Internet of Things and Artificial Intelligence—A Wining Partnership?

J. Semião, M. B. Santos, I. C. Teixeira, and J. P. Teixeira

Abstract Hardware/Software (hw/sw) systems changed the human way of living. Internet of Things (IoT) and Artificial Intelligence (AI), now two dominant research themes, are intended and expected to change it more. Hopefully, for the good. In this book chapter, relevant challenges associated with the development of a "society" of intelligent smart objects are highlighted. Humans and smart objects are expected to interact. Humans with natural intelligence (*people*) and smart objects (*things*) with artificial intelligence. The Internet, the platform of globalization, has connected people around the world, and will be progressively the platform for connecting "things". Will humans be able to build up an IoT that benefit them, while keeping a sustainable environment on this planet? How will designers guarantee that the IoT world will not run out of control? What are the standards? How to implement them? These issues are addressed in this chapter from the engineering and educational points of view. In fact, when dealing with "decision making systems", not only design and test should guarantee the correct and safe operation, but also the soundness of the decisions such smart objects take, during their lifetime. The concept of Design for Accountability (DfA) is, thus, proposed and some initial guidelines are outlined.

J. Semião
University of Algarve, Faro, Portugal
e-mail: jsemiao@ualg.pt

J. Semião · M. B. Santos · I. C. Teixeira (✉) · J. P. Teixeira
INESC-ID, Lisbon, Portugal
e-mail: isabel.teixeira@tecnico.ulisboa.pt

M. B. Santos
e-mail: marcelino.santos@tecnico.ulisboa.pt

J. P. Teixeira
e-mail: paulo.teixeira@tecnico.ulisboa.pt

M. B. Santos · I. C. Teixeira · J. P. Teixeira
IST, University of Lisboa, Lisbon, Portugal

M. B. Santos
Silicongate, Lisbon, Portugal

G. Mastorakis et al. (eds.), *Convergence of Artificial Intelligence and the Internet of Things*, Internet of Things,
https://doi.org/10.1007/978-3-030-44907-0_15

1 Introduction

The Internet has become the platform of human *globalization*. Internet allowed, unlike anything else before, a massive interconnection of people around the world. We can now say that there is an *Internet of People* (*IoP*), namely through social networks, wondering around and looking for education, business, pleasure and other more or less recommendable purposes. At present, humans are developing hardware/software (hw/sw) systems providing them a form of intelligence—the *Artificial Intelligence* (*AI*). Such hw/sw systems are now referred as smart objects, or "*things*". The Internet becomes progressively *the* platform for connecting "things". We can now say that there is an *Internet of Things* (*IoT*). In fact, a broader term for IoT has also been referred, the *Internet of Everything* (*IoE*), denoting physical devices and everyday objects connected to the Internet and outfitted with expanded digital features (like AI), extending the IoT emphasis on machine-to-machine (M2M) communications to describe more complex systems that also encompass people and processes.

Strangely, the Internet has relevantly modified *people's* values and way of thinking. For instance, *privacy*, once one of the most protected values, has been dropped out by social networks. In fact, in social networks, people expose themselves to the world in a way that would be unthinkable some decades ago. Hence, as some others aspects of modern life, that motivate people to share, not rationally explained, feelings and actions (e.g., the emotions associated to sports practices), the Internet is shaping the world. It gathers people in "one place", it modulates people's way of thinking.

History shows that modeling people's way of thinking according to some unique predefined pattern, can bring a dangerous and scaring outcome. The Babel tower is an interesting example of this, registered thousands of years ago in one of the most well preserved books in History. The evaluation of the situation was stated in the words: "now, nothing that they design to do will be impossible to them". Internet, like AI, does not differentiate good from bad. Its use for good or for bad greatly depends on the user. In fact, when a project falls in the hands of unscrupulous people, the outcome can be very negative. In reality, the Internet, being a valuable asset, can be used for excellent or for terrible ends. To avoid that this use depends entirely on the human user, some regulation is under way, by implementing Ethic values in the form of standards.

Now, that humans are developing smart objects, and connecting them through the Internet, the relevant question is: How to prevent the use of this global platform in IoT systems for bad purposes?

From the engineering point of view, the first goal is to guarantee the correct *functionality* of the hw/sw systems. However, in IoT systems driven by AI, it is important to make sure that these intelligent objects' *decisions* are also bounded by *Ethic values*. How can adequate Ethic values be implemented, regulated by Law, and monitored in the IoT world? One way to achieve this goal is by the establishment of *standards*. Design engineers already implement standards in their traditional system design. The establishment of standards for the IoT world is a must, particularly, with

the advent of the G5 technology [1]. Accordingly, and wisely, IEEE is working on a "Global Initiative on Ethics of Autonomous and Intelligent Systems" [2]. Similar effort is underway at the Montreal University [3], and in the European Union [4]. All these efforts make clear that these issues need to be taken into account in engineering education.

The availability of highly complex hw/sw systems with increased functionality, and simultaneously with the decrease of physical systems dimensions, allowed advanced semiconductor technologies to be widely used in commodities, and in the improvement of the potential of communication systems. It turned possible to build up one of the most powerful communication system men built, the Internet, and thus remotely control systems mission functionality with minor costs for the consumer.

At the beginning, electronics-based hardware subsystems were used only as the physical platforms to run the software, which drove programmed algorithms and heuristics to perform pre-specified functions. Artificial Intelligence (AI) introduced two additional features, that turned hw/sw systems into *smart objects* [5]. First, the ability *to learn*, by modifying Data Bases (DB) through Machine Learning (ML) algorithms [6, 7]. Second, the ability to make *autonomous decisions*, making these objects smart actors, progressively choosing and modifying their own role. With the Internet, the global platform allowing massive communication among smart objects, the IoT discipline emerged. It is being viewed as a new gold rush, not only for wealth generation, but also as a powerful mean to influence the decision centers of this world.

IoT introduced a new paradigm in computing. Reportedly, by 2020, the number of IoT devices could reach 24 billion [8]. IoT is stimulating the fourth industrial revolution, bringing significant benefits by connecting people, processes, knowledge and data [9–11]. The possibility of interconnecting a huge amount of smart objects, with increasing *local* AI, is opening new avenues of research. Innovative IoT applications are being developed across various markets, from smart cities [12] down to health systems [13], automotive applications [14], aerospace, e-government [15] and so on [16]. Simultaneously, but not surprisingly, *cybersecurity* has become a critical challenge in IoT systems [17, 18].

Therefore, IoT and AI are now two dominant research themes in our scientific world. Their possible synergies are stimulating the imagination of many, and driving new research and development applications.

Nevertheless, are there *risks* on the convergence of IoT and AI? Does the IoT-AI convergence always lead to a winning partnership? We should keep in mind that AI and Autonomous Intelligent Systems (A/IS) are assumed "to behave in a way that is *beneficial to people* beyond reaching functional goals and addressing technical problems", as stated in [19].

The purpose of this chapter is to highlight relevant *challenges* associated with the development of a "society" of intelligent smart objects, trained to perform Machine Learning (ML) [6, 7] and decision-making. Moreover, challenges associated with two intelligent communities—*people* and *things*—and their interactions are also highlighted, leading to the conclusion that an intelligent *strategy* for AI-driven IoT systems must be developed. For this purpose, two key areas are considered here:

engineering education, and smart object design, test and safe operation. The concept of Design for Accountability (DfA) is also introduced.

The chapter is organized as follows. Technology background is reviewed in Sect. 2. The basic characteristics of the existing intelligent society built of *people*, and how they influence the implementation of objects intelligence, is considered in Sect. 3. Key challenges in the convergence of IoT and AI, together with the interaction of two intelligent communities—*people* and *things*—are identified in Sect. 4. In order to devise a strategy for AI-driven IoT systems, Sect. 5 considers two key areas of urgent research: Sect. 5.1—Engineering education, in order that system designers may include smart solutions for eventual shortcomings, and Sect. 5.2—Design for Accountability (DfA), in such a way that smart objects and IoT systems may monitor safe operation, and evaluate the soundness of the decisions they make, along product lifetime. Finally, Sect. 6 summarizes the main conclusions of the work.

2 Technology Background—Five Decades that Changed the Electrical Engineering Paradigm

Let us take a brief look into electronics-based technology evolution during the 20th century. Also, some technology limitations are highlighted, in order to justify measures that must be taken into account in order to guarantee correct and safe systems operation, particularly, when objects (controlled by electronic systems) are expected to take decisions.

The first breakthrough has been reached when the ability of electronics-based systems to perform a required functionality was discovered. This feature has been driving innovation since the use of vacuum-based technologies down to solid-state technologies. Vacuum tubes, like diodes and triodes, were used since the yearly days of the 20th century. The solid-state Bipolar Junction Transistor (BJT) was discovered in 1948, after World War II. Metal-Oxide-Semiconductor (MOS) technology, although discovered decades before, only emerged as a dominant technology after the BJT technology.

Since we live in an analogue world, with physical entities assuming continuous, real values, we started with *analogue* electronics-based devices, circuits and systems, devised and implemented in domain applications. For instance, in order to generate audio signals strong enough to be heard by a large audience, amplification systems have been built, composed by sensors (microphones), a signal amplifier, and actuators (loudspeakers).

However, analogue signals can easily be influenced by *noise* signals, either captured by sensors, or generated by signal distortion, introduced by the physical system performing the desired functionality. Consequently, a simple measure of analogue

signal quality is the Signal-to-Noise (S/N) ratio. Often, S/N ratios in analogue physical systems are poor. However, what if we could transform continuous signal amplitude values into a set of discrete values? Using base 2, any integer value can be described by a sequence of two symbols: zero (0) and ones (1).

Hence, we have moved from analogue to *digital* data processing. The flexibility of the transistor, that can operate either in an analogue region, or as a switch, made everything easy. Using DC power supplies, ON transistors act like an almost short circuits (zero resistance), while OFF transistors operate like open circuits (infinite resistance). Circuit variables may easily assume '1' or '0' *logic* values, according to two discrete electric voltage levels, V_{DD} (the power supply voltage value) or 0 V (the Ground value). Digital electronics emerged, and digital data storage and processing can now be performed easily with virtually no noise. Digital electronics lead to the development of *computers*, in which a Central Process Unity (CPU) acts like the maestro in an orchestra, allowing software programs to be used to produce machine language, and to drive the underlying system hardware.

Now, having the ability to build digital systems, the next challenge was: how do we build complex systems, in a cost-effective way, so that hw/sw systems could reach the market at affordable costs?

Research on manufacturing technologies led to the development of simple Integrated Circuits (ICs). Progress in manufacturing IC technologies, namely in lithography, made possible to lay out and interconnect a large number of physical circuit elements. Very Large Scale Integration (VLSI) technologies emerged, as the move from micron-size to nano-size lithography made possible to build hardware silicon chips with millions of transistors. Gordon Moore, Intel co-founder, foresee in 1965 that computing would dramatically increase in power, and decrease in relative cost, at an exponential pace. The famous *Moore's Law* [20] hold for five decades, leading to extremely complex integrated systems, using semiconductor nanotechnologies. State-of-the-art IC technology goes down to 5 nm, i.e., the dimension of ten Si atoms! A positive feedback boost IC design, since increasingly sophisticated Electronic Design Automation (EDA) tools run in platforms that benefit from the progressive high performance achieved by recently designed chipsets, only possible to design with the support of EDA. Interconnection technologies, to interconnect complex devices, or modules, has also evolved from few modules, few interconnections to many, complex modules. Computer systems use *bus-based architectures*, in which control, and dialogue protocols define master-slave architectures. More complex interconnection architectures try to mimic the complexity of the interconnection among human brain's neurons, leading to the concept of *neural networks*.

A key concern in communication, as in digital operation, is power consumption. High speed processing (in the GHz range) means a significant power consumption. That is why, many portable smart objects require batteries with large autonomy. And *power management* becomes a key attribute of a good design.

Semiconductor Yield by existing manufacturing technologies is never 100%. Hence, there is a need to *test* individual components, which exhibit a spread in process parameters. Defining acceptable variable ranges, outliers must be discarded,

while "good" chips are sold. Such go-no-go test needs to be performed in a cost-effective way. Hence, not only short sequences of test vectors need to be applied to each manufactured component, but also the identification of defective parts need to be very accurate, as field returns are costly and erode customer's confidence. A functional test becomes inefficient. For instance, to test a 64-bit multiplier module, all combinations of two 64-bit inputs would have to be applied. Therefore, a *structural test* is the more adequate solution [21]. As an exhaustive test is prohibitive in time and cost, circuit and systems topology needs to be activated and checked for correct operation with high-quality, short test sessions.

Hence, two disciplines emerged, in connection to IC design: *Design for Testability* (*DfT*) and *Built-In Self-Test* (*BIST*).

DfT techniques aim at making the test of a given design more cost-effective. For instance, if the activation or the monitoring of the system response is difficult to perform, by a given system architecture, a different architecture may be chosen or additional controllability and/or observability nodes may be required. Moreover, system partition (in test mode), e.g., breaking complex modules in simpler submodules, may be introduced.

As the hardware silicon real estate has become cheap, test functionality may be designed on-chip, so the mission functionality may be tested by the component itself. Such BIST techniques are useful, not only for production test, but also during product lifetime, when the system is in the field, in its real-world application. This additional feature is very attractive, as physical systems suffer from *aging effects*, especially if the *things* (the smart objects) are supposed to operate for long periods of time, e.g., in automotive applications. Aging effects may induce a soft degradation (i.e., lower performance), or a hard failure (due to a physical defect that compromises correct operation) after a long period of time. Self-test and aging monitoring are especially important in *safety-critical applications*. Using the same example, automotive electronics failure may cause harm, namely, people's death. As many systems have idle periods of time, such time periods can be used to perform self-test, and inform the system manager if everything is OK, or if system maintenance is required, or even if the system must be shut down, as safe operation is at risk. Safety-critical applications [22] may justify adding redundancy and voting or the deployment of built-in sensors in IoT applications [23, 24].

It is true that design and test methodologies developed for traditional hw/sw systems, where the software is developed in a *deductive* way, are necessarily different from design and test methodologies of AI-driven objects and IoT systems, where the software is *inferred* from data and some rules. Yet, there are key aspects that must be preserved, in both application domains, and probably emphasized in IoT, namely, the controllability and observability of the systems inputs and outputs and internal critical nodes.

3 Human Background

Why do we dedicate a Section to the Human background? Basically, to identify the basic characteristics of *people's* (the natural, intelligent society) behavior, in order to compare it with the AI-driven IoT society, under construction.

Humans are a magnificent form of intelligent life. In fact, neurosciences are beginning to unveil the wonder of the complex functionality of the human brain, associated with other parts of the human body. Individuals have something like 100 billion neuron cells, each one with the possibility to establish thousands of connections (synapses) with other neuron cells, creating a neural network, allowing data storage (memory), data processing, reasoning processes, deep learning, and decision-making. Decisions are made, usually, using a rational process, basic in logic. However, and not less important, decisions, in the Human domain, are also influenced by another mysterious human characteristic: *feelings*. Either what we like it or not, feelings may significant influence on decisions, regardless of what we know it should be done.

"All men are born *equal*" [25], as quoted by Thomas Jefferson in the US Declaration of Independence, after Montesquieu. However, are they really born equal? All men and women are really born equal, in the sense that they belong to the same species. However, they start to differ in their genetics, namely in their DNA. *Each individual is unique in the Universe*. Differences tend to increase, as human life, experience and learning tend to progress, during their lifetime. Such differences are heavily influenced by their historical context, culture and society. Life span is a few decades. One generation tends to pass to the following generation their values and culture. Data, information and knowledge each individual grasps, together with a reasoning process, and feelings, are key attributes in their decision process. Communication is a key attribute. We are influenced by others, as this conditions our data gathering.

Humans are not born as adults. A baby has a magnificent way of *learning*. For instance, in 1–2 years, one or more human languages are captured, regarding sounds, words, their meaning, sentences, grammar in such a way that the baby's brain learns to talk, building his own sentences, and starting to communicate with other humans. Usually, parents love and protect their children, monitoring the learning process. *Decision-making* is progressively allowed, as they observe that teen agers are now able to start making their own, autonomous decisions. Hence, the freedom to make decisions is progressively allowed, so they learn how to behave in society, to make a positive input to it, and hopefully, to be happy.

Humans also have another interesting feature: *conscience*. In fact, we all have inherent to us, a sense of what is right and wrong, and a sort of moral/ethical judge inside us, that either praise us when we do good deeds, or condemn us when we do bad things. Conscience is also influenced by the social environment. A lovely baby can become a vicious thief, causing pain to those who love him (or her). Nevertheless, the decision-making process of each individual is also biased by his internal conscience, unless conscience is overwritten by mental training (brain washing).

Artificial Intelligence (AI) tries to mimic the human mindset, through the use of Machine Learning (ML) techniques applied to some data set. This is the reason why the following reflections are included in this chapter.

As a starting point, it is interesting to remember that the use of data sets to base decision-making is not a breakthrough of AI and ML. In fact, human knowledge is built through *inference*. Observation and data gathering is used with reasoning to devise a *theory*, a possible explanation of a given reality. *Abstract thinking* allows humans to identify a plausible *cause* for a given *effect*. This reflects the ability to identify patterns, things which are common to given data subsets, and the ability to analyze, rearrange, store and use available data. Then, each theory needs to be validated, or discarded. *Experiments* need to be carried out, so that the obtained results can be compared with the ones predicted by the theory. This is the base of the *scientific method*, so successfully applied to move forward, and to reach new levels of knowledge and understanding. If a theory is validated, then various *applications* are considered to take advantage of the new knowledge, and *decisions* are made. The first part of this process is referred as *creativity*. The second part is called *innovation*. Usually, innovation generates wealth and welfare in human society.

Humans have a *critic spirit*: when we analyze a given reality, we are usually able to reach a conclusion, whether a given outcome is good or bad, whether it accomplishes a desired goal, or not, whether the process needs to be improved or not.

Humans are organized as *societies* (families, tribes, nations). They build Cultures. Institutions. Law and Ethics. *Ethics* usually are identified as *fundamental principles* that rule life in society. One of such fundamental principles is the right to make individual or collective decisions. We call it *freedom*. Of course, individual freedom may collide with other individual's freedom. Human freedom is limited by Law. Humans may make decisions; however, not all decisions are legitimate. From time to time, or suddenly (e.g., traffic police stop operation) human's actions are under scrutiny, to see if people obey the Law. Laws usually are associated with penalties, or punishment, or restrictions, if the Law is violated. Humans also have inherent to them the concept of *discipline*, for continuous improvement.

Humans, individual or collectively, evaluate *the paths* they are pursuing, and either continue in that direction, or abandon the path and choose another path. By *trial and error*, incremental or significant improvements are made. Such improvements are reflected in the Laws humans formulate, to guide their paths in life. Interestingly, while the Physical world is deterministic and reliable (governed by immutable Physical and Chemical Laws), humans are non-deterministic, due to decision freedom. The outcome of their decisions is sometimes unpredictable (differences emerge, like between, e.g., election polls and results).

Nevertheless, *Humans are accountable*, no matter in what sphere they act: within the family, on the job, in their community, to the State, and so on. Sooner or later, they have to be made *responsible* for their acts, or for the lack of adequate action. Only small children are not accountable (their parents are for them), due to the fact that their autonomous decision-making, for obvious reasons, is very limited.

Looking at human History, *communication* evolved slowly, in small communities. Communication was basically *direct* among people, face to face. Knowledge was

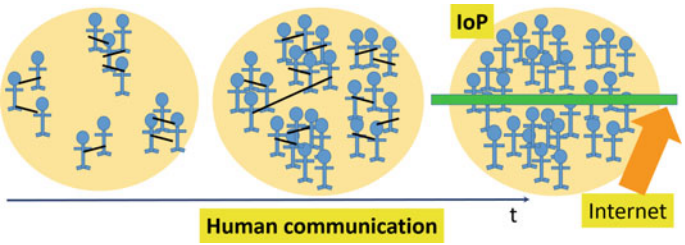

Fig. 1 Human communication along time, ending in the Internet of People (IoP)

scarce, and its dissemination was limited, from one generation to the following one (Fig. 1). Communication in *written* form was almost negligible.

The quantum, gigantic step was the discovery of the press, by Johannes Gutenberg. After 1432 AD, it was now possible to *print* as many copies of whatever knowledge people want. This made whatever knowledge accessible to large audiences, thus spreading human and scientific *knowledge*, and dramatically increasing human communication, in written form.

The knowledge cumulatively gathered in just the last few centuries lead to the spread of science and technology, and to the industrial revolution, as well as to many social transformations. *Education* became relevant in many parts of the world (although for limited subsets of the population). Basic human goals, such as how to get out of poverty, the pursuit of happiness, and the pursuit of wealth, became more visible and humans rush into science-based solutions to build a new society.

In the last decades, the Internet allowed communication among a wide set of humans, and the wide spread of education. Communication among people is carried out, more and more, through the Internet, rather than by direct, face to face communication. *Internet of People* (*IoP*) became a reality. Today, humans are dependent and addicted to science and technology, and flooded by huge amounts of information, not always reliable. Human society evolution moves very fast, sometimes making humans fragile.

As human society is going through an enormous transformation, one may wonder: is there a *strategy* to move from one model, to another model of society? Unfortunately, the answer seems to be *no*. When shining opportunities emerge, humans usually have no strategy—people just rush for gold. This was the case of 16th century Discoveries performed by the European nations, or the 19th century gold rush towards America's far west, or the industrial revolution, and now the AI. Competitiveness is the driving force.

4 Challenges of IoT-AI Partnership

Unlike humans, all smart objects belonging to a given product are born equal, as far as their hardware part is concerned. Nevertheless, even without AI, constant software

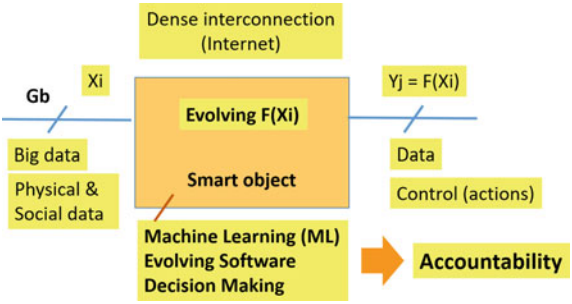

Fig. 2 Key characteristics of a single smart object

updates make clear that the functionality running on these hardware platforms is constantly changing.

As shown in Fig. 2, each complex smart object has relevant attributes. It receives and processes huge amounts of physical and/or social data, from neighboring objects, or from the Internet, eventually stored in a Cloud environment. It is as densely interconnected, as in a neural network. ML processes constantly examines and classifies data, identifying patterns which may modify internal Data Bases (DB). As a result, the software running in its hardware platform may be constantly evolving due to internal autonomous decisions made by the smart object, using AI techniques. The outcome, in form of an evolving software and functionality, and new data and control variables, may trigger actions, which hopefully are beneficial; however, they can cause harm, if not closely monitored. This is why each smart object must be accountable.

High-performance computing of Big data may require severe high power consumption. Often, large sets of smart objects are working in parallel, in order to achieve fast and accurate results, that may drive ML and decision making. For instance, a recent paper [26] considers model training processes for Natural-Language Processing (NLP), the subfield of AI that focuses on teaching machines to handle human language. They show that advanced techniques are computationally expensive—and highly *energy intensive*. The authors highlight that the carbon footprint required to fuel modern tensor processing hardware is severe, and thus a design constraint. When we think of 24 billion IoT devices, by 2020, the energy required can be awesome [8, 27].

Let us identify some key challenges in the new AI-driven society. As referred, ML and AI modify the software part of smart objects, based on autonomous decisions. Hence, a ***first challenge*** is (Fig. 3): ***how to test an evolving software, being modified along product lifetime?***

Human science is based on the capability *to infer* an analytical formal description of a given characteristic that is present in a huge amount of raw data. Hence, from a set of concrete data, it is possible to abstract some kind of relationship that is susceptible of being described by an equation or equation set.

Let us consider a simple example. From the simple verification that one orange with another orange are two oranges, and one apple with another apple are two

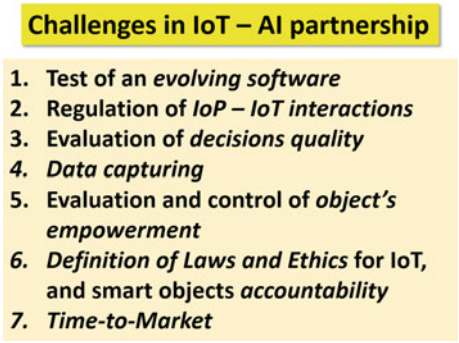

Fig. 3 Key challenges in IoT—AI partnership

apples, and so forth, we can conclude that one plus one (of the same object type) is always two. Then, we can create symbols describing this reality, namely $1 + 1 = 2$. In a more abstract level, we can say $a + b = c$. This is the basics of *arithmetic*.

In another domain, by observing the movement of planets around the Sun, Kepler, Galileo and Newton have derived formal Laws, that carry their names. This is also the way the Universe is being studied, from telescopic observation.

Therefore, the practice of inferring formal relationships, through data observation, has been the way scientists have gone through, and science has advanced. Providing IoT with AI, the goal is also to give these smart objects the ability to learn, to infer, and to make autonomous decisions, based on the knowledge they generate. This not only empowers individual *things*, but also can built an AI-driven *society*, as these objects communicate with each other using the Internet.

So far, the Human society has long been *the sole intelligent society* on this planet. We have the ability to observe, to analyze data, to learn, and to make decisions. Now, the development of AI-driven smart objects leads to the build-up of a *second intelligent society*. Two intelligent societies will simultaneously exist and interact.

Hence, a ***second challenge*** is ***how to regulate IoP—IoT interaction***. The Internet, already crowded with people, will be progressively filled with "*things*". Is the first society (*people*) prepared to build and control a healthy second society (*things*), that will benefit humans, while keeping a sustainable environment in this planet? What will be the *Ethic limits* of IoT and AI, as these two realities converge [2, 3]? How can adequate Ethic values be implemented, regulated by law, and monitored? Like in the area of biotechnology, a significant amount of AI-driven research is performed underground, due to industrial confidentiality. When research results reach the market, it may be too late to control unethical consequences.

Another aspect of IoP-IoT interaction is the fact that smart objects use an extraordinary ability to sense, to analyze Big data [28] in a time frame several orders of magnitude smaller than the human brain can perform. Thus, huge data analysis and high-speed data processing are key advantages of the new smart objects, as compared to the ability of humans. Decisions may be taken much faster than humans do; however, are these decisions better than the ones humans take? Therefore, a ***third challenge*** is: ***how to evaluate decisions' quality?***

A first constraint to the quality of the object decision is the "quality" of the trained *data* provided to these objects as starting point. Data collected by physical means, namely by sensors, can be erroneous if sensors have internal failures or environmentally-induced failures. Social data may also be misleading. Moreover, malicious data is being disseminated through the Internet, and can, maliciously, be introduced in the objects. Hence, a ***forth challenge*** for achieving sound decisions is ***data capturing***. Is the collected data *trustworthy*? Is the data sample *statistically significant*? How and who should evaluate this? May the imported data from an imperfect human society perpetuate its shortcomings, like discrimination? Is there a risk that decisions taken by IoT may be harmful, due to wrong or malicious data, or false information? How *safe* and *reliable* these IoT systems will be? Cybersecurity is already a key concern today [16, 18]. Moreover, similarly to what happens when chips and electronics reliability are analyzed, the identification of false positives and false negatives is relevant. For instance, too much false positives may often halt system performance, reducing system usefulness and customer's satisfaction. Too much false negatives may drive system operation into unsafe behavior. Hence, how to measure the rate of false positives and negatives?

As mentioned before, Artificial Intelligence (AI) tries to mimic the human mind-set. By using Machine Learning (ML) models, raw data is organized in order to derive *order* (or to identify *patterns*) of some interesting kind from even randomly acquired data. This allows smart objects to organize data in more sophisticated data bases, using rules that are inferred by the ML process. If the final conclusions are assessed by humans, the first society retains the control of AI-driven IoT systems. As smart objects move to become, more and more, *autonomous* objects, making (and assessing) their own decisions, how can we be sure that the decisions they make are appropriate? Any decision triggers an action, and actions have *consequences*. Therefore, a ***fifth challenge*** emerges: ***how to evaluate and control object empowerment*?**

A key concern is the fact that the AI embedded in these smart objects is created by people, i.e., by the *Natural Intelligence* (*NI*). As a byproduct of the existing NI, is not reasonable to assume that AI may be *inferior* to NI in the ability to reach the best possible solutions to problems facing life in this planet?

In fact, being a byproduct of NI, shouldn't AI-driven IoTs be inferior to NI-driven IoPs? If humans do not always make sound decisions, is it reasonable to be expect that human-inspired decisions, made by AI-driven IoTs, designed by humans, will make sound or even *better* decisions?

Moreover, behind every smart object and its artificial intelligence, there is a single or collective "*hidden mind*"—the mind of its designer(s), which may (or may not) target valuable, or good achievements. The outcome of their operation, in the field, will be beneficial to whom? Or, harmful to whom? Often, using NI, the outcomes in human societies are beneficial to a very limited minority, while they exploit the scarce resources of large communities…

This brings to a ***sixth challenge***: ***the definition of Laws and Ethics in the IoT domain, and object accountability***. Humans face Laws, and Ethics. Adult humans are considered responsible, in face of the Law. Humans are *accountable*. Human behavior is monitored. Should smart objects be facing also Law and Ethics? *Who*

should be defining such Ethics and Laws? As mentioned, efforts are under way to define Ethic values for IoT, like the ones carried out by IEEE, the University of Montreal or the European Union [2–4]. Panels of experts try to reach consensus on these issues. However, experts from a given community may be biased. In a recent paper [29], a Harvard Law professor warned that people should not "let industry write the rules for AI". Even unwillingly, they may be protecting corporate interests, not people's interests. AI is not supposed to be solely driven by the goal of maximizing company profit [3]. People (and a sustainable environment in the planet) are more important than companies.

Moreover, how will the design of such objects take into account these Laws, so that they are aware of the Laws, and eventually prevent Law violation? How do we make objects that know what *to do* (mission functionality), and what *to avoid doing*? Who will monitor, in real time, AI-driven IoT systems operation, and the consequences of the decisions and actions they take? Will autonomous smart objects be accountable, as humans are?

Similar to the human society, in which children are not accountable, but, as they grow up, they are progressively allowed, by adults, to take autonomous decisions, the society of *things* should be *progressively accountable*, as the scope of autonomous decisions is progressively enlarged by humans, while retaining control of the process. Hence, *Humans should define the boundaries of smart object autonomy* (either individually, or as part of an object community). This is crucial, in order for humans to keep control of this second intelligent society, making it a useful tool to serve the first intelligent society, Humanity.

So far, the first intelligent society has been "the master", and smart objects have been "the slave". Even the human society has *regulators*, people or Institutions that audit their actions. Shouldn't the AI-driven IoT society have also regulators? Or, should the AI-driven IoT society help humans to decide how AI itself is being progressively introduced? The second society is being developed to enrich people's lives, e.g., alleviating humans from tedious tasks or to perform huge data processing. Nevertheless, in the presence of two competing societies, which one will finally retain the leadership? Is there a risk of a dictatorship of the IoT society over the human society? Or, worst, is there a risk that some humans take control of the IoT society, in order to assume a dictator role over human society?

One way to constrain smart objects decisions, is *to limit the access to external data.* Moreover, *standards* must be established *for* guaranteeing the quality and scope of the *trained data seed,* or initial model with which objects start to "reason" as well as for the scope of the *ML algorithms*. In other words, individuals in human societies act under the Law. The Law establishes the limits of the human freedom. One way to establish this kind of limits in IoT is by limiting the access of each object to the data generated by any other object, or captured through the Internet. The Law also limits the actions each human can take. Similarly, *standards* should limit the scope of the ML engine.

A ***seventh challenge*** is ***Time-to-Market*** **(*TtM*)**. Time-to-Market, together with short period product lifetime, push hardware and software development to the edge. New products need to be introduced in the market as soon as possible, prior to

products introduced by competitors. Often, products are released without a thorough examination. This can be disastrous for safety-critical applications (automotive, aero spatial, medical applications, to name a few). For instance, it can cause an airplane crash (e.g., due to a flaw in the flight-control system) … Unfortunately, Murphy's Law is not a joke. Should we empower objects allowing that autonomous **critical** decisions will be made by them? Or, should AI and ML data help to regulate and decide the Time to-Market of products and Things?

5 Towards an Intelligent *Strategy* for AI-Driven IoT. Areas of Urgent Research

As referred, all identified challenges (and eventually others we fail to identify), lead to the conclusion that an intelligent *strategy* for AI-driven IoT systems must be developed. Instead of just letting running wild AI-driven IoT systems development, a set of guidelines and actions should be established. As a result, an improvement of the usefulness of IoT to mankind may be achieved.

5.1 Engineering Education

We are aware that two competing worlds already coexist—the *real world*, where humanity prevails, and the *virtual world*, which is assumingly being developed to help people to be happy, and to prosper. The virtual world is populated by hw/sw systems, which tend to be interconnected in arrays of smart objects, using the Internet as a giant platform. The virtual world, which already dominates all aspects of human society (Fig. 4), is rapidly capturing the keen interest of humans. Many people prefer to spend long hours surfing the Internet, touring social networks, posting irrelevant (and often, false) information, playing videogames, rather than to live their lives. Unfortunately, sometimes people look for a refuge in the virtual world, to escape (or forget) an unfortunate real world.

Beyond the social impact, note that *human activity* (e.g., economics, business, investment, government, education, science and so on) is rapidly being transferred and carried out in the virtual world. This may be considered very attractive, as it may be faster, and less expensive. However, such activity may become more difficult to manage and control.

Taking into consideration such impact of virtual world on the real world, it is mandatory that hw/sw systems engineers be prepared to deal with this fact. Hence, beyond the mission subjects of an engineering program, there is a large set of *multi-disciplinary knowledge* that must be acquired, if someone will reach a really "higher" education, that is adequate to the deal with present reality.

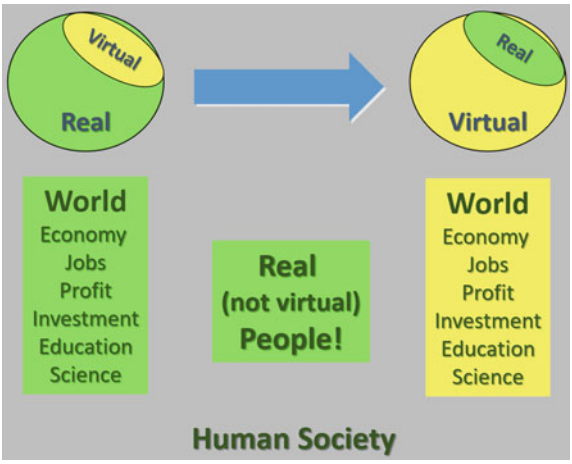

Fig. 4 The growing influence of the virtual world

In the past, people realized that engineering education should be complemented with management and business administration skills, so good ideas could be implemented in new products, and these could find a success path to the market and to the creation of wealth. Many engineers took MBA post-graduate courses, in order to be successful management professionals. Many become entrepreneurs, starting new companies.

More recently, as the planet resources are being wasted, and climate change is becoming a nightmare, there has been a growing awareness that a sustainable growth is mandatory. From the engineering education point of view, the concept of *green design* has emerged [30].

Now, the development of AI brought a new reality into stage. Hardware/Software system design has a new dimension, because smart objects are no longer time-invariant, in their architecture. As the ML procedures rearrange data bases, heuristics, algorithms and software code during product lifetime, the operating systems drift away from their original behavior.

Additionally, since new telecommunication technologies, such as 5G technologies [1, 10], are streaming into the market, and AI is continuously being implemented in complex hw/sw systems, making them ever smarter objects in IoT systems, there is a growing need to educate engineers to be aware of the impact that the "*things*" society will have on the human society.

The starting point is that education *professionals*, such as university professors, need to be aware and persuaded that a human society without common principles, Ethics and Law will not lead to a valuable outcome.

Hence, these professionals need to analyze the problem, foresee the consequences and be *convinced* that inserting these themes in a graduate program is a real asset. Moreover, as society is in rapid transformation, Continuing Education (CE) is becoming more and more urgent. Hence, from an engineering education perspective, *formal academic programs* and *CE courses* should be developed and presented to a

large audience, in order to guarantee that the message is rapidly spread, not only among engineer's professionals, but also among opinion-makers, decision-makers, managers, economists and social science experts.

From an engineering point of view, education professionals must also acknowledge that building AI-driven IoT systems that help and assist humans clearly influences these systems' *design process*.

In fact, the design variables and constraints are now expanded, and need to be clearly formulated. Problem analysis need to be deepened. Previously, the design process consisted in the design of *the mission functionality*. Then, we moved to DfT—Design for Testability. As a consequence, not only the design of the mission functionality is modified, to accommodate a more testable architecture, but also a test process is added - eventual embedding *self-testing* functionality (e.g., to allow architecture partitioning in test mode of operation) and defining the test stimuli and expected correct responses, so that the hardware of each design copy can be tested.

Software testing procedures must also be developed. At this stage, the design process of smart objects should be considering the ability of monitoring not only *correct operation* of their functional parts, but also *the soundness of the decisions* such smart objects take, during their lifetime, as well as the consequences of *the actions* which are triggered by the decisions taken. For achieving that purpose, Laws and Ethics are to be established for smart objects through the definition of new *standards*. The design process should allow the object, or an external auditor, to verify if these standards are embedded in the architecture of the smart object set. A large avenue of research is opened due to this new perspective. Results of such research should be fueled to academic curricula.

The professor's role on these issues remains essential, to make *students* aware of the challenges of the brave new world we live in, and how to improve the design process, taking into account these additional constraints. The awareness of such challenges, and how to overcome them, needs to be conveyed in graduate and post-graduate courses, and disseminated to industry and to society.

5.2 Design for Accountability (DfA)

Taking into consideration the impact that IoT systems will have in the human society, it is mandatory that smart objects are accountable for the *decisions* they make, and for the *actions* they trigger, since they may lead to people's wellbeing or, alternatively, to catastrophic results. This is why the concept of Design for Accountability (DfA) is introduced.

Accountability deals with being responsible, and being able to show what has been achieved, and how, which allows an external auditor to evaluate and judge the results of a given performance. In this context, by DfA we refer that IoT object and system design must take into consideration the possibility to assess the soundness ML processes, and of any object decision. This requires that the system (or an outside agent) must be able to evaluate the *consequences* of each object decision. This is not a

trivial task, particularly, if the final results are not completely known by the designer, which may be the case in various circumstances.

Depending on the *empowerment* given to a smart objects set, and on the domain of application of the IoT system, objects may be granted with the attribute to make *autonomous* decisions or not. If objects can make autonomous decisions, DfA deals with the design process to take into account decision soundness evaluation *after* system operation took place. If objects are not allowed to make autonomous decisions, decisions are first scrutinized by an external auditor, and then implemented or not, depending on the auditor's judgement. Of course, if the application of an IoT system is e.g., autonomous vehicles in urban traffic, objects must be allowed to make autonomous decisions, in real time. In other applications, in which there is time and resources to do so, it may be rewarding to first evaluate decision's soundness, and then applied them, or not.

A first example of object's accountability is antivirus software, in which a software module monitors the legitimacy of another software code to access and modify information in a given system. A second example is when an AI based system is used, to explore space. What is the range of values in the incoming data that is to be taken into consideration? What is the range of values that should be discarded, as outliers? The design engineer can provide an initial guess of that range of values. Yet, if the AI system detects a huge amount of data outside that range, should it be kept or discarded?

As stated, the above mentioned challenges in IoT-AI partnership (Fig. 3) should be seriously taken into account in *engineering education*, since they correspond to the problems new engineers will face. The comments provided in this sub-section are intended to provide some insight on the possible solutions to the identified challenges. Nevertheless, a significant R&D effort needs to be performed, to reach sound solutions.

Design challenges are addressed here from the *engineering* point of view. System designers need to turn out words or sentences into additional engineering requirements and specifications in the design of complex systems constituted by (eventually, among others) smart hw/sw objects, communicating through the Internet.

Prior to assess the soundness of object decisions, it is necessary to guarantee their correct operation, through the implementation of the *hardware test* procedure.

In order to make sure the hardware part is operating correctly, during product lifetime, it is mandatory that no external agent may be allowed to modify the hardware part. This is easy in non-programmable hardware. However, for programmable hardware, like the one using Field-Programmable Gate Arrays (FPGAs), it is mandatory that no external agent is allowed to reprogram the hardware functionality, unless some failure in the configuration memory (caused, e.g., by a Single Event Upset (SEU)) occurs. In such case, reconfiguration is carried out only to reestablish correct functionality.

In terms of the *software test* process, it is another story. As pointed out in [31, 32], the development paradigm of hw/sw systems using ML and AI techniques is completely different from the "traditional" hw/sw systems development. In fact, in the first, the functionality is known, the algorithms are implemented in order to

achieve a given functionality, and we know *what* to test, since, functionality is written in a *deductive* way, by writing down the rules as program code.

On the contrary, using ML techniques, rules are inferred from data and requirements are derived *inductively*. As the learning process modifies the rules, the outcome is not entirely known. This makes software test particularly difficult, given the fact that we do not have complete specifications (some of them may be inferred from data), the knowledge of the source code corresponding to inferred specification and yet less some of their critical behaviors. The problem is that, we may not know what to expect, and less *what* or *how* to test the unknown.

Although the ML process may present some opacity, there are some aspects that we know, namely, the *objectives* (or goals) we want the system to reach. Therefore, the decisions taken by smart objects should lead to the achievement of our goals. *Decision's quality* may be defined as its ability to achieve pre-specified goals. Consequently, the design process of such smart systems should consider the ability to monitor, not only the *correct operation* of their functional parts, but also the *soundness of the decisions* such smart objects take, during system operation. In this context, is there room for software self-test and for assessing decisions soundness?

To do so, designers need to establish a set of *metrics*. For instance, we may set (1) the range of values that we consider acceptable in terms of our goals, and (2) the maximum number of "iterations" the ML algorithm should run to reach the goal. *Merit functions* must be established, identifying relevant variables and their weighting factors. To be sure, these variables may now include those corresponding to ethical constraints. The variable set, and the weighting factors may be updated during product lifetime, as the object's learning process progresses.

In traditional hw/sw systems test, it is necessary to identify the input data, the system's functionality and the correct system's response. For the reasons presented in Sect. 4, in smart objects systems design, a key issue is to limit and identify *the data* each object is allowed to access, and to make sure that no malicious data is introduced in the system. The Internet provides an overwhelming set of data and information; however, the system under design is only allowed to access a very limited subset of that data, from other smart objects and/or from humans, navigating also in the Internet. The design should also take into consideration some sort of data evaluation, prior to processing. For instance, is unauthorized data collection taking place? If so, disregard the incoming data.

In terms of testing the software part of the smart object under design, the designer must now take into account that the testing process will include (1) *production* test, and (2) *lifetime* testing. Hence, the ability to test, during product lifetime, must include the ability of the object to test itself (*self-test*), or to be tested by another agent (*external test*), either another higher-hierarchy smart object, or a human. It may be necessary to test in real-time (*online testing*), or periodically (taking into advantage time periods in which the object is idle). In order to test correct operation, correct system responses should be observable. In order to test the soundness of object autonomous decisions, additional software must be embedded in-system.

In order to allow external test, as mentioned before, the external agent must know *what* to test. Hence, a key concept of DfA is that smart objects must *report the software modifications* introduced by the learning process, and by the decision process that may modify e.g., the internal data bases. This may not be easy to implement, in the design process. However, it makes possible for an external auditor to improve system test, especially software test. Unlike humans, that make changes along their lives and often try to hide their internal feelings and thoughts, smart objects can be designed in such a way that inner modifications are reported to the external world.

System designers must also take into account *design reuse*, and *legacy*. Will the next generation product be smarter than the previous one? This quests for taking advantage of the learning process of the previous generation; again, the *report* of software modifications will save time and money.

In a similar way, reusing design methodologies in a new context can be very rewarding. Software design methodologies, namely an extensive problem analysis, design specifications, architectural design and test, have been reused in hardware (and hw/sw) methodologies with success.

Moreover, hardware methodologies for safety-critical applications, using time or hardware *redundancy*, could be used in software design of complex, smart objects. Time redundancy involves executing a given functionality more than once, to make sure the operation is correct. Hardware redundancy techniques, like Triple Modular Redundancy (TMR) [21] triplicate the hardware module, apply the same input data to the three replicas and compare the outputs of the three modules. If outputs differ, a majority voting process takes place, and the correct output is applied to the following module. Can software redundancy techniques be applied in smart object design? For instance, if the designer does not know beforehand what are the best solutions for a given problem, a set of objects, running in parallel, may be programmed as seeds with different metrics, eventually different variables and/or weighting factors. By using the same input data, it is possible to compare the outputs of different learning and decisions processes, ascertaining what will be the "best" solution to implement in the field, after an experimental period of time.

As referred in Sect. 3, Humans have an inherent feature: *conscience*. The sense of right or wrong guides this mysterious inner moral judge to either praise us when we make good decisions, or find us guilty when we act in an erroneous way. Depending on the circumstances, the trial is made *before* or *after* the decision and action take place. When people face an unknown situation, and have to make fast decisions, usually the trial is made afterwards. On the contrary, if we have time to evaluate and decide, the trial is made in advance, and the moral sense, or the Ethics, are taken into account in the decision process.

Similarly, in smart object design, if acceptable Ethics are assumed as system requirements, and transformed into standards and rules, these ethic rules should be considered in the learning and decision-making processes, leading to more profitable IoT systems. How to do it, is a very interesting theme of research.

Nevertheless, system design needs to be performed in a way that *Ethic rules must be inerasable*, i.e., smart objects must not be allowed to decide without taking Ethics into consideration. As in the human society, where Ethics are changing with

time, ethic rules inserted in smart objects should be updatable. Embedded object "conscience" should be mandatory for product homologation, probably as a distinct object, operating as *master*, and should be viewed as the "regulator" in the human society.

Ethics standards will allow each smart object to analyze itself, ascertaining if a given decision or triggered action is (or is not) in compliance to the engraved *principles* embedded in itself, and in its follow objects. As in the human society, the trial may be performed either before, or after decisions and actions take place. Of course, the outcomes of such internal judgement may vary, whether the object itself, a master object or a human auditor will take action if the decisions broke any Law.

These questions led us to this innovative concept—***Design for Accountability* (*DfA*)**. As AI-driven IoT systems are allowed to make more and more autonomous decisions, they should be progressively accountable for their acts.

In order to evaluate the soundness of the decisions taken by the IoT system, often it takes a long time, allowing the cause-effect paradigm to take place. Hence, an Audit process may periodically take place, and be performed either internally, or more reasonably, externally by an independent agent. For instance, in medical diagnosis and subsequent treatment, medical trials and clinical evaluations require many patients, many years, and a significant amount of data. However, is the audit process that leads to progress in medical treatment and in the homologation of pharmaceutical products.

On the contrary, decision and action evaluation in safety-critical IoT systems must be performed as soon as possible, as some actions may be identified as harmful, potentially catastrophic. A collection of forbidden actions for each smart object should be identified, avoiding the object to autonomously make such decisions.

6 Conclusions

Humanity reached a new crossroad, with the advent of the Internet, and of AI-driven IoT systems. A new intelligent society, generated and empowered by humans, is under active development. It is our conviction that the creators must *master* their creation.

Relevant *challenges* in this path have been identified and discussed. How to test the smart object's modified software part, IoP-IoT interactions, the need to evaluate objects decisions quality, data capturing, control of objects empowerment, Ethics and Laws for IoT, smart objects accountability and Time-to-Market are the main challenges for which solutions need to be found.

A *strategy* for intelligent and controlled IoT society is, thus, mandatory, in order to achieve a harmonious, happy outcome. Such strategy should include *engineering education*, and *Design for Accountability (DfA)*. In reality, engineering education should precede a multi-disciplinary education, from academia down to the business environment, as additional actors in other domains need to be aware of the consequences of misusing the emerging AI-driven IoT society. This educational effort

should start by engineering education, due to the fact that engineers are the IoT system *designers*, who must develop ways to introduce root solutions for the above mentioned challenges.

As we saw, smart objects monitoring requires the ability of outside agents to be able to verify objects *correct operation* as well as the *soundness* of their *decisions* and triggered *actions*. Embedded functionality in these objects must drive each object to report ongoing transformations. Special care needs to be taken into account, when safety-critical applications are considered.

A fundamental issue is *the insertion of Ethics as standards in IoT* devices and systems, in an inerasable form. Embedded Ethics must be rooted to true Human values, to serve the human society, and act like an inner conscience of these smart objects. No privacy should be given to IoT devices and systems, in the sense that at any moment the inner values, embedded on them, must be accessible and verifiable from outside, eventually corrected if any malicious "values" are introduced in objects. The scope allowed by humans to smart objects to make *autonomous decisions* should be enlarged in a progressive way, as parents do with their children.

As a consequence, it becomes clear that a huge amount of R&D effort lays ahead of us. It must be pursued urgently, in order to obtain a wining partnership between IoT and AI.

References

1. G5: Moving to the next generation in wireless technology. https://www.sciencedaily.com/releases/2015/04/150430082723.htm
2. Ethically Aligned Design—Version II, Request for Input.: IEEE Global Initiative on Ethics of Autonomous and Intelligent Systems. https://standards.ieee.org/industry-connections/ec/autonomous-systems.html
3. Bengio, Y., Université de Montréal.: Montreal Declaration for a Responsible Development of Artificial Intelligence (2018). Available at https://www.montrealdeclaration-responsibleai.com/
4. High-Level Expert Group on Artificial Intelligence set up by the EU.: Ethic Guidelines for Trustworthy AI, European Commission, document made public in 8 Apr 2019. Available at https://ec.europa.eu/digital-single-market/en/news/ethics-guidelines-trustworthy-ai
5. Artificial Inteligence Tutorial. https://www.tutorialspoint.com/artificial_intelligence/
6. Nick McCrea—An Introduction to Machine Learning Theory and Its Applications: A Visual Tutorial with Examples in https://www.toptal.com/machine-learning/machine-learning-theory-an-introductory-primer
7. Xin, Y., Kong, L., Liu, Z., Chen, Y., Li, Y., Zhu, H., Gao, M., Hou, H., Wang, Ch.: Machine learning and deep learning methods for cybersecurity. IEEE Access **6**, 35365–35381 (2018)
8. Nouwens, M., Legarda, H.: China's pursuit of advanced dual-use technologies, Dec 2018. Available at https://www.iiss.org/blogs/analysis/2018/12/emerging-technology-dominance
9. Gubbi, J., Buyya, R., Marusic, S., Palaniswami, M.: Internet of Things (IoT): a vision, architectural elements, and future directions. Future Gener. Comput. Syst. **29**(7), 1645–1660 (2013)
10. Park, T., Abuzainab, N., Saad, W.: Learning how to communicate in the internet of things: finite resources and heterogeneity. IEEE Access: Optim. Emerg. Wirel. Netw.: IoT, 5G Smart Grid Commun. Netw. Spec. Session, IEEE Access **4**, 7063–7073 (2016)

11. Data Management for Artificial Intelligence. https://www.sas.com/en_za/whitepapers/data-management-artificial-intelligence-109860.html
12. Yu Shwe, H., King Jet, T., Han Joo Chong, P.: An IoT-oriented data storage framework in smart city applications. In: 2016 International Conference on Information and Communication Technology Convergence (ICTC), pp. 106–108 (2016)
13. Mohon Ghosh, A., Halder, D., Alamgir Hossain, S.K.: Remote health monitoring system through IoT. In: 2016 5th International Conference on Informatics, Electronics and Vision (ICIEV), pp. 921–926 (2016)
14. Yamauchi, T., Kondo, H., Nii, K.: Automotive low power technology for IoT society. In: 2015 Symposium on VLSI Circuits (VLSI Circuits), pp. T80–T81 (2015)
15. Nelson, G.S.: Getting started with data governance. In: Presented at the Annual Conference of the SAS Global Users Group, Dallas, TX, 28 Apr 2015
16. Ruan, J., et al.: An IoT-based E-business model of intelligent vegetable greenhouses and its key operations management issues. Neural Comput. Appl. 1–16 (2019)
17. He, H., et al.: The security challenges in the IoT enabled cyber-physical systems and opportunities for evolutionary computing & other computational intelligence. In: 2016 IEEE Congress on Evolutionary Computation (CEC), Vancouver, BC, pp. 1015–1021 (2016)
18. Xu, T., Wendt, J.B., Potkonjak, M.: Security of IoT systems: design challenges and opportunities. In: Proceedings of IEEE/ACM International Conference on Computer-Aided Design (ICCAD), pp. 417–423 (2014)
19. Ethically Aligned Design—A Vision for Prioritizing Human Well-being with Autonomous and Intelligent Systems, version 2. IEEE (2018). Available at http://theinstitute.ieee.org/resources/standards/ieee-releases-new-ethical-considerations-for-autonomous-and-intelligent-systems)
20. Mollick, E.: Establishing Moore's Law. IEEE Ann. Hist. Comput. **28**(3), 62–75 (2006)
21. Bushnel, M.L., Agrawal, V.D.: Essentials of Electronic Testing for Digital Memory and Mixed-Signal VLSI Circuits. Kluwer Academic Publishers (2000)
22. Will Safety-Critical Design Practices Improve First Silicon Success?. Mentor Graphics White Paper (2017). Available at http://s3.mentor.com/public_documents/whitepaper/resources/mentorpaper_102839.pdf
23. Semião, J., Cabral, R., Cavalaria, H., Santos, M.B., Teixeira, I.C., Teixeira, J.P.: Ultra-low-power strategy for reliable IoE nanoscale integrated circuits. In: Harnessing the Internet of Everything (IoE) for Accelerated Innovation Opportunities. IGI Global (2019)
24. Valdés, M., Freijedo, J., Moure, M.J., Rodríguez-Andina, J.J., Semião, J., Vargas, F., Teixeira, I.C., Teixeira, J.P.: Design and validation of configurable on-line aging sensors in nanometer-scale FPGAs. IEEE Trans. Nanotechnol. **12**(4), 508–517 (2013)
25. Charles de Secondat, Baron of Montesquieu "All men are born equal", The Spirit of the Laws, 1748. Available at https://americanart.si.edu/artwork/state-nature-indeed-all-men-are-born-equal-they-cannot-continue-equality-society-makes-them
26. Strubell, E., Ganesh, A., McCallum, A.: Energy and Policy Considerations for Deep Learning in NLP. University of Massachusetts at Amherst (2019)
27. Mavromoustakis, C.X., et al.: Socially oriented edge computing for energy awareness in IoT architectures. IEEE Commun. Mag. **56**(7), 139–145 (2018)
28. Big data needs a hardware revolution. Nature **554**, 145–146 (2018). https://doi.org/10.1038/d41586-018-01683-1
29. Benkler, Y.: Don't let industry write the rules for AI. Nature **569**(161), 2019 (2019). https://doi.org/10.1038/d41586-019-01413-1
30. Irimia-Vladu, M.: Green electronics: biodegradable and biocompatible materials and devices for sustainable future. Chem. Soc. Rev. **43**, 588–610 (2014)
31. Cagala, T.: Improving data quality and closing data gaps with machine learning. In: IFC National Bank of Belgium Workshop on Data Needs and Statistics Compilation for Macroprudential Analysis. Brussels, Belgium, 18–19 May 2017
32. Khomh, F., Adams, B., Cheng, J., Fokaefs, M., Antoniol, G.: Software engineering for machine-learning applications—the road ahead. IEEE Comput. Edge, pp. 21–24 (2019) (also in IEEE Software, vol. 35, no. 5, 2018)