

1

Introduction

What is a cyber-physical system? Why should I study it? What are its relations to cybernetics, information theory, embedded systems, industrial automation, computer sciences, and even physics? Will cyber-physical systems (CPSs) be the seed of revolutions in industrial production and/or social relations? Is this book about theory or practice? Is it about mathematics, applied sciences, technology, or even philosophy? These are the questions the reader is probably thinking about right now. Definitive answers to them are indeed hard to give at this point. During the reading of this book, though, I expect that these questions will be systematically answered. Hence, new problems and solutions can then be formulated, allowing for a progressive development of a new scientific field.

As a prelude, this first chapter will explicitly state the philosophical position followed in this book. The chapter starts by highlighting what is spontaneously thought under the term “CPS” to then argue the reasons why a general theory is necessary to build scientific knowledge about this object of inquiry and design. A brief historical perspective of closely related fields, namely control theory, information theory, and cybernetics, will also be given followed by a necessary digression of philosophical positions and possible misinterpretations of such broad theoretical constructions. In summary, the proposed demarcation can be seen as a risk management action to avoid mistakes arising from commonsense knowledge and other possible misconceptions in order to “clear the path” for the learning process to be carried out in the following hundreds of pages.

1.1 Cyber-Physical Systems in 2020

Two thousand and twenty is a remarkable year, not for the high hopes the number 20-20 brought, but for the series of critical events that have happened and affected everyone’s life. The already fast-pace trend of digitalization, which had started decades before, has boomed as a consequence of severe mobility restrictions

imposed as a response to the COVID-19 pandemics. The uses (and abuses) of information and communication technologies (ICTs) are firmly established and widespread in society. From dating to food delivery, from reading news to buying e-books, from watching youtubers to arguing through tweets, the *cyber* world – before deemed in science fiction literature and movies as either utopian or apocalyptic – is now very concrete and pervasive. Is this concreteness of all those practices involving computers or computer networks (i.e. cyber-practices) what defines *CPSs*? In some sense, yes; in many others, no; it all depends on how CPS is conceptualized! At all events, let us move step-by-step by looking at nonscientific definitions.

CPS is a term not broadly employed in everyday life. Its usage has a technical origin and is related to digitalization of processes across different sectors so that the term “CPS” has ended up being mostly used by academics in information technology, engineering, practitioners in industry, and managers. Such a broad concept usually leads to misunderstandings so much so that relevant standardization bodies have channeled efforts trying to establish a shared meaning. One remarkable example is the *National Institute of Standards and Technology* (NIST) located in the United States. NIST has several working groups related to CPS, whose outcomes are presented on a dedicated website [1]. In NIST’s own words,

Cyber-Physical Systems (CPS) comprise interacting digital, analog, physical, and human components engineered for function through integrated physics and logic. These systems will provide the foundation of our critical infrastructure, form the basis of emerging and future smart services, and improve our quality of life in many areas.

Cyber-physical systems (CPS) will bring advances in personalized health care, emergency response, traffic flow management, and electric power generation and delivery, as well as in many other areas now just being envisioned. CPS comprise interacting digital, analog, physical, and human components engineered for function through integrated physics and logic. Other phrases that you might hear when discussing these and related CPS technologies include:

- Internet of Things (IoT)
- Industrial Internet
- Smart Cities
- Smart Grid
- “Smart” Anything (e.g. Cars, Buildings, Homes, Manufacturing, Hospitals, Appliances)

As a commonplace when trying to determine the meaning of umbrella terms, the definition of CPS proposed by NIST is still too broad and vague (and excessively utopian) to become susceptible of scientific inquiry. On the other hand, such a




Figure 1.1 Illustration of a CPS. Sensors measure physical processes, whose data are transmitted through a communication network. These data are then processed to support decisions related to the physical process by either a human operator or an expert system.

definition offers us a starting point, which can be seen as the raw material of our theoretical investigation. A careful reading of the NIST text indicates the key common features of the diverse list of CPSs:

- There are physical processes that can be digitalized with sensors or measuring devices;
- These data can be processed and communicated to provide information of such processes;
- These informative data are the basis for decisions (either by humans or by machines) of possible actions that are capable of creating “smartness” in the CPS;
- CPSs are designed to intervene (improve) different concrete processes of our daily lives; therefore, they affect and are affected by different aspects of society.

These points indicate generalities of CPSs, as illustrated in Figure 1.1. In most of the cases, though, they are only *implicitly* considered when particular solutions are analyzed and/or designed. As a matter of fact, specific CPSs do exist in the real world without the systematization to be proposed in this book. So, there is an apparent paradox here: on the one hand, we would like to build a scientific theory for CPSs in general; on the other hand, we see real deployments of particular CPSs that do not use such a theory. The next section will be devoted to resolve this contradiction by explaining the reasons why a general theory for CPS is necessary while practical solutions do indeed exist.

1.2 Need for a General Theory

The idea of having a *general theory* is, roughly speaking, to characterize in a nonsubjective manner a very well-defined symbolic object that incorporates all the constitutive aspects of a class of real-world objects and therefrom obtain new knowledge by both symbolic manipulation and experimental tests. This generalization opens the path for moving beyond know-how-style of knowledge toward abstract, scientific conceptualizations, which are essential to assess existing objects, design new ones, and define their fundamental limits.

Example 1.1 *Chocolate cake.* You would like to prepare the chocolate cake you ate in your childhood, whose recipe you found in the Internet. This instructs you *how* to make it. If you follow the instructions, you will have the cake you wish but without knowing *why* that specific combination of elements and the processes of mixing and cooking can produce such a cake. If you decide to investigate the recipe, you will find that different elements are compatible for chemical reasons, and, after processing then in a specific way, such a combination will have the desired flavor and structural characteristic. Based on this abstraction, you can now (i) understand the reasons why the recipe works, (ii) propose modifications to the recipe to make the cake fluffier or moister, and (iii) create a vegan version of the cake by finding replacement for milk, butter, and eggs.

This example serves as a very simplified illustration of the difference between technical knowledge (know-how) and scientific knowledge. We are going to discuss sciences and scientific practice in more detail when pinpointing the philosophical position taken throughout this book. At this point, though, we should return to our main concern: the need for a general theory that conceptualizes CPSs. Like the particular chocolate cake, the existence of *smart* grids or cities, or even fully automated production lines neither precludes nor requires a general theory. Actually, their existence, the challenges in their particular deployments, and their specific operation can be seen as the necessary raw material for the scientific theory that would build the knowledge of CPSs as a symbolic (general, abstract) object. This theory would provide the theoretical tools for orienting researchers, academics, and practitioners with objective knowledge to analyze, design, and intervene in particular (practical) realizations of this symbolic object called CPS.

Without advancing too much too soon, let us run a thought experiment to mimic a specific function of *smart* meters as part of the *smart* electricity grid – one of the most well-known examples of CPS. Consider the following situation: the price of electricity in a household is defined every hour and the smart meter has access to this information. The smart meter also works as a home energy management system, turning on and off some specific loads or appliances that have flexibility in their usage as, for instance, the washing machine or the charging of an electric vehicle (EV). If there is no smartness in the system, whenever the machine is turned on or the EV is plugged in, they will draw electric energy from the grid. With the smart meter deciding when the flexible load will turn on or off based on the price, the system is expected to become smart overall: not only flexible loads could be turned on when there is a low price (leading to lower costs to the households) but this would also help the grid operation by flattening the electricity demand curve (which has peaks and valleys of consumption ideally reflected by the price).

This seems too good to be true, and it indeed is! The trick is the following. The electricity price is a *universal signal* so that all smart meters see the same value. By

facing the same price, the smart meters will tend to switch on (and then off) the loads at the same time, creating a collective behavior that would probably lead to unexpected new peaks and valleys in demand that cannot be predicted by the price (which actually reflects past behaviors, and not instantaneous ones). The adage *the whole is more than the sum of its parts* then acquires a new form: the smartness of the smart meters can potentially yield a stupid grid [2]. This outcome is surprising since individually everything is working as expected, but the system-level dynamic is totally undesirable. How to explain this?

The smart grid, as we have seen before, is considered a CPS where physical processes related to energy supply and demand are reflected in the cyber domain by a price signal that serves as the basis for the decisions of smart meters, which then modify the physical process of electricity demand by turning on appliances. However, the smart meters described above are designed to operate considering the grid dynamics as given so that they individually react to the price signal assuming that they cannot affect the electricity demand at the system level. If several of such smart meters operate in the grid by reacting to the same price signals, they will tend to have the same decisions and, consequently, coordinate their actions, leading to the undesirable and unexpected aggregate behavior. This is a byproduct of a segmented way of conceptualizing CPSs, which overestimates the smartness of devices working individually while underestimating the physical and logical (cyber) interrelations that constitute the smartness of the CPS.

Figure 1.2 illustrates this fact. The electricity hourly price is determined from the expected electricity demand following the indication of arrow (a). The smart meters would lead to a smart grid if the spikes in consumption would be flattened, as indicated by arrow (b). However, due to the unexpected collective effect

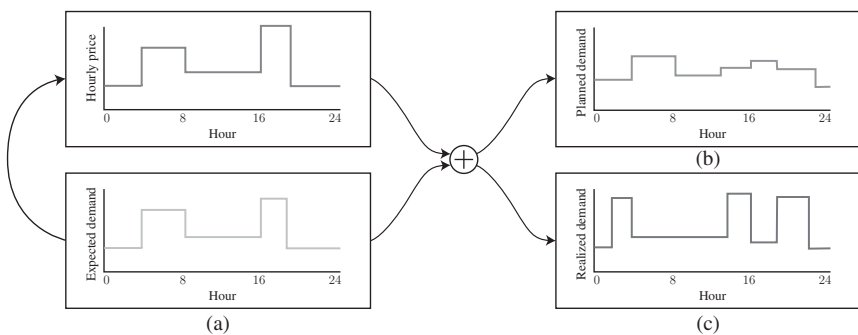


Figure 1.2 Operation of smart meters that react to hourly electricity price. Arrow (a) indicates that the hourly price is a function of the expected electricity demand. Arrow (b) represents the planned outcome of demand response, namely decreasing spikes. Arrow (c) shows the realized demand as an aggregate response to price signals, leading to more spikes.

of reactions to the hourly price, the realized demand has more spikes than before, creating a stupid grid. This is pointed by arrow (c).

At all events, coordination and organization of elements working together are old problems under established disciplines like systems engineering and operational research. However, those disciplines are fundamentally based on centralized decision-making and optimal operating points; this is usually called a top-down approach. In CPSs, there is an internal awareness and a possibility of distributed or decentralized decision-making based on local data and (pre-defined or learned) rules co-existing with hierarchical processes. Hence, CPSs cannot be properly characterized without explicit definition of how data are processed, distributed, and utilized for informed decisions and then actions. A general (scientific) theory needs to be built upon the facts that CPSs have internal communication-computation structures with specific topologies that result in internal actions based on potentially heterogeneous decision-making processes that internally modify the system dynamics. Some of these topics have been historically discussed within control and information theories, as well as cybernetics.

1.3 Historical Highlights: Control Theory, Information Theory, and Cybernetics

Control theory refers to the body of knowledge involving automatic mechanisms capable of self-regulating the behavior of systems. Artifacts dating back to thousands of years indicate the key idea behind feedback control loop. The oldest example is probably the *water clock* where inflow and outflow of water is used to measure time utilizing a *flow regulator*, whose function is [3]: “(...) to keep the water level in a tank at a constant depth. This constant depth yielded a constant flow of water through a tube at the bottom of the tank which filled a second tank at a constant rate. The level of water in the second tank thus depended on time elapsed.” Other devices based on the same principle have also been found throughout history. With the industrial revolution in the 1700s, different types of *regulators* and *governors* have appeared as part of the unprecedented technological development. It was desirable for the new set of machinery like windmills, furnaces, and steam engines to be controlled automatically. Hence, more and more solutions based on feedback and self-regulation started to be developed. At that point, the development was of a practical nature, by trial-and-error. In a groundbreaking work by J. C. Maxwell (also the “father” of electromagnetism) from 1868 [4], the first step to generalize feedback control as a scientific theory was taken. Maxwell analyzed systems that employ governors by linearized differential equations to establish the necessary conditions for stability based

on the system's characteristic equation. This level of generalization that states fundamental (mathematical) conditions that apply to all existing and potentially future objects of the same kind (i.e. feedback control systems) is the heart of the scientific endeavor of *control theory*.

A similar path can be also seen with information theory. Data acquisition, processing, transmission, and reception are biological facts, not only in humans but in other animals. Without the trouble to define what information is now, it is clear that different human societies have come up with different ways of sending informative data from one point to another. From spoken language to books, from smoke signals to pigeons, data can be transferred in space and time, *always facing the possibility of error*. With the radical changes brought by the industrial revolution, quicker data transmission across longer distances was desired [5]. The first telegraphs appeared in the early 1800s; the first transatlantic telegraph dates to 1866. By the end of the 19th century, the first communication networks were deployed in the United States. Interestingly enough, the development of communications networks depended on amplifiers and the use of feedback control, which have been studied in control theory.

Even with a great technological development of communication networks, mainly carried out within the Bell Labs, transmission errors had always been considered inevitable. Most solutions were focusing on how to decrease the chances of such events, also in a sort of highly complex trial-and-error fashion. In 1948, in one edition of *The Bell System Technical Journal*, C. E. Shannon published one amazing piece of work stating in a fully mathematical manner the fundamental limits of communication systems utilizing a newly proposed definition of information based on *entropy*; some more details will come later in Chapter 4. What is important here is to say that, in contrast to all previous contributions, Shannon created a mathematical model that states the fundamental limits of *all* existing and possible communication systems by determining the capacity of communication channels. One of his key results was the counter-intuitive statement that *an error-free transmission is possible if, and only if, the communication rate is below the channel capacity*. Once again, one can see Shannon's paper as the birth of a scientific theory of this new well-determined object referred to as *information*.

In the same year that Shannon published *A Mathematical Theory of Communication*, another well-recognized researcher—Norbert Wiener—published a book entitled *Cybernetics: Or Control and Communication in the Animal and the Machine* [6]. This book introduces the term *cybernetics* in reference to self-regulating mechanisms. In his erudite writing, Wiener philosophically discussed several recent developments of control theory, as well as preliminary thoughts on information theory. He presented astonishing scientific-grounded arguments to draw parallels between human-constructed self-regulating

machines, on the one side, and animals, humans, social, and biological processes, on the other side. Here, I would like to quote the book *From Newspeak to Cyberspeak* [7]:

Cybernetics is an unusual historical phenomenon. It is not a traditional scientific discipline, a specific engineering technique, or a philosophical doctrine, although it combines many elements of science, engineering, and philosophy. As presented in Norbert Wiener's classic 1948 book *Cybernetics, or Control and Communication in the Animal and the Machine*, cybernetics comprises an assortment of analogies between humans and self-regulating machines: human behavior is compared to the operation of a servomechanism; human communication is likened to the transmission of signals over telephone lines; the human brain is compared to computer hardware and the human mind to software; order is identified with life, certainty, and information; disorder is linked to death, uncertainty, and entropy. Cyberneticians view control as a form of communication, and communication as a form of control: both are characterized by purposeful action-based on information exchange via feedback loops. Cybernetics unifies diverse mathematical models, explanatory frameworks, and appealing metaphors from various disciplines by means of a common language that I call cyberspeak. This language combines concepts from physiology (homeostasis and reflex), psychology (behavior and goal), control engineering (control and feedback), thermodynamics (entropy and order), and communication engineering (information, signal, and noise) and generalizes each of them to be equally applicable to living organisms, to self-regulating machines, and to human society.

In the West, cybernetic ideas have elicited a wide range of responses. Some view cybernetics as an embodiment of military patterns of command and control; others see it as an expression of liberal yearning for freedom of communication and grassroots participatory democracy. Some trace the origins of cybernetic ideas to wartime military projects in fire control and cryptology; others point to prewar traditions in control and communication engineering. Some portray cyberneticians' universalistic aspirations as a grant-generating ploy; others hail the cultural shift resulting from cybernetics' erasure of boundaries between organism and machine, between animate and inanimate, between mind and body, and between nature and culture.

We can clearly see a difference between the generality of information and control theories with respect to their own well-defined objects, and the claimed universality of cybernetics that would cover virtually all aspects of reality. In this sense,

the first two can be claimed to be scientific theories in the strong sense. The last, despite its elegance, seems less a science but more a theoretical (philosophical) displacement or distortion of established scientific theories by expanding their reach towards other objects. This is actually a very controversial argument that depends on the philosophical position taken throughout this book, whose details will be presented next.

1.4 Philosophical Background

Science is a special type of formal discourse that claims to hold objective true knowledge of well-determined objects. Different sciences have different objects, requiring different methods to state the truth value of different statements. A given science is presented as a *theory* (i.e. a systematic, consistent discourse) that articulates different *concepts* through a chain of *determinations* (e.g. causal or structural relations) that are *independent of any agent (subject)* involved in the production of scientific knowledge. This, however, does not preclude the importance of scientists: they are the necessary *agents* of the scientific practice. Scientific practice can then be thought as the way to produce new knowledge about a given object, where scientists work on theoretical raw material (e.g. commonsense knowledge, know-how knowledge, empirical facts, established scientific knowledge) following historically established norms and methods in a specific scientific field to produce new scientific knowledge. In other words, *scientific practice is the historically defined production process of objective true knowledge*. Note that these norms, despite not being fixed, have a relatively stable structure since the object itself constrains which are the valid methods eligible to produce the knowledge effect.

Moreover, scientific knowledge poses general statements about its object. Such a generality comes with abstraction, moving from particular (narrow) abstractions of real-world, concrete objects to abstract, symbolic ones. Particular variations of a class of concrete objects can be used as the raw material by scientists to build a general theory that is capable of covering all, known and unknown, concrete variations of that class of objects. This general theory is built upon abstract objects that provide knowledge of concrete objects. However, this differentiation is of key importance since a one-to-one map between the concrete and abstract realities may not exist. Abstract (symbolic) objects as part of scientific theories produce a knowledge effect on concrete objects, understood as realizations of the theory, not as a reduction or special case. At any rate, despite the apparent preponderance of abstractions, the concrete reality is what determines in the last instance the validity of the theory (even in the “concrete” symbolic reality of pure mathematics, concreteness is defined by the foundational axioms and valid operations).

To illustrate this position, let us think about dogs. Although the concept of dog cannot bark, dogs do bark. Clearly, in the symbolic reality in which the concept of dog exists, it has the ability of barking. The concept, though, cannot transcend this domain so we cannot hear in the real world the barking sound of the abstracted dog. Conversely, we all hear real dogs barking, and therefore, any abstraction of dogs that assumes that they cannot bark shall not be considered scientific at all. This seems trivial when presented with this naive example, but we will see throughout this book the implications of unsound abstractions in different, more elusive domains. This is even more critical when incorrect abstractions are accompanied by heavily mathematized (therefore consistent) models. For instance, the fact that some statement is a true knowledge in mathematical sciences does not imply it is true in economics. Always remember: a mathematically consistent model is not synonymous with a scientific theory.

Philosophy, like science, is also a theoretical discourse but with a very important difference: it works by demarcating positions as correct or incorrect based on its own philosophical system that defines *categories* and their relations [8]. Unlike scientific proofs, philosophy works through rational argumentation to defend positions (i.e. *theses*), usually trying to answer universal and timeless questions about, for example, existence of freedom. In this case, philosophy has no specific (concrete) object as sciences do; consequently, it is not a science in the way we just defined. Philosophy then becomes its own practice: rational argumentation based on a totalizing system of categories defining positions about everything that exists or not. Following this line of thought, philosophy is not a science of sciences; it can neither judge the truth value of propositions internally established by the different sciences nor state *de jure* conditions for scientific knowledge from the outside.

Besides, scientific and philosophical practices exist among several other social practices. They are part of an articulated historical social whole, where different practices coexist and interfere with each other at certain degrees and levels of effectivity. As previously discussed, scientific practice produces general objective true knowledge about abstract objects, which very usually contradicts the commonsense ideas that are usually related to immediate representations arising from other social practices of our daily lives. This clearly leads to obstacles to scientists, who are both agents of the scientific practice and individuals living in society. The totalizing tendency of philosophy also plays a role: it either *distorts* scientific theories and concepts to fit in universal systems of philosophical categories or *judges* their truth value based on universal methodological assumptions. This directly or indirectly affects the self-understanding of the relation that the scientists have with their own practice, creating new obstacles to the science development [9, 10].

Example 1.2 *Differences between daily language, philosophical categories, and scientific concepts.* One word that exemplifies very well the difference is

time. We use word *time* in many ways in our daily lives: to discuss about our activities, plans, routine, and the like. However, in philosophy, the category *Time* has different roles depending on the philosophical system to be considered – this usually comes with the relation between other categories like Causality, Origin, and End. In sciences, *time* is also a concept in different disciplines. In physics, *time* is a very precise concept that has been changing throughout its history, changing (not without pain) from the classical definition that time is an absolute measure (i.e. the same everywhere) to today’s relativity theory where time is relative (and the speed of light is absolute). Such a scientific definition of the concept of time is not intuitive at all, and goes against most of our immediate use of the word. In this sense, scientists may find it difficult to operate with the scientific concept of time in relativity theory because of the other more usual meanings of the word. Besides, such a confusion between the scientific and the nonscientific may open philosophical questions and nonscientific interpretations of the scientific results.

In addition to this unavoidable challenge, the rationalization required by scientific theories appears in different forms. In this case, philosophical practice can help scientific practice by classifying the different types of rationality depending on the object under consideration. Motivated by Lepskiy [11] (but understood here in a different manner) and Althusser [8], we propose the following division.

- **Classical scientific rationality:** Direct observations and empirical falsification are possible for all elements of the theory, i.e. there is a one-to-one map between the physical and abstract realities.
- **Nonclassical scientific rationality:** Observations are not directly possible, i.e. the process of abstraction leads to nonobservable steps, resulting in a relatively autonomous theoretical domain.
- **Interventionist scientific rationality:** Active elements with internal awareness with objectives and goals exist, leading to a theory of the *fact to be accomplished* in contrast to theories of the *accomplished facts*.

By acknowledging the differences between these forms of rationality, sciences and scientific knowledge can be internalized as a social practice within the existing mode of production. Different from positivist and existentialist traditions in philosophy, this practice of philosophy attempts to articulate the scientific practice within the historical social whole, critically building demarcations of the correctness of the reach of scientific knowledge by rational argumentation [10].

Example 1.3 *Scientific efforts related to COVID-19.* In 2020, an unprecedented channeling of research activities was directed to combat the COVID-19 pandemics. These activities can be classified following the three aforementioned scientific rationalities. The **classical rationality** can be exemplified by the process

to test the effectiveness of vaccines following the historically defined norms. The **nonclassical rationality** incorporates the mathematical models for epidemics based on nonlinear dynamical processes over graphs where not all variables are observable. The **interventionist rationality** considers lockdown policies to control the virus propagation as a fact to be accomplished. A critical philosophical practice demarcates the reach of the three different scientific activities, determining both their interrelations and the articulation with other social practices. For example, a vaccine that is proved to work can be modeled by a mathematical model, which can be used to change the lockdown policy. However, within the capitalist mode of production, these activities are directly or indirectly determined by the economical reality – from the funds available to develop the vaccine and its respective property rights to the economic impact of lockdown policies and its justifications based on a wide range of epidemiological models. A critical philosophical practice acknowledges the autonomy of the results obtained through the scientific practice with respect to its object while it internalizes such a practice in the articulated social whole.

This philosophical practice goes hand in hand with the scientific practice by helping scientists to avoid overreaching tendencies related to their own theoretical findings. It also indicates critical points where other practices might be interfering in the scientific activity and vice versa. Although a deep discussion of the complex relations between scientific practice and other practices are far beyond our aim here, we will throughout this book deal with one specific relation: how scientific practice is related to the technological development. We have seen so far that the practical development of techniques does not require the intervention of (abstract) scientific rationalization. On the other hand, the knowledge produced by the sciences has a lot to offer to practical techniques. The existence of the term *technology*, referring to techniques developed or rectified by the sciences, indicates such a relation. More than what this definition might suggest, technology cannot be simply reduced to a mere application of scientific knowledge; it can indeed create new domains and objects subject to a new scientific discourse.

The aforementioned control and information theories perfectly exemplify this. New technological artifacts had been constructed using the up-to-date knowledge of physical laws to solve specific concrete problems, almost in a trial-and-error basis to create know-how-type of knowledge pushed by the needs of the industrial revolution. At some point, these concrete artifacts were conceptualized as abstract objects toward a scientific theory with its own methods, proofs, and research questions, constituting a relatively autonomous science of specific technological objects. The new established science not only indicates how to improve the efficiency of existing techniques and/or artifacts but also (and very importantly) defines their fundamental characteristics, conditions, and limits.

Example 1.4 Information theory. Although large-scale communications systems had already been deployed for some decades before the 1940s, the engineers considered that errors in transmission were somehow inevitable. This commonsense practical knowledge was scientifically proven false when Claude Shannon published *A Mathematical Theory of Communication* [12] formulating the concept of *information entropy* and *mutual information*. Using these concepts, Shannon mathematically proved the existence of a code that leads to error-free communication if, and only if, the coding rate is below the channel capacity. This theory proposed in 1948 opened up a new field of theoretical research and also oriented practical deployments by giving an absolute indication of how far from the fundamental limit specific technologies are. It is noteworthy that, although Shannon had mathematically proven the existence of capacity-achieving codes, he has not indicated how to practically design them. For many years, researchers and engineers have pushed the technological boundaries and have developed different coding schemes. Only with the new millennium, feasible solutions have been proposed (or rediscovered) and, currently, the turbo codes and low-density parity-check (LDPC) codes are feasible options to reach a performance close to Shannon's limit. These high-performance techniques are used for example in cellular networks and satellite communications. The fundamental limit proposed by Shannon, though, cannot be surpassed by any existing or future technologies. A similar development happened in physics when the fundamental laws and limits of thermodynamics; firstly motivated by the development of thermal engines, the thermodynamic laws imposed fundamental limits of all existing or future engines [10].

An important remark is that sciences as theoretical discourses are historical and objective, holding a truth value relative to what is scientifically known at that time considering limitations in both theoretical and experimental domains. In this sense, scientific practice is an open-ending activity constituted by historically established norms. These norms, which are not the same for the different sciences and are internally defined through the scientific practice, determine the valid methodologies to produce scientific knowledge. Once established, this knowledge can then be used as raw material not only for the scientific practice from where it originates but also it can be (directly or indirectly) employed by other practices. As demonstrated in, for example, Noble [13], Feenberg [14], the scientific and technical development as a historical phenomenon cannot be studied isolated from the society and its articulation with the social whole becomes necessary.

From this perspective, this book will pedagogically construct a scientific foundation for CPSs based on existing scientific concepts and theories without distorting and displacing their specific objects. The resulting general theory will then be used to explain and explore different particular existing realizations of

the well-defined abstract scientific object called *CPS* following the three proposed scientific rationalities. We are now ready to discuss the book structure and its rationale to then start our theoretical tour.

1.5 Book Structure

This book is divided into three main parts with ten core chapters, plus this introduction and the last chapter with my final words. The first part covers Chapters from 2 to 6, and focuses on the key concepts and theories required to propose a new theory for CPSs, which is presented in Chapters 7 and 8 (the second part of this book). The third part (Chapters 9, 10, and 11) deals with existing enabling technologies, specific CPSs, and their social implications

Part 1 starts with *systems* – the focus of Chapter 2, where we will revisit the basis of system engineering and then propose a way to demarcate particular systems following a cybernetic approach. Chapter 3 focuses on how to quantify *uncertainty* by reviewing the basis of probability theory and the concept of random variable. In Chapter 4, we will first define the concept *information* based on uncertainty resolution and then discuss its different key aspects, which includes the relation between data and information, as well as its fundamental limits. Chapter 5 introduces the mathematical theory of graphs, which is applied to scientifically understand interactions that form a *network* structure, from epidemiological processes to propagation of fake news. *Decisions* that determine *actions* are the theme of Chapter 6 discussing different forms of decision-making processes based on uncertainty, networks, and availability of information. Since decisions are generally associated with actions, *agents* are also introduced, serving as a transition to the second part.

Part 2 is composed of two dense chapters. Chapter 7 introduces the concept of *CPS as constituted by three layers*, which are interrelated and lead to a self-developing system. In Chapter 8, such a characteristic is further explored by introducing different approaches to model the *dynamics of CPSs*, also indicating performance metrics and their possible optimization, as well as vulnerabilities to different kinds of attacks. With these scientific abstractions, we will be equipped to assess existing technologies and their potential effects, which is the focus of the third part of this book.

Part 3 then covers concrete technologies and their impacts. Chapter 9 presents the *key enabling ICTs* that are necessary for the promising widespread of CPSs. Chapter 10 aims at different real-world applications that, following our theory, are conceptualized as realizations of CPSs. Chapter 11 is devoted to aspects *beyond technology* related to governance models, social implications, and military use.

At the end of each chapter, a summary of the key concepts accompanied by the most relevant references are presented followed by exercises that are proposed for the readers to actively learn how to operate with the main concepts.

1.6 Summary

In this chapter, we highlighted the reasons why a general scientific theory of CPSs is needed. We have briefly reviewed the beginnings of two related scientific fields, namely control theory [4] and information theory [12], contrasting them with the more philosophically leaned cybernetics as introduced by Wiener in [6]. To avoid potential threats of theoretical displacements of scientific theories, we have explicitly stated the philosophical standpoint taken in this book: science is a formal discourse holding true objective knowledge about well-defined abstract objects, which produces a knowledge effect on particular concrete objects. Scientific theories are then the result of a theoretical practice that produces new knowledge from historically determined facts and knowledges following a historically determined normative method of derivation and/or verification, which depends on the science/object under consideration. This leads to a philosophical classification that identifies three broad classes of scientific rationalities, helping to avoid misunderstandings of scientific results. The philosophical position taken here follows the key insights introduced by L. Althusser [8, 9], and I. Prigogine and I. Stengers [10]; the classification of different scientific rationalities is motivated by the work of Lepskiy et al. [11, 15] (although I do not share their philosophical position).

Exercises

- 1.1 Daily language and scientific concepts.** The idea is to think about the word *power*.
- Check in the dictionary the meaning of the word *power* and write down its different meanings.
 - Compare (a) with the meaning of the concept as defined in physics: ***power is the amount of energy transferred or converted per unit time.***
 - Write one paragraph indicating how the daily language defined in (a) may affect the practice of a scientist working with the physical concept indicated in (b).
- 1.2 Scientific rationalities.** In Example 1.3 in Section 1.4, an example of the three scientific rationalities was presented. It is your turn to follow the same steps.
- Find an example of scientific practices that can be related to the three rationalities.
 - Based on (a), provide one example of each scientific rationality: (i) *classical*, (ii) *nonclassical*, and (iii) *interventionist*.
 - Articulate the scientific practices defined in (a) and (b) with other practices present in the social whole.

- 1.3 Alan Turing and theoretical computer sciences.** The seminal work of Claude Shannon was presented in this chapter as the beginning of information theory. Alan Turing, a British mathematician, is also a well-known scientist considered by many as the father of theoretical computer sciences. Let us investigate him.
- (a) Read the entry “Alan Turing” from The Stanford Encyclopedia of Philosophy [16].
 - (b) Establish the historical background of Turing’s seminal work *On computable numbers, with an application to the Entscheidungsproblem* [17] following Example 1.4 in Section 1.4. This is the first paragraph of the text: *The ‘computable’ numbers may be described briefly as the real numbers whose expressions as a decimal are calculable by finite means. Although the subject of this paper is ostensibly the computable numbers, it is almost equally easy to define and investigate computable functions of an integral variable or a real or computable variable, computable predicates, and so forth. The fundamental problems involved are, however, the same in each case, and I have chosen the computable numbers for explicit treatment as involving the least cumbersome technique. I hope shortly to give an account of the relations of the computable numbers, functions, and so forth to one another. This will include a development of the theory of functions of a real variable expressed in terms of computable numbers. According to my definition, a number is computable if its decimal can be written down by a machine.*
 - (c) Discuss the relation between the work presented in (b) with the more philosophically leaned (speculative) *Computing Machinery and Intelligence* [18]. This is the first paragraph of the text: *I propose to consider the question, “Can machines think?” This should begin with definitions of the meaning of the terms “machine” and “think”. The definitions might be framed so as to reflect so far as possible the normal use of the words, but this attitude is dangerous. If the meaning of the words “machine” and “think” are to be found by examining how they are commonly used it is difficult to escape the conclusion that the meaning and the answer to the question, “Can machines think?” is to be sought in a statistical survey such as a Gallup poll. But this is absurd. Instead of attempting such a definition I shall replace the question by another, which is closely related to it and is expressed in relatively unambiguous words.*

References

- 1 National Institute of Standards and Technology. Cyber-physical systems; 2020. Last accessed 2 October 2020. <https://www.nist.gov/el/cyber-physical-systems>.

- 2 Nardelli PHJ, Kühnlenz F. Why smart appliances may result in a stupid grid: examining the layers of the sociotechnical systems. *IEEE Systems, Man, and Cybernetics Magazine*. 2018;4(4):21–27.
- 3 Lewis FL. *Applied Optimal Control and Estimation*. Prentice Hall PTR; 1992.
- 4 Maxwell JC. On governors. *Proceedings of the Royal Society of London*. 1868;16(16):270–283. 10.1098/rspl.1867.0055.
- 5 Huurdeman AA. *The Worldwide History of Telecommunications*. John Wiley & Sons; 2003.
- 6 Wiener N. *Cybernetics or Control and Communication in the Animal and the Machine*. MIT press; 2019.
- 7 Gerovitch S. *From Newspeak to Cyberspeak: A History of Soviet Cybernetics*. MIT Press; 2004.
- 8 Althusser L. *Philosophy for Non-Philosophers*. Bloomsbury Publishing; 2017.
- 9 Althusser L. *Philosophy and the Spontaneous Philosophy of the Scientists*. Verso; 2012.
- 10 Prigogine I, Stengers I. *Order Out of Chaos: Man's New Dialogue with Nature*. Verso Books; 2018.
- 11 Lepskiy V. *Evolution of Cybernetics: Philosophical and Methodological Analysis*. Kybernetes; 2018.
- 12 Shannon CE. A mathematical theory of communication. *The Bell System Technical Journal*. 1948;27(3):379–423.
- 13 Noble DF. *Forces of Production: A Social History of Industrial Automation*. Routledge; 2017.
- 14 Feenberg A. *Transforming Technology: A Critical Theory Revisited*. Oxford University Press; 2002.
- 15 Umpleby SA, Medvedeva TA, Lepskiy V. Recent developments in cybernetics, from cognition to social systems. *Cybernetics and Systems*. 2019;50(4): 367–382.
- 16 Hodges A., Zalta EN, editor. Alan Turing. *Metaphysics Research Lab, Stanford University*; 2019. Last accessed 20 October 2020. <https://plato.stanford.edu/archives/win2019/entries/turing/>.
- 17 Turing AM. On computable numbers, with an application to the Entscheidungsproblem. *Proceedings of the London Mathematical Society*. 1937;2(1):230–265.
- 18 Turing AM. Computing machinery and intelligence. In: *Parsing the Turing Test*. Robert E., Gary R., Grace B., editors. Springer, Dordrecht; 2009. pp. 23–65.

