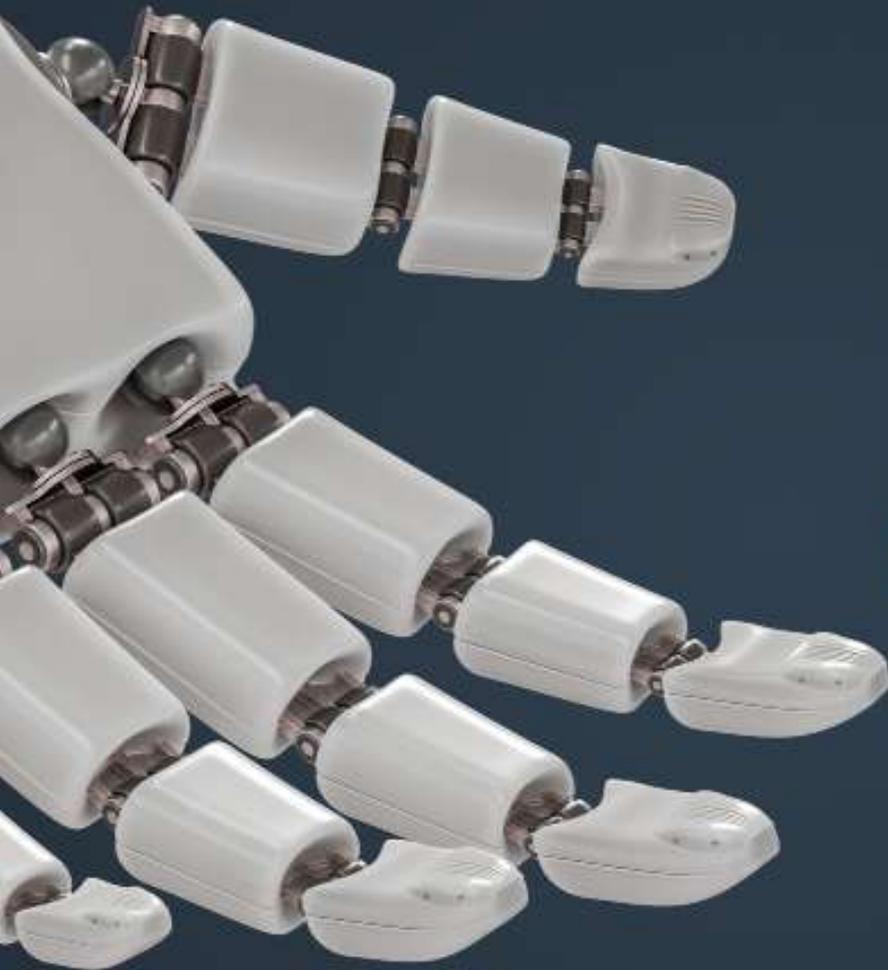




4 PEOPLE



ON GOOD AI GOVERNANCE

14 PRIORITY ACTIONS,
A S.M.A.R.T. MODEL OF
GOVERNANCE, AND A
REGULATORY TOOLBOX

A B S T R A C T

AI4People's second year of activities has focused on applying – concretely, in real world scenarios and through appropriate governance – those ethical principles of AI announced by AI4People in 2018. The 2019 White Paper gives shape to – whilst establishing priorities and critical issues – **14 Priority Actions**, a **Model of S.M.A.R.T. Governance** and a **Regulatory Toolbox**, to which governments and businesses alike can refer to – immediately and efficiently.

To conceive the aforementioned, we examine current initiatives and debates on the governance of AI, and consequently provide:

- A definition of the notion of governance and the principles that are at stake in this context
- **14 Priority Actions** that can be undertaken immediately, existing within three new groups of priority: (i) forms of engagement; (ii) no-regrets actions; and (iii) coordination mechanisms for the governance of AI
- A **S.M.A.R.T. Model of Governance**, for both governments and businesses, adequate for tackling the normative challenges of AI, while being **Scalable, Modular, Adaptable, Reflexive, and Technologically-savvy**. We call for specific forms of governance that are neither bottom-up, nor top-down, but that are in-between, and argue that neither co-regulatory models of AI governance – nor forms of self-regulation, nor its variants, such as ‘monitored self-regulation’ – are adequate
- A **Regulatory Toolbox**, illustrating how the model of **S.M.A.R.T.** governance works



TABLE OF CONTENTS

INTRODUCTION	5
1. DEFINITIONS	10
1.1 Three Levels of Legal Regulation, Co-Regulation, and Self-Regulation	10
1.2 Law And Ethics: Three Kinds of Interaction	11
2. 14 PRIORITY ACTIONS	16
2.1 No-regrets Actions – a Top-Down Approach	17
2.2 Forms of Engagement – a Bottom-Up Approach	18
2.3 Coordination Mechanisms – a “Middle-Out” Approach	19
3. A S.M.A.R.T. MODEL OF GOVERNANCE	22
3.1 Function of the Model	23
3.2 Requirements of the Model.....	23
3.3 The GDPR Model	24
3.4 The Limits of the GDPR Model	25
3.5. On S.M.A.R.T. Coordination	26
3.5.1 The "Middle-out" Approach to AI	27
3.5.2 A New Coordination at Work	29
3.5.3 Scalability	30
3.5.4 Granularity	31
4. A REGULATORY TOOLBOX	34
4.1 Function of the Toolbox.....	34
4.2 Requirements of the Toolbox.....	35
4.3 The Tools.....	36
CONCLUSIONS	39
REFERENCES	40

Authors

Ugo Pagallo¹, Paola Aurucci², Pompeu Casanovas³, Raja Chatila^{4, 5}, Patrice Chazerand⁶, Virginia Dignum^{7, 8}, Christoph Luetge⁹, Robert Madelin¹⁰, Burkhard Schafer¹¹ and Peggy Valcke^{12, 13}.

¹Department of Law, University of Turin, Turin, Italy.

²Department of Management, University of Turin, Turin, Italy.

³La Trobe University Law School, Melbourne, Australia.

⁴French National Center of Scientific Research, France.

⁵Institute of Intelligent Systems and Robotics at Pierre and Marie Curie University, Paris, France.

⁶Digital Europe, Brussels, Belgium.

⁷Digital Europe, Brussels, Belgium.

⁸University of Umeå, Umeå, Sweden.

⁹Delft Design for Values Institute, Delft University of Technology, Delft, the Netherlands.

¹⁰Centre for Technology and Global Affairs, University of Oxford, Oxford, United Kingdom.

¹¹University of Edinburgh Law School, Edinburgh, United Kingdom.

¹²Centre for IT & IP Law, Catholic University of Leuven, Flanders, Belgium.

¹³Bocconi University, Milan, Italy.



A I 4 P

I N B R I E F

AI4People is a multi-stakeholder forum, bringing together all actors interested in shaping the social impact of new applications of AI, including the European Commission, the European Parliament, civil society organisations, industry and the media.

Launched in February 2018 with a three year roadmap, the goal of AI4People is to create a common public space for laying out the founding principles, policies and practices on which to build a “good AI society”. For this to succeed we need to agree on how best to nurture human dignity, foster human flourishing and take care of a better world. It is not just a matter of legal acceptability, it is really a matter of ethical preferability.



EXECUTIVE SUMMARY

This White Paper reports the findings of AI4People, an Atomium – EISMD initiative, whose 2019 activities have aimed to giving shape to tangible, actionable and immediately available recommendations for Good AI Governance. Drawing on the ethical principles published in last year’s White Paper -*AI4People’s Ethical Framework for a Good AI Society* - AI4People’s second year of activities has focused on applying concretely, in real world scenarios and through appropriate governance, those ethical principles of AI announced by AI4People in 2018. To introduce this year's recommendations of the Scientific Committee, the White Paper is divided into four parts that regard definitions, 14 different priorities actions, a model of S.M.A.R.T. governance, and its corresponding regulatory toolbox.

First, we introduce the notion of governance and the principles that are at stake in this context. The focus is on (i) different normative systems in competition, such as ethics, the forces of the market, and of social norms; (ii) different kinds of legal regulation, as a key ingredient of every governance model; and, (iii) three ways in which law and ethics may interact. On this basis, we stress a substantial convergence of today's ethical and legal debate on the governance of AI.

Second, drawing on this convergence of today's debate, we present 14 priority actions that can be undertaken immediately. We present them in accordance with three new groups of priority that regard (i) no-regrets actions; (ii) forms of engagement; and (iii) coordination mechanisms for the governance of AI. Such group of priorities illustrate three different kinds of legal regulation, namely, top-down (no regrets actions); bottom-up (forms of engagement); and, in between such forms of top-down and bottom-up regulation (coordination mechanisms). Rather than a simple list of priorities, this approach should thus be understood in connection with a corresponding new model of governance that properly tackles the complexity of current moral and legal issues of AI.

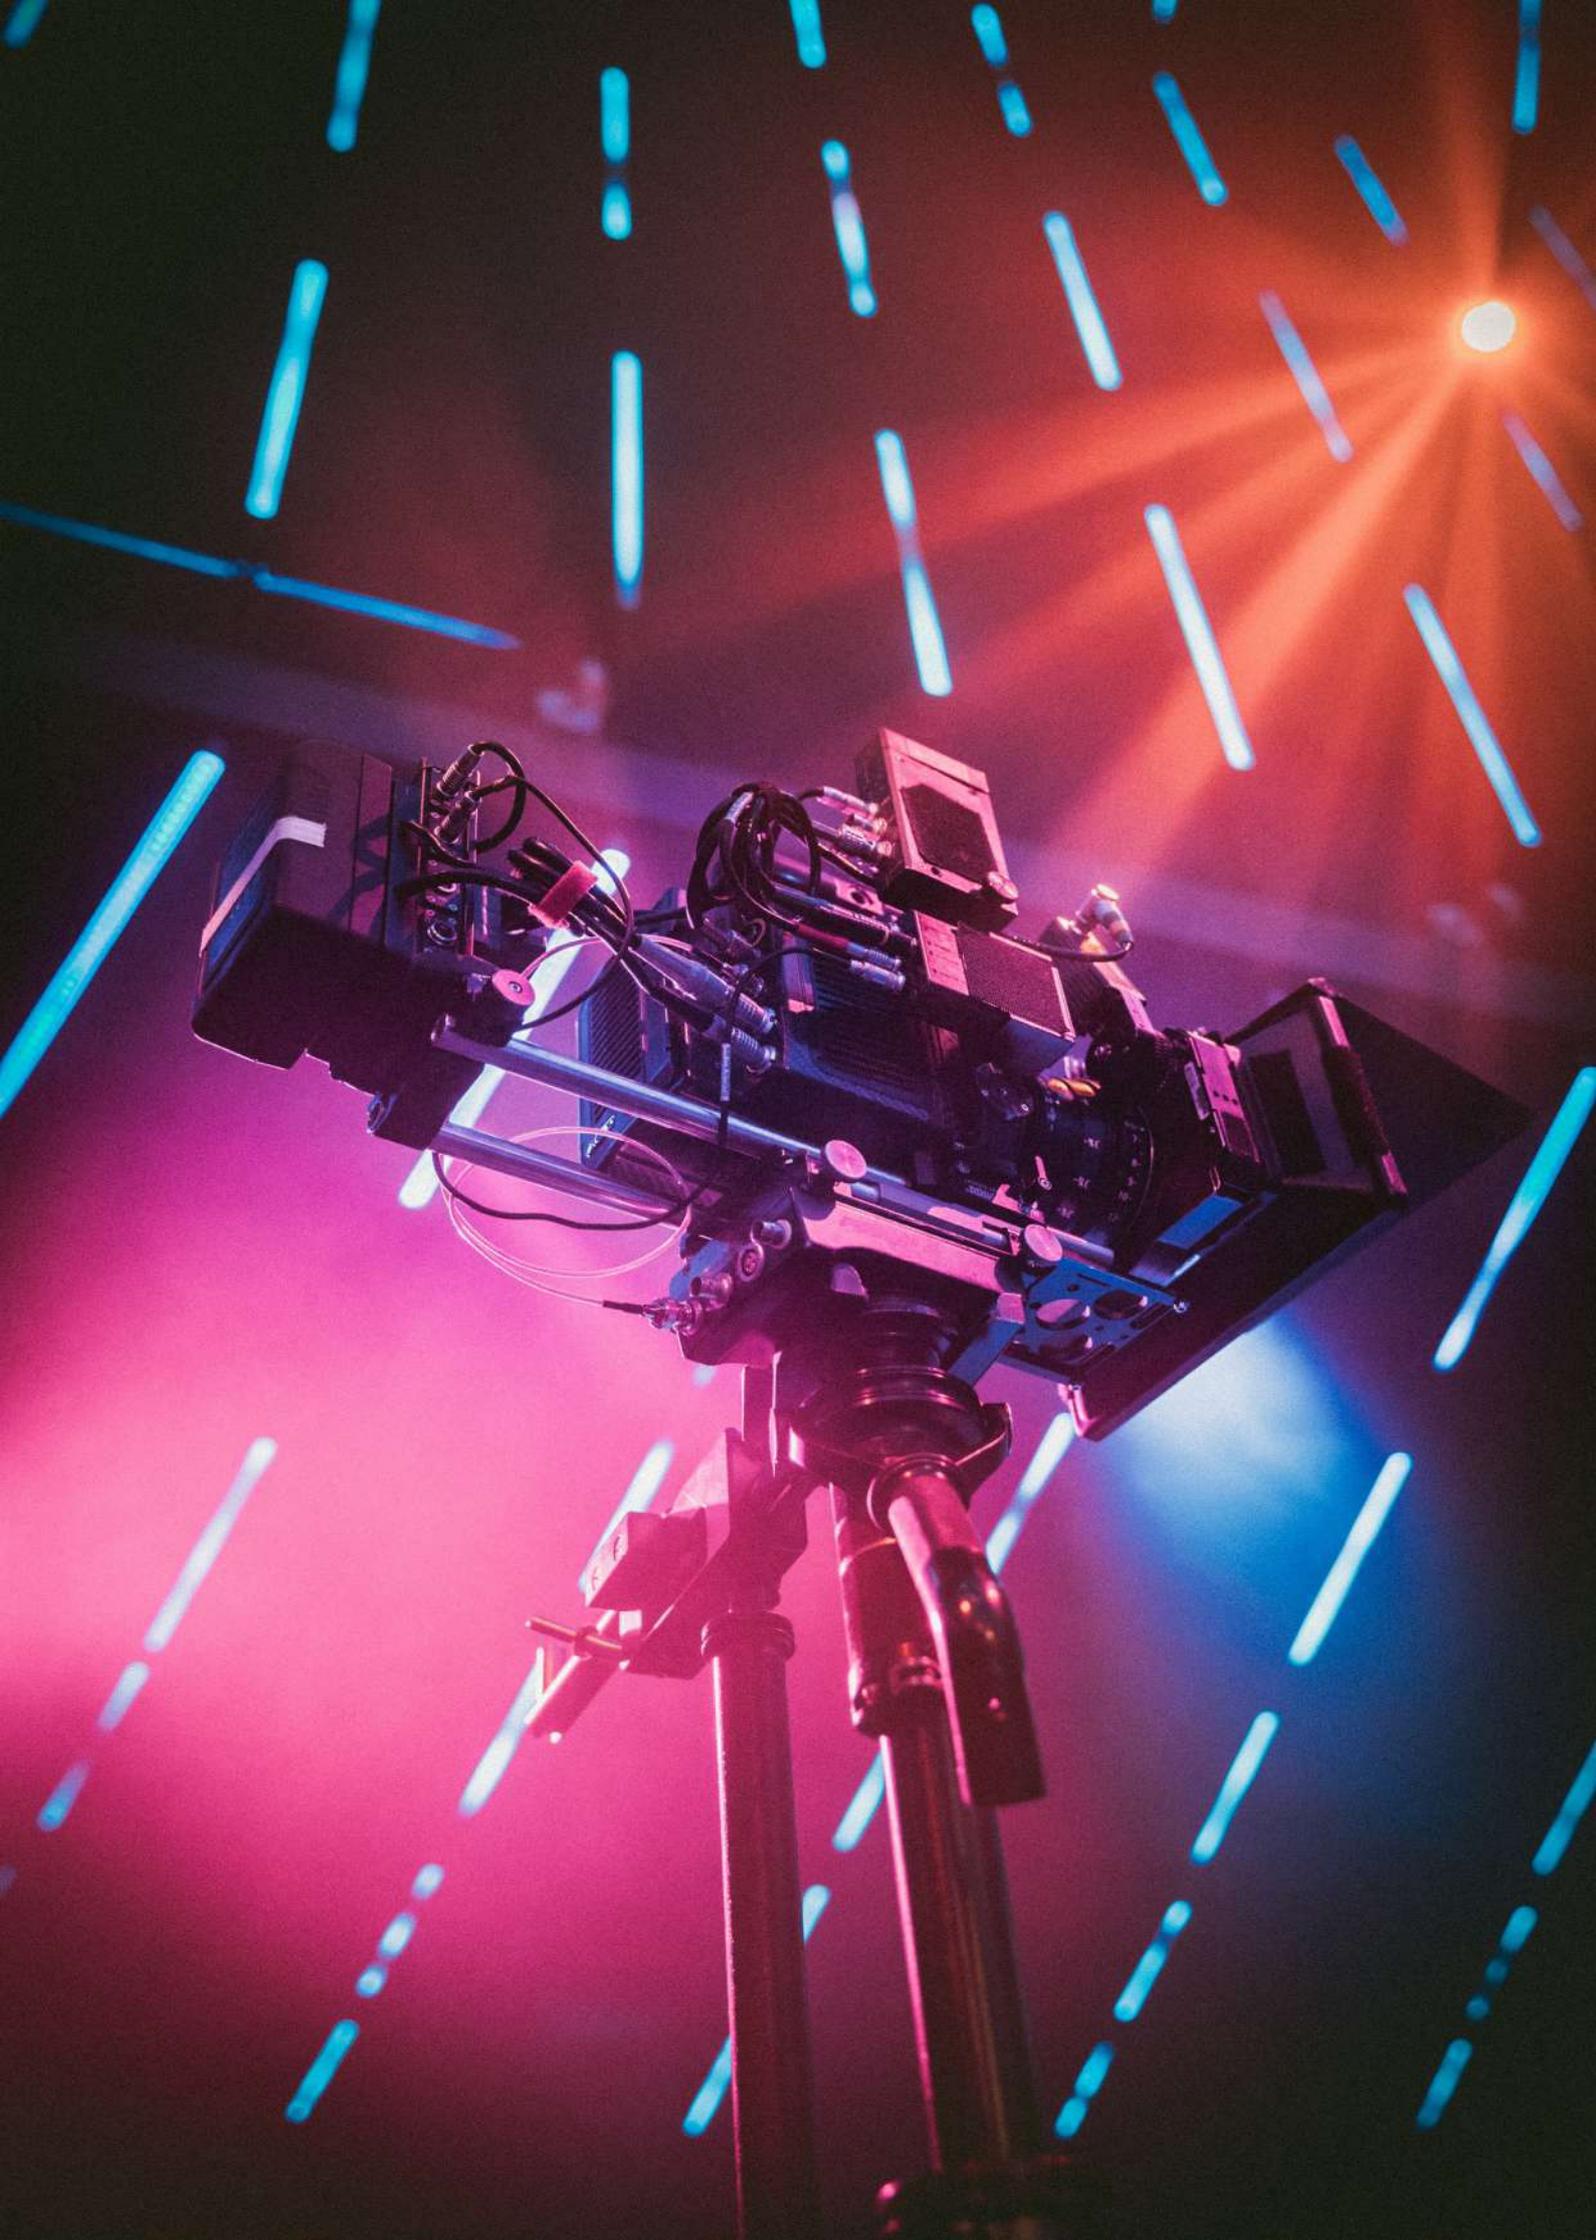
Third, the complexity of current moral and legal issues of AI calls for a specific form of governance that is neither bottom-up nor top-down, but it is in between forms of co-regulation and every variant of self-regulation. This is what we call the “middle-out” approach. As technological challenges grow increasingly complex, both top-down and



bottom-up approaches will become less fruitful. Correspondingly, we scrutinize this “middle-out” layer of the analysis, which should be able to strike the balance between technology, ethics, market, and social norms. In the current EU legal framework, this “middle-out” layer is mostly associated with forms of co-regulation, such as the GDPR. In the GDPR the mixed approach to governance revolves around the ‘principle of accountability’ that through a mix of primary and secondary legal rules aims to strike a balance between guaranteeing compliance with both the principles and the top-down rules of the system, while leaving room for self-regulatory measures. We illustrate the limits of this co-regulatory model of governance for the field of AI and provide for a more adequate model of S.M.A.R.T coordination, i.e. a model which should be **Scalable, Modular, Adaptable, Reflexive, and Technologically-savvy**; Scalability, Modularity and Adaptability of the model can be suitably addressed by adding resources to the mechanisms set up for coordination, and complemented by further procedures of Reflexivity, i.e. sound scrutiny, assessment and evaluation of risks, and Technological-savvy oversight. This model fits neatly with recent EU policies on better and smart regulation, and is consistent with the stance on the rule of law taken by standardisation agencies and governance models in the business field.

Fourth, the White Paper complements this model of S.M.A.R.T coordination with a regulatory toolbox. The latter is developed in accordance with European and national laws, much as case-based law and comprises all components for the good AI governance stressed throughout the paper, namely, (i) principles and legal schemes; (ii) digital innovation hubs; (iii) innovation deals; (iv) compliance through design and by design; (v) regulatory sandboxes; (vi) online and alternative dispute resolution; (vii) regulatory computing architectures; (viii) regulatory ontologies, interoperability & ontology design patterns; (ix) monitoring functions; (x) standardisation, verification and impact assessment mechanisms; (xi) oversight agency. The aim is to accommodate the uncertainties of innovation and, at the same time, to capture expanding understanding with increasing regulatory granularity.





INTRODUCTION

The governance of artificial intelligence (AI) is one of the hottest topics in contemporary institutional debate. In the past years, around the globe, institutions, governments, non-profits and businesses alike have analysed, questioned and attempted to answer what ought to be the principles, guidelines and ethical principles of a “Good AI Society”.

AI4People, an Atomium – EISMD initiative, published one of the world’s most relevant of such guidelines in November 2018. *The AI4People’s Ethical Framework for a Good AI Society* laid out the foundations for a “Good AI Society”, and its guidelines were later adopted by the EU’s High-Level Expert Group on Artificial Intelligence (AI HLEG), in the first draft of their AI ethics guidelines, published in December 2018.

Although important, initiatives such as the one undertaken by the EU - as well as other governments and institutions (outlined in a separate box) - have mainly focused on the ethical principles that should undergird the adoption of AI, firstly outlined in Europe in *AI4People’s 2018 Ethical Framework*.

In 2019, the report “On Good AI Governance: 14 Priority Actions, a S.M.A.R.T. Model of Governance, and a Regulatory Toolbox” has taken a step towards tangible, actionable and immediately available recommendations for Good AI Governance for both governments and businesses alike.

To introduce these recommendations, the White Paper is structured in the following manner:

1. DEFINITIONS

This section **introduces the notion of governance and the principles** that are at stake in this context. First, attention is drawn to the formal and informal rules that govern the public arena, in which the **different regulatory systems** of ethics, the forces of the market, or of social norms **interact**. Second, focus is shifted on a specific set of formal rules, specifically the **norms of legal regulation**, which we distinguish between top-down regulation, co-regulation, and self-regulation. Third, the section examines how legal regulation and governance relate to the **ethics and the moral principles** outlined in *AI4People’s 2018 Ethical Framework*.



2. 14 PRIORITY ACTIONS

Section 2 scrutinises the **substantial convergence** of today's ethical and legal debate **on the governance of AI**, highlighting 14 priorities - that can be undertaken immediately – which exist within three new groups of priority:

(i) **No-regrets actions** concern the achievement of sustainable development goals, such as capacity building in a good AI society; an interoperable AI strategy between the EU and Member States; a support for the capacity of corporate boards of directors to take responsibility for the ethical implications of companies' AI technologies; strategies of inclusive innovation; the creation of educational curricula around the impact of AI and a coherent European AI research environment. This set of no-regrets actions includes both tools of hard law and soft law, corresponding to a **top-down approach** to legal regulation;

(ii) **Forms of engagement** including, amongst others, cross-disciplinary and cross-sectorial cooperation - and debate - on the issues of AI, the creation of an European observatory for AI, and of legally deregulated special zones, or living labs, for AI empirical testing - and development for a better interaction between scientists and laymen. By taking into account today's limited understanding of the stakes of AI, the creation of new type of forums for collective consultation and discussion becomes a priority. Such forms of engagement correspond to a **bottom-up approach to legal regulation**;

(iii) **Coordination mechanisms** that represent a sort of interface between top-down and bottom-up approaches, that is, between the different forms of engagement and the set of no-regrets actions. These coordination mechanisms include participatory procedures for the alignment of societal values and understanding of public opinion, upstream multi-stakeholder mechanisms for risk mitigation, systems for user-driven benchmarking of marketed AI offerings, cross-disciplinary and cross-sectorial cooperation, and a European observatory for AI to consolidate these forms of coordination.

Rather than a simple list of 14 priorities, this threefold approach should be understood in connection with a corresponding new model of institutional design. The complexity of current moral and legal issues in AI calls for a **more adequate model of coordination that should be able to accommodate the uncertainties of innovation and, at the same time, to capture expanding understanding with increasing regulatory granularity**. The aim is to formalise specific forms of governance that are neither bottom-up, nor top-down, but in-between, as occurs for example, with the current EU model of governance and co-regulation in the field of data protection – the GDPR.



3. A S.M.A.R.T. MODEL OF GOVERNANCE

Section 3 examines current debates on the governance of AI and its legal regulation, in order to provide an actionable model for the field. The section is divided into five parts. First, focus is on the function of the model, namely, what the latter is supposed to do. Second, attention is drawn to the requirements of the model, i.e. what the model is meant to be. Third, the section examines **how such models of governance work in practice** with the set of provisions and procedures from the EU's 2016 regulation on data protection, the GDPR. Fourth, we illustrate the limits of this co-regulatory model of governance for the field of AI. Fifth, we provide for a more adequate model **in between current forms of co-regulation, e.g. the GDPR, and self-regulation**. The provided model is S.M.A.R.T., i.e. **Scalable, Modular, Adaptable, Reflexive, and Technologically-savvy**; Scalability, Modularity and Adaptability of the model can be suitably addressed by adding resources to the mechanisms set up for coordination, and complemented by further procedures of Reflexivity, i.e. sound scrutiny, assessment and evaluation of risks, and Technologically-savvy oversight. The model fits neatly with the recent EU policies on better and smart regulation, and is consistent with the stance on the rule of law taken by standardisation agencies and governance models in the business field. This convergence is unsurprising. **As technological challenges grow increasingly complex, both top-down and bottom-up approaches will become less fruitful.**

4. A REGULATORY TOOLBOX

Section 4 complements such model of S.M.A.R.T. governance with a **regulatory toolbox**. This entails leaning on some practical mechanisms to bridge the relationship between ethical principles and legal norms, much as the relation between top-down and bottom-up approaches to regulatory toolboxes. The latter should be developed **in accordance with European and national laws, much as case-based law**—i.e. *acquis communautaire*—in the Single Digital Market strategy. A regulatory toolbox can accordingly be used as a key enabler to enhance this legal content, and foster both security and trust. **The toolbox comprises all components for the good AI governance stressed throughout the paper**, namely, (i) principles and legal schemes; (ii) digital innovation hubs; (iii) innovation deals; (iv) compliance through design and by design; (v) regulatory sandboxes; (vi) online and alternative dispute resolution; (vii) regulatory computing architectures; (viii) regulatory ontologies, interoperability & ontology design patterns; (ix) monitoring functions; (x) standardisation, verification and impact assessment mechanisms; (xi) oversight agency.



AI ETHICS AND GOVERNANCE: AN INTERNATIONAL CONCERN

This White Paper has taken in consideration the work of various institutions and governments.

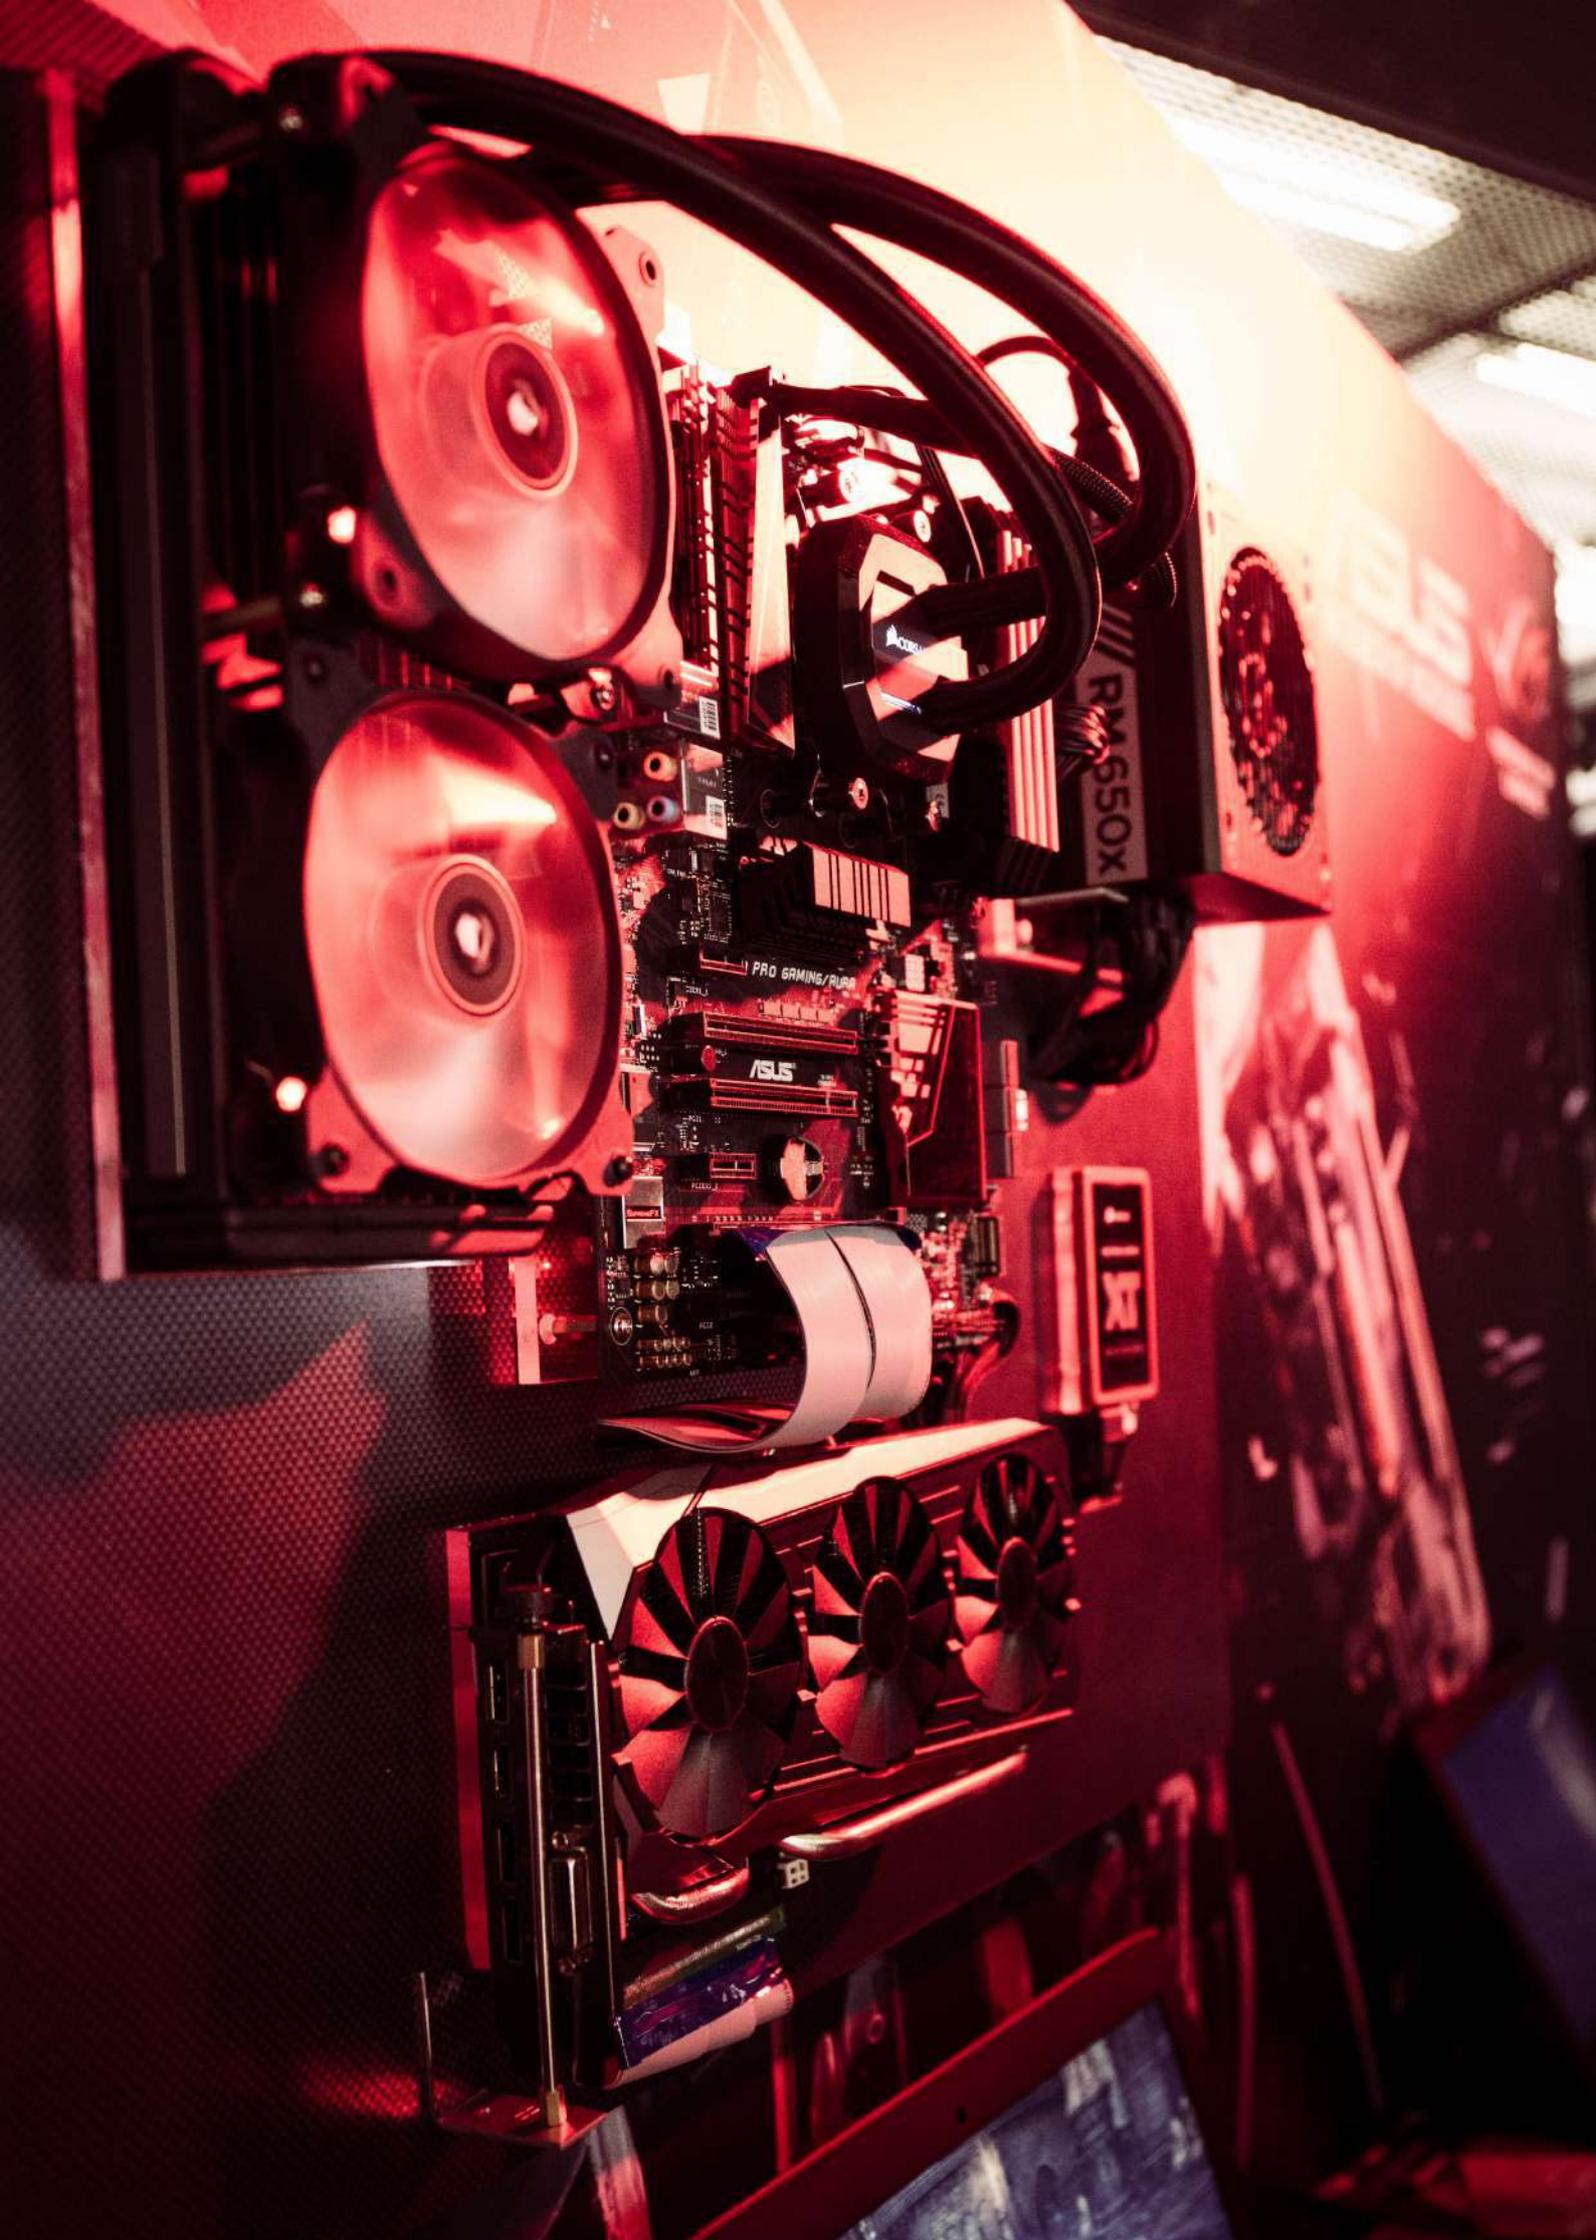
AI ethics and governance have been a recurrent topic of interest and discussion in the past years, with international governments such as the USA, the EU, China and the UK, as well as institutions and organizations such as the World Economic Forum and the United Nations joining the discourse.

Three years ago, in 2016, the White House Office of Science and Technology Policy (OSTP) conducted a series of public workshops on questions of AI and policy, culminating in a report that addressed the many ethical issues in AI, such as fairness, accountability, and social justice, which need to be addressed when aiming to increase transparency. The Trump administration and the Pentagon released similar documents in early 2019. Meanwhile, China declared its ambition to become the AI world leader by 2030, and, in September 2017, Russia's president Vladimir Putin, predicted that the nation becoming the world's leader in AI will be 'the ruler of the world'.

European institutions have been making significant contributions of their own, such as the European Parliament's Resolution from February 2017, the Economic and Social Committee's Opinion on AI from May 2017, the European Commission's AI Strategy and the Work of the High-Level Expert Groups on AI in 2019, down to the work of the Council of Europe. More recently, in July 2019, the European Parliament nominated the first woman to be President-elect of the European Commission, Ursula von der Leyen who released her own political guidelines¹. Remarkably, in the third section of this document on 'A Europe fit for the digital age,' special attention is drawn to the challenges and opportunities brought about AI. In the wording of the President-elect, 'it may be too late to replicate hyperscalers, but it is not too late to achieve technological sovereignty in some critical technology areas.' In addition to a new digital services act and a joint cyber unit, the aim of 'my first 100 days in office [is to] put forward legislation for a coordinated European approach on the human and ethical implications of Artificial Intelligence.'

To this work in progress can be added further initiatives. The list includes the United Nations' 'AI for Good Global Summits', OECD Reports, and numerous partnerships and institutional proposals at the international level ranging from the Global Initiative on Ethics of Autonomous and Intelligent Systems of the Institute of Electrical and Electronics Engineers (IEEE), the Committee on Professional Ethics and Public Policy Council of the Association for Computing Machinery (ACM), the World Economic Forum's Centre for the Fourth Industrial Revolution, the Future of Life Institute's Asilomar Principles, OpenAI, the Partnership on AI, the Software and Information Industry Association (SIIS), and the work of the AI4People project (Floridi et al. 2018).

¹https://ec.europa.eu/commission/sites/beta-political/files/political-guidelines-next-commission_en.pdf



1

DEFINITIONS

The wider framework of both the formal and informal rules that govern the public arena is commonly summed up under the term 'governance.' For example, in the words of the World Bank, governance refers to 'the process and institutions through which decisions are made and authority in a country is exercised' (Grindle 2007). Others define governance as 'the formation and stewardship of the formal and informal rules that regulate the public realm, the arena in which states as well as economic and societal actors interact to make decisions' (Grindle 2007).

Thus, a crucial feature of governance has to do with different levels of legal regulation and how they may relate to ethics and moral theories.

1.1

THREE LEVELS OF REGULATION: LEGAL REGULATION,
CO-REGULATION, AND SELF-REGULATION

Different levels of legal regulation are traditionally presented as a tripartite classification:

- i. **Legal regulation** – associated with a top-down approach - can be understood as a set of rules or instructions for the determination of every legal subject of a system. These are the rules that aim to directly govern social and individual behaviour, and mainly hinge on the threat of physical (and financial) sanctions as a means of social control (Kelsen 1949);
- ii. **Self-regulation** – associated with a bottom-up approach - refers - in this paper - to any kind of bottom-up approach with limited accountability and legal framing. The AVMSD describes this kind of regulation as 'a type of voluntary initiative which enables economic operators, social partners, non-governmental organisations and associations to adopt common guidelines amongst themselves and for themselves.' A good example is the Pan-European Game Information (PEGI) system, which is the only harmonised ratings system for digital content available in Europe. Self-regulation-originated ratings are here combined with technology-driven parental control tools built into game consoles, affording parents a fool-proof shield to protect children – while delivering on the game industry's leading purpose of setting up an ambitious model to match the unprecedented 'moral panic' running rampant at the time. As PEGI exemplifies, self-regulation can indeed be effective. It can, however, lack resilience, especially at times of greater societal division on what are the problems - and what are the possible solutions – for the governance of AI. Too often, 'self-regulation' is taken to define a single and uniform set of practices: as a matter of fact, there is a rich variety of practices,



some even with a long history, in addition to several theoretical frameworks for self- and co-regulation, either under the latter terms, or under - for example - the terms of the ‘Ethical Business Regulation’ a model of co-regulation – based on behavioural science and ethics - which proposes an effective co-operation between business practitioners and regulators. (Hodges and Steinholtz 2018).

iii. Co-regulation – a type of “middle-out” approach - refers to how legal regulation and self-regulation interact. According to Recital 44 of the 2010 Audiovisual Media Services Directive (AVMSD), namely, Directive 2010/13/EU, co-regulation provides in its minimal form ‘a legal link between self-regulation and the national legislator... In co-regulation, the regulatory role is shared between stakeholders and the government or the national regulatory authorities or bodies’. A good recent example of a legally framed but adaptive and agile framework is the Regulation 2019/1150 on platform-to-business (P2B) trading practices, where hard law is the basis for corporate codes of conduct and dispute resolution mechanisms, under public sector surveillance and analysis of outcomes, with updates of the hard law requirements as a possible step at any time.

In between legal regulation and self-regulation, scholars have provided different kinds of scales and taxonomies – and different kinds of co-regulations, all aiming to understand how top-down forms of legal regulation and bottom-up solutions can be mixed (Tirole 2018). Within this broad mix, ethics and law can interact in a variety of ways; we explore these kinds of interaction in the next section.

1.2

LAW AND ETHICS: THREE KINDS OF INTERACTION

In light of current declarations and guidelines on the principles of AI, such as the Asilomar AI Principles¹, the Montreal Declaration for Responsible AI², the Tenets of the Partnership on AI³, the AI HLEG's *Ethics Guidelines for Trustworthy AI*⁴ or the recent OECD Principles on AI⁵, we may wonder how these initiatives fit different forms of legal regulation and governance. Three options are at hand:

First, the ethical principles of such declarations and guidelines can be understood as the principles that will be embraced and enforced through the top-down tools of hard law. However, some of these ethical principles—the principle of transparency, for example—may clash with current provisions of the law on, say, the protection of intellectual property and national security. How should we deal with multiple jurisdictions and different legal traditions? Should we consider the principles of today's declarations

¹Asilomar AI Principles, available at: <https://futureoflife.org/ai-principles/>

²Montreal Declaration for Responsible AI, available at: https://docs.wixstatic.com/ugd/ebc3a3_c5c1c196fc164756afb-92466c081d7ae.pdf

³Tenets of the Partnership on AI, available at: <https://www.partnershiponai.org/tenets/>

⁴AI HLEG's Ethics Guidelines for Trustworthy AI, available at: <https://ec.europa.eu/digital-single-market/en/news/ethics-guidelines-trustworthy-ai>

⁵OECD Principles on AI, available at: <https://www.oecd.org/going-digital/ai/principles/>



as being on a higher rank than current legal regulations?

A **second way** to grasp the interaction between the ethical principles of current guidelines and today's legislation is to interpret their relation as such that no conflict would follow as a result of a further work of clarification. However, should we distinguish, according to the jargon of international lawyers, between absolute principles, which admit no exception, and relative principles, which protection should be balanced against further principles of the system? How should we strike such a balance, and what room is left for self-regulation, in this second scenario?

Third, we could imagine that economic operators, social partners, non-governmental organisations and associations determine their own moral values, and that such values go hand-in-hand with the moral principles of guidelines and declarations. As we underlined in the principles published in last year's White Paper - *AI4People's Ethical Framework for a Good AI Society* - the 21st century sees a deepening of societal values in particular around the fundamentals of human flourishing for all individuals in all their experiences (as citizens, workers, consumers, members of family and community), as well as around the uncertain future of the planet and the changing world of work. It follows from this that no ethical check-list can be regarded as closed, and any list will require adaptation to lived realities. This invites a deeper reflection, too, as to what kind of governance and what kind of legal regulation would fit such a period of change? How should we set the bar between different forms of co-regulation and self-regulation? Are these regulatory models of governance related to the difference between absolute principles and relative principles? If so, how might we formalise such a correlation?

Against this threefold option, we need to debunk a popular misconception in the debate on AI Governance. Going back to the initiatives of the European Commission on AI, some have interpreted the work of the HLEG on the ethics of AI as if the aim were to identify the (moral) basis for legal regulation.

The discussions of such experts, however, are not about what should or should not be done against current legislation, or despite it. Rather, the guidelines revolve around how to complement and strengthen the existing regulation. In the phrasing of the *Ethics Guidelines for Trustworthy AI* from April 2019, 'trustworthy AI should respect all applicable laws and regulations, as well as a series of requirements.'

All in all, a panoply of legal rules and conventions already regulates the multifaceted world of AI. This does not mean there are no legal loopholes, but such rules and conventions should indeed be taken into account, following a threefold perspective:

First, attention should be drawn to the boundaries established by the principles of both the ECHR and the EU Charter of Fundamental Rights (CFR). The current work of the HLEGs deals with the fields of applied ethics (i.e. HLEG on AI and ethics); and with secondary - as opposed to primary - legislation (i.e. the two HLEGs on AI and the product liability directive and responsibility for new technologies formation).



Second, multiple hard law provisions already govern several aspects of AI innovation in Europe today, such as the processing of personal data through AI systems. Every model of AI governance must therefore accommodate this body of existing law.

Third, many of these legal provisions, as well as several ethical principles of AI, are context-dependent (see separate box for more info). These top-down features of current regulation do not suggest that the corresponding model of governance should exclude bottom-up forms of self-regulation, or variants in between, such as co-regulation. Relevant to this, the EU model of co-regulation with the General Data Protection Regulation, for example, is norm-rich, but offers a list of six enforceable principles in Art. 5(1): the principles of (i) lawfulness, fairness, and transparency; (ii) purpose limitation; (iii) data minimisation; (iv) accuracy; (v) storage limitation; and, (vi) integrity and confidentiality. These principles are available as a framework for the developers of, for example, binding corporate rules within GDPR.

However, it is important to keep the complexity of today's legal framework from leading to another kind of misconception. The authors of this paper admit that AI raises a whole set of new moral and legal challenges, such as the possibly reducing human control, removing human responsibility, devaluing human skills, eroding human self-determination, and enabling human wrongdoing (Floridi et al. 2018). Such challenges are real, but they are emergent, so that they are twice as difficult to tackle.



CONTEXT - DEPENDENCY

Several current rules on AI are context-dependent. In addition to the rules on data protection in the EU, for example, consider the current legal regulation in the fields of self-driving cars, drones, e-health, financial services, and more. The testing and use of self-driving cars in the EU, for example, has to contend with a complex legal network involving three directives and one regulation¹. The testing and use of drones hinge on one regulation, two European Commission's implementing and delegated acts, in addition to several Opinions and guidelines of the European Aviation Safety Agency (EASA)². Whilst the e-health sector raises several issues of contractual and tortious liability that mostly concern national regulations of the EU member states, 'sandboxing' (a security mechanism for separating running programs, usually in an effort to mitigate system failures or vulnerabilities) has been a popular mechanism for governments and technology to collaborate, in particular, in the financial sector, and address current legal loopholes at both national and EU level³. On top of this complex network of legal provisions, current debates and initiatives on both the ethics and law of AI, such as the three HLEGs set up by the EU Commission in 2018, have to be attentive to a long-standing tradition in Europe, which is defined by the 1950 ECHR and the 2000 EU Charter of Fundamental Rights⁴. This common ground makes it clear why the HLEGs on AI & ethics and on AI & the law have worked independently over the past year.

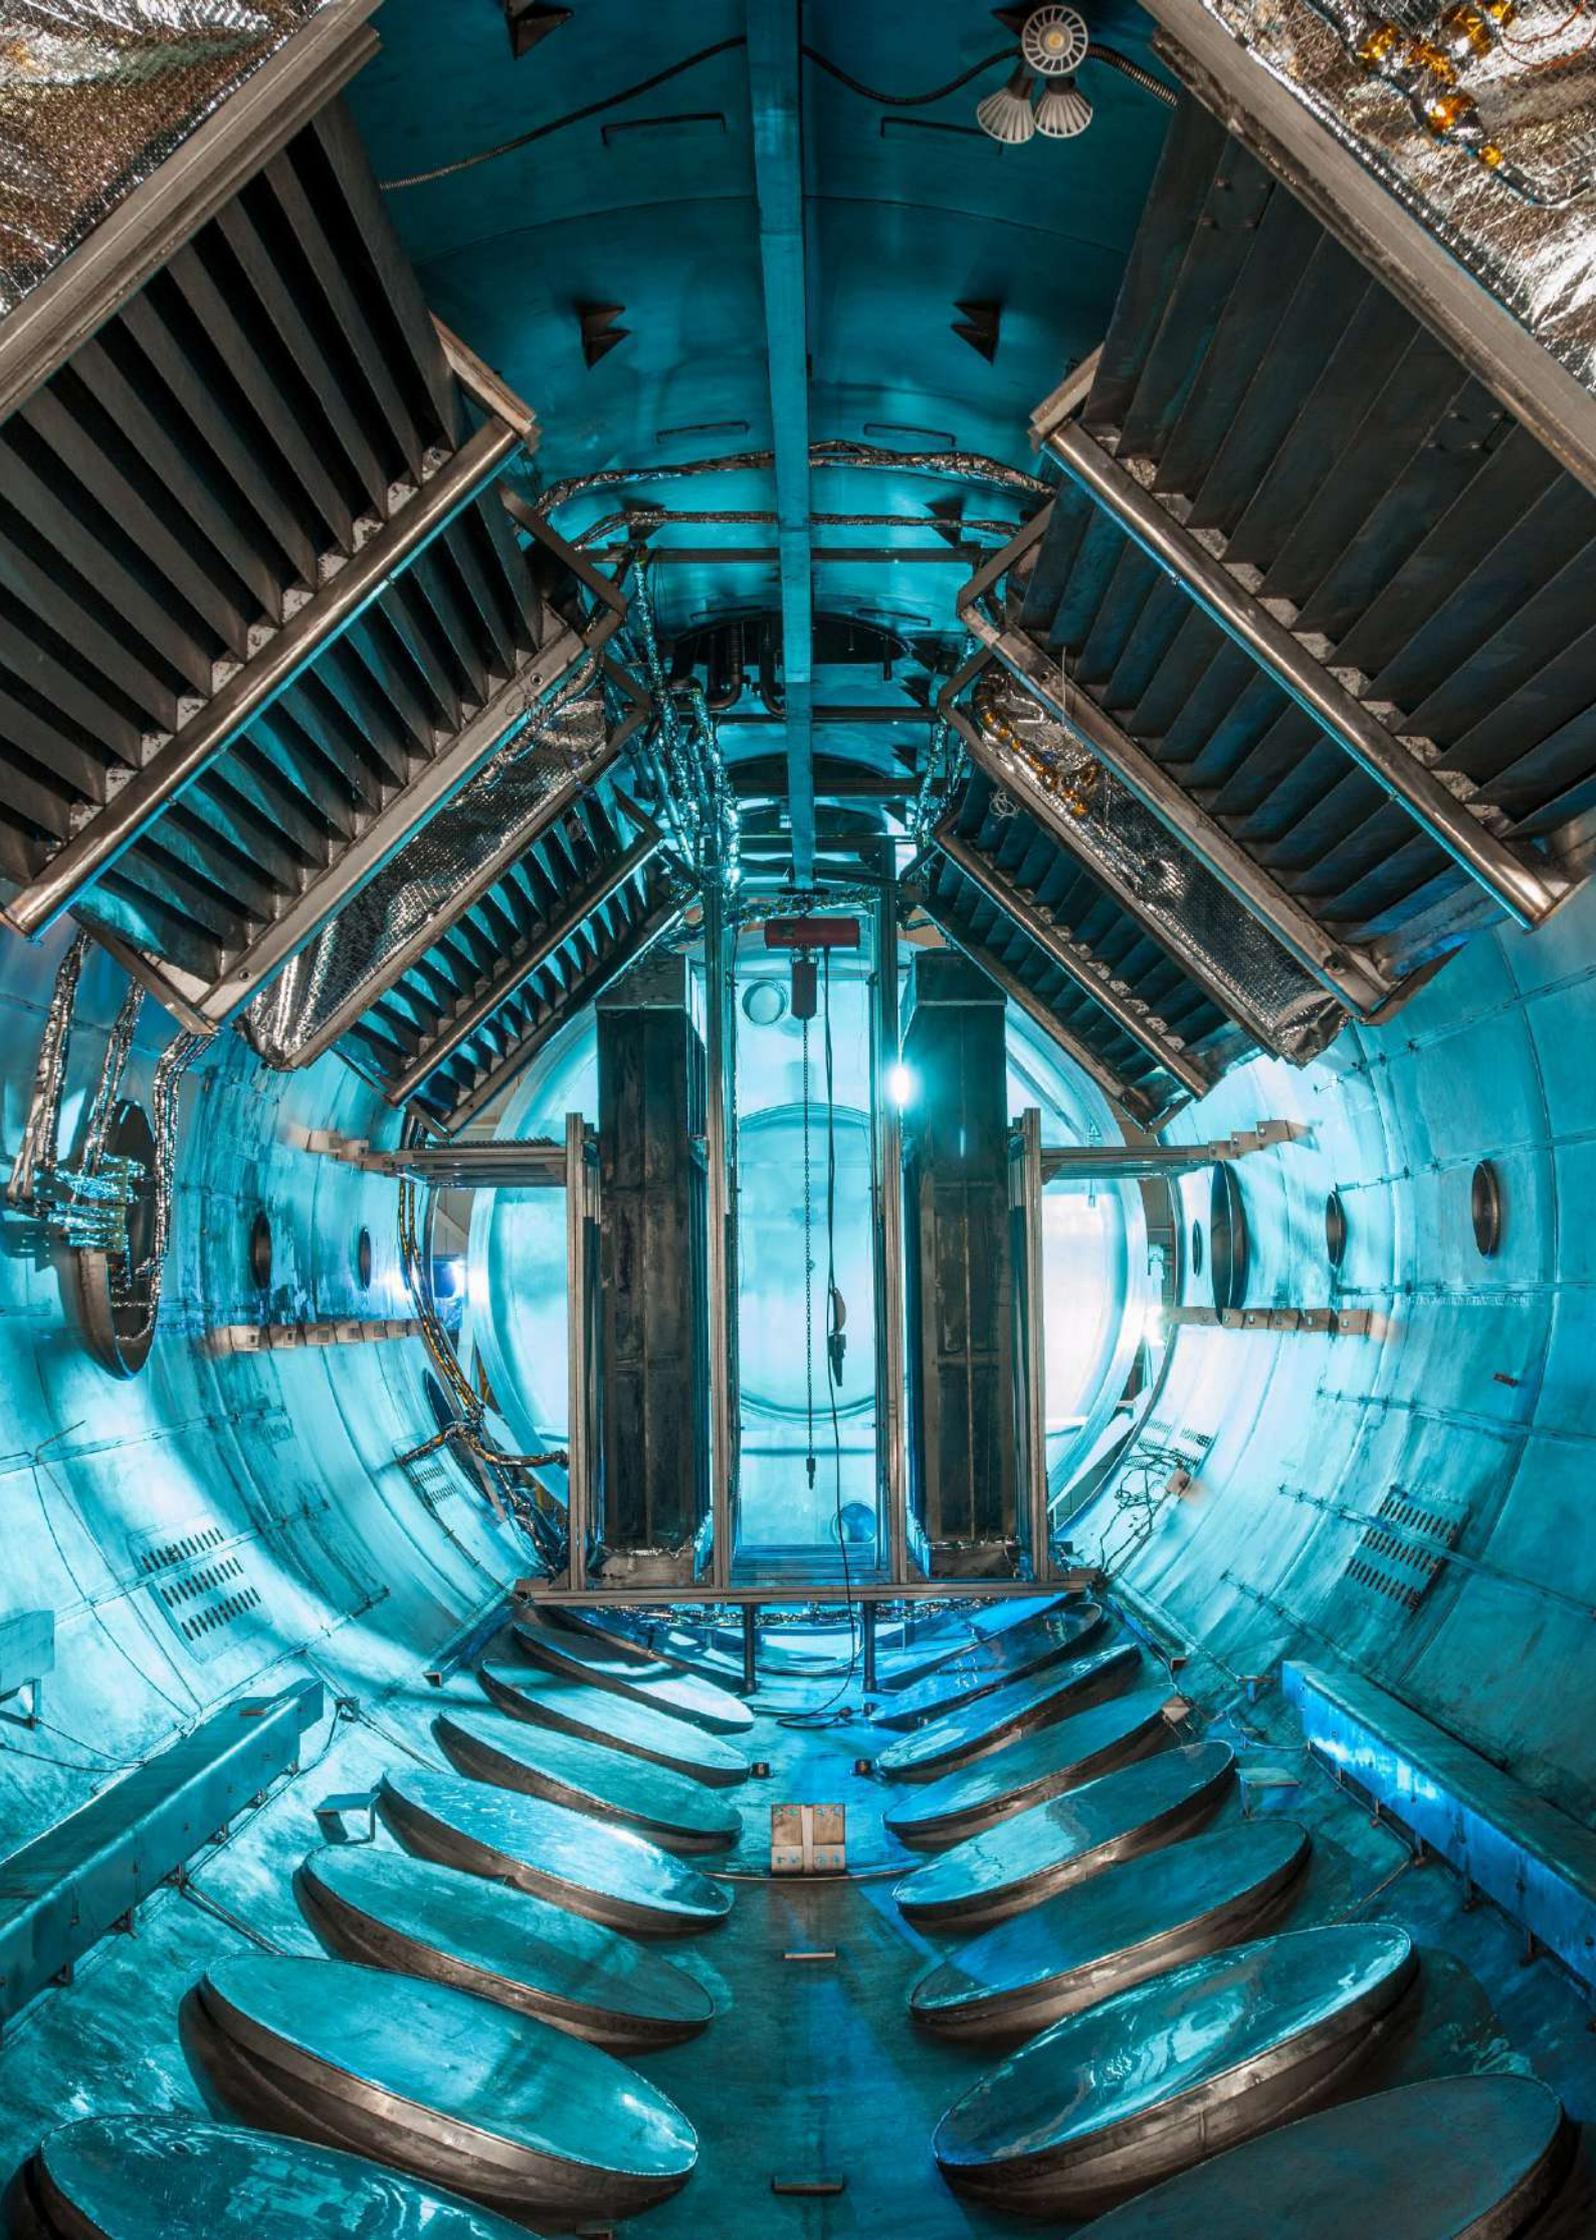
¹The list concerns, (i) Council Directive 85/374/EEC on the approximation of the laws, regulations and administrative provisions of the Member States concerning liability for defective products; (ii) Directive 1999/44/EC on certain aspects, such as repair and replacement, price reduction and termination, of the sale of consumer goods and associated guarantee; (iii) Directive 2009/103/EC relating to insurance against civil liability in respect of the use of motor vehicles, and the enforcement of the obligation to insure against such liability; (iv) Regulation 2018/858 on the approval and market surveillance of motor vehicles and their trailers, and of systems, components and separate technical units intended for such vehicles.

²See Reg. (EU) 2018/1139 on common rules in the field of civil aviation and establishing a European Aviation Safety Agency. On this basis, the European Commission has adopted so far Delegated Regulation 2019/945 and the Implementing Regulation 2019/947.

³See for instance at a national level "Regulatory Sandbox", *FCA Website*, 14 February 2018, <https://www.fca.org.uk/firms/regulatory-sandbox>, and "FinTech Regulatory Sandbox", *Moneyart Authority of Singapore Website*, 1 September 2017, <http://www.mas.gov.sg/Singapore-Financial-Centre/Smart-Financial-Centre/FinTech-Regulatory-Sandbox.aspx>.

⁴On 25 April 2018, the EU Commission set up three different groups of experts on (i) the ethics of AI; (ii) whether and to what extent to amend the directive on liability for defective products; and, (iii) liability and new technologies formation. See the Commission's document on 'Artificial Intelligence: a European approach to boost investment and set ethical guidelines' (doc. IP/18/3362).





14 PRIORITY ACTIONS

This section of the White Paper outlines 14 Priority Actions for good AI governance.

We intend as “Priority Actions” that which we can reasonably conceive as good, right, or lawful (Floridi et al. 2018). Most importantly, we focus on those actions which are immediately actionable and applicable in real world scenarios.

A different – and less immediate - kind of priority would be that which we can reasonably judge to be good, right, or lawful - but would require some time to implement. As an illustration of this second-step priority, consider for the sake of argument the 2019 list of key findings and recommendations by the EU’s HLEG on liability for new technologies¹. The list includes principles of law-making and legal personality, mechanisms on burdens of proof for causation and fault, new standards and duties of care, product and strict liability, compensation funds, and more. Even if we accepted all of HLEG’s findings and proposals, the regulatory process for the amendment of the EU legal system would take some time -around three or four years. It’s unclear what should we do in the meantime.

The set of actions, mechanisms, and procedures we advocate in this section draws on our previous work, *The AI4People’s Ethical Framework for a Good AI Society* (Floridi et al. 2018). This ethical framework has gained increasing traction in today’s debate, with some of our recommendations and action points being advanced in other documents and initiatives, such as those by the HLEGs on ethics & AI.

Some of these recommendations and action points are presented here in accordance with the three types of legal regulation discussed above, i.e.:

- i. **Top-down, or traditional legal regulation**, which includes both hard law and soft law tools, such as the Opinions of the Art. 29 Working Party and, nowadays, of the European Data Protection Board (EDPB) in the field of data protection;
- ii. **Bottom-up**, or forms of **self-regulation** in all its variants;
- iii. **“Middle-out”**, that is, everything that lies in between the top-down and bottom-up approaches. As pointed out in Section 1, the “middle-out” stance incorporates both elements of top-down legal framing and bottom-up empowerment of individual actors. More importantly, it sheds light on some of the governance models adopted by the EU, such as the EU Better Regulation scheme for interoperability (TOGAF 2017), and the current European Commission Better Regulation Toolbox 1, Principles, Procedures & Exceptions².

Against this backdrop, we can now group different sets of priorities for AI under the heading of institutional design (Fig. 1). Section 2.1 focuses on five no-regrets actions that correspond to top-down governance. Section 2.2 illustrates the bottom-up approach with four different forms of engagement, such as new types of forums, for example. Section 2.3

¹See above note 4.

²European Commission, Better Regulation Toolbox 1, Principles, Procedures & Exceptions, available at: https://ec.europa.eu/info/sites/info/files/file_import/better-regulation-toolbox-1_en_0.pdf



presents five different forms of coordination mechanisms to demonstrate the middle-out approach.

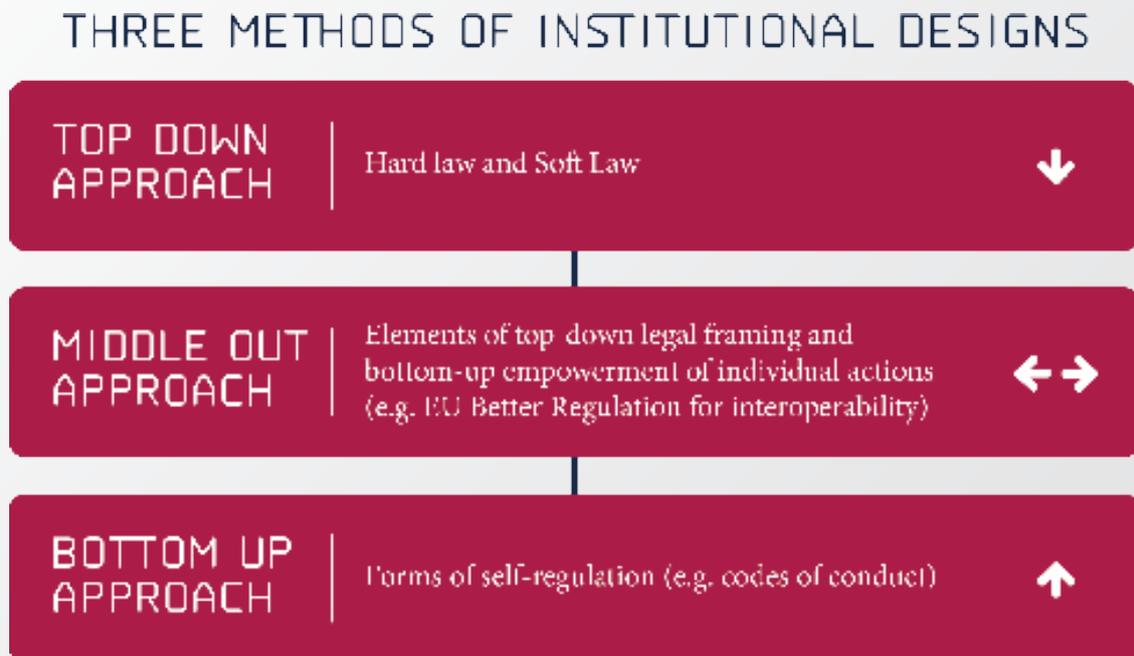


Figure 1: Setting the Priorities for AI.

2.1 NO-REGRETS ACTIONS – A TOP-DOWN APPROACH

The concept of ‘no-regrets action’ was first popularised in the field of environmental policies and legal regulation in the 1990s. As highlighted in our 2018 paper *AI4People’s Ethical Framework for a Good AI Society*, ‘of all areas of applied ethics, bioethics is the one that most closely resembles digital ethics in dealing ecologically with new forms of agents, patients, and environments’ (Floridi et al. 2018). In addition to the set of core principles commonly used in bioethics, such as beneficence, non-maleficence, autonomy and justice, we included a fifth principle of explicability, in order to properly tackle the normative challenges of AI. Explicability should incorporate both intelligibility of AI - as well as accountability for its uses.

As in the field of environmental law, however, those responsible for policy making and governance in the field of AI must couple precaution with no-regrets actions. In the jargon of the international community³, such no-regrets actions concern the achievement of sustainable development goals, such as capacity building - the process by which individuals and organizations improve their skills and knowledge - in a good AI society. There are five different kinds of no-regrets actions, which we already examined in our 2018 White Paper. These sets of actions correspond to top-down governance and to both forms of soft and hard law:

³See the ITU’s AI for Good Global Summit initiative, for example: <https://aiforgood.itu.int/>



- i. The EU's expectation of developing national AI strategies by the Member States, with the positive vision that Europe can be a global leader in the field, and that the field can be utilized to support EU's goals and human flourishing;
- ii. The creation of educational curricula and public awareness activities around the impact of AI, involving schools, academia, qualification programmes in business and the public at large;
- iii. A sustained, increased and coherent European research effort, which provides for the integration of ethical, legal and social considerations in AI research projects, together with research about public perception and understanding of AI and its applications (an extensive and international public consultation of the public's opinion on AI will be conducted in 2020 by ECAI⁴);
- iv. The idea of 'inclusive innovation' and smooth transition to new kinds of jobs via rewarding human-machine collaboration, while preserving not only a human-centric design but also a design able to keep human-human interaction going, for instance, by way of a European AI Adjustment Fund taking cues from the current European Globalisation Adjustment Fund;
- v. Support for the capacity of corporate boards of directors to take responsibility for the ethical implications of companies' AI technologies, both within the firm and across the corporate footprint.

This non-exhaustive set of actions, could indeed increase in accordance with the growing convergence of today's debate on the principles, both moral and legal, of AI. What is relevant here, however, does not concern whether we have prioritised all possible types of no-regrets actions. Rather, attention is drawn to that which can be put into action immediately.

2.2 FORMS OF ENGAGEMENT – A BOTTOM-UP APPROACH

The bottom-up approach can be illustrated with new forums for collective consultation and discussion, such as 'designing-by-debate' (Ausloos et al. 2018), or 'cooperative responsibility' (Helberger et al. 2018). The currently limited understanding of the stakes of AI favours these kinds of initiatives. Four different types of engagement are of particular note in this context:

- i. The use of **participatory mechanisms** to ensure alignment with societal values and understanding of public opinion through an ongoing dialogue between all stakeholders;
- ii. A proposed **European observatory for AI** to provide a forum to: nurture debate, consensus and awareness activities; to drive the development of real-time vigilance

⁴The European Consultation on Artificial Intelligence (ECAI) -coordinated by Atomium-EISMD - was set up to organize an extensive and first of its kind consultation of the public-at-large on Artificial Intelligence, across twelve European Countries. Engaging citizens in this crucial debate - and collecting their opinions on one of the most relevant subject matters of our time - is an utmost priority in order to create policies considerate of the public's opinion on AI, its Ethics and its Governance. The results of this extensive and international public consultation will be presented by AI4People in July 2020.



- and monitoring tools, expert assessment of incoming data on AI deployment; and to deliberate as to whether the resulting new knowledge calls for fresh deliberation and action;
- iii. A scheduled **cross-disciplinary and cross-sectorial cooperation** and debate around the intersections between technology, social issues, legal studies, and ethics;
 - iv. The interaction between scientists and laymen in **legally deregulated special zones**, or living labs, for AI empirical testing and development. Here, researchers and the larger public can understand whether AI systems fulfil their task specifications in ways that are acceptable and comfortable to humans, while increasing our understanding of how the future of the human-AI interaction could turn out. This form of engagement should be complemented with the development of a framework, in which individuals have the ability 'to obtain a factual, direct, and clear explanation of the decision-making process, especially in the event of unwanted consequences' (Floridi et al. 2018).

These bottom-up forms of engagement are indeed not the only ones possible. But they are ones that can be immediately - and concretely - realised.

2.3 COORDINATION MECHANISMS – A “MIDDLE-OUT” APPROACH

Between the bottom-up solutions and top-down approaches, a series of coordination mechanisms represent the middle layer of our priorities. These function as a sort of interface between the aforementioned set of no-regrets actions and four different forms of engagement. These coordination mechanisms – or “middle-out” approach -include forms of participatory and multi-stakeholder procedures, such as the aforementioned European observatory for AI, and communications networks or platforms made available to the public for sharing experiments led in a variety of sectors or places. The lessons learnt from these specific experiences about the direct impact of AI deployment or its possible effects on the social side may provide helpful inspiration for other experiments, and build not only a unique database on AI-related developments, but also a space wherein to engage with the teams conducting these experiments. More particularly, the coordination mechanisms relate to:

- i. Participatory procedures for the alignment of societal values and understanding of public opinion;
- ii. Upstream multi-stakeholder mechanisms for risk mitigation, e.g. unwanted consequences of human-AI interaction;
- iii. Systems for user-driven benchmarking of marketed AI offerings, allowing trust in products and services, as well as providers to be measured and made transparent;
- iv. Cross-disciplinary and cross-sectorial cooperation and debate incentivisation;



- v. A European observatory for AI to consolidate these forms of coordination, in addition to the forms of engagement stressed above in Section 2.2.

Such coordination mechanisms are crucial, since they can help us tackle current limits on any clear understanding of the stakes of AI, while consolidating new forums for collective consultation and discussion and developing new standards and mechanisms for good AI regulation. This coordination will be especially important in ensuring there is as close a consensus as feasible between local regulators and administrations, so that innovation can enjoy the benefits of a cross-EU level playing field. Among desirable areas for standards, we would identify the development of new social standards to address societal concerns in environmental and other fields, in addition to technological standards, such as the IEEE's P7000 standards series⁵, or legal standards. One of the key findings of the HLEG's 2019 document on liability for new technologies concerns the development of these new social standards with duties of care.

⁵The IEEE P7000 series includes to date 14 standards under development in the framework of the IEEE Global Initiative on Ethics of Autonomous and Intelligent Systems, <https://ethicsinaction.ieee.org>





A S.M.A.R.T. MODEL OF GOVERNANCE

The complexity of – as well as the time needed to create or update today's legal regulation - has gradually revealed the limits of traditional top-down approaches. Trust in indirect and less accountable forms of self-regulation has eroded of late, too. Scholars and lawmakers have progressively focused on that which lies in between these top-down and bottom-up approaches, and how could such top-down forms of legal regulation, e.g. the hard accountability of legislation, be blended with the agility and adaptability typical of the more effective self-regulatory options.

In the EU legal framework, this middle-out layer is mostly associated with forms of co-regulation, as defined above in section 1, pursuant to Recital 44 of the 2010 AVMSD. Such governance and regulatory options concern the balance between multiple regulatory systems and how different kinds of rules are combined. Depending on where we set the bar between the top-down and bottom-up approaches, different models of governance and legal regulation follow as result (Fig. 2) (Pagallo et al. 2019).

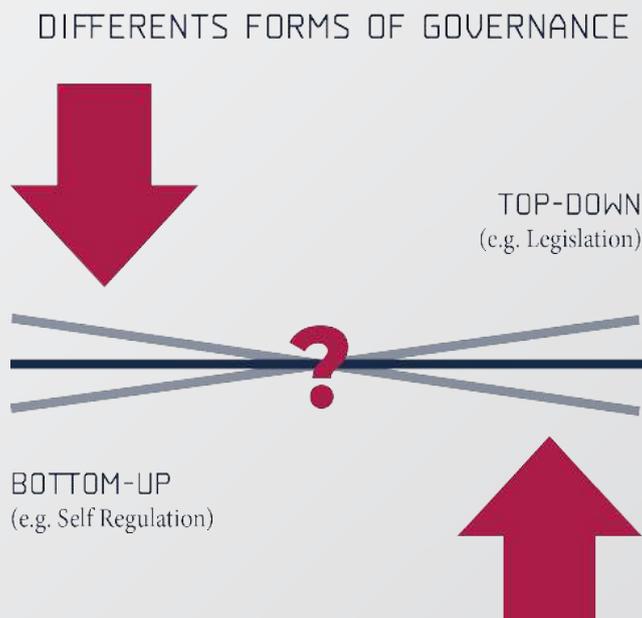


Figure 2: Setting the Level of Abstraction.

The analysis of this section is divided into five parts, and aims to present a new model for the three sets of priorities illustrated in the previous section (2.1 No-Regrets Actions - a Top-Down Approach; 2.2 Forms of Engagements - a Bottom-Up Approach and 2.3 Coordination Mechanisms - a "Middle-Out" Approach)

Section 3.1 examines what the model is called on to do, namely, the function of the middle-out approach. Section 3.2 sheds light on its requirements, that is, what the interface of the model is supposed to be vis-à-vis with the constraints of institutional design. Section 3.3 illustrates how this middle-out approach works in practice with the set of provisions and procedures of the EU's 2016 regulation on data protection, the GDPR. Section 3.4 sheds light on the reasons why the GDPR's model of legal regulation offers no magic bullet for today's governance of AI. We therefore need to imagine new "middle-out" approaches for the field of AI; Section 3.5 presents one such new S.M.A.R.T. model (Scalable, Modular, Adaptable, Reflexive, and Technologically-savvy).

3.1 FUNCTION OF THE MODEL

There are multiple regulatory systems out there. In addition to the law, consider the role of ethics, the forces of the market, or of social norms as well as the normative constraints of technology. Every regulatory system aims to govern individual and social behaviour through its own means. Regulatory systems may thus compete, or indeed render the claim of another regulatory system superfluous, or in other cases strengthen each other. The function of a model concerns what such model is called on to do.

The “middle-out” approach draws attention on how the balance is struck between multiple competing regulatory systems. New forms of co-regulation, as defined in the 2010 AVMSD, should prevent the risk of both self-regulatory fragmentation and legal failure, as happened with the EU e-money directive 46 from 2000. Soon after its implementation, new forms of payment, such as PayPal, forced the legislators in Brussels to intervene, amending themselves with a new directive nine years later (Directive n. 110 from 2009).

Thus the model of governance for AI has to strike a balance between competing regulatory systems. The “middle-out” approach specifies what balance is remaining, once forms of top-down legislation and pure bottom-up solutions are discarded.

3.2 REQUIREMENTS OF THE MODEL

The remaining balance between multiple regulatory systems emerges as a mix of different kinds of legal rules. In addition to the distinction between principles (e.g. those of the 1950 ECHR and the 2000 EU Charter of Fundamental Rights) and rules (e.g. those specific norms in the fields of self-driving cars, drones, e-health, financial services etc. existing within the boundaries established by the principles) (e.g. Dworkin 1985), we have to distinguish between primary and secondary rules of the law (Hart 1961):

The primary rules of the law are an illustration of the top-down approaches that aim to directly govern individual and social behaviour also (but not only) through the threat of physical and financial sanctions.

The secondary rules of the law include rules of recognition, of adjudication, and of change, i.e. rules that allow the creation, modification, and suppression of the primary rules on people's conduct (Pagallo 2017a). The role and function of the secondary rules of the law and, more specifically, the rules of adjudication and change, can hardly be overestimated. The complexity of technological regulation has increasingly recommended the adoption of this kind of rule (Pagallo et al. 2019). For example, in the case of the EU's GDPR, there are four different types of secondary rules, namely, (i) mechanisms of delegation of power; (ii) mechanisms of legal coordination; (iii) procedures for a preemptive approach to data protection; and, (iv) procedures for effective judicial remedies (Pagallo 2017a)

Regulatory options of governance and multiple scales of self- and co-regulation can



thus be properly represented as a mix of primary and secondary rules, namely, norms that aim to govern social and individual behaviour (primary rules), and the secondary rules that set up judicial remedies, mechanisms of legal coordination and of delegation of power and the norms allowing amending the primary rules of the law. This mix determines how the resulting eco-system establishes its own interaction with further regulatory systems, such as the technology industry, as an example.

Out of all the possible combinations and taxonomies among the different regulatory systems and mixes of primary and secondary rules of the law, the model proposes a “middle-out” functionality: an intersection between top-down (hard law) and bottom-up options, coming to life as a network of rules that strikes the balance between technology, ethics, market, and social norms.

3.3 THE GDPR MODEL

An example of a “middle-out” approach is a new model of legal governance that EU lawmakers have adopted in the field of data protection, the GDPR. According to the “middle-out” functionality of this model, top-down legislation establishes the principles that must be followed by data controllers - as well as the goals that they should strive for - though it is in large part up to data controllers how they should attain such outcomes, under the supervision of public guardians.

This mixed approach to governance revolves around what the GDPR calls the ‘principle of accountability.’ Recital 85 of the GDPR refers to accountability in connection with the responsibility of data breaches as well as the risks of the data subjects’ rights and freedoms. Article 5(2) mentions the ‘principle of accountability’ as the way in which data controllers shall be able to prove compliance with the six ‘principles relating to the processing of personal data’ pursuant to Art. 5(1). It is then up to data controllers to prove compliance with this six-fold set of principles, by assessing the level of risk triggered by their own data processing, in accordance with the provisions of Articles 24(1) and 25(1).

In particular, data controllers shall define how to approach the prevention of risk through technological and corporate organisational measures for the protection and security of personal data processing. In addition, Art. 32 sets up further duties and self-regulatory measures of data controllers, in regards to the ‘security of processing,’ ‘the pseudonymisation and encryption of personal data’; ‘the ability to ensure the ongoing confidentiality, integrity, availability and resilience of processing systems and services’; and so forth.

The aim of the legislator is not only to set the conditions for the fair processing of personal data. The intent is also to prevent possible risks brought about by inadequate systems and services of data processing.

This complex interaction of principles and legal rules can be illustrated with another figure (Fig.3). This sheds light on the GDPR's principle of accountability as the middle-out layer of the model. Art. 5(2) of the GDPR sets the legal interface between



top-down principles and legal rules of Art. 5(1), or Art. 12 on transparency, and the bottom-up organisational and technical measures of Art. 24(1), 25(1) and 32 (see Fig. 3) (Pagallo et al. 2019).



Figure 3: The GDPR's middle-out approach.

The GDPR's model of legal regulation should thus be complemented by a number of secondary rules. These frame the dynamics of institutional relations. Understandably, many such secondary rules are rules of adjudication, such as in Articles 77 ff. of the GDPR. Risks of fragmentation that depend on multiple jurisdictions of national supervisory authorities are tackled with further secondary rules on coordination and procedural regularity, as in Articles 60, 61, 75(4) and 97(2)(b). The model comprises a more complex mix of primary and secondary rules for some specific cases, such as impact assessments and the protection of corporate rights, i.e., Articles 36 and 80 of the GDPR.

Admittedly, the primary and secondary rules of the GDPR have ignited a hot debate on whether or not they provide a good model of legal governance, for example with the "territorial scope" of Art. 3 of the GDPR which also applies to non-EU companies that target both EU organizations and EU residents by profiling or proposing products or services.

In this context, the focus is on whether the accountability principle can offer the foundation for a middle-out approach to today's AI governance. Although the GDPR is already valid law for AI processing of personal data, it is an open question whether this model of regulation can, or should, be transplanted into the entire domain of AI.

3.4 THE LIMITS OF THE GDPR MODEL

The GDPR represents an intermediate system (or "middle-out" approach), somewhere between the pure bottom-up and top-down approaches, albeit one step closer



to top-down - and detailed norming - than to more agile 'adaptive' self-regulations exemplified, for example, in the recently agreed Regulation on platform-to-business transparency. The GDPR's 'principle of accountability' strikes a balance between the regulatory claims of the forces of the market, of social norms, and technology, through a mix of primary and secondary legal rules. The mechanism is set up to guarantee compliance with both the principles and the top-down rules of the system, while leaving room for self-regulatory measures, both technical and organisational, on the part of economic operators. Whether and to what extent this model can adequately deal with the challenges of data protection remains an open issue. Yet, it seems fair to admit, this model hardly fits the challenges of AI (Pagallo et al. 2019). This is for a variety of reasons:

First, we lack a list of principles to be enforced through forms of co-regulation in all fields of AI, as occurs with Art. 5(1) of the GDPR.

Second, the legal regulation of AI concerns not only personal data issues but also the protection of other fundamental rights and non-discrimination law.

Third, current AI regulation is already context-dependent (see box on page 14): in addition to the rules on data protection, there are a plethora of rules in the fields of self-driving cars (Lütge 2017), drones (Bassi 2019), e-health (Blasimme and Vayena 2019), financial services (Turner 2019), and so on (Barfield and Pagallo 2019).

Fourth, self-regulatory measures for AI corporate governance are still in the early phases and need to be strengthened in accordance with the fifth of our no regrets-actions outlined in section 2.1 (i.e. Support for the capacity of corporate boards of directors to take responsibility for the ethical implications of companies' AI technologies, both within the firm and across the corporate footprint.)

As a result, we need another model of good AI governance, which brings us back to Fig. 2 in this paper. How should we imagine the middle-out approach of this model? We provide one such example in the next section 3.5

3.5 ON S.M.A.R.T COORDINATION

The governance of AI and its ethical principles provide a rich test bed and important application field for new models of legal governance that depend on how its "middle-out" layer is designed. Both forms of top-down regulation and bottom-up approaches fall short in coping with the challenges of the field. This section is divided into three parts: section 3.5.1 addresses the limits of current proposals on the governance of AI and, correspondingly, what the middle-out layer of the model should look like. Section 3.5.2 illustrates how the model works in practice. Section 3.5.3 explains the scalability of the model. Finally, Section 4 provides the adequate – and corresponding - regulatory toolbox for EU regulators, national governments and multi-stakeholders.



3.5.1

THE “MIDDLE-OUT” APPROACH TO AI

The need for new governance models for AI is a popular claim in today's debate. Some propose a model of ‘monitored self-regulation’ (CEPS 2019). This means that principles, organisational measures, and technological solutions for AI are up to the forces of the market and of social norms, while the role of public institutions would be to monitor of these activities. Others, like the AI Now Institute, recommend new approaches to governance that include internal accountability structures for the industry and the implementation of independent monitoring and transparency efforts. By taking into account further proposals, such as the argument for an ‘agile and comprehensive’ governance for AI (e.g., Wallach and Marchant 2019), which side should we take in today's debate?

First of all, as mentioned in section 1.2, we need to avoid misconstruing the current debate. We have underlined the fact that most legal systems already govern the field of AI in a context-dependent way (see box on page 14): these sets of norms on self-driving cars, on drones, on e-health systems and more, represent the current legal constraints for any model of AI governance. Present discussions on moral principles and whether - or to what extent - should we amend current legal systems, should not however overlook a substantial convergence on some normative issues brought about by AI.

14 Priority Actions, divided into three sets of actions and mechanisms were presented in section 2. There, the aim was to stress the priorities of today's good AI governance, namely, that which we can reasonably conceive as good, right, or lawful and immediately actionable. **Here, the intent is to determine how the priorities presented in section 2 may provide a model for the field.** The complexity of today's legal regulation recommends dwelling once again on the middle-out layer of the analysis, so as to (i) ascertain what kind of balance is struck between regulatory systems according to (ii) a certain mix of primary and secondary legal rules, the result of which (iii) provides a framework for our three sets of priorities.

Looking at the balance between the top-down and bottom-up regulatory systems, attention should be drawn to the current limits of a clear understanding of the stakes of AI. For instance, in the field of autonomous vehicles, the lack of data on the probability of events, consequences, and costs, makes it hard to determine the level of risk in this sector of AI innovation (Robert 2019). These limits are crucial to determine how we intend to set the bar between the top-down and bottom-up approaches.

We have already stressed that today's regulatory frameworks do not provide a set of legal and moral principles to be enforced either through forms of top-down regulation, or intermediate methods of co-regulation, in all fields of AI. This vacuum seems to argue for the adoption of different forms of engagement and of participatory mechanisms as illustrated in Section 2.2 above.

Yet these forms of self-regulation with their possible variants - such as the proposal of a ‘monitored self-regulation’ (e.g. CEPS 2019) - should take into account current



applicable top-down regulations and their possible amendments, either as a form of soft law, or of hard law. Consider the role of the European observatory for AI, as mentioned in Section 2.3(v), and that which we proposed in our 2018 paper, namely, the development of ‘a new EU oversight agency responsible for the protection of public welfare through the scientific evaluation and supervision of AI products, software, systems, or services’ (Floridi et al. 2018). Both recommendations (monitored self-regulation and an EU oversight agency) fit hand in glove, since **the sets of priorities proposed in Section 2 of this paper on no-regrets actions, forms of engagement, and coordination mechanisms, can and sometimes should be complemented by top-down forms of intervention.** Such is the case with a new EU oversight agency that plays the role of meta-regulator, helping the vertical agencies, e.g. the data protection boards, to do AI right.

Both the existence and limits of current regulatory frameworks recommend a middle-out solution in between the top-down and bottom-up approaches and in between the models of co-regulation and self-regulation.

The set of coordination mechanisms introduced in Section 2.3 represents this middle-out layer of the model.

Such coordination mechanisms have a twofold aim: vis-à-vis bottom-up solutions, where the goal is to consolidate new forums for collective consultation and discussion; vis-à-vis top-down actions, where the purpose is to develop new standards and procedures for good AI regulation. See Fig. 4 below.

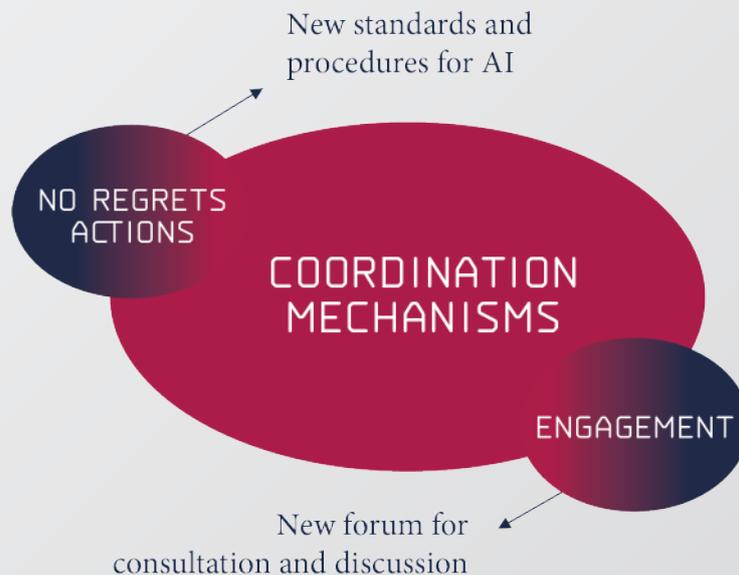


Figure 4: A middle-out approach for AI.



3.5.2

A NEW COORDINATION AT WORK

Contrary to models of self-regulation, the middle-out layer in Fig. 4 is not incompatible with forms of top-down intervention. We have examined these forms of intervention in this paper, specifically with:

- i. Existing context-dependent regulations of AI, e.g. the GDPR;
- ii. Possible improvements of such regulations via, e.g., a new EU oversight agency;
- iii. The set of no-regrets actions illustrated above in 2.1;
- iv. New forms of legal regulation, such as the P2B paradigm and its adaptive rules for emerging innovation.

In the latter example of top-down intervention (iv), the law frames the use of a toolbox that can be adapted and customised for each AI-driven sector, through the adoption of codes of conduct, alternative dispute resolution, independent ongoing expert analysis, real-time monitoring and fast-to-market rule-making - and adjustment.

These forms of legal experimentation go hand-in-hand with the creation of lawfully de-regulated special zones, a sort of living labs for the empirical testing and development of AI and robotics. Over the past two decades, it's worth noting the Japanese government has set up a number of special zones, or *Tokku*, to improve the understanding of how AI systems may react in specific contexts - and satisfy human needs. These special zones have concerned the fields of road traffic laws (at Fukuoka in 2003), radio law (Kansai 2005), data protection (Kyoto 2008), safety governance and tax regulation (Tsukuba 2011), road traffic laws in highways (Sagami 2013), and more. Meanwhile, a regulatory sandbox has also been set up by the National Environmental Agency in Singapore¹. Likewise in Europe these experiments with lawfully de-regulated special zones have been particularly popular in the fields of self-driving cars and drones. In 2016, Sweden sponsored the then world's first large-scale autonomous driving pilot project, while Germany allowed a number of tests with various levels of automation on highways. The first special zone for the test of drones in open labs was established in the harbour of Antwerp, in January 2019.

Generally these forms of experimentation through lawfully de-regulated special zones represent the legal basis upon which to collect empirical data and sufficient knowledge to make rational decisions for a number of critical issues. Different forms of coordination mechanisms – including, but not only de-regulated special zones - can improve our understanding of how AI systems may react in various contexts and satisfy human needs. On this basis, we can better comprehend risks and threats brought about by possible losses of control of AI systems, in order to keep them in check. The aim of the special zones is thus to develop theoretical frameworks that allow us to better appreciate the space for potential systems that avoid undesirable behaviours. Additionally, by developing theoretical frameworks, we can rationally address the legal aspects of this experimentation, 'covering many potential issues raised by the next-generation AI systems and managing

¹ NEA, Regulatory Sandbox, available at: <https://www.nea.gov.sg/industry-transformation-map/regulatory-sandbox>



such requirements, which often represent a formidable obstacle for this kind of research, as public authorizations for security reasons, formal consent for the processing and use of personal data, mechanisms of distributing risks through insurance models and authentication systems, and more' (Pagallo 2017b).

The coordination mechanisms shown in Fig. 4 are thus a mix of primary and secondary rules of the law, as defined above in Section 3.2, that work as a set of adaptive norms for emerging innovation, while striking the balance between the different regulatory systems of the market, of social norms and of technology.

In particular, the bar is set in between the models of self-regulation and co-regulation, since the approach takes into account both the existence and limits of current regulatory frameworks, as examined above in Sections 1 and 3.3.

Three important outcomes follow as a result of this stance:

First, we prevent the chilling effect of some ideological debates on the ethics of AI and the complexity of the legal environment. A middle-out approach for coordination mechanisms tackles the normative challenges of AI dynamically, as the proper interface amongst the top-down and the bottom-up approaches.

Second, the stance is pragmatic, for the coordination mechanisms of the model aim to implement the sets of priorities illustrated previously in Section 2.

Third, the action plan is scalable, both in the sense that growing workload can be suitably addressed by adding resources to the middle-out layer of the model, as well as the possibility of complementing the model by further mechanisms of systemic oversight.

The next section explores the scalability aspect of the model separately.

3.5.3 SCALABILITY

So far, our analysis has focused on the horizontal facet of the figures discussed in this paper. The aim has been to flesh out that which lies in between the top-down and bottom-up approaches and in between the forms of co- and self-regulation as the proper level of abstraction in order to govern the complexity of today's developments in AI - and to implement the priorities discussed in Section 2.

This stance on the institutional layers and procedural mechanisms for good AI governance can be integrated with further models of governance that dwell on the vertical facet, i.e., from top-down to bottom-up solutions, of the models discussed so far. The aim of this integration is to capture our expanding understanding of AI with increasing granularity.

For example, when dealing with the dimensions and layers of the Web of Data and the complexities of linked data systems, the middle-out layer of the analysis is structured to accommodate three additional layers. Moving from bottom-up forms of self-regulation to the hard-law tools of hetero-regulation, these layers concern (i) organic



decentralisation and intermediate conceptualisation; (ii) systemic interdependence and coordinate agency; (iii) semantic interoperability and abductive reasoning.

These layers, in addition, can be understood in accordance with the desirable features of such stratified interaction between bottom-up and top-down approaches. Other domains of technological innovation further illustrate how the “middle-out” approach works in practice. In (Blasimme and Vayena 2018), the components of the model regard oversight structures and processes for data-intense fields such as biomedicine and AI. Their features concern (i) Adaptivity; (ii) Flexibility; (iii) Inclusiveness; (iv) Reflexivity; (v) Responsiveness; and, (vi) Monitoring. All the coordination mechanisms of Section 2.3 can be specified in accordance with the ways in which we should address the features of this ‘A.F.I.R.R.M. model’ (Blasimme and Vayena 2018). In particular,

- i. The coordination mechanisms should be adaptable and flexible enough to accommodate the uncertainties of innovation through appropriate forms of oversight, such as the European AI observatory and the European AI Meta-regulator mentioned above in 2.2 and 3.5.1;
- ii. Such flexibility entails the capacity to treat different AI types depending both on their data sources and on their actual use. As stressed in our 2018 paper (Floridi et al. 2018), many of the normative challenges of AI are context-dependent, from both a moral and legal point of view;
- iii. Inclusiveness refers to the engagement of all affected parties in deliberations and decision-making practices about the use of AI. The mechanisms of adaptive regulation should be clear enough to impose society's preferences on emerging innovation;
- iv. Reflexivity concerns the sound scrutiny, assessment and evaluation of risks;
- v. In contrast, responsiveness refers to the aim of mitigating the effects of unintended consequences, such as unauthorised access to personal data;
- vi. Monitoring views regards the regular scrutiny of AI activities with their effects. The intent is to anticipate the emergence on new vulnerabilities and undesirable outcomes. We may opt for fixed or adaptive monitoring, in real time or predictive, automated or used as a recommender system.

3.5.4 GRANULARITY

Another way of increasing the granularity of the model is by looking at its compatibility with further administrative and corporate governance models (Poblet, Casanovas, Rodríguez-Doncel, 2019). The middle-out approach actually contributes to bridging these, fleshing out their interoperative, coordinative and participatory dimensions.

Furthermore, our stance fits neatly with the recent EU policy on better and smart regulation as well as with some technical developments of the EU Better Regulation scheme for interoperability².

² See above in page 16 (TOGAF 2017).



Likewise, our approach is consistent with the stance on the rule of law taken by standardisation agencies, e.g. NIST-800-53 from 2013 and NIST-800-63C from 2016, together with ISO/IEC 27002 and 27001 on security and privacy controls for Federal Information Systems and Organizations.

In addition, the “middle-out” approach is in line with such governance models in the business field, as for example the recent COBIT2019 framework launched by ISACA, and the Enterprise Architecture model which aims to align management information systems with business interests.

This convergence is unsurprising. As technological challenges grow increasingly complex, both top-down and bottom-up approaches will become less fruitful, thus demanding we pay greater attention to the “middle-out” level of the approach.





A REGULATORY TOOLBOX

According to the EU's Better Regulation scheme for interoperability¹ mentioned in the previous section, the EU Commission should ensure that (i) decision-making is open and transparent; (ii) citizens and stakeholders can contribute throughout the policy and law-making process; (iii) EU actions are based on evidence and understanding of the impacts; and, (iv) regulatory burdens on businesses, citizens or public administrations are kept to a minimum². A toolkit for better regulation has already been provided, launching a better regulation planning throughout the whole European policy cycle³.

Drawing on this agenda, the aim is now to develop our own toolkit of good AI governance. This should align with further EU initiatives on the governance of technological innovation, as well as the S.M.A.R.T. Model of Governance and the 14 Priority Actions illustrated in the previous sections of this article.

Thus, Section 4.1 scrutinizes the functions of the toolbox, namely, what the latter is called on to do. Then, Section 4.2 examines its requirements, that is, what the toolbox is meant to be. Section 4.3 provides for the tools of such regulatory governance box.

4.1 FUNCTIONS OF THE TOOLBOX

A Communication from the EU's Commission dating 23 March 2017 has defined the strategy for governance and interoperability across the Member States⁴. The EU has adopted a relational view to foster citizen participation, transparency, public monitoring and control, considering interoperability as a prerequisite 'for enabling electronic communication and exchange of information between public administrations' and 'for achieving a digital single market' [COM (2017) 134].

Although the Commission's strategy is not tailored to AI in particular, it represents a pillar of today's EU data governance that ought to be taken into account. This strategy establishes what our toolkit is called on to do, namely, the function of the toolbox. More particularly, the EU provides a set of principles and recommendations⁵ to promote electronic communication across administrations and private parties, distinguishing four layers of interoperability:

¹ See above in page 16 (TOGAF 2017).

² https://ec.europa.eu/info/law/law-making-process/planning-and-proposing-law/better-regulation-why-and-how_en

³ https://ec.europa.eu/info/sites/info/files/better-regulation-toolbox_2.pdf

⁴ *European Interoperability Framework – Implementation Strategy*. Brussels, 23.3.2017 COM (2017) 134 final. https://eur-lex.europa.eu/resource.html?uri=cellar:2c2f2554-0faf-11e7-8a35-01aa75ed71a1.0017.02/DOC_1&format=PDF

⁵ Underlying principles for public administration are citizen- and user-centred: (i) subsidiarity and proportionality, (ii) openness, (iii) transparency, (iv) reusability, (v) technological neutrality and data portability, (vi) user-centricity, (vii) inclusion and accessibility, (viii) security and privacy, (ix) multilingualism, (x) administrative simplification, (xi) preservation of information, (xii) assessment of effectiveness and efficiency.



- i. Legal (ensuring that organisations operating under different legal frameworks, policies and strategies are able to work together, setting interoperability checks to identify legal barriers). In particular, legislation should not impose unjustified barriers to the reuse of data in different policy areas;
- ii. Organisational (the relationship between service providers and service consumers). This layer includes the request of formal agreements on the conditions applicable to cross-organisational interactions;
- iii. Semantic (developing vocabularies and schemata to describe data exchanges in the same format). The use of common descriptions of exchanged data represents one of the more effective ways to address current data/semantic concerns;
- iv. Technical (the applications and infrastructures linking systems and services). The aim of this layer is to set up the necessary information systems environment, in order to allow an uninterrupted flow of bits and bytes [COM (2017) 134]. As assessed in the document *Artificial Intelligence for Europe*⁶ [COM (2018) 237 final], AI-based systems can either be purely software-based (e.g. voice assistants); or embedded in hardware devices (e.g. autonomous cars); or be fed with collected data from extended data sets or from the Web (i.e. machine learning); or trained to improve their performance (deep learning).

Information data flows geared through AI should be monitored and regulated. Stemming from the Better Regulations Agenda and this interoperable perspective, a middle-out toolbox for good AI governance could also be set and put into practice. This entails leaning on some practical mechanisms to bridge the relationship between ethical principles and legal norms, much as the technical middle ground and the social meso-level, i.e. the relation between top-down and bottom-up approaches to regulatory toolboxes.

4.2 REQUIREMENTS OF THE TOOLBOX

We already stressed that models of governance should be scalable, modular, and decomposable. Focus here is drawn on what the toolbox is meant to be (its requirements), rather than what it is called on to do (its function).

There are three of such requirements:

First, toolboxes should be developed in accordance with European and national laws, much as case-based law—i.e. *acquis communautaire*—in the Single Digital Market strategy;

Second, a regulatory toolbox should accordingly be used as a key enabler to enhance this legal content and foster both security and trust;

Third, a middle-out toolbox should comprise all components for a good AI governance stressed so far in this article.

⁶ <https://ec.europa.eu/transparency/regdoc/rep/1/2018/EN/COM-2018-237-F1-EN-MAIN-PART-1.PDF>



4.3 THE TOOLS

Drawing on the set of functions and requirements of the toolbox, as illustrated above in Sections 4.1 and 4.2, our kit for good AI governance comprises 11 different tools. They comprise all the components stressed throughout this document. In particular, the middle-out toolbox for good AI governance includes:

- I. **Principles and Legal schemes:** conceptual kernels and systematic plans of action for AI, analysed in Sections 1, 2 and 3.5.2. These parts of our White Paper stressed the moral and legal constraints of the model of S.M.A.R.T. coordination for the governance of AI;
- II. **Digital Innovation Hubs:** Sections 2.1 and 3.5.2 insisted on research efforts, open labs and further forms of legal experimentation, as key tools for making rational choices in the field of technological innovation through the test and use of AI;
- III. **Innovation Deals:** Sections 2.3 and 3.5.2 illustrated some of such deals with benchmarking and adaptable legal rules as in the P2B paradigm;
- IV. **Compliance through Design and by Design:** that is, the aim to embed legal constraints into technology. Sections 3.3 and 3.5.3 examined this approach with the rules of the GDPR as valid law for AI with its limits, and further examples of enforcement by-design through e.g. NIST standards;
- V. **Regulatory Sandboxes:** Section 3.5.2 dwelt on such lawfully de-regulated experiments as in e.g. finance and with the creation of special zones, or Tokku;
- VI. **Online & Alternative Dispute Resolution:** Section 3.5.2 stressed new ways to make deals and facilitate the resolution of disputes between parties, in accordance with e.g. the P2B model;
- VII. **Regulatory Computing Architectures:** Sections 3.5.3 and 4.2 insisted on the crucial role of linked data systems for the governance of AI, drawing on the Communications of the European Commission on data governance and coding;
- VIII. **Regulatory Ontologies, Interoperability & Ontology Design Patterns:** Sections 3.5.3 and 4.2 drew attention to the necessary information system environment for the governance of AI in the Web of Data. Additional layers for the middle-out layer of the analysis took into account the European Commission's strategy for governance and interoperability across the Member States;
- IX. **Monitoring Functions:** Sections 3.5.2 and 3.5.3 focused on the regular scrutiny of AI activities and a proactive approach to the risks brought about by AI through innovation hubs and deals;
- X. **Standardisation, Verification and Impact Assessment Mechanisms:** Sections 2.2, 3.3 and 3.5.3 concerned the sound scrutiny, assessment and evaluation of risks through the creation of new standards, much as the aim to mitigate the effects of unintended consequences in the use of AI through participatory mechanisms;



XI. Oversight Agency: Sections 3.5.1-3.5.3 regarded the protection of public welfare through the scientific evaluation and supervision of AI, which includes e.g. a new EU oversight agency that should play the role of meta-regulator.

All these elements are knowledge-based and should be legally driven. Accordingly, they can be ordered into normative clusters to facilitate their management via hard law, policies, soft law and ethics. They also can be focused and developed into more elaborated complex regulatory and legal systems, partly automated. It's worth noting that these regulatory systems are hybrid—semi-automated—and require both coding and human intervention to be designed, re-designed, and monitored (Pagallo and Durante 2016).

Identifying principles and legal schemes, that is, conceptual kernels and systematic plans of action on a particular topic, internally provided in a statute or regulation, is the first step to turn them into legal data-based instruments.





88SS9174-BLD2

PFV6590A.1

1203 B1P

TW



C07F

3R3

3R3

3R3

3R3

CONCLUSION

The paper has dwelt on today's debate and institutional initiatives on the principles that should undergird the adoption of AI. Drawing on our 2018 paper, *The AI4People's Ethical Framework for a Good AI Society* (Floridi et al. 2018), we fleshed out that which can be deemed as good, right, lawful and moreover, is immediately doable. 14 Priority Actions, a S.M.A.R.T. Model of Governance and the corresponding Toolkit followed as a result. The aim has been to incorporate both elements of top-down legal framing and bottom-up empowerment of individual actors, in order to accommodate the uncertainties of innovation and, at the same time, to expand and strengthen our understanding of the stakes of AI. The outcome was summed up as S.M.A.R.T., because Scalability, Modularity and Adaptability can be suitably addressed by adding resources to the mechanisms set up for coordination, and complemented by further procedures of Reflexivity, i.e. sound scrutiny, assessment and evaluation of risks, and Technological-savvy oversight. By drawing the attention to that which is immediately actionable, leaving thus aside second-order priorities that will require some time to implement (i.e. the development of new standards), five different outcomes followed as a result of this approach to a good AI governance.

First, the sets of priorities for the field of AI recommend a **model of governance that lies in between** current forms of EU co-regulation and multiple variations of self-regulation, such as 'monitored self-regulation.' Our model of coordination shows the ways in which **the middle-out approach** works at both institutional levels through, e.g., the European AI observatory and oversight agency, and in accordance with further models of governance, like the EU Better Regulations Agenda.

Second, the model takes into account how the various sets of priorities may **evolve**, and how regulators should approach this evolution through procedures for engagement and systemic oversight. The model is here smart because it is pro-active and dynamic.

Third, the model is **scalable** and suggests a meta-model, according to which we may also determine how good AI governance could evolve in the next future.

Fourth, it describes possible technical support for the **integration** of such directions, including standardisation, verification and impact assessment mechanisms.

Fifth, the model is **decomposable** and thus properly addresses crucial differences in the field. After all, a considerable amount of current legal provisions, as well as several ethical principles of AI, are already context-dependent.

Our model of smart coordination pinpoints how we can be **flexible enough to accommodate the uncertainties of innovation** and, at the same time, **clear enough to impose society's preferences** on emerging innovation as a matter of priority.

The model is sufficiently agile to capture our expanding understanding of AI with increasing regulatory granularity.



REFERENCES

- Ausloos, J. Heyman, R., Bertels, N., Pierson, J. and P. Valcke (2018) Designing-by-Debate: A Blueprint for Responsible Data-Driven Research & Innovation, in F. Ferri et al. (eds), International Conference on Responsible Research and Innovation in Science, Innovation and Society (RRI-SIS2017), Edition 1: 47 - 64;
- Barfield, W. and U. Pagallo (eds.) (2019) Research Handbook on the Law of Artificial Intelligence, Elgar: Cheltenham, UK - Northampton, MA.;
- Bassi, E. (2019) European Drones Regulation: Today's Legal Challenges. 2019 International Conference on Unmanned Aircraft Systems (ICUAS), 443-450;
- Blasimme, A. and E. Vayena (2019), The Ethics of AI in Biomedical Research, Patient Care and Public Health, Oxford Handbook of Ethics of Artificial Intelligence, Forthcoming. Available at SSRN: <https://ssrn.com/abstract=3368756> or <http://dx.doi.org/10.2139/ssrn.3368756> Available at SSRN: <https://ssrn.com/abstract=>;
- CEPS (2019), Artificial Intelligence: Ethics, governance and policy challenges, edited by A. Renda, Brussels;
- Dworkin, R. (1985) A Matter of Principle. Oxford University Press: Oxford;
- Floridi, L., Cowls, J., Beltrametti, M., Chatila, R., Chazerand, P., Dignum, V., Luetge, Ch., Madelin, R., Pagallo, U., Rossi, F., Schafer, B., Valcke, P. and E. Vayena (2018) AI4People - An Ethical Framework for a Good AI Society: Opportunities, Risks, Principles, and Recommendations, *Minds and Machines*, 28(4): 689-707;
- Grindle, M. (2007) Good Enough Governance Revisited, *Development Policy Review*, 25(5): 533-574;
- Hart, H. L. A. (1961) The Concept of Law. Oxford: Clarendon;
- Helberger, N., Pierson, J. and Th. Poell (2018) Governing Online Platforms: From Contested to Cooperative Responsibility, *The Information Society*, 34(1): 1-14;
- Hodges, C and Steinholtz, R. (2018) Ethical Business Practice and Regulation. A Behavioural and Values-Based Approach to Compliance and Enforcement, Civil Justice Systems, Hart/Beck: London
- Kelsen, H. (1949) General Theory of the Law and the State. Harvard University Press: Cambridge (Massachusetts);
- Lütge, C. (2017) The German Ethics Code for Automated and Connected Driving, *Philosophy & Technology* 30(4): 547-558.
- Pagallo, U. (2017a) The Legal Challenges of Big Data: Putting Secondary Rules First in the Field of EU Data Protection, *European Data Protection Law Review*, 3(1): 34-46;
- Pagallo, U. (2017b) LegalAIze: Tackling the Normative Challenges of Artificial Intelligence and Robotics Through the Secondary Rules of Law, in Corrales M., Fenwick M. and N. Forgó (eds.), *New Technology, Big Data and the Law. Perspectives in Law, Business and Innovation*, Springer: Singapore, Edition 1: 281-300;
- Pagallo, U., Casanovas, P. and R. Madelin (2019) The Middle-out Approach: Assessing Models of Legal Governance in Data Protection, Artificial Intelligence, and the Web of Data, *The Theory & Practice of Legislation*, DOI: 10.1080/20508840.2019.1664543;
- Pagallo, U. and M. Durante (2016) The Pros and Cons of Legal Automation and its Governance, *European Journal of Risk Regulation*, 7(2): 323-334;
- Poblet, M., Casanovas, P. And V. Rodríguez-Doncel (2019) *Linked Democracy. Foundations, Tools, Applications*, Springer Briefs in Law. <https://www.springer.com/gp/book/9783030133627>
- Robert, L.P. (2019) Are automated vehicles safer than manually driven cars?, *AI & Society*. DOI: <https://doi.org/10.1007/s00146-019-00894-y>;
- Tirole, J. (2018) *Economics for the Common Goods*, Princeton University Press: Princeton (New Jersey);
- Turner, J. (2019) *Robot Rules: Regulating Artificial Intelligence*, Palgrave MacMillan;
- TOGAF, 2017. An introduction to the European Interoperability Reference Architecture (EIRA®) v2.1.0. Available at https://joinup.ec.europa.eu/sites/default/files/distribution/access_url/2018-02/b1859b84-3e86-4e00-a5c4-d87913cdcc6f/EIRA_v2_1_0_Overview.pdf;
- Wallach, W. and G. Marchant (2019), Toward the Agile and Comprehensive International Governance of AI and Robotics, *Proceedings of the IEEE*, 107(3): 505-508.



 **ATOMIUM**
EUROPEAN INSTITUTE
FOR SCIENCE, MEDIA AND DEMOCRACY



From left to right: Valéry Giscard d'Estaing, Jean-Claude Juncker and Michelangelo Baracchi Bonvicini.

Atomium-European Institute for Science, Media and Democracy (EISMD), convenes leading European universities, media, businesses, governments and policymakers to increase the exchange of information and interdisciplinary collaboration, to develop innovative collaborative initiatives and to encourage frontier thinking about science, media and democracy.

Atomium-EISMD was launched publicly by the former President of France Valéry Giscard d'Estaing, Michelangelo Baracchi Bonvicini and by the leaders of the institutions engaged during the first conference on the 27 November 2009 at the European Parliament in Brussels.

AI4PEOPLE

SCIENTIFIC COMMITTEE

- **Ugo Pagallo**, Chairman, AI4People Scientific Committee; Professor of Jurisprudence at the Department of Law, University of Turin, faculty at the Center for Transnational Legal Studies (CTLS) London, faculty fellow at the Nexa Center for Internet and Society at the Politecnico of Turin.
- **Pompeu Casanovas**, Research Professor at La Trobe University Law School; Director of Advanced Research and Professor of Law and Legal Studies at Autonomous University of Barcelona; founder and Head of the UAB Institute of Law and Technology (IDT-UAB).
- **Raja Chatila**, Director of Research at the French National Center of Scientific Research, Director of the Institute of Intelligent Systems and Robotics at Pierre and Marie Curie University in Paris, Director of the Laboratory of Excellence “SMART” on human-machine interaction.
- **Patrice Chazerand**, Director in charge of Digital Economy and Trade Groups at Digital Europe.
- **Virginia Dignum**, Professor of Social and Ethical Artificial Intelligence, Umeå University; Scientific Director of Wallenberg Artificial Intelligence, Autonomous Systems and Software Program - Humanities and Society (WASP-HS)
- **Christoph Luetge**, Chair of Business Ethics, Technical University of Munich; Director of the Institute for Ethics in Artificial Intelligence, Technical University of Munich
- **Robert Madelin**, Visiting Research Fellow at the Centre for Technology and Global Affairs, University of Oxford, and Fipra International.
- **Burkhard Schafer**, Professor of Computational Legal Theory, University of Edinburgh Law School.
- **Peggy Valcke**, Research Professor, Centre for IT & IP Law – IMEC, KU Leuven; Visiting Professor Tilburg University & Bocconi University Milan; Member Scientific Committee CMPF and FSR (EUI Florence).



With the contribution of:



INTESA  SANPAOLO



facebook

FUJITSU

Johnson & Johnson

Microsoft

Ada-AI



EurAi