewed outcomes* [online], Proceedings of the ACM on Human-Computer Interaction 2019, Retrieved from <https://arxiv.org/abs/1904.02095>
2. Anderson, P. (2013), *Passages from Antiquity to Feudalism*, London, Verso, Original work published 1974
3. Ansary, T. (2009), *Destiny Disrupted: A History of the World Through Islamic Eyes*, New York PublicAffairs
4. BBC (2018), *The country where Facebook posts whipped up hate* [online], BBC, Retrieved from <https://www.bbc.com/news/blogs-trending-45449938>
5. Bonnefon, J., Shariff, A., Rahwan, I. (2016), *The social dilemma of autonomous vehicles* [online], Science vol. 352, Retrieved from <https://arxiv.org/abs/1510.03346>
6. Bozdag, E., van den Hoven, J. (2015), *Breaking the filter bubble: democracy and design* [online], Retrieved from <https://doi.org/10.1007/s10676-015-9380-y>
7. Busby, M. (2018), *Corbyn supporters attack Labour moderates for 'using targeted Facebook ads to trick him about own general election campaign'* [online], UK Independent, Retrieved from <https://www.independent.co.uk/news/uk/home-news/labour-hq-jeremy-corbyn-targeted-facebook-ads-2017-election-a8448036.html>
8. Cadwalladr, C., Graham-Harrison, E. (2018), *Revealed: 50 million Facebook profiles harvested for Cambridge Analytica in major data breach* [online], UK The Guardian, Retrieved from <https://www.theguardian.com/news/2018/mar/17/cambridge-analytica-facebook-influence-us-election>
9. Collewet, M., Sauermann, J. (2017), *Working hours and productivity* [online], Labour Economics, Retrieved from <https://doi.org/10.1016/j.labeco.2017.03.006>
10. Collins, K. (2020), *As the GDPR turns 2, Big Tech should watch out for big sanctions* [online], cnet, Retrieved from <https://www.cnet.com/news/as-the-gdpr-turns-2-big-tech-should-watch-out-for-big-sanctions/>
11. Corbett, S. (2019) *Computer game licences: The EULA and its discontents*, NL Elsevier, Retrieved from <https://doi.org/10.1016/j.clsr.2019.03.007>
12. Cox, J. (2020), *How the U.S. Military Buys Location Data from Ordinary Apps* [online], Vice, Retrieved from <https://www.vice.com/en/article/jgqm5x/us-military-location-data-xmode-locate-x>
13. Creemers, R. (2018), *China's Social Credit System: An Evolving Practice of Control*, SSRN, Retrieved from <https://dx.doi.org/10.2139/ssrn.3175792>
14. Croft, J., Venkataramakrishnan, S. (2020), *Police use of facial recognition breaches human rights law, London court rules* [online], Financial Times, Retrieved from <https://www.ft.com/content/b79e0bee-d32a-4d8e-b9b4-c8ffd3ac23f4>
15. Data Protection Commissioner v Facebook Ireland Limited and Maximillian Schrems (European Court of Justice 2020), Retrieved from <https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX:62018CA0311>
16. Delcker, J. (2018), *Europe's AI ethics chief: No rules yet, please* [online], Politico.eu, Retrieved from <https://www.politico.eu/article/pekka-ala-pietila-artificial-intelligence-europe-shouldnt-rush-to-regulate-ai-says-top-ethics-adviser/>
17. DiResta, R. (2018), *Free Speech Is Not the Same As Free Reach* [online], Wired, Retrieved from <https://www.wired.com/story/free-speech-is-not-the-same-as-free-reach/>
18. Duncan, M. (2012), *The History of Rome* [Audio podcast] Retrieved from https://thehistoryofrome.typepad.com/the_history_of_rome/
19. European Commission, *Adequacy Decisions* [online], Retrieved from https://ec.europa.eu/info/law/law-topic/data-protection/international-dimension-data-protection/adequacy-decisions_en [Accessed 3 December 2020]

20. European Commission (2020), *COMMUNICATION FROM THE COMMISSION TO THE EUROPEAN PARLIAMENT AND THE COUNCIL Data protection as a pillar of citizens' empowerment and the EU's approach to the digital transition - two years of application of the General Data Protection Regulation* [online], Retrieved from <https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX%3A52020DC0264>
21. European Commission (2020), *Data protection as a pillar of citizens' empowerment and the EU's approach to the digital transition - two years of application of the General Data Protection Regulation* [online], Retrieved from <https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX%3A52020DC0264>
22. European Commission (2020), *Opening remarks by Vice-President Jourová and Commissioner Reynders at the press point following the judgment in case C-311/18 Facebook Ireland and Schrems* [online], Retrieved from https://ec.europa.eu/commission/presscorner/detail/en/statement_20_1366
23. European Commission (2020), *Proposal for a REGULATION OF THE EUROPEAN PARLIAMENT AND OF THE COUNCIL on a Single Market For Digital Services (Digital Services Act) and amending Directive 2000/31/EC* [online], Retrieved from <https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX:52020PC0825>
24. European Data Protection Board (2019), *EDPB Annual Report 2018* [online], Retrieved from https://edpb.europa.eu/sites/edpb/files/files/file1/edpb_annual_report_2018_-_digital_final_1507_en.pdf
25. European Data Protection Board (2020), *EDPB Annual Report 2019* [online], Retrieved from https://edpb.europa.eu/sites/edpb/files/files/file1/edpb_annual_report_2019_en.pdf.pdf
26. European Parliament and Council of European Union (2016) Regulation (EU) 2016/679, Retrieved from <https://eur-lex.europa.eu/legal-content/EN/TXT/HTML/?uri=CELEX:32016R0679&from=EN>
27. European Union Agency for Fundamental Rights (2018), *#BigData: Discrimination in data-supported decision making* [online], Retrieved from <https://fra.europa.eu/en/publication/2018/bigdata-discrimination-data-supported-decision-making>
28. Exec. Order No. 13925, 85 FR 34079-34083 (2020). <https://www.federalregister.gov/documents/2020/06/02/2020-12030/preventing-online-censorship>
29. Facebook, *Data Policy* [online] <https://www.facebook.com/privacy/explanation/> [accessed 22 Oct 2020]
30. Falch M., Henten A., Tadayoni R., Windekilde I. (2009), *Business Models in Social Networking* [online], CMI International Conference, Retrieved from https://www.researchgate.net/publication/242178725_Business_Models_in_Social_Networking
31. Fukuyama, F. (1989), *The End of History?* [online], The National Interest, Retrieved from <http://www.wesjones.com/eoh.htm>
32. Glancy, D. J. (1974), *The Invention of the Right to Privacy*, US Ariz. L. Rev, Retrieved from <https://digitalcommons.law.scu.edu/facpubs/317/>
33. Google, *Privacy Policy - Privacy & Terms* [online] <https://policies.google.com/privacy?hl=en-US> [accessed 22 Oct 2020]
34. Green, D. (2017), *Big Brother Is Listening to You: Digital Eavesdropping in the Advertising Industry* [online], US Duke L. & Tech. Rev., Retrieved from <https://scholarship.law.duke.edu/dltr/vol16/iss1/12>
35. Grill, G. (1997), *GLCM Guidance System* [online], Association of Air Force Missileers, Retrieved from <http://afmissileers.com/newsletters/NL1997/Dec97.pdf>
36. Grother, P., Ngan, M., Hanaoka, K. (2019), *Face Recognition Vendor Test (FRVT) Part 3: Demographic Effects* [online], National Institute of Standards and Technology, Retrieved from <https://doi.org/10.6028/NIST.IR.8280>

37. Guariglia, M. (2020), *Police Will Pilot a Program to Live-Stream Amazon Ring Cameras* [online], EFF, Retrieved from <https://www.eff.org/deeplinks/2020/11/police-will-pilot-program-live-stream-amazon-ring-cameras>
38. Hampton, M. (2010), “The Fourth Estate Ideal in Journalism History”, *The Routledge Companion to News and Journalism* [online], UK Routledge, Retrieved from http://chinhnghia.com/The_Routledge_Companion_to_News_and_Journalism.pdf#page=48
39. Heywood, A. (1992) *Political Ideologies*, UK Palgrave Macmillan
40. High-Level Expert Group on Artificial Intelligence (2019), *A definition of Artificial Intelligence: main capabilities and scientific disciplines* [online], Retrieved from <https://ec.europa.eu/digital-single-market/en/news/definition-artificial-intelligence-main-capabilities-and-scientific-disciplines>
41. High-Level Expert Group on Artificial Intelligence (2019), *Ethics Guidelines for Trustworthy AI* [online], Retrieved from <https://ec.europa.eu/digital-single-market/en/news/ethics-guidelines-trustworthy-ai>
42. High-Level Expert Group on Artificial Intelligence (2019), *Policy and investment recommendations for trustworthy Artificial Intelligence* [online], Retrieved from <https://ec.europa.eu/digital-single-market/en/news/policy-and-investment-recommendations-trustworthy-artificial-intelligence>
43. High-Level Expert Group on Artificial Intelligence (2020), *Assessment List for Trustworthy Artificial Intelligence (ALTAI) for self-assessment* [online], Retrieved from <https://ec.europa.eu/digital-single-market/en/news/assessment-list-trustworthy-artificial-intelligence-altai-self-assessment>
44. High-Level Expert Group on Artificial Intelligence (2020), *Sectoral Considerations on Policy and Investment Recommendations for Trustworthy AI* [online], Retrieved from <https://futurium.ec.europa.eu/en/european-ai-alliance/document/ai-hleg-sectoral-considerations-policy-and-investment-recommendations-trustworthy-ai>
45. Hill, C. (1980), *The Century of Revolution*, UK Routledge
46. Hill, K. (2020), *Wrongfully Accused by an Algorithm* [online], The New York Times, Retrieved from <https://www.nytimes.com/2020/06/24/technology/facial-recognition-arrest.html>
47. Hobbes, T. (1651) *Leviathan or The Matter, Forme and Power of a Common-Wealth Ecclesiasticall and Civil*, UK n.p.
48. Holstein, T., Dodig-Crnkovic G., Pelliccione P. (2018), *Ethical and Social Aspects of Self-Driving Cars* [online], Gothenburg, ARXIV’18, Retrieved from <https://arxiv.org/abs/1802.04103>
49. Information Commissioner’s Office, *How do we comply with the cookie rules?* [online], Retrieved from <https://ico.org.uk/for-organisations/guide-to-pecr/guidance-on-the-use-of-cookies-and-similar-technologies/how-do-we-comply-with-the-cookie-rules/> [Accessed 1 December 2020]
50. Isaak, J., Hanna M. J. (2018), *User Data Privacy: Facebook, Cambridge Analytica, and Privacy Protection* [online], US IEEE, Retrieved from <https://doi.org/10.1109/MC.2018.3191268>
51. Iyengar, R. (2020), *The antitrust case against Facebook: Here's what you need to know* [online], CNN Business, Retrieved from <https://edition.cnn.com/2020/12/11/tech/facebook-antitrust-lawsuit-what-to-know/index.html>
52. Kates, S. (2006), *Literacy, Voting Rights, and the Citizenship Schools in the South, 1957-1970* [online], US NCTE, Retrieved from <https://www.jstor.org/stable/20456898>
53. Kymlicka, W. (2005) *Contemporary Political Philosophy. An Introduction* (G. Molyvas, Trans.) GR Polis (Original work published 2002)
54. Lallement, M. (2004) *Histoire des idées sociologiques*, (H. Agkiranopoulou, Trans.) GR Metaixmio (Original work published 2000)

55. Landau, S. (2013) *Making Sense from Snowden: What's Significant in the NSA Surveillance Revelations* [online], US IEEE, Retrieved from <https://doi.org/10.1109/MSP.2013.90>
56. Lenin, V. I. (1918) "Bourgeois And Proletarian Democracy", *The Proletarian Revolution and the Renegade Kautsky* [online], Retrieved from <https://www.marxists.org/archive/lenin/works/1918/prrk/democracy.htm>
57. Levien, R., Maron, M. E. (1964), *Cybernetics and its Development in the Soviet Union* [online], Santa Monica, CA, Rand Corporation, Retrieved from <https://apps.dtic.mil/sti/citations/AD0602705>
58. Liang, F., Das, V., Kostyuk, N., Hussain, M. (2018), *Constructing a Data-Driven Society: China's Social Credit System as a State Surveillance Infrastructure* [online], US Wiley, Retrieved from <https://doi.org/10.1002/poi3.183>
59. Lowen, M. (2015), *Greek debt crisis: What was the point of the referendum?* [online], BBC, Retrieved from <https://www.bbc.com/news/world-europe-33492387>
60. Lyons, J. (2017), *Bay Area start-up deploys robots to harass homeless* [online], Salon, Retrieved from <https://www.salon.com/2017/12/13/bay-area-start-ups-robots-employed-to-harass-the-homeless/>
61. Mangan, D, Breuninger, K. (2020), *Twitter fact-checks Trump, slaps warning labels on his tweets about mail-in ballots* [online], CNBC.com, Retrieved from <https://www.cnbc.com/2020/05/26/twitter-fact-checks-trump-slaps-warning-labels-on-his-tweets-about-mail-in-ballots.html>
62. Manheim, K., Kaplan, L. (2019), *Artificial Intelligence: Risks to Privacy and Democracy* [online], 21 Yale Journal of Law and Technology 106, Loyola Law School, Retrieved from <https://ssrn.com/abstract=3273016>
63. Massad, D. [@badnetworker] (2019), *1999: there are millions of websites all hyperlinked together 2019: there are four websites, each filled with screenshots of the other three.* [Tweet], Twitter, Retrieved from <https://twitter.com/badnetworker/status/1133363823728091136> [accessed 9 January 2021]
64. Metal Gear Solid 2: Sons of Liberty (PlayStation 2 version) [video game]. (2001). Tokyo, Japan: Konami
65. Metz, R. (2019), *These people do not exist. Why websites are churning out fake images of people (and cats)* [online], CNN Business, Retrieved from <https://edition.cnn.com/2019/02/28/tech/ai-fake-faces/index.html>
66. Miles, T. (2018), *U.N. investigators cite Facebook role in Myanmar crisis* [online], Reuters, Retrieved from <https://www.reuters.com/article/us-myanmar-rohingya-facebook-idUSKCN1GO2PN>
67. Mitrou, L., Piskopani, A., Tassis, S., Karyda, M., Kokolakis, S. (2013), *Facebook, Blogs και Δικαιώματα*, GR Sakkoula
68. Mitrou, L. (2018), *Data Protection, Artificial Intelligence and Cognitive Services: Is the General Data Protection Regulation (GDPR) 'Artificial Intelligence-Proof'?* [online], Retrieved from <https://ssrn.com/abstract=3386914>
69. Montesquieu, C.d.S., B.d. (1899), *The Spirit of Laws* (T. Nugent, Trans.), New York, The Colonial Press (original work published 1748)
70. Mueller, R. S. (2019), *Report On The Investigation Into Russian Interference In The 2016 Presidential Election* [online], Washington D.C., US Department of Justice, Retrieved from <https://www.justice.gov/storage/report.pdf>
71. Mullan, P. (2020), "The Anti-democratic Roots of Neoliberalism", *Beyond Confrontation: Globalists, Nationalists and Their Discontents* [online], Emerald Publishing Limited, Retrieved from <https://doi.org/10.1108/978-1-83982-560-620200005>
72. Newton, C. (2019), *Why Twitter has been slow to ban white nationalists* [online], The Verge, Retrieved from <https://www.theverge.com/interface/2019/4/26/18516997/why-doesnt-twitter-ban-nazis-white-nationalism>

73. Oracle, *What is Big Data?* [online], Retrieved from <https://www.oracle.com/big-data/what-is-big-data.html> [accessed 31 Oct 2020]
74. Papadopoulos, T., Charalabidis, Y. (2020), *What do governments plan in the field of artificial intelligence? Analysing national AI strategies using NLP* [online], Athens ICEGOV 2020, <https://doi.org/10.1145/3428502.3428514>
75. Pariser, E. (2011). *The Filter Bubble: What the Internet Is Hiding from You*, US Penguin Press
76. Pencavel, J. (2014), *The Productivity of Working Hours* [online], The Economic Journal, Retrieved from <https://doi.org/10.1111/econj.12166>
77. Peters, B. (2016), *How Not to Network a Nation: The Uneasy History of the Soviet Internet*, Cambridge, Massachusetts, The MIT Press
78. Horowitz, J., Igielnik, R., Kochhar, R. (2020), 1. *Trends in income and wealth inequality* [online], Pew Research Center, Retrieved from <https://www.pewsocialtrends.org/2020/01/09/trends-in-income-and-wealth-inequality/>
79. Popkin, J. D. (2011) *A Concise History of the Haitian Revolution*, US Wiley
80. Popper, K. (1945), *The Open Society And Its Enemies* [online], London, Routledge, Retrieved from <https://archive.org/details/in.ernet.dli.2015.59272>
81. R. Neisse, G. Baldini, G. Steri and V. Mahieu (2016), *Informed consent in Internet of Things: The case study of cooperative intelligent transport systems*, US IEEE, Retrieved from <https://doi.org/10.1109/ICT.2016.7500480>
82. Raghavan, M., Barocas, S., Kleinberg, J., Levy, K. (2020), *Mitigating bias in algorithmic hiring: evaluating claims and practices* [online], FAT* '20: Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency, Retrieved from <https://doi.org/10.1145/3351095.3372828>
83. Reuters Staff (2020), *Twitter, Facebook outline action on posts claiming early U.S. election victory* [online], Retrieved from <https://www.reuters.com/article/us-usa-election-twitter/twitter-facebook-outline-action-on-posts-claiming-early-u-s-election-victory-idUSKBN2711UX>
84. Roth, Y., Achuthan, A. (2020), *Building rules in public: Our approach to synthetic & manipulated media* [online], Twitter, Retrieved from https://blog.twitter.com/en_us/topics/company/2020/new-approach-to-synthetic-and-manipulated-media.html
85. Roth, Y., Pickles, N. (2020), *Updating our approach to misleading information* [online], Twitter, Retrieved from https://blog.twitter.com/en_us/topics/product/2020/updating-our-approach-to-misleading-information.html
86. Russell, S., Norvig, P. (2005), *Artificial Intelligence A Modern Approach* (Refanidis, G., trans.) Athens, Kleidarithmos (original work published 2003)
87. SAS, *Big Data What is it And Why it Matters* [online], Retrieved from https://www.sas.com/en_us/insights/big-data/what-is-big-data.html [accessed 31 Oct 2020]
88. Savage, C. (2019), *Trump Can't Block Critics From His Twitter Account, Appeals Court Rules* [online], The New York Times, Retrieved from <https://www.nytimes.com/2019/07/09/us/politics/trump-twitter-first-amendment.html>
89. Schechner, S. (2020), *Twitter Fined for Breaking EU Privacy Law in First for U.S. Tech Firm* [online], The Wall Street Journal, Retrieved from <https://www.wsj.com/articles/twitter-fined-546-000-in-first-cross-border-gdpr-case-for-u-s-tech-firm-11608027373>
90. Silver, L., Smith, A., Johnson, C., Jiang, J., Anderson, M., Rainie, L. (2019), *Mobile Connectivity in Emerging Economies* [online], Pew Research Center, Retrieved from <https://www.pewresearch.org/internet/2019/03/07/mobile-connectivity-in-emerging-economies/>
91. Sprenger, P. (1999), *Sun on Privacy: 'Get Over It'* [online], Retrieved from <https://www.wired.com/1999/01/sun-on-privacy-get-over-it/>

92. Strick, B. (2020), *West Papua: New Online Influence Operation Attempts to Sway Independence Debate* [online], Retrieved from <https://www.bellingcat.com/news/2020/11/11/west-papua-new-online-influence-operation-attempts-to-sway-independence-debate/>
93. Technological Determinism (2021, January 13), In *Wikipedia*, https://en.wikipedia.org/w/index.php?title=Technological_determinism&oldid=1000026818
94. Telecommunications Act of 1996, S.652, 104th Cong. (1996). <https://www.congress.gov/bill/104th-congress/senate-bill/652>
95. Tessian (2020), *11 Biggest GDPR Fines of 2020 (So Far)* [online], Retrieved from <https://www.tessian.com/blog/biggest-gdpr-fines-2020/>
96. The Onion (2012), *Report: Every Potential 2040 President Already Unelectable Due To Facebook* [online], The Onion, Retrieved from <https://politics.theonion.com/report-every-potential-2040-president-already-unelecta-1819595196>
97. Thorley, J. (1996), *Athenian Democracy*, UK Routledge
98. Trotsky, L. (1938), *Freedom of the Press and the Working Class* [online], MX Clave, Retrieved from <https://www.marxists.org/archive/trotsky/1938/08/press.htm>
99. Twitter, *Privacy Policy* [online] <https://twitter.com/en/privacy> [accessed 22 Oct 2020]
100. Twitter (2018), *Update on Twitter's review of the 2016 US election* [online]. Retrieved from https://blog.twitter.com/en_us/topics/company/2018/2016-election-update.html
101. Twitter Safety [@TwitterSafety] (2021), *Future violations of the Twitter Rules, including our Civic Integrity or Violent Threats policies, will result in permanent suspension of the @realDonaldTrump account.* [Tweet], Twitter, Retrieved from <https://twitter.com/TwitterSafety/status/1346970432017031178> [accessed 7 January 2021]
102. Twitter Safety [@TwitterSafety] (2021), *After close review of recent Tweets from the @realDonaldTrump account and the context around them we have permanently suspended the account due to the risk of further incitement of violence.* [Tweet], Twitter, Retrieved from <https://twitter.com/TwitterSafety/status/1347684877634838528> [accessed 9 January 2021]
103. Uggen, C., Larson, R., Shannon, S. (2016), *6 Million Lost Voters: State-Level Estimates of Felony Disenfranchisement, 2016* [online], Sentencing Project, Retrieved from <https://www.sentencingproject.org/publications/6-million-lost-voters-state-level-estimates-felony-disenfranchisement-2016/>
104. Veale, M. (2019), *A Critical Take on the Policy Recommendations of the EU High-Level Expert Group on Artificial Intelligence* [online], Retrieved from https://discovery.ucl.ac.uk/id/eprint/10084302/7/Veale_HLEG_preprint_revised.pdf
105. Vinocur, N. (2019), *'We have a huge problem': European tech regulator despairs over lack of enforcement* [online], Politico, Retrieved from <https://www.politico.com/news/2019/12/27/europe-gdpr-technology-regulation-089605>
106. Wagner, K. (2017), *Twitter says Donald Trump's tweets are newsworthy, which might explain why he hasn't been suspended* [online], Vox, Retrieved from <https://www.vox.com/2017/9/25/16364054/president-donald-trump-twitter-north-korea-tweet-guidelines-nuclear-war>
107. Warner, M. (2018), *Statement of U.S. Sen. Mark R. Warner on Cambridge Analytica* [Press Release], Retrieved from <https://www.warner.senate.gov/public/index.cfm/2018/3/statement-of-u-s-sen-mark-r-warner-on-cambridge-analytica>
108. Wille, M. (2020), *New Internet sex bill could hurt sex workers the most* [online], InputMag, Retrieved from <https://www.inputmag.com/culture/new-internet-sex-bill-could-hurt-sex-workers-the-most>