Introduction: From intelligent machines to self-driven organisations

One of the sources of human 'competitive advantage' over other animals stems from our capacity for communication and abstract thought. This allows us not only to collectively perform certain tasks but also to transfer years' worth of cumulative experience and knowledge to our children and other members of the community. This knowledge constitutes a key factor in our development: the more we acquire and learn to utilise, the faster we are able to grow.

Thousands of years of history allowed humans to develop many effective methods of transferring knowledge, particularly by teaching it to others. Still, those methods are not without certain inherent limitations: they are time-consuming and highly dependent on the aptitude and motivation of the respective people involved. But even despite this fact, the dynamics of societal development, be it in the technological, economic (welfare) or social dimension (level of education, health, value systems) remains staggeringly high.

But let us imagine structures that would (1) have the ability to gather and process considerably larger amounts of information from a far broader spectrum of stimuli than that available to humans, (2) be able to extract knowledge therefrom and learn from the experiences gained and (3) have the capacity to instantaneously and reliably transfer the same to all members of their community. As soon as any one of the thousands of elements comprising such a structure were to learn a certain skill, all other members thereof would also gain access to the same. And if a certain mistake were made and then successfully remedied in a given situation, all members would learn the ability to avoid it in the future. How quickly would such a structure be able to develop? And what would be the direction of its evolution?

For years, artificial intelligence (AI) has excited the imaginations of science fiction writers and business managers alike. On the one hand, there are many hopes for potential future benefits, and on the other, exaggerated fears of less than welcome repercussions. Optimists see it as a chance for accelerated growth, lower production costs or improved safety. Pessimists fear the expected disappearance of job opportunities, excessive dependence on technology, deterioration of human cognitive abilities, a world controlled by elites with exclusive access to intelligent technologies and ultimately ripped from our grasp by AI itself. Meanwhile, researchers are busy looking for new

methods and value generation structures applicable to organisational and networked context alike.

As follows from the author's experience, many managers and owners of medium and large enterprises are open (often eagerly so) to the possibility of implementing such solutions and have the necessary resources at their disposal, but find themselves incapable of developing a portfolio of projects that would rationally utilise the potential of AI in generating tangible business value in short-, medium- and long-term perspectives. At the same time, local and global markets are rife with providers offering such systems (or components necessary to create them) to both mature companies and newly established start-ups. The situation is further complicated by the rapid growth of industries, methods and technologies fundamental to the development of autonomous solutions. The situation is not unlike that observed in the context of integrated information technology systems several dozen years ago: both the technologies themselves and the related demand existed, but there were still no reliable methods of building viable business cases or experienced teams capable of effectively implementing such solutions.

In order to discover and effectively utilise the potential of a system, one needs to have at least some intuitive grasp, if not full understanding, of the way in which it operates. This provides an insight into both the abilities (explicitly stated in e.g. the provider's offer or derived from application scenarios, as well as still uncovered) and limitations of a given solution. With his knowledge, one is able to better project the particulars of implementation, necessary resources (human, financial, infrastructural) and feasible time frames. A more in-depth intuition, coupled with industry-specific experience and knowledge, will also be helpful when determining the potential impact that a given technology can have on markets and future trends.

The aim of Chapter 1 will be to familiarise the reader with the basic methodologies that constitute the foundation of autonomous systems. First, the general landscape of AI and machine learning will be presented. We will discuss the key methods and techniques involved, describe their underlying principles and suggest possible applications. The next sections will explore in greater detail the ways in which intelligent systems are able to learn through interactions with their environment. A set of such methodologies, jointly described as reinforcement learning, is fundamental to many contemporary intelligent systems – an intuitive understanding of the basic ideas related to this area will prove very useful when trying to discuss the potential and limitations inherent in autonomous solutions. Chapter 1 will close with a presentation of various definitions of autonomy applicable to currently employed systems, thus naturally paving the way for the deliberations included in the subsequent parts of the book.

As already mentioned, this publication is an attempt to describe a fairly universal model of value generation in organisations utilising the potential of autonomous technologies. Decision makers (leaders, managers, owners) in companies that collaborated with the author typically identified two primary needs: to identify the advanced technologies that offer the greatest potential for facilitating them in the achievement of already defined strategic goals and, in the case of those more versed in the topic, to modify their existing strategies with the view to better utilising the potential of AI and preparing for imminent changes (e.g. in their markets or competitive milieu). In either case, even with the knowledge of various models of value generation gained from e.g. scientific publications, anyone attempting to satisfy both those needs would have to be familiar with the viable uses of intelligent technologies, if not all of them then at least in the context of organisations similar to one's own.

This is one of the reasons why such a large part of this book (all of Chapter 2 in fact) will be devoted to real-life examples of practical AI applications. First, the key technologies employed in intelligent systems will be presented. Subsequently, dozens of AI applications in various functional areas of organisations will be discussed: e.g. in logistics, manufacturing, sales or customer service. Finally, the possible impact on various industries will also be considered.

This compiled factographic material will provide the basis for the subsequent identification of best practices applicable to intelligent system implementations, a categorisation of the possible values attainable and the proposal of an implementation strategy developed through analogy with processes inherent in human resources management. Those conclusions, presented in Chapter 3, will hopefully contribute to achieving the primary purpose of this publication: to aid managers in reliably and effectively generating value from intelligent systems and solidifying the resulting changes in their organisations.

Chapter 4 will be dedicated to future trends: the possible impact of autonomous technologies on market sectors, the concepts of distributed AI and cognitive networks, as well as some particularly fascinating studies conducted at research centres and corporate laboratories. Some of those ideas may set the trends for future technological development and help to create a new niche for systems that are not only autonomous but also capable of continuous learning and self-improvement.

BIBLIOGRAPHICAL NOTE

Anyone interested in gaining a deeper insight into AI methods and technologies would benefit greatly from studying, above all, some of the great textbooks pertaining to the topic. One of the best and most comprehensive analyses was offered by Russell et al. (2010). This admittedly rather bulky

publication (nearly 1,000 pages) provides a solid mathematical foundation before proceeding to discuss the methods of both so-called *Good Old Fashioned AI* and the most recent advances in the areas of machine learning or deep learning. Valuable insights into the concept of machine learning can also be derived from publications by Géron (2019), Witten (2017) and Goodfellow et al. (2016). The foundations of reinforcement learning methodologies were analysed in considerable detail by Sutton and Barto (2017), while practical implementations (including algorithms and software fragments, as well as some really intuitive insights) were discussed in the books by Morales (2020) and Lapan (2020).

Most of the content included in Chapter 2 was based on case studies. Initially, over 500 companies were identified that either utilised or marketed intelligent, autonomous systems. The search was facilitated by services such as index.co or Crunchbase as well as various market reports (e.g. Gartner.com as well as major advisory companies). Subsequently, a preliminary analysis of the identified enterprises allowed the selection of the 186 most pertinent examples (cf. Appendix). The same were then analysed in terms of the employed technologies, spectra of applications and values generated, with the view to classifying them under respective categories (functional areas in organisations and industry sectors). At the stage of summarising and classifying the values generated by intelligent systems (Chapter 3), information published by data science providers such as DataRobot, Rapid Miner, H2O or Dataiku proved very helpful.

Chapter 4 makes numerous references to scientific research conducted in the area of AI. This was primarily facilitated by results presented in the course of major conferences in this field, e.g. *Neural Information Processing Systems: NeurIPS*. References were also made to research results published by companies such as DeepMind (2020) in *Nature*, *Science* and other periodicals.

REFERENCES

DeepMind (2020), accessed 22 March 2020 at www.deepmind.com.

Géron, A. (2019), *Hands-on Machine Learning with Scikit-Learn, Keras, and TensorFlow Concepts, Tools, and Techniques to Build Intelligent Systems*, Newton, MA: O'Reilly.

Goodfellow, I., Y. Bengio and A. Courville (2016), *Deep Learning*, Cambridge, MA: MIT Press.

Lapan, M. (2020), *Deep Reinforcement Learning Hands-On: Apply Modern RL Methods to Practical Problems of Chatbots, Robotics, Discrete Optimization, Web Automation, and More*, 2nd ed., Birmingham: Packt Publishing.

Morales, M. (2020), 'Grokking Deep Reinforcement Learning', accessed 6 December 2019 at www.manning.com:443/books/grokking-deep-reinforcement-learning.

Russell, S. J., P. Norvig and E. Davis (2010), *Artificial Intelligence: A Modern Approach*, 3rd ed., Upper Saddle River, NJ: Prentice Hall.

Sutton, R. S. and A. G. Barto (2017), *Reinforcement Learning*, 2nd ed., Cambridge, MA: MIT Press.
Witten, I. H. (ed.) (2017), *Data Mining: Practical Machine Learning Tools and Techniques*, 4th ed., Amsterdam: Elsevier.