Preface

From the early days of vacuum tube computers to the current generation of supercomputers which operate in the exaflop computing range much progress has been made in the field of artificial intelligence (AI). But perhaps the most striking example of progress in AI has been the advances in the design of algorithms and analytical techniques which allow AI systems to learn from experience, to be creative, and to perform autonomously from humans. However, while these aspects of AI are impressive, current systems operating with AI have begun to challenge established areas of law and have resulted in the beginnings of an important conversation among legal scholars in Europe, the U.S., and Asia on how to regulate AI.

Looking back to the beginnings of AI, in the summer of 1956 the Dartmouth Summer Research Project on AI was attended by a group of imminent scientists that were to become the pioneers in the then fledgling field of AI; among others, these included Marvin Minsky, John McCarthy, Claude Shannon, Allen Newell, and Herbert Simon. The Dartmouth project lasted from six to eight weeks, and was essentially an extended brainstorming session on how to create AI that could perform tasks much like a human. Interestingly, the proposal for the project which was written over 60 years ago still has relevance for AI today:

The study is to proceed on the basis of the conjecture that every aspect of learning or any other feature of intelligence can in principle be so precisely described that a machine can be made to simulate it. An attempt will be made to find how to make machines use language, form abstractions and concepts, solve kinds of problems now reserved for humans; and improve themselves. We think that a significant advance can be made in one or more of these problems if a carefully selected group of scientists work on it together for a summer.

While the time-period of one summer proved insufficient to make major advances in AI, six decades later AI has clearly “come of age.” That isn’t to say that AI has reached the broad level of intelligence shown across the board by humans as predicted would happen fairly soon by the attendees of the Dartmouth conference, but in narrow domains AI is now superior to the performance shown by human experts. From the perspective of law, an interesting consequence of AI gaining in intelligence is that in some cases humans are beginning to be taken completely out of the decision-making loops of the system such that they are unaware of the solution(s) used by the AI in solving a problem. As AI learns from experience and from the vast amount of data found on the internet, “new forms” of AI are beginning to write their own algorithms and importantly are becoming more autonomous from humans in their actions. This increased level of intelligence for AI has raised significant challenges to established areas of law that have traditionally looked to humans as legal actors in disputes ranging from issues of constitutional law to violations of the criminal law. In fact, it is due to the increased autonomy and intelligence of AI, that AI has raised significant questions for legal scholars and judges to consider, such as: “When AI controls a system (for example, a robot), what happens when humans are harmed, or property is damaged, and there is no human in the system to hold

responsible?” Holding “someone” responsible is a cornerstone to allocating criminal or tort liability, yet current versions of AI lacking legal personhood status, cannot be held responsible for their actions.

Based on the definitions of AI proposed by leading researchers in the field, the question of exactly what comprises AI is not only difficult to state but surprisingly open to question; from the perspective of law this is an important observation because a regulatory body must be able to define a technology to effectively regulate it. In this regard, one can ask: is AI the code, the algorithms, or the analytical techniques which lead to creativity and novel solutions to problems, or is AI the architecture of the operating system, the wireless transmission protocols, or the design of the semiconductor chips processing information? Alternatively, is AI simply the act of solving problems as humans would regardless of the decision-making technique, the computer architecture (e.g., von Neumann architecture), or wireless protocol used? To further complicate the complexity of regulating AI, there are many kinds of intelligence, some more amenable to regulation than others. This will require that legislators wade through the difficult task of defining intelligence in its various forms as they propose regulations for AI.

Adding yet another wrinkle to the level of complexity associated with regulating AI is the realization that even though the legal issues associated with machines approaching human-like levels of intelligence are complex, it is possible that future forms of AI may not be human-like at all, the goals and aspirations of such systems may be alien to humans, even incomprehensible to our intellect. Clearly humanity must decide what is to be the future of our interaction with intelligent entities as AI continues to improve and to operate more autonomously from us. Will AI end up being an existential threat to humanity, or our partner in solving major threats to our survival, such as hunger, war, or disease, or will we simply be ignored by a super intelligence that has its own interests? Different scenarios, of course, require different legal approaches to protect the rights of the parties involved.

No matter what form AI takes, its role in society will need to be regulated as will its relationship with individual humans; however, as-of-yet, no major regulatory scheme for AI exists even though AI is involved in almost all aspects of society. In fact, we are at the beginnings of such a discussion now (e.g., with drones, autonomous vehicles, and robotics) focusing on how to regulate AI systems that are not only gaining in intelligence but autonomy. Contributing to that discussion the “Research Handbook on the Law of Artificial Intelligence” consists of 25 chapters written by leading scholars in the emerging field of law and artificial intelligence. This book is also a testament to how far we have come since the first conference on AI was held at Dartmouth in 1956. As detailed in this Research Handbook, antitrust law, criminal law, corporate law, constitutional law, and other areas of law are now being challenged by increasingly smart and autonomous forms of AI. And as AI continues to evolve and get smarter, more issues of law will be raised. This will require new laws and regulatory schemes designed to encourage innovation in AI, while also protecting the rights of humans in a world of increasingly smart forms of AI. Eventually, we may even have to consider rights for AI entities, not the least of which could be granting AI legal personhood status (which would solve some problems but introduce others).

Artificial intelligence will not be “content” to stay within the geographical boundaries of any particular jurisdiction, or nation state for that matter, therefore to be effective, the

regulatory approach to AI will have to be international in scope. For that very reason, the coeditors of this “Research Handbook on the Law of Artificial Intelligence” recruited legal scholars from the U.S., the European Union, and Asia to provide a comprehensive coverage of current thinking on how to apply existing law to developments in AI and to begin the discussion of regulating AI from a global perspective. The chapters in this research handbook taken together show that as AI becomes increasingly smart and autonomous, many areas of law established for humans and for those entities which have been granted the fiction of legal personhood becomes relevant. It’s as if the more “human-like” AI becomes, the more the law for humans is implicated. But if one considers the expansive body of law regulating human behavior, this suggests that the regulation of AI will be immensely difficult. For example, consider the common law or canon law, while both legal systems took centuries to develop, in some ways the capabilities of AI are increasing exponentially implying that the timeframe to regulate AI is decades or less.

Artificial intelligence will continue to advance and will produce amazing intellectual entities in the future (some will have a physical embodiment, others will not), but when that happens will AI continue to serve as a tool for human use or will AI surpass our ability to comprehend its goals, aspirations, and decision making? Either way, the law has a major role to play in how our future develops, and how we eventually co-exist with what may very likely be a more intelligent being “living” amongst us. In response to a future which will surely involve increasingly smart and autonomous entities, some legal systems have already adopted forms of de-regulation and legal experimentation, in order to address the normative challenges of AI. For example, over the past 15 years, the Japanese government has worked out a way to address such challenges through the creation of special zones for the empirical testing and development of AI, namely, a form of living lab, or Tokku. In the field of autonomous vehicles, several EU countries have endorsed similar experimentation: Sweden has sponsored the world’s first large-scale autonomous driving pilot project, in which self-driving cars use public roads in everyday driving conditions; so too has Germany allowed a number of tests with various levels of automation on highways, e.g. Audi’s tests with an autonomously driving car on highway A9 between Ingolstadt and Nuremberg; and the U.S. is a major testbed for the development and regulation of autonomous cars. And as the Committee on Legal Affairs of the European Parliament stressed in the recommendations to the EU Commission on civil law rules for robotics in 2016, “testing robots in real-life scenarios is essential for the identification and assessment of the risks [that AI & robots] might entail, as well as of their technological development beyond a pure experimental laboratory phase.”

Such forms of experimentation through lawfully de-regulated special zones are often complemented with further legal techniques and mechanisms that aim to properly address further potential issues raised by the next-generation AI systems. Forms of experimental federalism and multiple techniques to govern technological innovation, such as the principle of implementation, can indeed improve our understanding of how AI systems may react in various contexts and satisfy human needs. In the first case, the idea of an experimental federalism is to flesh out the content of the rules that shall govern social and individual behavior through a beneficial competition among legal systems. This is what occurs nowadays in the field of self-driving cars in the U.S., where at the time of this writing, seven states have enacted their own laws for this kind of technology. At its best possible light, the same policy will be at work with the EU regulation in the field of data

protection. As to the principle of implementation neutrality, lawmakers can pass regulations that are specific to that technology, e.g. AI, and yet do not favor one or more of its possible implementations, or non-compliant implementations can be easily modified to become compliant. Since the regulation of AI will be challenging, the aim is to employ the legal rules of change that create, modify, or extinct the primary rules of the law, in order to collect empirical data and sufficient knowledge to make rational decisions for a number of critical issues.

Importantly, we must improve our understanding of how AI systems may react in various contexts and satisfy human needs. We shall better appreciate risks and threats triggered by possible losses of control of AI systems, so as to keep them in check. New theoretical frameworks should allow us to better appreciate the space of potential systems that avoid undesirable behaviors, whilst new and old authorities shall coordinate the management of some requirements, which often represent a formidable obstacle for research in AI, such 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. By harnessing the legal rules of change, such as forms of legal experimentation, experimental federalism, techniques of technological governance, risk assessment, and more, how basic concepts, standards, and principles of the law may realign in the foreseeable future will become clearer in constitutional law, criminal and business law, and so forth.

However, as the grandfather of legal automation and AI and law, the German polymath Wilhelm Leibniz, used to say, “every mind has a horizon in respect to its present intellectual capacity but not in respect to its future intellectual capacity.” Once we admit there being AI machines capable of autonomous decisions similar in all relevant aspects to the ones humans make, the next step would be to acknowledge that the legal meanings of intent, of negligence, of consciousness and self-consciousness, and even of legal person, will radically change. Although nobody knows where this scenario may ultimately lead, current debate represents an opportunity to properly take the challenge seriously, namely, what new environment we may wish. The *Research Handbook on the Law of Artificial Intelligence* is a significant step forward in presenting the current framework for thinking about law and AI, with a discussion of how the future direction of law and AI may develop by leading scholars in the field.

Woodrow Barfield, Chapel Hill, North Carolina, USA
Ugo Pagallo, Turin, Italy