MONDO DIGITALE

No. 106 – 2026

JULY 2026

ISSUE 106

Computer Science and Digital Culture in the New National Guidelines for Schools

CONTENTS

In This Issue

4 **INTRODUCTION**
Antonio Piva, President of AICA

5 **EDITORIAL**
Paolo Ciancarini

8 **OPEN LETTER**
A Realistic View of Artificial Intelligence
Enrico Nardelli and Walter Quattrociocchi

Articles

10 **Computer Science Education in the New National Guidelines**
Toward a School System Capable of Understanding, Designing, and Governing the Digital World
Vincenzo Vespri, University of Florence

53 **The Fourth Writing**
Conversation, Illusions, and the Future of Education
Paolo Ciancarini, University of Bologna

59 **Toward a Harmonious Relationship Between Education and Technology**
Christian Fiori, AICA

68 **A Few Sentences, One Discipline**
Computer Science in the New National Guidelines for the First Cycle of Education
Michael Lodi, University of Bologna · Violetta Lonati, University of Milan

81 **Calibra: An AI-Based Framework for Personalizing Learning Pathways**
Nicola Illuzzi, Claudio Vergini, and Giorgio Tullio De Negri, Calibra Project

Columns

CERTIFICATIONS AND REGULATIONS
99 **ICDL Artificial Intelligence: Italy's First Certification in AI Literacy**
Antonella Fancello, AICA

COMPUTER ETHICS
116 **Building a "Better Rifle"**
Palantir, Artificial Intelligence, and the Ethics of the "Technological Republic"
Christian Fiori, AICA

INTERNATIONAL PROJECTS
124 **AICA's International Activities**
Marina Cabrini, CEPIS – Member of the Board

ECOLOGICAL IMPACT OF COMPUTING
129 **Digital Ecology**
Giovanni Franza, AICA

HISTORY OF COMPUTING
135 **LIFO avant la lettre**
Pierangelo Indolfi, Polytechnic University of Bari

No. 106 – 2026

Introduction

by **Antonio Piva, *President of AICA***

It is with particular enthusiasm and satisfaction that I present the new direction of *Mondo Digitale*, AICA's scientific and cultural journal, which, since its inception in 2002, has served as a point of reference for the dissemination of computer science culture and digital innovation in our country.

Over more than twenty years of activity, *Mondo Digitale* has kept pace with the evolution of technology, offering authoritative contributions from the worlds of academia, research, the professions, and business. Today, in a historical period characterized by increasingly rapid and profound transformations, we believe that this space for discussion and in-depth analysis is more necessary than ever.

Mondo Digitale is fully aligned with the mission that AICA has been pursuing for over sixty years: promoting digital culture as a driver of growth, innovation, and informed citizenship. For this reason, I would like to thank the new Scientific Director, Paolo Ciancarini, who has agreed to lead this project with the expertise and authority that set him apart. I also extend my heartfelt thanks to Professor Viola Schiaffonati for her invaluable dedication over the course of a decade at the helm of *Mondo Digitale*, having made a decisive contribution to the journal's growth and consolidation.

Current events present us with decisive challenges. The new National Guidelines for schools, the growing focus on digital skills, and the impact of Artificial Intelligence call for forums for informed reflection—ones capable of combining scientific rigor with an attention to society's needs. In this context, the work carried out by AICA on digital skills, syllabi dedicated to Artificial Intelligence, and certification programs represents a concrete contribution to helping citizens, students, teachers, and professionals understand and navigate this change.

Mondo Digitale aims to continue being an authoritative voice in this debate, fostering dialogue among research, education, institutions, and the business world. It is a space open to the ideas, experiences, and best practices that are shaping our country's digital future.

To all those who will contribute to this new phase of the magazine, I extend my sincerest thanks and invite you to actively participate in this journey.

MONDO DIGITALE

No. 106 – 2026

Editorial

by **Paolo Ciancarini**

With this issue, *Mondo Digitale* embarks on a new editorial phase. The magazine aims to strengthen its role as a meeting place for research, technological innovation, education, institutions, and the professional world, focusing in particular on the new era of digital transformation ushered in by conversational Artificial Intelligence and related technologies. It is not merely a matter of following a technological evolution, but of helping to interpret it, guide it, and make it the subject of critical and systematic reflection.

For me, this role also holds personal significance. As a doctoral student, I had the opportunity to publish an article in *the* historic *Rivista di Informatica*, the predecessor of *Mondo Digitale*, and later contributed other works to the new journal. In the years that followed, I was appointed to the scientific board. Returning today to contribute to its relaunch is therefore not only a scientific responsibility but also a source of particular personal satisfaction. It is also an opportunity to give back, through my editorial commitment, for all that this community has contributed to my education and personal growth.

In recent years, digital technology has ceased to be merely a set of tools supporting organizational processes: it is becoming a cognitive, productive, and social environment that profoundly changes the way we learn, work, communicate, and make decisions. Digital technologies are no longer limited to automating existing activities; rather, they actively participate in the construction of knowledge, the definition of operational contexts, and the transformation of relationships between individuals and organizations. In this scenario, the emergence of conversational AI systems introduces unprecedented forms of interaction, in which dialogue with machines becomes part of everyday learning and work practices.

Conversational generative AI currently represents the most visible aspect of this systemic transformation, but it is not the only relevant element. Cybersecurity, data science, digital platforms, interoperability, automation, and new forms of human-machine interaction are redefining the relationship between technologies, skills, and governance. This gives rise to new, crucial issues: the quality and reliability of information, the transparency of algorithms, data management, the ethics of innovation, and the impact on work and skills.

Mondo Digitale aims to contribute to understanding and critically discussing these changes, with a scientifically rigorous approach that remains open to dialogue with the applied and institutional worlds. The journal

will therefore strengthen its academic profile by promoting scientific articles in English, surveys, position papers, and applied contributions, while maintaining a strong focus on qualified outreach and the transferability of results to schools, organizations, and society.

It is no coincidence that the first issue of this new series is dedicated to **"Digital Culture and Computer Science in the New National Guidelines for Schools**." The new ministerial guidelines represent an important milestone for the Italian educational system: for the first time, specific objectives related to computer science and digital culture are being introduced in an explicit and structured manner, starting in elementary school. This recognition marks a significant change, as it attributes to computer science not only instrumental value but also an autonomous cultural and educational significance.

This change goes far beyond the introduction of new curriculum content. Bringing computer science into elementary school means promoting the development of computational thinking, the ability to model complex problems, an understanding of how digital systems work, and the capacity to interact thoughtfully with increasingly pervasive technologies. It also means fostering mature digital citizenship, which includes understanding the mechanisms underlying the production, management, and circulation of data, as well as the ethical and social implications of technology.

In this context, reference to European frameworks for digital competencies provides an important anchor. In particular, the **DigComp** framework for citizens defines a comprehensive set of competencies ranging from information and data literacy to communication and collaboration, from digital content creation to security, and to problem-solving. Similarly, the **DigCompEdu** framework provides specific guidance for developing teachers' digital competencies, emphasizing the role of technology not only as a subject of instruction but also as a lever for innovating teaching and learning processes.

This evolution also encompasses the emerging topic of **AI literacy**, which is gradually being incorporated into the most recent European frameworks. One possible way to interpret this is through the lens of the "three generations" of information technology: a first phase centered on computation and automation; a second phase characterized by the spread of networks and digital platforms; and a third phase—the current one—in which systems based on artificial intelligence—particularly conversational AI—actively participate in the production of content, knowledge, and decisions.

From this perspective, this transition can be summarized as follows: **from computing that executes, to computing that connects, to computing that dialogues and co-constructs knowledge**. It is at this stage that conversational AI enters learning and work processes as an active interlocutor, making it necessary to develop new forms of critical awareness.

While in the first two phases digital literacy could be traced primarily to the informed use of tools and an understanding of their mechanisms, in the current phase it becomes essential to develop more advanced skills: understanding the limitations and biases of AI systems, evaluating the reliability of outputs, responsibly managing interactions with conversational AI systems, and integrating these technologies without delegating cognitive control to them.

The inclusion of computer science in the National Guidelines also raises the question of striking a balance between theoretical and practical aspects. It is not simply a matter of teaching how to use tools, but of providing the analytical frameworks that enable students to understand what happens “behind the scenes” of digital technologies.

The implications for the education system are profound. First, the issue of teacher training emerges as a key priority; it must be rethought in terms of both initial training and ongoing professional development, particularly in light of European guidelines and the new requirements related to AI literacy. Second, there is the issue of instructional materials and methodologies, which must align with these new objectives.

A further aspect concerns the integration of computer science with other disciplines. The new Guidelines pave the way for a cross-curricular approach, in which computer science concepts interact with other areas of knowledge, helping to build richer and more interconnected educational pathways.

In this context, the journal aims to offer a concrete and informed contribution to the scientific debate and educational practice, providing critical and knowledgeable guidance on the innovation processes shaping the education system within the broader context of digital transformation.

I would like to extend my sincere thanks to the outgoing editor-in-chief, **Viola Schiaffonati** of the Politecnico di Milano, for her work over the past years and for her contribution to the journal’s cultural growth. Special thanks also go to AICA President Antonio Piva for the trust he has placed in me in assuming scientific responsibility for this new editorial direction.

Finally, I would like to thank the entire editorial staff for their commitment to revitalizing the journal and, in particular, **Pierfranco Ravotto**, who, with his expertise, experience, and consistency, serves as a fundamental pillar for *Mondo Digitale*.

The goal of this new phase is ambitious: to build a journal capable of connecting research, practice, and policy, contributing to the development of a digital culture that is informed, critical, and scientifically grounded. In a context where technological innovation is evolving rapidly, it becomes essential to have authoritative, open, and interdisciplinary forums for reflection. *Mondo Digitale* aims to be one of these forums.

MONDO DIGITALE

No. 106 – 2026

Open Letter

by **Enrico Nardelli and Walter Quattrociocchi**

The open letter to society titled *"A Realistic View of Artificial Intelligence"* https://www.petizioni.com/visione_realistica_intelligenza_artificiale was initiated by Enrico Nardelli and Walter Quattrociocchi in response to a specific need: to bring the discussion on AI—often distorted by idealized and anthropomorphized portrayals—back to a clear and realistic foundation.

The context that prompted the authors to take action is a public debate that is increasingly detached from technical reality. Recently, an interview by Walter Veltroni in *Corriere della Sera* with Claude, Anthropic's AI model, reignited the discussion among scholars about how far off track the debate on AI has gone: the concern centers on Italians' lack of literacy on the subject and the attribution to AI systems of characteristics they do not possess. There is talk of machines that "think," of artificial consciousness, and of the imminent replacement of humans—a narrative fueled, according to the letter's authors, by hundreds of billions in investments still awaiting a significant economic return. Public discourse on the matter is rife with anthropomorphic metaphors that overlook the real issues and portray technology as magic, reverting to a pre-scientific approach. Instead, we must talk about statistics and probability, architectural and infrastructural constraints, and specific human responsibilities.

One of the key concepts to reflect on is that of *episteme*, introduced by Quattrociocchi—that is, the illusion of knowledge linked to the purely statistical answers provided by large-scale language models. This illusion is at the heart of the problem: when a generative AI system produces a fluent and convincing response, it is neither understanding nor reasoning—it is calculating which output is statistically most consistent with the context. When it appears to know something, it generates the most probable sequence of words, without direct access to the world and without autonomous verification mechanisms. The technical results achieved are impressive, but they do not equate to understanding in the human sense of the term.

A large language model does not construct knowledge in the same way a human does. It does not perceive the surrounding world and does not accumulate lived experience. It generates plausible sequences of words based on statistical associations between data. This is not a technical subtlety reserved for specialists but the premise upon which a clear understanding of the problem depends. Moreover, highlighting the technical and infrastructural aspects of AI allows us to focus on what really matters: who sets the rules of the game (including control over data) and who ensures compliance, who benefits from the prevailing narrative, and who bears the costs.

The letter, which was signed by more than 335 scientists and researchers in late May, is a call from the academic community to bring the discussion back within the Galilean approach to science and to reject the media-driven approach based on pundits debating the ethics of algorithms with little scientific or technological knowledge.

The letter also clarifies where these systems reveal their structural weaknesses: they work well when the problem resembles situations already present in the training data; they become unreliable when information changes rapidly, when the topic is controversial, or when there is a significant deviation from what the model has already encountered. In these situations, they can produce responses that appear consistent but are actually incorrect. The point is not that they sometimes make mistakes: it is that they have no way of realizing it.

As for AGI—artificial general intelligence, capable of understanding and reasoning autonomously like a human being—the authors are explicit: current generative systems do not possess these characteristics, and the increase in data and computing power improves performance without changing the nature of the underlying mechanism.

The real risk, according to the signatories, is not the technology itself, but the way we interpret it. If we treat the fluidity of language as proof of knowledge, if we delegate judgment without understanding how the tool works, we risk altering—often without realizing it—the criteria by which we recognize what constitutes knowledge. And this has concrete cultural, social, and economic consequences.

If we confuse the narrative surrounding artificial intelligence with its technical and scientific reality, then we open the door to social control through technology.

The solution outlined in the letter is literacy. Understanding how these systems work, what their limitations are, and how to use them responsibly is now a fundamental skill for all citizens. The appeal is directed at the academic community—with an explicit call to computer scientists, but also to scholars in every other discipline—to actively contribute to this effort at clarification. Explaining precisely what artificial intelligence really is—and what it is not—is identified as an integral part of the work of researchers and educators: not a secondary task, but a public responsibility.

Computer Science Education in the New National Guidelines

Toward a School System Capable of Understanding, Designing, and Governing the Digital World

Vincenzo Vespri, University of Florence

Abstract

This article addresses the role of computer science in the new National Guidelines, drawing on direct experience gained while coordinating the Commission tasked with developing the guidelines for STEM disciplines. The central thesis is that truly introducing computer science into schools does not mean adding a technical subject, but rather redefining the cultural model of education in the age of algorithmic society and generative artificial intelligence. The text distinguishes between instrumental digital literacy and critical understanding of digital technology, dispels the misconception that reduces computer science to coding, and proposes viewing it as a new form of thinking and a cultural discipline on par with mathematics. By comparing international models (Estonia, Finland, the United Kingdom, the United States, and China) with the European framework (DigComp, AI Act, digital sovereignty), the article outlines an Italian cultural approach to computer science education, whose goal is not to train efficient users but to foster citizens capable of understanding, designing, and critically governing digital society.

Keywords: computer science, school, National Guidelines, artificial intelligence, computational thinking, coding, digital citizenship, algorithmic society, pedagogy

Introduction

This article stems from concrete experience, not from purely theoretical reflection. It arises from the months spent coordinating the Commission tasked with drafting the new National Guidelines for STEM disciplines and, above all, from the discussions—often intense, sometimes even contentious—that accompanied the attempt to place computer science at the heart of the Italian school curriculum.

At first glance, the issue seemed straightforward. Everyone, at least formally, agreed that the Italian school system could no longer ignore the role of computer science and artificial intelligence in contemporary society. No one truly disputed that the world is changing rapidly and that schools must address this transformation.

But the consensus ended right there. Behind the term "computer science," in fact, lay profoundly different—and at times incompatible—ideas about what schools should become in the 21st century. There were mathematicians concerned that the introduction of computer science would end up further squeezing the space devoted to mathematics, which had already been reduced by years of pedagogical simplifications. There were teachers convinced that computer science in schools should focus primarily on coding or the practical use of digital tools. Still others proposed opposite scenarios, envisioning the introduction of artificial intelligence, genetic algorithms, machine learning, and advanced robotics as early as elementary school. Alongside these positions, an even more concrete problem emerged: Who should teach this new discipline? With what training? With what skills? And above all: Should computer science become a standalone subject or a cross-curricular skill distributed across all disciplines?

Little by little, it became clear that the debate over computer science was not merely about a new school subject. Rather, it concerned the very cultural model of the school itself. Because truly introducing computer science inevitably means redefining the relationship between memory and knowledge, the role of the teacher, the meaning of learning, the relationship between humans and machines, and, ultimately, the cultural function of the school in the digital society. Many of the conflicts that emerged during the Commission's work stemmed precisely from this. On the surface, the discussions were about curriculum hours, coding, labs, or programs; in reality, they were about something much deeper: what conception of the human being, of knowledge, and of citizenship should guide the school of the future.

The emergence of generative artificial intelligence has made all of this even more evident. In the space of just a few years, systems like ChatGPT have radically transformed the educational landscape. For the first time in history, a machine has become capable of producing coherent texts, solving problems, generating images, summarizing books, translating languages, and even writing code. Many activities traditionally used by schools to assess learning have suddenly become automatable. Faced with this transformation, schools have often reacted defensively: by banning, monitoring, restricting, and being suspicious. But a technology of this magnitude cannot be addressed solely with disciplinary measures or

administrative tasks. Artificial intelligence is not simply a new piece of software to be used in classrooms; it alters the very relationship between knowledge, cognitive work, and cultural production. For this reason, the issue of computer science in the new National Guidelines cannot be reduced to the question of how many hours of computer science to include in the curriculum. The real question is far more radical: what kind of school can still make sense in a society where machines now participate directly in cognitive processes? This question runs throughout this entire contribution.

In recent years, the public debate on digital technology has often oscillated between two opposing extremes. On the one hand, there has been a naive, technocratic enthusiasm, according to which simply filling schools with tablets, platforms, and devices would be enough to drive educational innovation. On the other hand, there has been a nostalgic and defensive reaction that views technology solely as a threat to traditional knowledge. Both of these positions now appear insufficient. The central issue is not introducing more technology into schools, but understanding the cultural significance of the digital revolution.

For many years, computer science education has been interpreted almost exclusively as technological literacy: using software, browsing the web, creating documents, and using platforms. Today, this approach is no longer enough. A society that knows how to use technologies without understanding them is a culturally dependent society. The algorithms that organize information, guide consumption, influence political opinions, and govern digital platforms are not neutral tools. They embody cultural models, economic interests, and power structures. For this reason, computer science must become part of the new critical literacy of the 21st century. Understanding the digital world today means understanding the relationship between data and power, how algorithms work, the limits of automation, the significance of artificial intelligence, the transformation of cognitive labor, and the role of platforms in contemporary democracy. Schools cannot limit themselves to training efficient users of digital tools. They must educate citizens capable of critically understanding algorithmic society. And this is where the deepest conflict emerges.

The 20th-century school system was designed for a relatively stable world, characterized by slow-evolving knowledge, long-lasting professions, predictable technologies, and limited access to information. Contemporary society, on the other hand, is dominated by continuous acceleration, cognitive automation, constant information flows, and the ceaseless transformation of skills. Artificial intelligence brings this historical divide into sharp relief. If a machine can produce complex content, then the specifically human value of education can no longer be equated with the mere reproduction of information. Critical thinking, authentic creativity, the ability to formulate problems, interdisciplinary understanding, and ethical responsibility become central. Computer science thus takes on a new role. It is not merely a technical discipline. It becomes one of the fundamental languages of the contemporary world, on par with mathematics and the sciences.

This contribution stems precisely from this awareness. It is not intended to be a technical manual on teaching computer science nor a naive celebration of technology. Rather, it aims to reflect

on a question that is decisive for the future of European education: How can we educate free, critical, and aware human beings within a society increasingly governed by algorithms, platforms, and intelligent systems? The issue certainly concerns computer science. But it concerns the future of cultural democracy even more.

Chapter 1 — The End of the Twentieth-Century School

1. A School Built for the Industrial World

The school system we live in today is, to a large extent, a product of the 20th century. Its underlying structures—the organization of academic disciplines, the structure of the school day, the centrality of the lecture-style class, the division by age, standardized curricula, and the hierarchical relationship between teacher and student—emerged within European industrial society between the late 19th century and the 20th century.

That school system had a clear mission: to educate ever-larger segments of the population, build shared national identities, and train workers capable of integrating into relatively stable production systems. It was a school system designed for a society in which knowledge evolved slowly, professions changed little, and access to information was limited. For decades, that model worked. It fostered social mobility, economic growth, and cultural democratization. In many European countries, the 20th-century school system served as the primary vehicle for social mobility in the postwar era. What was taught in school largely coincided with the knowledge needed to participate in economic and civic life. Today, however, many of the assumptions upon which that model was built are rapidly disappearing.

The digital revolution has not merely introduced new technological tools; it has altered the very relationship between human beings, knowledge, and power. The internet, smartphones, global platforms, and artificial intelligence have created an informational environment radically different from the one in which the modern school system was born. For the first time in history, access to information is no longer the main problem. Today, every student has, via a smartphone, access to an immensely greater amount of information than was available in a school library just a few decades ago. The problem is no longer finding information, but understanding it, selecting it, verifying it, and interpreting it critically. This change challenges the very core of the traditional, knowledge-transmission-based school.

If the primary function of school was to transmit content, what happens when that content becomes immediately accessible? If machines can produce texts, translations, images, and summaries in a matter of seconds, what meaning does learning still hold? The issue is far more radical than is often realized. It concerns not only teaching methods or the introduction of technology into classrooms. It concerns the cultural model of the school itself.

2. The Rise of the Algorithmic Society

Over the past twenty years, society has become increasingly algorithmic. A growing part of our lives is mediated by automated systems that collect data, process information, and guide behavior. Algorithms decide what news we read, what content we see on social media, what products we buy, which routes we take in cities, what ads we receive, and, increasingly, what opportunities are offered to us. The most surprising feature of this transformation is its invisibility.

The great machines of the Industrial Revolution were physically present: factories, engines, assembly lines. Algorithms, on the other hand, operate silently. They operate within seemingly neutral platforms, hiding behind simple and intuitive interfaces. Most people use extremely sophisticated computational systems on a daily basis without having any idea how they work. This situation creates a massive cultural divide. In the 20th century, literacy meant teaching people to read, write, and do arithmetic. Today, that is no longer enough. A citizen unable to grasp at least the fundamental principles of the digital world risks becoming dependent on systems that they are unable to either control or interpret.

The issue is not limited to employment or economic competitiveness. It concerns the very quality of democracy. A society in which collective decisions are increasingly influenced by opaque platforms and automated systems requires citizens capable of understanding the relationship between data, algorithms, and power. And this is where the school system faces a crisis. The 20th-century school system was designed to impart relatively stable knowledge. The algorithmic society, however, is constantly changing languages, tools, and cognitive models. Content quickly becomes obsolete. Technical skills transform within a few years. Consequently, the educational challenge can no longer be limited to the accumulation of knowledge. The ability to understand complex systems, interpret change, and learn continuously becomes central.

3. From Memory to Understanding

For centuries, schooling has been based on the centrality of memory. Memorizing dates, formulas, definitions, and procedures was an essential component of learning. In a society with limited access to texts and information, memory was a fundamental tool for cultural participation. The digital revolution is profoundly changing this situation. Today, an enormous amount of knowledge is instantly available. Generative artificial intelligence even makes it possible to obtain summaries, explanations, and complete texts in a matter of seconds. This does not mean that memory becomes useless. It does mean, however, that its role is changing.

Educational value is gradually shifting from the mere retention of information to the ability to critically understand it. It is becoming more important to know how to ask questions than to remember.

automatically provided answers. Paradoxically, it is precisely this abundance of information that makes deep cultural literacy even more essential. In an environment dominated by a constant flow of content, the main risk is not ignorance, but confusion. Students must learn to distinguish between information and knowledge, between correlation and explanation, and between genuine authority and mere algorithmic visibility. Schools should therefore transform from places of transmission into places of interpretation. This change inevitably alters the role of the teacher as well. The teacher is no longer merely the one who possesses information that students lack. Instead, the teacher becomes a cultural guide, a critical mediator, and a designer of cognitive pathways. It is a difficult transformation. Many teachers feel that digital technology is undermining the traditional authority of the school. This is partly true. But precisely for this reason, the educational function can become even more important. In a society dominated by commercial platforms and cognitive automation, the school remains one of the few public spaces dedicated to the critical construction of knowledge.

4. Artificial Intelligence Changes the Meaning of Learning

The arrival of generative artificial intelligence likely marks the definitive breaking point for the traditional school model. For many years, digital technology was perceived primarily as an operational tool. Computers were mainly used to write texts, conduct research, create presentations, or use specialized software. Artificial intelligence radically alters this perspective because it penetrates directly into the heart of cognitive processes. Today, a machine is capable of writing an essay, translating a text, summarizing a book, producing images, generating code, solving mathematical problems, and creating complex content in a matter of seconds. Many activities traditionally used by schools to assess learning thus suddenly become automatable. The most immediate reaction is often defensive: to ban, control, restrict, and be suspicious. But it would be a historic mistake to interpret artificial intelligence solely as a threat.

In fact, AI forces schools to question their deepest purposes. If a machine can produce complex content, then the specifically human value of education can no longer be equated with the mere reproduction of information. Critical thinking, authentic creativity, the ability to construct meaning, interdisciplinary reasoning, ethical understanding, and the formulation of complex problems become central. Artificial intelligence does not render schools useless. Rather, it renders obsolete a certain conception of school: one built primarily on the transmission and repetition of knowledge.

5. Education in the Age of Cognitive Automation

The great educational challenge of the 21st century likely lies in preventing technological automation from leading to cultural passivity. Every powerful technology tends, in fact, to outsource certain human abilities. The automobile shortens the journey, the GPS weakens spatial orientation, and artificial intelligence risks reducing cognitive autonomy. If students limit themselves to using intelligent systems

without understanding them, schools risk producing individuals who are technologically competent but culturally dependent. For this reason, the issue of computer science cannot be reduced to the mere use of technologies. The real issue is understanding the digital world.

Understanding the digital world means grasping how algorithms work, how data is used, what logic governs platforms, what limitations automation has, and what relationships exist between technology, the economy, and democracy. Computer science thus takes on a new role. It is no longer merely a technical discipline but a form of critical literacy for the contemporary world. The school of the future must therefore educate not merely users of digital tools, but citizens capable of understanding, designing, and critically governing algorithmic society. And it is precisely this need that gives rise to the necessity of radically rethinking the role of computer science in the new National Guidelines.

Chapter 2 — Whoever Controls the Algorithms Controls Society

1. Data, Algorithms, and New Centers of Power

To understand why computer science has become a central educational issue, we must first understand how the structure of power in contemporary society has changed. For centuries, economic and political power was based primarily on control over land, natural resources, productive infrastructure, or financial capital. The digital revolution has introduced a new strategic element: control over information. Today, large global technology platforms collect immense amounts of data on individual behavior. Every online search, every purchase, every trip, every social media interaction, and every piece of content viewed generates information that is recorded, processed, and transformed into predictive models.

Data has become the new fuel of the digital economy. But true power lies not only in the collection of data. It lies above all in the ability to interpret it through increasingly sophisticated algorithms. Algorithms today represent the invisible heart of contemporary society. They curate content, guide consumption, filter information, suggest connections, predict behavior, and influence collective decisions. Their power also stems from their invisibility. Digital platforms often appear to be neutral and easy-to-use tools. In reality, behind those intuitive interfaces lie extremely complex computational systems designed to maximize user attention, engagement, and retention. The fundamental problem is that most people use these systems daily without understanding how they work. For the first time in history, a growing portion of our social, cultural, and political experience is mediated by algorithmic infrastructures controlled by global private entities. This profoundly alters the relationship between citizens, knowledge, and power.

In the 20th century, cultural power flowed primarily through recognizable institutions: schools, universities, newspapers, political parties, and public television. Today, a growing portion of the formation

of public opinion flows through private platforms governed by opaque algorithmic logic. It is an enormous transformation, and one that is often underestimated. When an algorithm decides which news stories to display, which content to make visible, or which topics to bring to the forefront of public debate, it inevitably exercises a form of cultural and political power. The issue is not merely about technology. It concerns the very structure of contemporary democracy.

2. Democracy in the Age of Platforms

For a long time, the internet was envisioned as a space of widespread freedom, universal access to knowledge, and the democratization of information. In part, this promise has been fulfilled. Never in human history have so many people had access to such a vast amount of content, tools, and communication possibilities. However, the maturation of the digital ecosystem has produced far more ambiguous effects. Contemporary platforms do more than simply host information. They organize collective attention. They decide which content to prioritize, which relationships to make visible, and which messages to amplify. The digital society is not simply a society with more information; it is a society in which access to information is constantly mediated by algorithms.

This profoundly changes the functioning of the public sphere. In the classical democratic model, pluralism was based, at least in part, on the multiplicity of information sources and the relative transparency of editorial processes. In platform society, however, much of the selection of information occurs through automated mechanisms that are difficult to understand even for specialists. The dominant logic is not that of cultural quality or democratic relevance, but that of attention optimization. Content is selected because it generates emotional engagement, extended time spent online, and continuous interaction.

This has profound consequences. The polarization of public debate, the spread of disinformation, the speed of manipulative campaigns, and the success of fake news are not merely accidental side effects. They are also the result of an algorithmic ecosystem designed to maximize the intensity of interaction. In this context, democratic citizenship requires entirely new skills. It is no longer enough to know how to read and write in the traditional sense of the term. It is necessary to understand how platforms work, what economic logic governs them, how data is collected and used, and how algorithms influence the formation of public opinion. Digital literacy thus takes on a political dimension. Schools cannot ignore this transformation. If students spend an increasing portion of their lives within algorithmic environments, then understanding those systems becomes an essential part of democratic education.

3. Invisible Algorithms and the Manipulation of Attention

One of the most disturbing aspects of algorithmic society is the gradual invisibility of control mechanisms. In the past, forms of power were generally recognizable. Political,

economic, and cultural institutions had relatively visible structures. Today, however, a growing proportion of the decisions that influence daily life is delegated to opaque automated systems. The algorithms that govern social networks, search engines, video platforms, and advertising systems operate silently. The user perceives only the final result: the content that appears on the screen, the suggestion received, the news item displayed. But behind that apparent simplicity lie extremely sophisticated mathematical models designed to predict and steer behavior.

The true raw material of digital society is human attention. Platforms are constantly competing to capture cognitive time. Every click, every view, and every interaction becomes data that can be used to further refine predictive models. This creates a cycle in which algorithms progressively learn about users' desires, fears, preferences, and vulnerabilities.

This brings about a significant anthropological transformation. Contemporary human beings live immersed in a constantly personalized informational environment. Two different people no longer necessarily see the same digital world. Algorithms construct different informational experiences based on the data they collect. The consequence is a progressive fragmentation of the public sphere.

Schools should openly address these issues—not to demonize technology, but to foster critical awareness. Students must understand that platforms are not merely neutral tools; they are economic and cultural ecosystems designed to steer attention and behavior. Without this understanding, there is a risk of producing citizens who are formally connected but essentially passive.

4. Why Computer Science Is a Political Issue

For many years, computer science was perceived primarily as a technical discipline. It was thought to concern mainly computers, software, programming, and technological infrastructure. Today, this view seems inadequate. Computer science has become a political issue in the deepest sense of the term, because digital infrastructure now organizes a growing part of collective life. It influences work, information, communication, consumption, education, and even democratic participation. Algorithms are not neutral. Every computational system inevitably incorporates design choices, economic priorities, and cultural models. Artificial intelligence also reflects this problem. AI systems learn from enormous amounts of historical data and therefore tend to reproduce cultural structures already present in society. Bias, discrimination, and imbalances are not automatically eliminated by technology; they are often amplified. For this reason, understanding how digital systems work cannot be considered a skill reserved for specialists. It becomes a form of citizenship. Schools should help students understand the relationship between technology and power in the same way they teach how democratic institutions, economic systems, or political history function. In the

the absence of this awareness, society risks progressively delegating fundamental decisions to automated systems that few truly understand.

5. Digital Sovereignty and Citizenship

In recent years, the issue of digital sovereignty has emerged with increasing force. Much of the technological infrastructure used daily in Europe is, in fact, owned by large non-European corporations. Search engines, social networks, cloud platforms, operating systems, and artificial intelligence models are controlled primarily by private entities based in the United States or, to an increasing extent, China. This situation creates a new form of cultural and technological dependence. The issue is not limited to the market or industrial competitiveness. It concerns a society's ability to understand and govern the tools that shape its collective life.

A population unable to understand how digital systems work inevitably risks passively submitting to them. For this reason, schools now play a strategic role. Educating students in computer science also means fostering cultural autonomy. It means training citizens capable of understanding the digital world rather than merely consuming it. What is truly at stake is not simply the acquisition of technical skills. It is the possibility of maintaining a space for critical freedom within a society increasingly organized by algorithms, data, and intelligent systems. Computer science education should therefore not be limited to teaching the use of digital tools. It should contribute to the formation of a new democratic citizenship, capable of understanding the relationship between technology, power, and freedom in contemporary society.

Chapter 3 — Artificial Intelligence and the Crisis of Academic Knowledge

1. ChatGPT and the End of the Traditional Assignment

The arrival of generative artificial intelligence systems has brought about one of the greatest cultural disruptions in schools in recent decades. For the first time, tools accessible to anyone have become capable of producing coherent texts, solving problems, constructing arguments, translating languages, summarizing content, and even simulating complex dialogues. Many teachers immediately grasped the significance of this transformation. In just a few months, school activities that had been considered essential for years—essays, reports, summaries, exercises, and presentations—have become easily automatable. Today, a student can produce a formally correct paper in a matter of seconds without actually possessing the skills needed to create it independently.

This situation inevitably leads to confusion. Some in the education system are reacting by trying to defend the traditional model through bans, monitoring, or attempts to limit the use of AI. But the problem cannot be addressed solely in disciplinary terms. Artificial intelligence is not merely a technological shortcut. It represents a profound change in the relationship between human beings

and knowledge. For centuries, schools have equated learning with the ability to produce content. Writing a text, solving a problem, or summarizing a topic were relatively reliable indicators of specific cognitive skills. Today, this equivalence is being challenged. If a machine can produce formally convincing results, then the educational challenge no longer consists simply of verifying the final product. It becomes necessary to understand the cognitive processes underlying the production of knowledge. This is an enormous transformation. Twentieth-century schooling was built largely on the scarcity of information and the difficulty of cultural production. Artificial intelligence, on the other hand, introduces the opposite condition: an almost unlimited abundance of automatically generatable content. This does not mean that studying loses its value. It does mean, however, that the meaning of learning changes radically.

2. If Machines Write, What Should Schools Do?

The question currently facing the world of education is as simple as it is unsettling: if machines can write, what should schools do? To grasp the depth of this question, we must avoid two opposing errors. The first is to view artificial intelligence as an absolute threat destined to destroy human learning. The second is the naive enthusiasm of those who imagine that AI can automatically replace complex educational processes. Both of these positions stem from a superficial understanding of the problem. Machines do not truly understand the meaning of the content they produce. Modern generative systems operate through statistical correlations learned from enormous amounts of data. They effectively simulate human language, but they lack worldly experience, intentionality, or semantic awareness. Yet this philosophical distinction does not eliminate the educational problem. Because, from a practical standpoint, the results produced by AI are often convincing enough to render many traditional school assessment tools inadequate.

Schools therefore face a historic choice. They can attempt to artificially preserve educational models designed for a world prior to the algorithmic revolution, or they can redefine their fundamental purposes. In reality, artificial intelligence brings to light a problem that has existed for some time. Many school activities had gradually become mechanical and repetitive long before the arrival of ChatGPT. In many cases, students were trained primarily to reproduce information, apply standard procedures, or produce formally correct assignments. AI brutally exposes the limitations of this model. If a machine can automatically perform purely reproductive tasks, then the specifically human value of education must shift toward other dimensions: interpretation, creativity, critical thinking, the ability to formulate authentic questions, interdisciplinary understanding, and ethical responsibility. Paradoxically, artificial intelligence might force schools to rediscover their deepest cultural function.

3. Artificial Intelligence and Creativity

One of the most surprising aspects of generative AI concerns the relationship between automation and creativity. For a long time, it was thought that machines could primarily replace repetitive and standardized tasks, while creativity would remain an exclusively human prerogative. Today, this belief appears less solid. Artificial intelligence systems are capable of producing artistic images, literary texts, music, computer code, and narrative content that is often highly sophisticated.

This phenomenon inevitably raises concerns. If a machine can generate seemingly creative works, what still distinguishes human creation? And above all: what role can schools play in a world where content production can be automated? In reality, the issue extends beyond art or writing. It concerns the very meaning of contemporary creativity. Authentic creativity does not simply consist of combining linguistic or stylistic elements. It involves experience, intentionality, the ability to attribute meaning to the world, an awareness of historical context, and interpretive responsibility. Machines can produce sophisticated variations on existing content, but they lack an existential relationship with what they produce. Nevertheless, the effectiveness of generative AI inevitably changes human intellectual work.

Many creative professions are already changing. Designers, journalists, programmers, translators, illustrators, and authors find themselves collaborating daily with intelligent systems. Schools should prepare students for this new reality, avoiding both a nostalgic rejection of technology and the illusion that human creativity can simply be replaced. Education will likely need to focus increasingly on the ability to make sense of things, interpret critically, and construct original worldviews.

4. The Risk of Cognitive Dependence

Every powerful technology inevitably alters the way humans think. Writing reduced the role of oral memory. The printing press transformed our relationship with knowledge. The Internet changed access to information. Artificial intelligence now risks directly transforming the relationship between humans and cognitive processes. The main risk is not so much the complete replacement of human intelligence as the gradual delegation of fundamental cognitive functions. If students become accustomed to constantly using automated systems to write, summarize, translate, analyze, or argue, they may gradually weaken their intellectual autonomy. The immediate availability of automated answers risks reducing the time available for reflection, in-depth analysis, and the personal construction of knowledge. The speed of artificial intelligence, in fact, clashes with the slow pace of human learning. Truly understanding a problem often requires uncertainty, error, cognitive effort, and personal reworking. Machines, on the other hand, tend to eliminate precisely these areas of friction, offering immediate and seemingly complete results. Schools should therefore openly address the

issue of cognitive dependence. It is not a matter of preventing the use of artificial intelligence, but of teaching students to use it critically. Students must understand when automation serves as a useful aid and when it risks replacing essential cognitive processes. Without this awareness, there is a risk of producing individuals who are capable of using sophisticated tools but are progressively less able to think independently.

5. Schools as Places for Critical Thinking

Just as digital platforms and artificial intelligence are transforming our relationship with knowledge, schools could become one of the few social spaces dedicated to the critical construction of knowledge. This role seems more important today than ever before. We live immersed in a constant stream of information, images, and automatically generated content. In this environment, the challenge is no longer merely accessing knowledge, but distinguishing between truth and manipulation, between information and propaganda, and between authentic knowledge and mere algorithmic plausibility. Schools should therefore assume a new cultural centrality—no longer merely a place for the transmission of information, but a space for critical interpretation, rational debate, and the development of awareness.

This also implies a transformation of the teacher's role. The teacher can no longer be merely the one who possesses information that students lack. Instead, the teacher becomes a cultural guide, a critical mediator, and a designer of cognitive pathways capable of fostering intellectual autonomy. Artificial intelligence does not eliminate the need for school. Instead, it makes school even more necessary. Because in a society governed by algorithms and automated systems, the real educational challenge is not to produce more information, but to educate human beings capable of understanding its meaning.

Part II — What Computer Science Really Is

Chapter 4 — The Misconception of Coding

1. How Coding Has Changed Schools

Over the past fifteen years, coding has become one of the most visible symbols of educational innovation. Few terms have spread so rapidly in the Italian and European educational debate. Workshops, coding weeks, educational robots, visual environments, and digital platforms have gradually made their way into schools at all levels, often accompanied by the idea that programming represents the new literacy of the 21st century. This phenomenon has not been without positive effects. For many years, schools had viewed digital technology almost exclusively as a passive tool. Students mainly learned to use software designed by others: word processors, presentations, spreadsheets, and online platforms. Coding, on the other hand, introduced a different idea: the computer not as a mere tool to be consumed, but

as an environment to design and control. For many children and teenagers, writing a program—even a very simple one—has meant discovering that digital technology is not magic, but a human creation. Through environments such as Scratch or educational robotics, many students have begun to develop important insights into logic, the sequential nature of processes, and the relationship between instructions and machine behavior.

Coding has also had the merit of reintroducing a hands-on dimension—one that is often overlooked—into schools. Programming means experimenting, identifying errors, correcting procedures, and building solutions. In this sense, coding has helped promote more active and participatory teaching methods. It has also held strong symbolic value. In a society dominated by digital technologies, coding has appeared to many as a tool for cultural empowerment—not merely using applications built by others, but learning, at least in part, to understand their underlying logic. However, the very success of coding has given rise to a profound misunderstanding: the idea that coding is synonymous with computer science. And this is where the problems begin.

2. Why Coding Isn't Enough

Reducing computer science to coding means reducing a complex scientific discipline to a single technical activity. It's a bit like equating mathematics with simple arithmetic or literature with spelling. Programming is certainly an important component of computer science, but it by no means exhausts the discipline's cultural and scientific nature. In fact, many coding activities in schools risk remaining superficial. Students learn to move characters on the screen, build simple games, or assemble logic blocks without developing a real understanding of computational systems. In these cases, coding becomes a sort of new, simplified technical education—often closer to technological entertainment than to scientific training. The problem does not lie with coding itself, but with the way it is interpreted. When coding is presented as the ultimate goal of computer science education, schools risk losing sight of much deeper questions: What is an algorithm? How does data work? How do platforms organize information? What is the relationship between automation and society? How does artificial intelligence work? Contemporary computer science is not just about writing code. It is about building models, managing information, simulating reality, the functioning of networks, cybersecurity, artificial intelligence, and the relationship between data and power.

The cultural risk is enormous. A school that reduces computer science to coding might produce students capable of using basic programming tools but unable to truly understand the digital world in which they live.

This misunderstanding is even more evident today, in the age of generative artificial intelligence. If a machine can automatically generate computer code, then educational value can no longer be equated simply with learning the syntax of a programming language. Instead, it becomes

essential to understand problems, design systems, interpret results, and critically evaluate the behavior of technologies.

3. Programming Does Not Mean Understanding

There is a fundamental difference between knowing how to program and understanding computer science.
Many students today are able to build small programs without having any idea of how computational systems actually work. Similarly, millions of people drive a car every day without knowing the mechanics that make it work. Schools should avoid this same oversimplification when it comes to digital technology. Programming can certainly develop logical and design skills. But if coding remains confined to procedural exercises, it risks producing an extremely limited form of literacy.
Understanding computer science, on the other hand, means understanding that behind every digital system there are abstract models, data structures, algorithms, networks, computational architectures, and logics of information representation. It means understanding that platforms are not simply neutral tools, but complex ecosystems that incorporate economic, cultural, and political choices. Computer science is a science of information representation and processing. Programming is merely one of the languages through which this science expresses itself. The central educational challenge, therefore, is not simply to teach how to write code, but to develop a mindset capable of understanding complex processes, relationships, and systems. This profoundly shifts the pedagogical perspective. Coding should not be presented as an end in itself, but as a potential introductory tool toward a broader understanding of the computational paradigm.

4. The Risk of Reducing Computer Science to a Lab Activity

In recent years, many schools have addressed the topic of computer science through extracurricular activities, occasional workshops, or special projects. To some extent, this has been inevitable, given the historical absence of a truly structured presence of computer science in Italian curricula. However, this approach has produced a significant side effect: computer science has often been perceived as something separate from the school's general culture. An afternoon coding lab, a week dedicated to robotics, or a few activities on online safety are certainly useful initiatives, but on their own they are not enough to build a true computer science education. The risk is that digital technology will become a sort of technological appendix to the curriculum, without truly impacting the cultural model of education. Computer science cannot be confined to a marginal or purely laboratory-based space. The digital revolution has transformed the very nature of contemporary knowledge. It now permeates mathematics, physics, biology, economics, communication, art, and the social sciences. For this reason, computer science should become an integral part of general cultural education. This does not mean that every discipline must

turn into computer science. Rather, it means recognizing that the computational paradigm is now one of the fundamental languages of our time. The Italian school system is still struggling to make this transition. Very often, digital technology is treated as a mere technological update, whereas the true transformation concerns the very way in which society produces knowledge, organizes information, and builds social relationships.

5. From Syntax to Understanding Systems

The real educational challenge, therefore, lies in moving from syntax to an understanding of systems. For many years, computer science education has focused primarily on tools: learning a programming language, using a platform, or mastering a software application. But in contemporary society, tools change rapidly. Programming languages and applications become obsolete within a few years. What remains constant, however, is the ability to understand deep structures. Understanding a digital system means grasping how data, algorithms, networks, and decision-making processes interact. It means understanding that every platform represents a complex ecosystem in which technological, economic, and cultural dimensions are constantly intertwined. This approach also radically changes the meaning of computational thinking. Computational thinking should not be interpreted as mere programming training, but as the ability to interpret the complexity of the world through models, relationships, and processes. From this perspective, computer science naturally engages with mathematics, the sciences, philosophy, economics, languages, art, and civics. The school of the future will likely need to teach less syntax and more systemic understanding. Because the central challenge will no longer be simply using digital technologies, but understanding a society increasingly organized by invisible computational systems. And it is precisely here that coding reveals both its value and its limitations. It can serve as an initial gateway to digital logic, but it cannot be mistaken for computer science itself. Computer science, today, is much more than that.

Chapter 5 — Computer Science: A New Form of Thought

1. Algorithms, Data, and Models

To truly understand what computer science is today, we must dispel a widespread misconception: the idea that it is primarily synonymous with computers. Computers are merely the physical manifestation of something much deeper. Contemporary computer science, in fact, concerns the way information is represented, organized, transformed, and used. At the heart of this paradigm lie three fundamental concepts: data, algorithms, and models. Data constitutes the raw material of the digital society. Every human activity now produces information that can be transformed into data: movements, communications,

purchases, images, text, preferences, social relationships. Contemporary society generates vastly greater amounts of data than any previous era.

But data alone means nothing. To acquire meaning, it must be organized and interpreted through algorithms. An algorithm is a procedure capable of transforming initial information into final results through a defined sequence of operations. Algorithms are not exclusive to computer science: they have existed for centuries in mathematics, logic, and even in everyday activities. However, the digital revolution has transformed algorithms into the primary mechanism for organizing contemporary reality. They determine what content we see on social networks, suggest products, classify information, manage financial flows, guide logistics systems, and control critical infrastructure.

The third fundamental element is models. Contemporary computer science is not limited to processing data; it constructs abstract representations of reality. A computational model is a simplification of the world designed to be processed by a machine. Through these models, it is possible to simulate physical phenomena, biological systems, social networks, economic markets, and cognitive processes. Artificial intelligence represents the most advanced stage of this evolution today. Machine learning systems build mathematical models capable of recognizing correlations and generating predictions based on enormous amounts of data. The true cultural revolution in computer science stems precisely from this: the ability to represent the world through computational structures. Schools should help students understand this transformation. It is not simply a matter of learning to use digital technologies, but of understanding the new language through which contemporary society organizes knowledge.

2. Computational Thinking

In recent years, the concept of computational thinking has become central to the international educational debate. However, it is often interpreted in a reductive way, almost as a synonym for coding or basic programming. In reality, computational thinking is something much broader. It represents a way of tackling complex problems through processes of analysis, decomposition, modeling, and the construction of procedures. It means learning to recognize structures, relationships, and patterns within complexity. Computational thinking is not the same as programming. Programming is merely one of the possible applications of a more general mindset. A student can learn the syntax of a programming language without developing true computational thinking. Conversely, it is possible to develop deep computational insights even without directly using a computer.

The true essence of computational thinking lies in the ability to abstract problems, identify relationships, build models, analyze systems, and design processes. This way of thinking has enormous cultural significance because it extends far beyond computer science. It naturally intersects with mathematics, physics, biology, economics, philosophy, and even the humanities. Contemporary society is characterized by increasingly complex and interconnected systems. Climate change, social networks,

the global economy, artificial intelligence, and cybersecurity are problems that cannot be understood through simple linear models. Computational thinking offers tools to address this complexity. Schools should therefore avoid reducing it to a mere specialized technique. Rather, it represents a new form of cognitive literacy capable of integrating logic, creativity, and systems thinking.

3. Simulation and Representation of Reality

One of the most revolutionary aspects of contemporary computer science is the ability to simulate reality. For centuries, scientific knowledge has been based primarily on observation, experimentation, and mathematical formalization. The computational paradigm introduces a new dimension: the possibility of building executable models capable of dynamically reproducing complex phenomena. This profoundly changes the very way science is conducted. Through computational simulations, it is now possible to study climatic phenomena, biological systems, cosmological evolutions, economic networks, or social dynamics that would be impossible to observe directly. Simulation also alters the relationship between theory and practice. A computational model is not merely an abstract description; it is a system capable of producing observable behaviors.

Schools should place much greater emphasis on this aspect. Students should not limit themselves to memorizing formulas or static definitions. Instead, they could explore dynamic phenomena through simulations, visualizations, and computational models. This approach fosters systemic understanding and interdisciplinarity. Computer science, in fact, allows us to connect mathematics, physics, biology, economics, and the social sciences within shared simulation environments. Contemporary reality appears less and less divisible into isolated disciplines and increasingly interpretable as a network of interconnected complex systems. In this sense, computer science is not merely a new technology. It is a new way of representing the world.

4. Computer Science and Mathematics

One of the most significant conflicts that emerged during the debate on the new National Guidelines concerns the relationship between mathematics and computer science. Many mathematicians fear that the introduction of computer science may further reduce the space devoted to mathematics in schools. This concern is not without merit, especially in an educational system where curriculum time is limited.

But the problem is often framed incorrectly. Computer science and mathematics are not opposing disciplines. On the contrary, they share deep-seated structures and have historically developed in close mutual relation. Computer science stems largely from mathematics. Algorithms, logic, computation theory, formal structures, and modeling all have very deep mathematical roots. At the same time, computer science also profoundly changes the way mathematics is done. Computational simulation, data analysis, symbolic computation, and artificial intelligence are transforming numerous areas of contemporary mathematical research

. Schools should therefore move beyond the artificial dichotomy between these disciplines. The goal should not be to replace mathematics with computer science, but to foster a deeper dialogue between mathematical thinking and the computational paradigm. Today, both represent fundamental languages of the modern world. Mathematics provides rigor, abstraction, and the ability to prove. Computer science adds dynamism, simulation, and the management of complexity. Strictly separating them risks impoverishing both.

5. Computer Science as a Cultural Discipline

Perhaps the most important point is to understand that computer science is not merely a technical discipline. It represents one of the great cultural transformations of the contemporary modern era. Just as mathematics has changed the way we interpret the physical world and the experimental sciences have transformed our relationship with nature, computer science is changing the relationship between human beings, knowledge, and reality. The computational paradigm now permeates the economy, communication, medicine, art, scientific research, politics, and social organization. Today, failing to understand computer science means failing to understand a growing part of the contemporary world.

For this reason, schools should recognize computer science as a discipline with full cultural significance. It is not merely a matter of training future ICT specialists or programmers. Rather, it is about providing all students with cognitive tools capable of helping them understand the algorithmic society in which they live. Computer science should therefore be introduced into schools not as a mere technical update, but as a new form of critical literacy for the contemporary world. The real educational challenge is not teaching students how to use increasingly sophisticated machines. It is teaching them to understand the cultural significance of the world those machines are building.

Chapter 6 — Computer Science at the Intersection of Science, Philosophy, and Society

1. The Problem of Artificial Intelligence

For many years, computer science was viewed primarily as a technical discipline linked to automatic computation, information management, and software development. The advent of artificial intelligence, however, has profoundly altered this perception. AI is no longer just about machines. It directly concerns the very meaning of human intelligence, knowledge, and creativity. For the first time in history, artificial systems have become capable of producing texts, images, music, translations, and complex solutions—often with results indistinguishable from those produced by humans. This phenomenon compels schools, philosophy, and society as a whole to grapple with questions that, until just a few years ago, seemed to belong almost exclusively to the realm of science fiction.

What does it really mean to understand? Can a machine be considered intelligent? What is the difference between simulation and thought? And above all: what role remains for human beings in a society in which more and more cognitive activities are being automated? Artificial intelligence has brought questions that directly concern the nature of human beings back to the center of the educational and cultural debate. For this reason, computer science can no longer be taught as a mere set of operational techniques. It now straddles the boundary between science, philosophy, ethics, and social organization. Schools should help students understand that AI represents not just a new technology, but a profound anthropological and cultural transformation.

2. What Does "Understanding" Mean?

One of the most fascinating problems posed by artificial intelligence concerns the very concept of understanding. When a system like ChatGPT generates a coherent text, translates a poem, or solves a math problem, can we say that it "understands" what it is doing? Or is it simply a sophisticated statistical simulation? This question now permeates not only scientific research but also the world of education. For centuries, human knowledge has been associated with experience, consciousness, and the ability to attribute meaning to the world. Artificial intelligence systems, on the other hand, operate through mathematical correlations learned from enormous amounts of data. They produce often impressive results without possessing direct experience of reality. A machine can speak of the sea without ever having seen the sea. It can write about pain, love, or fear without ever having felt emotions.

This creates a culturally new situation. Contemporary society is, in fact, faced with systems capable of effectively simulating human language even without possessing consciousness or intentionality in the traditional sense of the term. Schools should address this issue openly. Not to turn computer science education into abstract philosophy, but because the relationship between humans and intelligent systems inevitably changes the way students perceive knowledge, intelligence, and even creativity. Understanding these differences becomes part of the new contemporary literacy.

3. Machines That Generate Knowledge

Artificial intelligence is no longer limited to executing instructions. It produces content, organizes information, and participates ever more directly in the processes of knowledge construction. This phenomenon profoundly alters the relationship between knowledge and cultural authority. For centuries, the production of knowledge has been associated with recognizable human figures such as scholars, teachers, researchers, journalists, and artists. Today, however, a growing portion of the information circulating in society is filtered, synthesized, or even generated by automated systems. Many students already use artificial intelligence as a study tool, writing aid, personal tutor, advanced search engine, summarization system, or full-fledged

cognitive assistant. The traditional distinction between “human knowledge” and “technical tools” is thus becoming increasingly blurred.

This has enormous educational implications. Schools can no longer limit themselves to conveying content that is readily accessible through automated systems. Instead, they must help students develop interpretive, critical, and reflective skills that no machine can fully replace. Furthermore, a particularly delicate epistemological problem arises: how can we assess the reliability of automatically generated content? Generative systems can, in fact, produce convincing but erroneous texts, nonexistent citations, distorted interpretations, or false information with a high degree of linguistic confidence. Formal authority no longer necessarily coincides with truth. For this reason, critical thinking becomes even more important. Students will have to learn not only to use artificial intelligence but also to critically question what it produces.

4. Automation and Freedom

Every technological revolution changes the relationship between human beings and freedom. Industrial machines transformed physical labor. Digital technologies are now transforming cognitive work. An increasing portion of decisions, analyses, and organizational processes is gradually being delegated to automated systems. This phenomenon yields enormous benefits in terms of speed, efficiency, and processing capacity. However, it also raises a fundamental question: to what extent can a society automate its decision-making processes without reducing human autonomy?

This issue directly concerns education. If students become accustomed to continually delegating activities such as writing, memorization, summarization, translation, design, or problem-solving to intelligent systems, they risk gradually losing part of their cognitive autonomy. Freedom, in fact, consists not only in the ability to choose, but also in the capacity to critically understand the world. A society that uses intelligent systems without understanding them risks becoming a technologically advanced but culturally dependent society. For this reason, computer science education inevitably takes on an ethical and political dimension. Schools must educate citizens who are capable not only of using sophisticated tools, but also of maintaining independent judgment within increasingly automated ecosystems.

5. Computer Science as a New Form of Literacy

In the 20th century, literacy meant teaching people to read, write, and do arithmetic. Today, these skills remain fundamental, but they are no longer sufficient for full participation in contemporary society. The digital revolution has introduced a new form of literacy: the ability to understand the computational paradigm. This does not mean turning all students into programmers or ICT specialists. Rather, it means providing the cultural tools to interpret a society

in which data, algorithms, and intelligent systems organize an ever-growing part of collective life. Computer science should therefore be taught as a cultural discipline capable of connecting logic, mathematics, philosophy, science, ethics, politics, and social organization.

Today, it represents one of the primary languages of the contemporary world. Not understanding computer science means not fully understanding how the contemporary world works. For this reason, schools should recognize that computer science plays a role analogous to that which mathematics played during the modern scientific revolution—not merely a specialized technique, but a fundamental way of interpreting reality. The true educational challenge of the 21st century, therefore, does not lie in teaching students how to use increasingly sophisticated digital tools. It lies in educating human beings capable of critically understanding the cultural, social, and ethical significance of the algorithmic revolution.

Part III — How Schools Are Changing

Chapter 7 — What Kind of Computer Science Should We Really Teach?

1. A Question More Difficult Than It Seems

One of the most widespread illusions in the public debate on digital education is the belief that the main problem is deciding “how much” computer science to introduce into schools. In reality, the decisive question is much more complex: what kind of computer science should be taught? Behind this question lies a profound cultural conflict. For some, computer science in schools should primarily consist of coding and programming. For others, it should be a form of digital literacy tied to the informed use of technology. Then there are those who envision a highly scientific discipline, closely related to mathematics and systems theory, and those who instead place greater emphasis on the social, ethical, and cultural aspects of the digital world. All of these perspectives contain valid elements. The problem arises when one of them claims to exhaust the meaning of contemporary computer science on its own.

The digital revolution has, in fact, brought about a transformation far broader than can be represented by a single technical discipline. Today, computer science cuts across scientific research, economics, communication, medicine, art, politics, and social organization. Consequently, schools must also develop a more nuanced and mature vision of how to teach it. The goal cannot simply be to train future programmers. Instead, schools must educate citizens capable of understanding how algorithmic society functions. This implies an important consequence: computer science in schools cannot be organized as mere technological training. It must become a cultural discipline capable of developing logic, systemic thinking, critical understanding, and awareness of the relationship between technology and society.

2. Elementary School: Intuition, Logic, and Discovery

In elementary school, the main challenge is not teaching specialized content, but building fundamental cognitive insights. One of the most common mistakes is trying to introduce complex topics or unnecessary technical details too early. Computer science in elementary school should not become either a form of professional "mini-programming" or mere digital entertainment. Its goal should be much deeper: to develop children's ability to reason through relationships, sequences, rules, and processes. At this age, computational thinking can be introduced even without computers. Logic games, hands-on activities, problem-solving paths, building activities, simulations, and collaborative activities help children develop fundamental insights into the structure of processes and the representation of information.

Digital technology should be introduced in elementary school primarily as a creative and exploratory language. Visual environments like Scratch can be highly valuable if used not merely as technical training but as expressive tools through which children create stories, simulations, games, and representations of the world. At this stage, it is essential to maintain a balance between hands-on experience and technological mediation. Elementary school should not foster an early dependence on screens but rather help children gradually understand that behind the functioning of technologies lie rules, models, and design decisions. Even the first elements of digital citizenship can be introduced at this age, through discussions about respecting others online, managing information, and using technology in a balanced way. The true goal is not to train young programmers but to foster curiosity, logical thinking, and critical awareness.

3. Secondary School: Systems, Data, and Networks

In secondary school, the curriculum can become progressively more structured. Students develop greater abstract thinking skills and can begin to understand how digital systems work in a more rigorous way. It is at this stage that computer science should cease to be viewed as merely a technology lab and begin to be presented as a cultural and scientific discipline. Students should gradually become familiar with fundamental concepts such as algorithms, data, networks, information representation, modeling, and automation. Programming can certainly play an important role, but as a tool for understanding processes and building models—not as an end in itself.

The topic of data becomes particularly important. Contemporary society relies increasingly on the collection, interpretation, and management of enormous amounts of information. Students should therefore understand how search engines, social media platforms, recommendation systems, and predictive algorithms work. The topics of cybersecurity and digital citizenship also emerge at this stage. Young people now spend a significant part of their lives online. Schools should therefore explicitly address issues such as digital identity, misinformation, algorithmic manipulation, privacy, technology addiction, and information security. Secondary schools should also foster dialogue between

computer science and other disciplines. Mathematics can engage with algorithms and modeling; the sciences with simulations and data; philosophy with the ethics of artificial intelligence; and the humanities with the relationship between digital platforms and identity formation. Computer science should gradually become a tool for understanding the complexity of the contemporary world.

4. High School: AI, Simulation, and Complexity

In upper secondary school, computer science should fully assume the role of a cultural discipline of the contemporary world. Students now possess sufficient cognitive tools to tackle abstract problems, epistemological reflections, and complex systems. It is precisely at this stage that schools should explicitly introduce the relationship between computer science, artificial intelligence, and the organization of society. The goal should not simply be to teach technologies destined to become obsolete within a few years. The real challenge lies in providing enduring conceptual frameworks. Students should understand the significance of computational models, the functioning of intelligent systems, the relationship between data and prediction, the limits of automation, the role of algorithms in social life, and the transformation of cognitive work.

Artificial intelligence should become a permanent part of the school curriculum—not as a technological fad, but as a central cultural issue of the 21st century. High schools should also place much greater emphasis on simulation, modeling, and interdisciplinarity. Today, the computational paradigm allows us to address climate, economic, biological, and social problems through complex dynamic systems. Understanding these models means developing cognitive abilities that go far beyond the simple memorization of information. From this perspective, computer science becomes a tool for interpreting the contemporary world.

5. What Should Not Be Taught

Two opposing errors often emerge in the public debate on digital education. The first consists of reducing computer science to the operational use of technologies. In this case, schools risk becoming training centers for the use of commercial platforms that are destined to change rapidly. The second error, on the other hand, consists of introducing abstract technical concepts too early that lack real educational significance. Schools should not constantly chase the latest technological trends. It makes no sense to build curricula centered on specific programming languages, temporary applications, or tools destined to become obsolete quickly. Similarly, it would be a mistake to turn computer science education into a form of early professional specialization. The goal is not to prepare hyper-specialized technicians as early as twelve or thirteen years of age. It is to develop cognitive skills, critical understanding, and cultural flexibility. Schools should avoid both technical jargon and superficiality. Computer science cannot be reduced to mere recreational coding or to

sterile abstract theory. Instead, it must help students understand the cultural significance of the digital transformation.

6. Why Progressive Curricula Are Needed

One of the most obvious problems in Italian schools is the fragmentation of digital experiences. Many computer science-related activities still depend on the individual initiative of teachers or on temporary projects lacking real educational continuity. This approach is no longer sufficient. The digital transformation requires progressive, coherent, and vertical curricula capable of accompanying students throughout their entire educational journey. Computer science cannot be introduced through isolated episodes. It must become a structural part of contemporary education. A truly progressive curriculum should gradually develop logical intuition in elementary school, an understanding of systems in middle school, and critical thinking skills along with more advanced forms of modeling in high school. Above all, however, it should maintain strong cultural continuity. The true goal is not to teach individual technologies, but to build a new form of literacy capable of helping students live consciously within algorithmic society. And this is precisely the great challenge of the new National Guidelines: not simply to introduce a new subject, but to rethink the relationship between school, knowledge, and technological transformation.

Chapter 8 — Teachers Facing the Digital Revolution

1. The Real Problem Isn't Computers

In the public debate on digital education, the topic is often addressed almost exclusively in technological terms. The discussion centers on computers, internet connections, platforms, devices, labs, and software. For many years, it was believed that the main obstacle to innovation was the lack of technological tools in schools. Of course, infrastructure and devices are important. A school without adequate internet connections or even basic technological resources cannot seriously tackle the digital transformation. However, the main problem has never been merely technical. The real issue concerns the relationship between teachers and the cultural shift brought about by the digital revolution. Many teachers today are facing a rapid transformation that is simultaneously changing tools, languages, cognitive models, and even the very meaning of teaching. For a significant portion of the teaching staff, this change inevitably leads to disorientation.

The difficulty does not stem from mere conservative resistance, as is sometimes superficially claimed. Rather, it arises from the speed with which the digital paradigm is altering the relationship between knowledge, cultural authority, and learning. For centuries, teachers have been one of the primary sources of access to knowledge. Today, students live immersed in information ecosystems in which teachers

no longer holds a monopoly on information. The internet and artificial intelligence have radically altered this balance. Many teachers therefore perceive a crisis in their traditional role. But the problem cannot be addressed by pitting technology against schools. The central issue is understanding how the teacher's educational role can be transformed within a society increasingly organized by digital systems and artificial intelligence. The key is not to introduce more computers into classrooms. It is to help teachers redefine their cultural role in an algorithmic society.

2. Is Teacher Training Impossible?

One of the arguments most frequently raised against the introduction of computer science in the new National Guidelines concerns teacher training. Many argue that the Italian school system currently lacks a sufficient number of teachers prepared to seriously address this transformation. The problem is real. For decades, computer science has remained a marginal subject in the initial training of most Italian teachers. Many teachers currently in service were trained in a cultural context that predates the contemporary digital revolution entirely. But the issue is even more complex.

It is not simply a matter of teaching teachers how to use new technological tools. The real challenge lies in building a new pedagogical culture capable of integrating computer science, artificial intelligence, and critical thinking into the school curriculum. Many digital refresher courses have failed precisely because they were reduced to mere technical training. Learning to use a platform or a software program does not mean understanding the cultural significance of the digital transformation. Teacher training should instead address much deeper issues, such as the relationship between artificial intelligence and learning, the meaning of computational thinking, the role of data and algorithms in society, the transformation of cultural authority, the problem of cognitive dependence, and the development of critical thinking in the age of platforms. This inevitably makes training much more difficult. There are no quick or purely administrative solutions. Schools are undergoing a historic transformation that requires time, cultural reflection, and a redefinition of professional identities. For this reason, simply talking about "digital training" for teachers risks being misleading. The real challenge is to develop a new type of teacher capable of navigating the algorithmic society critically.

3. The Fear of Computer Science

Behind much of the resistance to introducing computer science in schools lies a deeper factor: fear. It is not merely a fear of technology itself, but a fear of losing established cultural reference points. In fact, computer science and artificial intelligence are transforming our relationship with knowledge, the role of the expert, the value of memory, the meaning of writing, and even the very concept of learning. For many teachers, this inevitably creates a sense of instability. Traditional schooling was built around a relative stability of knowledge and language. Today, however,

technologies, platforms, and tools are constantly changing. This creates the perception of a cultural landscape that is increasingly beyond our control. There is also another factor that is often underestimated: many teachers fear that digital technology may impoverish thought, reduce concentration, and foster cognitive superficiality.

These concerns are not without merit. The contemporary digital environment does indeed tend to prioritize speed, fragmented attention, and the rapid consumption of information. Even artificial intelligence can lead to cognitive dependence if used passively. But precisely for this reason, schools cannot simply reject technology. The only way to critically engage with the digital transformation is to understand it. Excluding computer science and artificial intelligence from schools would mean leaving students completely on their own within digital ecosystems designed by commercial platforms. Instead, schools should become the place where students learn to critically understand these transformations.

4. New Professional Roles in Schools

Digital transformation does not merely change the content of instruction. It also changes the very organization of schools. Educational institutions of the future will likely need new professional roles capable of combining technological, pedagogical, and organizational skills. For many years, digital initiatives in schools have relied almost exclusively on the goodwill of individual teachers with a particular interest in technological innovation. This model is no longer sustainable. Contemporary computer science, in fact, requires scientific, organizational, ethical, methodological, and infrastructural skills that are unlikely to be concentrated in a single professional role. Schools will therefore need to develop new forms of collaboration among teachers, technologists, cybersecurity experts, specialists in educational artificial intelligence, and professionals capable of integrating technical expertise with instructional design. School administrators will also need to gradually assume new responsibilities in managing digital ecosystems. The school of the future will not simply be a school with more computers. It will be a cultural and organizational environment profoundly transformed by the computational paradigm. This will inevitably require new professional roles and new forms of educational cooperation.

5. The Teacher as a Critical Guide

Paradoxically, just as the internet and artificial intelligence seem to render the traditional role of the teacher obsolete, the teacher's cultural role could become even more important. In a society where information is universally available, the real educational challenge is no longer merely conveying content. It becomes helping students interpret it critically. The 21st-century teacher can no longer be merely a transmitter of facts. They must become a cultural guide, a critical mediator, and a designer of cognitive pathways. This means helping students distinguish between

information and knowledge, to recognize algorithmic manipulation, to assess the reliability of sources, to understand the limitations of artificial intelligence, and to develop independent judgment.

Teachers thus take on a new role. They do not compete with machines in terms of speed or the quantity of information. Their uniquely human value lies in the ability to attribute meaning, build cultural connections, and foster critical thinking. Artificial intelligence does not eliminate the need for teachers. Rather, it makes the presence of educators capable of guiding students through the complexities of digital society even more essential. The real challenge for contemporary schools, therefore, is not to replace teachers with technology, but to redefine the teacher's role within a world transformed by technology. And it is precisely on this redefinition that the future of education in the algorithmic society will largely depend.

Chapter 9 — Labs, Data, and Smart Environments

1. Why Digital Technology Isn't Just an Interactive Whiteboard

For many years, the digital transformation of schools has been identified almost exclusively with the introduction of technological devices into classrooms. Interactive whiteboards, tablets, online platforms, and electronic gradebooks have become the visible symbols of educational innovation. But all too often, this innovation has remained superficial. In many cases, digital technology has merely replaced older tools without truly changing the cultural model of teaching. The interactive whiteboard has taken the place of the traditional blackboard, digital documents have replaced paper ones, and online platforms have transferred essentially unchanged teaching practices to the screen. The problem is not the presence of technology, but the way it is interpreted. The digital revolution is not simply about using new tools. It changes the very way knowledge is produced, organized, and shared. For this reason, schools cannot limit themselves to introducing technologies into educational structures that have remained essentially 20th-century in nature.

Digital technology is not just a more sophisticated blackboard. It represents a completely new cognitive environment, in which data, simulations, networks, computational models, and artificial intelligence transform the relationship between students, teachers, and knowledge. Schools should therefore stop viewing digital technology as a mere technological tool and begin to conceive of it as a cultural environment for learning. This inevitably implies a transformation of school spaces, teaching methodologies, and educational relationships.

2. Simulation and Learning

One of the most innovative aspects of educational computing concerns the possibility of using computational simulations to transform the way students learn.

For centuries, school instruction has been based primarily on the static transmission of knowledge. Students studied theoretical descriptions of natural, economic, or social phenomena that were often difficult to visualize concretely. Computational technologies now make something different possible. Through dynamic simulations, students can directly observe the behavior of complex systems. It is possible to simulate climatic phenomena, biological processes, social networks, economic markets, energy systems, or physical models that are difficult to reproduce in a traditional laboratory. This profoundly changes the relationship with knowledge. Students are no longer limited to memorizing abstract descriptions. They can modify parameters, observe consequences, formulate hypotheses, and understand dynamic relationships between variables. Simulation thus fosters a form of exploratory learning that is much closer to the actual workings of contemporary scientific research.

The very concept of a laboratory is also changing. The laboratory of the future will not merely be a physical space dedicated to traditional experiments, but a hybrid environment in which hands-on experience and computational modeling are in constant dialogue. This approach also fosters interdisciplinarity and systems thinking. Many contemporary problems—from climate change to energy management, from the spread of epidemics to the global economy—can only be understood through complex models and dynamic simulations. Computer science therefore offers schools new tools not only for conveying information but also for developing a deep understanding of contemporary systems.

3. Educational Robotics

Educational robotics represents one of the most interesting examples of integration between computer science, creativity, and experiential learning. Robots possess a pedagogically powerful feature: they make the relationship between software and the physical world visible. Students see firsthand how instructions, algorithms, and sensors produce observable behaviors. This profoundly transforms their relationship with digital technology. Many contemporary technologies, in fact, operate invisibly. Algorithms, neural networks, and digital platforms operate silently behind seemingly simple interfaces. Robotics, on the other hand, makes the computational process tangible and observable. Students can build systems that move, react, measure, avoid obstacles, and interact with their environment. Educational robotics also fosters collaborative learning, design, and problem-solving. Building a robot means tackling concrete problems, correcting errors, testing hypotheses, and integrating diverse skills. In this sense, robotics provides an ideal environment for developing computational thinking and design skills. But here, too, we must avoid oversimplifications. Robotics should not be reduced to a mere spectacle or competitive activity. Its educational value lies not in the technological effect itself, but in the opportunity to understand the relationship between information, automation, and the organization of systems. The

schools should use robotics as a tool to develop a critical understanding of smart technologies, not just operational skills.

4. Cloud, Cybersecurity, and Infrastructure

Contemporary schools now depend increasingly on complex digital infrastructures. Cloud platforms, online environments, digital archives, communication systems, and artificial intelligence services organize a growing part of school life. However, very often these infrastructures remain invisible to students and even teachers. Schools use digital technology without truly understanding its architecture. This creates a new form of cultural and technological dependence. Contemporary educational platforms are not simply neutral tools. They collect data, organize information, define how content is accessed, and influence the very functioning of educational processes. For this reason, schools should develop greater infrastructure awareness. Students should understand at least the fundamental principles of how networks, data management, cloud systems, and cybersecurity work.

The issue of digital security has also taken on strategic importance today. Contemporary society is increasingly dependent on infrastructure that is vulnerable to cyberattacks, manipulation, and data breaches. Schools cannot ignore these issues. Educating students about cybersecurity does not mean training military specialists or security technicians, but rather developing critical awareness of the relationship between technology, data, and the protection of individual freedom. Privacy, authentication, information protection, and the responsible management of digital identities should become an integral part of modern citizenship.

5. Schools as Digital Ecosystems

The digital transformation inevitably changes the way we conceive of the school as an institution. The school of the future will not simply be a traditional school enhanced by technological tools. It will gradually become a complex digital ecosystem, in which computational infrastructure, simulation environments, collaborative platforms, and intelligent systems will continuously interact with the educational dimension. This change requires a new cultural vision of the school. For many years, digital technology was perceived as an add-on to the traditional school system. Today, it is increasingly clear that it is, in fact, transforming the entire learning environment. Schools will need to become spaces where students and teachers learn not only to use technology but also to critically understand the digital ecosystems in which they live. This implies rethinking spaces, labs, libraries, modes of collaboration, the organization of knowledge, and the relationship between physical presence and virtual environments. The challenge is not to build schools dominated by technology, but schools capable of culturally governing technology. Because the decisive issue is not having more digital tools. It is

to educate citizens capable of understanding the human, social, and political significance of the computational environments in which an increasing part of contemporary life will take place.

Part IV — Europe, the World, and the Italian Challenge

Chapter 10 — Europe Between Rights and Algorithms

1. Europe Facing the Digital Revolution

In recent years, the European Union has gradually come to understand that the digital transformation is not merely a technological or economic issue, but a cultural and political challenge destined to redefine the future of democratic societies. For a long time, Europe viewed the digital revolution primarily as an industrial phenomenon dominated by major U.S. platforms and, more recently, by China's technological growth. However, the expansion of artificial intelligence, algorithmic systems, and global platforms has made it clear that the issue concerns not only the market but also the cultural and democratic sovereignty of contemporary societies. Schools are directly at the heart of this scenario. European educational institutions, in fact, face a dual challenge. On the one hand, they must prepare students to live and work in an advanced digital society. On the other, they must defend fundamental values of the European democratic tradition, such as cultural pluralism, individual freedom, transparency, and public accountability. It is precisely this tension that gives rise to the distinctive European approach to digital technology. Unlike other international models, Europe tends to view technology not only as a tool for economic innovation but also as an ethical, political, and cultural issue. This approach is clearly evident in European education policies of recent years.

2. DigComp and the Limits of Digital Competence

One of the most influential documents produced by the European Union is the DigComp framework, created with the goal of defining the fundamental digital competencies for European citizens. DigComp has played an important role because it has helped overcome the notion that digital literacy is simply about the technical use of tools. In fact, the framework introduces topics related to information, security, online communication, content creation, and problem-solving. However, the very success of DigComp has also highlighted some limitations of the European approach. The very concept of "digital competence" risks remaining ambiguous. In many educational contexts, it has been interpreted primarily as an operational skill: using platforms, browsing online, producing digital content, or using collaborative software. But today, all of this is no longer enough. Artificial intelligence and the algorithmic society have shown that the real issue is not merely using technologies, but understanding how they work and their

cultural implications. A society can be highly competent from an operational standpoint and at the same time completely dependent on systems it does not truly understand. For this reason, digital education cannot be limited to the development of technical skills. It must become an education that fosters critical thinking about algorithmic society. Europe today seems to be moving progressively in this direction, even if the path still appears incomplete.

3. The AI Act and the Issue of Liability

The adoption of the European AI Act likely represents one of the most significant and courageous steps toward building a digital policy grounded in democratic principles. For the first time, a major geopolitical region has attempted to comprehensively regulate artificial intelligence not only in economic terms but also in ethical, social, and political terms. Europe has consciously chosen to assert that not everything that is technologically possible must necessarily be accepted without limits. This choice carries profound cultural significance. The AI Act implicitly recognizes that intelligent systems are not neutral. They can lead to discrimination, manipulate information, influence behavior, and directly impact individual rights. The issue of algorithmic transparency thus becomes central. Many artificial intelligence systems, in fact, operate as “black boxes”: they produce results without users being able to truly understand the internal decision-making processes. This creates an enormous democratic problem. How can a citizen fully exercise their rights within a society governed by opaque automated systems? Europe has decided to address this question while also taking on an economic risk. The AI Act represents a painful yet highly civilized choice: accepting the possibility of having artificial intelligence that is less competitive commercially in order to defend values considered inviolable. It is the choice of a society that consciously decides to limit part of its economic efficiency in the name of protecting human dignity, transparency, and fundamental rights. This approach is also reminiscent, in part, of other major contemporary European initiatives, such as the Green Deal. In that case as well, Europe accepted higher energy costs and competitive disadvantages to uphold a vision of society founded on environmental sustainability and responsibility toward future generations. In both cases, the same underlying idea emerges: economic wealth cannot be the sole criterion by which to measure a civilization’s progress.

From an educational standpoint, this element is fundamental. Understanding artificial intelligence does not simply mean learning how to use it. It means understanding how models are built, what data they use, what biases they may incorporate, what limitations they have, and what social effects they produce. AI education thus becomes an integral part of the new democratic citizenship. Through the AI Act, Europe appears to recognize that the future of artificial intelligence cannot be left solely to market forces or global technological competition. This awareness should also profoundly influence the way schools teach computer science.

4. European Digital Sovereignty

One of the most important concepts to emerge in recent years is that of digital sovereignty. Much of the digital infrastructure used daily in Europe belongs to large non-European platforms. Search engines, social networks, cloud systems, and artificial intelligence models are predominantly controlled by U.S. or Chinese companies. This situation creates a new form of cultural, economic, and technological dependence. The problem is not limited to the economic control of these platforms. It concerns a society's very ability to understand and govern the tools that shape its collective life. A population unable to understand how digital technologies work inevitably risks passively submitting to them. For this reason, education plays a strategic role. Schools are, in fact, one of the few places where it is possible to develop a critical understanding of the digital transformation independent of the commercial logic of these platforms. But European digital sovereignty is not just about culture. It also represents an attempt to support and protect the European technology industry in a global competition dominated by foreign giants with vastly superior resources. Europe is aware that, without an autonomous industrial and regulatory policy, the European technology sector risks gradually succumbing to pressure from major U.S. and Chinese platforms.

This choice, too, entails inevitable contradictions. Defending a European digital ecosystem may, in fact, mean accepting services that are less efficient or more expensive than those of global competitors. This could penalize part of the European industry itself, which would be forced to use technological infrastructure that is less efficient and less cost-effective.

Europe, however, seems to view this sacrifice as the necessary price to pay for preserving strategic autonomy and cultural independence. Once again, a choice between civilizations emerges: giving up some immediate efficiency to defend the ability to maintain democratic control over the fundamental infrastructure of contemporary society. And this is where computer science and education once again become deeply intertwined. A society's true digital autonomy does not depend solely on its technological infrastructure. It depends above all on the level of cultural awareness among its citizens. Without citizens capable of critically understanding how digital systems work, no technological sovereignty can ever be truly democratic**.**

5. Education and Democracy in the Algorithmic Society

Today, the digital revolution is forcing Europe to redefine the relationship between education and democracy. For centuries, European schools have been tasked with shaping citizens capable of participating in public life through literacy, critical thinking, and knowledge of democratic institutions. The algorithmic society is profoundly altering this landscape. Today, a growing part of democratic life takes place on digital platforms governed by automated systems. Political information, opinion formation

Public life, social communication, and even electoral processes are influenced by algorithms. In this context, democratic citizenship requires entirely new skills. It is not enough to know political history or how institutions function. One must understand the relationship between data, platforms, algorithms, and power. European schools should therefore take on a new cultural mission: to educate citizens capable of engaging critically with algorithmic society. This inevitably implies a redefinition of the role of computer science in the National Curriculum Guidelines. Computer science can no longer be considered merely a technical or professional skill. It now represents a fundamental component of contemporary democratic education. What is truly at stake is not merely Europe's technological competitiveness. It is the very possibility of maintaining spaces of critical freedom within a society increasingly organized by automated systems and global platforms.

Chapter 11 — International Models

1. Why Look Beyond Italy

In recent years, many countries have undertaken profound transformations of their educational systems to address the digital revolution and the growing impact of artificial intelligence. Observing these experiences does not mean automatically importing foreign models, but rather understanding how different political cultures and school systems are attempting to address common challenges. The issue of computer science in schools is not merely about technological innovation. It concerns the relationship between education, economic development, citizenship, and cultural sovereignty. Each country interprets this transformation according to its own history and strategic priorities. Some education systems emphasize technological competitiveness and STEM education above all else. Others place greater emphasis on digital citizenship, creativity, or the ethics of artificial intelligence. An international comparison, however, reveals a common trend: almost everywhere, computer science is gradually moving beyond its marginal or purely technical role to become an integral part of general education.

Italy is arriving at this transition somewhat late. For many years, our school system has approached digital technology primarily as technical literacy or as an administrative tool. Today, however, the transformation of algorithmic society makes a much more profound rethinking inevitable. Looking at international experiences can therefore help us understand not only what works but also the mistakes to avoid.

2. Estonia and Education for the Digital Society

Among European countries, Estonia is probably the most emblematic example of the integration of digital transformation and the education system. After gaining independence from the Soviet Union, the country invested heavily in digital technology as a lever for economic, administrative, and cultural modernization. School

became one of the central tools of this strategy. Estonia realized very early on that the digital revolution is not simply about introducing computers into classrooms, but about building a new national culture founded on the ability to understand and navigate the digital world. Today, most Estonian public services are digitized. Digital identity, online government services, and electronic voting systems are part of citizens' daily lives. In this context, computer science education inevitably takes on a civic dimension. Estonian schools introduce elements of computational logic, educational robotics, and algorithmic thinking from the earliest grades. However, the Estonian model avoids reducing computer science to mere coding. Significant attention is also devoted to cybersecurity, data management, and digital citizenship. The relationship between schools and society is particularly interesting. Students grow up in an environment where digital technology is not seen as an artificial or supplementary element, but as an integral part of contemporary social organization. This creates a significant cultural difference compared to school systems where digital technology is still perceived as an extracurricular activity or a mere technological update.

3. Finland and Computational Thinking

Finland represents one of the most influential educational models internationally, and its approach to computer science fully reflects the pedagogical philosophy that characterizes the Finnish school system. In the Finnish national curriculum, computational thinking was introduced as a cross-curricular competency as early as elementary school. However, Finland has chosen a very different approach compared to more technically focused models. The main goal is not to teach specialized programming, but to develop cognitive skills related to problem-solving, modeling, and understanding processes. Computational thinking is integrated into various disciplines and is not confined to a separate space. Mathematics, science, art, and the humanities constantly interact with computational tools and digital environments. This choice reflects a key conviction: digital transformation is not just about technology, but about the very way knowledge is organized. Finland also places enormous importance on teacher training. The digital teacher is not viewed as a mere user of technological tools, but as a professional capable of critically integrating digital elements into complex educational processes. Particular emphasis is also placed on students' cognitive well-being. The Finnish school system avoids indiscriminate digitization and insists on a balanced relationship between technology, learning, human relationships, and psychological development. The Finnish model thus demonstrates that computer science and the humanities are not necessarily in conflict. On the contrary, the computational paradigm can be integrated into an educational vision centered on the holistic development of the individual.

4. The United Kingdom and Computer Science as a Scientific Discipline

The United Kingdom represents one of the most radical cases of transformation in European school computer science. In 2014, the British government replaced the previous ICT curriculum with a new subject called "Computing." This decision marked a very significant cultural shift. The old approach had been criticized for being too focused on the operational use of commercial software and lacking real scientific depth. Students learned to use digital tools without understanding the fundamental principles of computer science. The new curriculum, however, introduced a conception of computer science as an autonomous scientific discipline. Students engage with algorithms, logic, programming, data representation, and systems architecture as early as elementary school. The goal is to develop a conceptual understanding of how computational systems work. The British model places great importance on computer science as part of contemporary general knowledge. However, this transformation has also revealed some significant challenges, particularly regarding teacher training. Many teachers came from the old ICT model and lacked adequate preparation to teach computer science as a scientific discipline. The British case thus illustrates both the potential and the challenges of a radical reform of computer science education.

5. The United States and China: Two Opposing Models

Outside Europe, profoundly different models are emerging. In the United States, computer science education is strongly influenced by a culture of technological innovation and entrepreneurship. Major Silicon Valley companies have played an important role in the spread of computer science in schools. The American approach emphasizes creativity, startups, technology design, and economic competitiveness. Computer science is often viewed as a tool for preparing future generations of innovators. However, the U.S. system also exhibits significant regional and social inequalities. Educational opportunities often depend on the financial resources of individual school districts. China, on the other hand, represents a completely different model. There, digital education is viewed as part of a national geopolitical strategy. The Chinese government considers artificial intelligence, big data, and advanced technologies to be decisive tools for global competition in the 21st century. Computer science has been progressively introduced throughout the school system with a strong emphasis on AI, robotics, and programming. But the Chinese model also raises very sensitive questions about the relationship between technology, social control, and individual freedom. Digital technologies are, in fact, integrated into systems of surveillance and social management that differ greatly from the European democratic tradition. A comparison between the United States, China, and Europe thus shows that the issue of computer science education cannot be separated from the model of society one intends to build.

6. What Italy Can Learn

Italy cannot simply copy foreign models. The Italian school system has a unique cultural history, strongly tied to the humanistic tradition and the centrality of the disciplines. However, this very tradition could represent an important resource. The challenge for Italy does not lie in rapidly transforming schools into a replica of Silicon Valley or highly technocratic Asian systems. Rather, it lies in building a European and culturally grounded approach to computer science education. Italy can learn a great deal from international experiences. It can understand the importance of vertical curricula, teacher training, digital citizenship, and computer science as a cultural discipline. But it should avoid two opposing mistakes. On the one hand, the risk of superficial digitization reduced to the use of platforms. On the other, the risk of excessive technicality that transforms schools into mere centers for technological training. Instead, computer science should become a tool for strengthening critical thinking skills, an understanding of complexity, and democratic awareness. The real challenge for Italy is not simply to introduce more technology into schools. It is to build a culture of computer science capable of engaging with Europe's scientific, philosophical, and humanistic traditions. And it is precisely on this balance that the future of the new National Guidelines will depend.

Chapter 12 — The Italian Question

1. A School System Still Rooted in the Twentieth Century

The Italian school system is facing the digital revolution today while still carrying a cultural structure that is deeply rooted in the 20th century. This does not mean that the Italian school system is stagnant or incapable of innovation. In recent years, many schools have experimented with interesting programs related to coding, educational robotics, artificial intelligence, and digital citizenship. There are highly qualified teachers and extremely dynamic schools. However, the overall model of the Italian school system remains built around a relatively static conception of knowledge—one that is compartmentalized by subject area and heavily focused on the transmission of content. The digital revolution challenges precisely these assumptions. Contemporary society is characterized by continuous acceleration, the constant transformation of skills, cognitive automation, and a growing dependence on digital platforms. The Italian school system still struggles to fully engage with this new reality. Very often, digital technology is treated as a mere technical update, whereas the transformation underway concerns the very relationship between knowledge, learning, and social organization. In many cases, schools continue to use cultural tools designed for a world that predates the algorithmic revolution. The issue is not merely about introducing computer science as a subject. It concerns the education system's overall capacity to rethink its role within the digital society.

2. The Issue of Academic Disciplines

One of the most difficult issues that emerged during the debate on the new National Guidelines concerns the relationship between computer science and the disciplinary organization of the Italian school system. The Italian school system, in fact, has a strong disciplinary tradition. Mathematics, physics, literature, philosophy, history, and science constitute not only areas of instruction but also well-established cultural and professional identities. The introduction of computer science inevitably creates tension within this balance. Many fear that computer science will take away from traditional disciplines. Others, however, argue that it should not become a standalone subject but rather a cross-curricular competency integrated throughout the curriculum.

Both positions capture real aspects of the problem. Computer science certainly has a very strong interdisciplinary dimension. It naturally interacts with mathematics, physics, economics, biology, philosophy, and the social sciences. But it is precisely this cross-disciplinary nature that sometimes risks making it invisible or marginal. When a responsibility belongs to everyone, it often ends up belonging to no one at all. The Italian school system must therefore address a delicate cultural issue: how to integrate computer science without undermining existing disciplines and without reducing it to mere technical support. The problem cannot be solved solely through a redistribution of curriculum hours. Rather, we need to rethink the relationship between traditional disciplines and the computational paradigm. Computer science should not be perceived as a discipline competing with mathematics or the sciences. Instead, it should become one of the languages through which we interpret and connect contemporary complexity. This inevitably requires a cultural transformation far deeper than a simple administrative reform.

3. Computer Science or Digital Literacy?

In the Italian debate, there remains significant ambiguity between computer science and digital literacy. For many years, education policies have interpreted digital literacy primarily as a set of practical skills—such as using software, browsing the internet, creating presentations, or using collaborative platforms. This approach certainly had its historical utility. In the early stages of the digital revolution, it was necessary to ensure at least minimal access to technological tools.

Today, however, this is no longer enough. Artificial intelligence, algorithmic systems, and the centrality of data have made it clear that the issue is not merely using technologies, but understanding how they work and their cultural implications. A school system that teaches only how to use tools risks producing citizens who are technologically competent but culturally dependent. Computer science cannot, therefore, be reduced to mere education in the use of digital tools. It should introduce students to an understanding of computational models, algorithms, networks, data, intelligent systems, and the relationship between technology and power. This shift is fundamental. True contemporary literacy consists not only in using digital tools but in understanding the algorithmic society in which we live. And it is precisely here that the Italian school system still finds itself in a phase of incomplete transition.

4. Cultural Resistance

Many of the difficulties in introducing computer science do not stem solely from organizational or infrastructural problems. There are also deep-seated cultural resistances. A segment of the educational community still perceives digital technology as a threat to traditional culture. There is a fear that technology will impoverish language, reduce concentration, foster cognitive superficiality, and replace deeper forms of learning. These concerns are not entirely unfounded. The contemporary digital environment does indeed tend to prioritize speed, fragmented attention, and the rapid consumption of information. Even artificial intelligence can lead to cognitive dependence if used passively. But precisely for this reason, schools cannot simply reject digital transformation. Ignoring computer science and AI would mean leaving students completely exposed to the logic of digital platforms without the critical tools to understand them. The real choice is not between traditional schooling and technology. It is between a society that understands the digital world and one that passively endures it. There is also another, more subtle form of resistance. Many Italian teachers and intellectuals continue to view computer science as purely technical knowledge, lacking the cultural dignity traditionally attributed to the humanities or mathematics. This view seems increasingly untenable today. Computer science is not merely technical. It is one of the great cultural forms of our time. Failing to understand it means failing to understand a growing part of the modern world.

5. Why This Reform Is Crucial

The issue of computer science in the new National Guidelines is not simply one curricular reform among many. It concerns the cultural future of Italian schools. For the first time in many decades, the educational system is facing a transformation that simultaneously alters the relationship with knowledge, the role of the teacher, the very structure of knowledge, the meaning of learning, and even the relationship between humans and machines. Artificial intelligence makes this transformation even more radical. Schools can no longer limit themselves to conveying information that is easily accessible through automated systems. Instead, they must develop critical thinking skills, an understanding of complexity, independent judgment, and an awareness of the relationship between technology and democracy. For this reason, the inclusion of computer science in the National Guidelines represents a strategic choice. This is not simply a matter of teaching new technologies. It is about deciding what kind of citizen the Italian school system wants to shape in the 21st century. A society that does not understand how digital systems work inevitably risks losing its cultural autonomy and democratic capacity. What is truly at stake, therefore, is not merely the modernization of schools. It is the very possibility of shaping citizens capable of critically understanding the world in which they live. And it is precisely from this awareness that a new vision of computer science education should emerge: not as a technical add-on to the curriculum, but as an essential part of contemporary cultural education.

Conclusions

Not Training Users, but Citizens

The issue of computer science in the new National Guidelines is not simply about introducing a new school subject. It concerns the way in which the Italian school system intends to address one of the most profound cultural transformations of contemporary modernity. For many years, the digital realm was perceived primarily as a set of technical tools. Schools have sought to adapt by introducing computers, platforms, coding, and practical skills. All of this has had its historical utility, but today it appears insufficient. The algorithmic revolution has altered the very relationship between human beings, knowledge, and power. Artificial intelligence, recommendation systems, digital platforms, and the centrality of data are transforming cognitive work, the formation of public opinion, the organization of the economy, and even the functioning of democracy. In this context, the educational challenge can no longer be reduced to the use of technology. The real issue is understanding the cultural significance of the digital transformation. Understanding means questioning how algorithms work, the relationship between data and power, the limits of automation, the impact of artificial intelligence, and the new forms of cognitive dependence produced by digital ecosystems. The 21st-century school cannot limit itself to training efficient users of technological platforms designed by others. It must educate citizens capable of critically understanding the world in which they live. And it is precisely here that computer science takes on a new role. It is no longer merely a technical or professional discipline. It becomes a form of critical literacy for the contemporary world, one of the fundamental languages through which to interpret algorithmic society. This inevitably requires a transformation of the school system. The digital revolution does not merely alter the content of instruction. It changes the role of the teacher, the meaning of learning, the relationship with memory, the structure of knowledge, and the very organization of educational environments. Artificial intelligence makes all of this even more evident. If a machine can produce texts, summaries, images, and complex solutions, then the specifically human value of education can no longer be equated with the mere reproduction of information. Critical thinking, the ability to construct meaning, interdisciplinary understanding, independent judgment, and ethical responsibility become central. Schools must therefore increasingly become places for the critical interpretation of reality and not merely for the transmission of knowledge. This change also requires a new conception of the teacher. Teachers do not compete with machines in terms of speed or the quantity of available information. Their value lies in their ability to guide students through the complexity of the contemporary world, helping them distinguish between information and knowledge, between authority and manipulation, and between the passive use of technology and a critical understanding of systems. For this reason, teacher training likely represents the most important challenge for future educational policies. It will not be enough to teach

teachers how to use new digital platforms or tools. It will be necessary to build a new pedagogical culture capable of integrating computer science, artificial intelligence, and democratic citizenship. Europe, too, seems to have gradually come to understand this transformation. Documents such as DigComp and, above all, the AI Act demonstrate a growing awareness that the issue is not merely technological, but deeply political and cultural. A society that uses intelligent systems without understanding them inevitably risks losing democratic autonomy and cultural sovereignty. Schools therefore bear a historic responsibility. For the first time in many decades, the education system is facing a transformation that simultaneously reshapes knowledge, communication, cognitive work, and social organization.

The new National Guidelines therefore represent a decisive opportunity—not to chase technological trends or superficially introduce new practical skills, but to rethink the relationship between education and the digital society. The real challenge does not lie in having more computers in schools or more platforms in classrooms. It lies in shaping citizens capable of understanding the human, social, and ethical significance of the algorithmic revolution. Because the future of computer science education is not just about technology. It concerns the very future of cultural freedom in contemporary society.

Bibliography

Nardelli, Enrico (ed.), *Coding e oltre. L'informatica nella scuola*, Lisciani Scuola, Teramo, 2020.

Nardelli, Enrico, *La rivoluzione informatica. Conoscenza, consapevolezza e potere nella società digitale*, Themis, Rome, 2022.

Nardelli, Enrico, "Perché insegnare informatica nella scuola," in *Nuova Secondaria*, vol. XLIII, no. 3, 2025, pp. 69–74.

Vespri, Vincenzo, Papini, Paolo, Tortoriello, Francesco Saverio, *La scuola indivisa. Insegnare Dante con la Matematica (e viceversa)*, Edizioni Studium, 2026.

Italian Institutional Documents

Ministry of Education, University and Research, *Indicazioni nazionali per il curricolo della scuola dell'infanzia e del primo ciclo d'istruzione*, Ministerial Decree No. 254 of November 16, 2012, Rome, 2012.

Ministry of Education, University and Research, *Indicazioni nazionali e nuovi scenari*, document of the National Scientific Committee for the Curriculum, Rome, 2018.

Ministry of Education and Merit, *Linee guida per le discipline STEM*, Ministerial Decree No. 184 of September 15, 2023, Rome, 2023.

Ministry of Education and Merit, *Piano Nazionale Scuola Digitale (PNSD)*, Rome, 2015.

Ministry of Education and Merit, *Atto di indirizzo politico-istituzionale per l'anno 2025*, Rome, 2025.

Ministry of Education and Merit, *Indicazioni Nazionali per il curricolo – Scuola dell'infanzia e Scuole del Primo ciclo di istruzione*, Ministerial Decree No. 221 of December 9, 2025, published in the Official Gazette No. 21 of January 27, 2026, Rome, 2025.
Ministry of Education and Merit, *Le nuove Indicazioni Nazionali per i Licei*, draft for public consultation, Rome, 2026.

European and International Documents

European Commission, *DigComp 2.2. The Digital Competence Framework for Citizens*, Publications Office of the European Union, Luxembourg, 2022.
European Parliament and Council of the European Union, *Regulation (EU) 2024/1689 laying down harmonised rules on Artificial Intelligence (AI Act)*, Official Journal of the European Union, 2024.
European Commission, *Ethics Guidelines for Trustworthy AI*, High-Level Expert Group on Artificial Intelligence, Brussels, 2019.
UNESCO, *Guidance for Generative AI in Education and Research*, Paris, 2023.
OECD, *Artificial Intelligence and the Future of Skills*, OECD Publishing, Paris, 2021.

Computer Science, Computational Thinking, and the Digital Society

Wing, J. M., "Computational Thinking," *Communications of the ACM*, vol. 49, no. 3, 2006, pp. 33–35.
Papert, Seymour, *Mindstorms: Children, Computers, and Powerful Ideas*, Basic Books, New York, 1980.
Floridi, Luciano, *La quarta rivoluzione. Come l'infosfera sta trasformando il mondo*, Raffaello Cortina Editore, Milan, 2017.
Han, Byung-Chul, *Infocrazia. Le nostre vite manipolate dalla rete*, Einaudi, Turin, 2023.
Postman, Neil, *Technopoly. La resa della cultura alla tecnologia*, Bollati Boringhieri, Turin, 1993.
Morozov, Evgeny, *L'ingenuità della rete*, Codice Edizioni, Turin, 2011.
Harari, Yuval Noah, *Homo Deus. Breve storia del futuro*, Bompiani, Milan, 2017.

AI, Education, and Cognitive Transformation

Selwyn, Neil, *Should Robots Replace Teachers? AI and the Future of Education*, Polity Press, Cambridge, 2019.
Luckin, Rose, *Machine Learning and Human Intelligence*, UCL Institute of Education Press, London, 2018.
Williamson, Ben, *Big Data in Education: The Digital Future of Learning, Policy and Practice*, Sage, London, 2017.



Vincenzo Vespri

vincenzo.vespri@unifi.it

Vincenzo Vespri is a Full Professor of Mathematical Analysis at the University of Florence. His research focuses on mathematical analysis, partial differential equations, and their applications. In recent years, he has expanded his interests to include blockchain, cryptography, artificial intelligence, and data science, developing interdisciplinary applications of mathematics. He is the author of over 170 scientific publications and has held institutional positions dedicated to promoting STEAM disciplines and innovation in teaching. He is a corresponding member of the Turin Academy of Sciences and the Peloritana Academy.

The Fourth Writing

Conversation, Illusions, and the Future of Education

Paolo Ciancarini, University of Bologna

Abstract

Drawing on Clarisse Herrenschmidt's Les trois écritures, which interprets Western history through three major forms of writing—linguistic, monetary-arithmetic, and computer code—this article argues that generative language models reveal a fourth form of writing. In this framework, conversation becomes the general interface of computation: it is the machine that presents itself in the user's language, and not the other way around. Each form of writing produces its own illusion; the fourth produces the illusion of shared understanding, encapsulated by the concept of episteme—the comforting sensation of knowledge generated solely by the linguistic plausibility of a response. The article discusses the implications for schools, which must shift from teaching writing to teaching critical conversation with machines that write, and from evaluating the product alone to observing the process of judgment.

Keywords: fourth writing, generative AI, language models, epistemia, writing, school, critical literacy, assessment, Herrenschmidt

Clarisse Herrenschmidt is a French anthropologist and a member of the prestigious Collège de France. She belongs to that rare group of scholars who do not limit themselves to a single field but are able to connect different eras, multiple disciplines, and various forms of knowledge. A specialist in ancient Persia and Greece, a historian of writing, an archaeologist, a linguist, a semiologist, and an anthropologist, she has, with her book *Les trois écritures. Langue, nombre, code (Gallimard 2007),* a work that spans more than fifty-three centuries of the history of signs, from Sumerian clay tablets to computer codes. Her book demonstrates that writing is never merely a simple tool for recording information, but a powerful and symbolic technology that transforms the way human beings think, remember, measure, govern, and imagine the world.

The core of her thesis is that the West—and subsequently the globalized world—has been shaped by three major forms of writing. The first is linguistic writing, which allows us to communicate "in writing": poems,

stories, bills, and shopping lists; the second is monetary-arithmetic writing, which allows us to do calculations, record income and expenses, and solve a cubic equation; the third is computational and networked writing, which gave rise to the digital transformation. These forms of writing do not cancel each other out: they layer upon one another, incorporate one another, and mutually reinforce one another. Today, however, in light of large language models (such as ChatGPT), this thesis must be expanded.

Not to contradict Herrenschmidt, but to remain faithful to her. Generative models, in fact, reveal a **fourth** form **of writing**.

The first form of writing emerged when language became visible. The second emerged when numbers, through currency and then accounting, became a socio-technical environment of abstraction. The third emerged when computer code made the sign executable and transformed writing into an operational instruction for a machine, whether real or virtual. The fourth form of writing lies beyond—and at the same time within—all these previous forms. In fact, the generative model does not write like a scribe, does not measure like a coin, and does not execute like a program in the conventional way—instruction by instruction. A generative model generates text based on probability. It uses tokens, context, statistical correlations, and vector spaces to produce a linguistically plausible response. At this historical juncture, the thesis that **"conversation is the computer**" takes on its full meaning. The computer does not disappear: it changes its interface. It no longer asks only for commands; it asks for dialogue. In the traditional computing paradigm, the user had to learn the language of the machine—sequential and logical-mathematical. In the paradigm of models, it is the machine that presents itself in the user's language. This does not eliminate computation; it makes it conversational. Conversation thus becomes the general interface of the new computational environment.

This transformation is both technical and anthropological. Technical, because natural language becomes an operational layer capable of activating complex functions. Anthropological, because speaking with a machine alters our relationship with knowledge, with the authority of the text, and with the very act of writing. Writing no longer appears merely as the individual production of an author, but as the outcome of a dialogic relationship between human intention and automatic generation.

The Illusions of Writing

Every great form of writing has generated its own illusion—powerful and, to a large extent, unconscious. These illusions are not accidental errors. They are structural effects of new systems of signs, and over time they have also shaped the way societies have conceived of learning and the transmission of knowledge.

The **first form of writing**—that of languages—produced the illusion of presence. The text seemed to contain the voice of the author, the teacher, and sometimes even a god. To read meant to encounter an authoritative word that was already

complete. From this stems a long scholastic tradition based on commentary, gloss, and listening to the text as a living and normative presence.

The **second** form of **writing**—monetary and arithmetic—generated the illusion of exact measurement. Numbers gave the impression of being able to objectively capture the value of things, performance, and even people. In modern times, this illusion has fostered a culture of quantification that has also entered the school system: grades, class rankings, rankings, indicators, merit metrics, and benchmark averages.

The **third form** of **writing**—that of code—has produced the illusion of total control. All it took was writing the right instruction for the machine to execute it. In education, this logic has often fueled a procedural view of learning: input, processing, output; objective, exercise, assessment; competence, measurement, performance. It has been a powerful simplification, but also a reduction. Computer scientists say: garbage in, garbage out.

The **fourth form** of writing—generative writing—is producing an even more subtle illusion: the illusion of shared understanding. The model speaks so fluently, contextually, and convincingly that we are led to believe not only that it knows, but that it understands. This is where the concept of **episteme** comes into play, now used to describe the comforting illusion of knowledge produced by the linguistic plausibility of LLMs. This is not simply a matter of error or misinformation: it is the tendency to mistake a well-constructed form for genuine understanding on the part of the artificial interlocutor, who appears understanding and friendly.

It is important to distinguish between the terms *"episteme"* and *"epistemia*." *"Episteme"* is a word that belongs to the philosophical tradition of well-founded knowledge, justified knowledge, and truth that requires argumentation and proof. *"Epistemia*," on the other hand, is a recent neologism that refers to the cognitive short circuit whereby a well-written, elegant, and coherent response is perceived as reliable even when it is not. In the case of language models, this illusion is structural. The model does not understand in the human sense of the term, but it produces utterances that resemble those of someone who does understand. It simulates judgment, reflectiveness, an expert tone, and continuity of reasoning. From this stems one of the most important issues in contemporary digital culture: it is no longer enough to ask whether a text is grammatically correct or rhetorically persuasive; we must ask what relationship exists between its fluency and its truth.

Schools in the Fourth Writing

In schools, this transition is particularly delicate. For centuries, education has been based on the idea that writing was a human act of responsibility, memory, and thought. Writing meant processing, organizing, and laboring over language to clarify concepts and judgments. Today, a growing number of students can delegate the production of formally flawless texts—or at least texts convincing enough to pass many standard assessment thresholds—to generative models. The main risk is not merely that students "fail to learn how to write." The deeper risk is that they lose the experience of **thinking through writing**. If

writing is experienced solely as an output to be submitted, rather than as a process of conceptual clarification, then schools risk losing one of their most important historical functions: transforming language into a space for the formation of judgment.

This issue is already at the center of many educational reflections on generative AI. Several studies emphasize that AI-assisted writing can undermine the practice of cognitive effort, revision, independent argumentation, and critical literacy if it is adopted as a replacement rather than as a support. In other words, the fourth writing does not merely impose new tools; it demands a redefinition of educational work.

For this reason, the school's task can no longer be merely to teach writing in the traditional sense. It must teach students to engage responsibly with "fourth writing." This implies at least four new competencies, which should become standard components of contemporary literacy.

- **The rhetoric of the prompt**—not as a technical trick, but as the art of asking well-formulated questions, contextualization, and argumentative intent.
- **Epistemic verification**—that is, verifying sources, comparing answers, and the ability to distinguish between textual plausibility and epistemic reliability.
- **Text curation**, understood as critical review, editing, and taking responsibility for the generated output.
- **Distinguishing between fluency and comprehension**, to recognize when a text "sounds good" but doesn't actually convey meaning.

In this sense, schools must shift from teaching writing to teaching **critical dialogue** with writing machines. This does not mean abandoning the subject of writing. It means elevating it to a higher level of cultural responsibility.

This change also calls for a reevaluation of assessment practices. While the second writing introduced the temptation of blanket measurability into schools, the fourth writing makes the limitations of an assessment focused solely on the final product even more evident. A well-crafted text is no longer sufficient to demonstrate authentic learning.

We must therefore shift the focus from performance alone to the process: tracking drafts, versions, revisions, and choices; observing how a student questions a model, verifies an answer, corrects an error, or justifies an intervention. In this sense, the education of the future requires less emphasis on ready-made answers and more practice in exercising judgment. The decisive skill is not generating text, but knowing how to assess its validity.

The easiest mistake would be to divide the field between doomsayers and optimists. On one side are those who see generative AI as the end of school and writing. On the other are those who see it merely as a neutral accelerator of efficiency. Both positions miss the essential point: every new form of writing redefines the conditions of what it means to be human, and thus calls for new institutions of knowledge.

Herrenschmidt's lesson helps precisely here. Forms of writing are never merely technical tools; they are anthropological environments. The fourth form of writing, then, does not simply ask us to make good use of a digital product or service. It asks us to redefine what knowledge, authority, responsibility, and education mean in a world where conversation has become computational infrastructure.

Conclusions

Read today, *Les trois écritures* appears not as a concluded and definitive book, but as the best prologue for understanding the threshold we have crossed and the world we have entered. If the first writing made language visible, the second made value visible, and the third made the executable command visible, the fourth makes the generation of meaning visible in the form of dialogue.

The great educational challenge of our time is not to prevent students from using generative models, but to prevent them from becoming passive tools of a fourth writing whose constitutive illusion they do not understand. For this reason, it is no longer enough today to teach writing and arithmetic, with or without a machine. We must teach students to recognize the illusions of writing, to govern epistemia, and to critically navigate the new conversational environment of computation.

Key References

1. Clarisse Herrenschmidt, *Les trois écritures. Langue, nombre, code*, Gallimard, 2007.
2. *"Epistemia" in the Vocabolario Treccani and the related popular article.* https://www.treccani.it/vocabolario/neo-epistemia_(Neologismi)/



Paolo Ciancarini

paolo.ciancarini@gmail.com

Paolo Ciancarini is a full professor of Computer Science at the University of Bologna. His research focuses on software engineering, distributed systems, and software architectures. Throughout his career, he has held numerous academic and institutional positions, including the directorship of the National Interuniversity Consortium for Computer Science (CINI), contributing to the development of computer science research in Italy.

Toward a Harmonious Relationship Between Education and Technology

Christian Fiori, AICA

Abstract

This article argues that the school system's top priority for investment should not be technology-based instruction, but rather the cultivation of a critical and informed perspective on digital technologies. Building on Ivan Illich's distinction between convivial and non-convivial tools, the article traces the decline of the paradigm that assigned schools the task of chasing innovation, and documents its reversal between 2023 and 2026: from measures limiting minors' access to social media to declassified court documents, which reveal that young people's addiction was a deliberate part of the business model of major platforms. Drawing on the work of Marco Gui, Antonio Calvani, and David Buckingham, the article argues for the priority of media education over educational technology: it is not enough to know how to use the tools; one must understand the commercial, algorithmic, and representational logics that govern them. The theoretical framework is provided by Gert Biesta: while technology-based instruction falls almost entirely within the realm of skill development, media education addresses subjectivation—that is, the formation of a subject rather than the training of a user. In this light, the article analyzes the new National Guidelines that took effect in 2026 and the European guidelines of March 2026, outlining the conditions necessary to ensure they do not result in yet another wave of instrumental training.

Keywords: media education, educational technology, digital competencies, 2026 National Guidelines, subjectivation, conviviality, non-neutrality of platforms, digital well-being

A Paradigm Under Strain

When Ivan Illich published *Tools for Conviviality* in 1973, he proposed a distinction that has since permeated every serious discussion of technology. *Convivial* tools are those that anyone can master for their own purposes, capable of expanding individual autonomy and creativity without ever exceeding the limits of human

understanding. *Non-convivial tools*, on the other hand, are those that, behind a facade of efficiency, diminish these capacities, generate dependence, and concentrate control over the many in the hands of the few. Mass education, Illich warned, is always tempted to slip into the latter category; and it is in light of this tension that the title of these pages calls for a "harmonious"—and not merely instrumental—relationship between education and technology. Fifty years later, his question—which tools truly deserve a place in a project that aims to shape individuals, rather than train users—is once again pressing, and the following reflection stems from it.

For at least twenty years, public discourse on schools and digital technology has revolved around a catchphrase, "innovation," and an implicit conviction: that young people, as early adopters of technology, point the way to the future—almost as natural vanguards of change—and that it falls to schools to *keep up with them* by adopting the most innovative tools as soon as possible. This framework falls almost point by point within Illich's category of non-convivial tools, and among the Italian scholars currently investigating this topic is Marco Gui, whose analyses fully inspire these pages. Already in *Il digitale a scuola* (2019), Marco Gui—a media sociologist at the University of Milan-Bicocca and coordinator of the "Benessere Digitale" research center—effectively demonstrated the limitations of this approach: "The race for the latest technology has fueled a sense of inadequacy within educational institutions vis-à-vis the outside world, without thereby strengthening students' ability to engage critically with the digital realm."

The tide, after all, has turned. Gui speaks of a veritable "backlash" that has overturned the debate in recent years: between 2023 and 2025, smartphones and social media shifted from a promise of empowerment to an educational problem, amid high-impact documentaries, parent petitions, ministerial circulars for elementary schools, and growing institutional restrictions. To cite just a few notable examples, consider the lawsuits filed in 2023 by forty-two state attorneys general in the U.S. District Court for the Northern District of California; the Australian ban on social media for those *under 16*, effective December 10, 2025, and already cited as a model by the rest of the world; or the announcement in which the British government declared its intention to adopt this framework and even go further, with a ban on social media for those *under 16* expected to take effect in the spring of 2027 and additional restrictions on *livestreaming*, communication with strangers, and infinite *scrolling*. This shift has thus gained alarming momentum and, in 2026, is verging on scandal. On January 25, 2026, the Tech Oversight Project published a report with the telling title: *"Unsealed Court Documents Show Teen Addiction Was Big Tech's 'Top Priority*.'" Internal documents unsealed in the MDL (Multi-District Litigation) proceedings before the federal courts of Northern California and North Dakota show that *Meta*, *Google*, *Snap,* and *TikTok* deliberately designed their platforms to induce addiction in minors, and that this "mass addiction" among young people was considered the very core of their business model.

An internal TikTok document explicitly admits that "compulsive use of the platform is rampant" and that younger users, who are "particularly sensitive to social reward reinforcement," have "limited ability to self-regulate." The report also reveals Big Tech's coordination with industry-funded third-party organizations (from the National PTA to the Family Online Safety Institute) to control the public discourse and shape the narrative surrounding minors' use of social media. In the weeks that followed,

further court documents have revealed Meta's explicit strategy to "hook teenagers" as a priority objective, aiming to reach the entire family through them.

The evidence is emerging from courtrooms, but the substance of it is that of a scientific investigation: engineered addiction, awareness of the harm, and systematic concealment. It is this body of research and legal implications that is now guiding the actions of lawmakers, from the European *Digital Services Act* to individual U.S. state regulations for the protection of minors online.

Educating About Media Before Educating with Media

Given these premises, Gui's fundamental thesis—which serves as the backbone of this book—is that the priority should not be to introduce technology into teaching, but rather to *first* cultivate a critical and informed relationship with technological tools. This distinction was clearly articulated by Antonio Calvani, a leading figure in Italian pedagogy, a longtime full professor at the University of Florence, and scientific director of *Form@re*: while *education technology* concerns teaching that uses digital tools as a support for instruction, *media education* concerns the development of a critical perspective on media, understood as cultural languages and artifacts. And the order in which these concepts are presented is significant.

Among the empirical evidence supporting this assertion is the EYES UP project presented by Gui's research group in February 2025, which conducted the first longitudinal study dedicated to *the age at which students first gain access* to digital devices. By tracking over 6,600 high school students in Lombardy, the study found that those who began using social media before the age of twelve achieved lower scores in Italian and mathematics compared to those who started around the age of fourteen, with a particularly marked impact on boys. Not only that: early exposure tends to widen socioeconomic gaps, because it is families with greater cultural resources who are better able to manage the timing and manner of access. The same research group has also shown, on another front, that what protects young people is not so much technical familiarity as a less naive attitude toward how platforms operate: it is not enough to "know how to use" the tools; one must develop a critical eye.

In other words, the distinction that David Buckingham—one of the most authoritative voices in the field internationally—makes with great clarity in his *Manifesto for Media Education* between *media literacy* reduced to functional training in the use of tools (online safety, computer literacy, *coding*) and *media education* in the full sense, capable of questioning the mechanisms, commercial logic, and forms of representation of the world conveyed by the media. For decades, Buckingham has warned against two mirror-image yet equally harmful narratives: that of the "digital native"—the idea that young people, simply by using media, spontaneously develop the critical skills they need—and the narrative that treats digital literacy as an individualistic solution to structural social problems, shifting collective responsibilities onto the individual. Critical media literacy, he argues, is not a spontaneous process resulting from mere immersion, but requires a coherent educational path structured around four conceptual pillars: media languages, their representation of reality, the conditions of production (who governs

the platforms, with what interests and within what power dynamics), and their audiences. This competency must be conceived as a "fundamental right for all students," to be guaranteed from the earliest years of schooling, and not as a belated option reserved only for certain fields of study.

The conclusion that follows, therefore, concerns not the methods or tools of the school, but *the focus* of its attention: its historic mission to enable students to critically interpret the world in which they live today depends first and foremost on the media through which that world is portrayed. From this perspective, digitization certainly constitutes an unprecedented revolution for schools, but in purely pedagogical terms, it must be viewed as a process of incremental improvement and enrichment. As Sonia Livingstone—who has spent years at the London School of Economics studying the relationship between young people, schools, and digital technology—points out, the educational value of technologies does not stem from their mere presence, but from their integration into a thoughtful pedagogical framework: technologies can thus become a valuable addition to the teacher's toolkit.

Qualifying or Educating Individuals? Biesta's Lesson

Gert Biesta, one of the most influential international voices in the philosophy of education, clarifies—on a theoretical level—why media education *must precede* technology-based instruction. He argues that education operates through three inseparable functions: *qualification* (providing knowledge and skills), *socialization* (integration into traditions and practices), and *subjectivation*—that is, the emergence of *the self* into the world as the subject of one's own life and "not as an object of forces or desires originating elsewhere." His sharpest criticism concerns what he calls *"learnification"*: the tendency to reduce every pedagogical issue to the neutral language of "learning," which strips education of its questions of meaning and purpose. Linked to this is his critique of the "age of measurement," obsessed with producing measurable outcomes at the expense of forming individuals capable of "standing in the world on their own." Biesta reminds us that what matters most is also what is least amenable to measurement.

The point is that *educational technology*, on its own, falls almost entirely within the realm of skill development: it teaches how to use tools and navigate digital environments; it is necessary but not sufficient. It is the most easily measurable type of objective, and precisely for this reason it has been able to monopolize the school agenda. Media education, understood in its strongest sense, however, touches on subjectivation: it concerns a young person's ability to resist the pitfalls of the attention economy, to discern, and not to let themselves be defined by a platform's visibility metrics. It is the formation of an autonomous subject, not the training of a user. In his essay *World-Centred Education*, Biesta builds on these premises to formulate a clear proposal: education should be "centered on the world," not on the student or the curriculum, and the task of teaching cannot be reduced to transmitting knowledge but must redirect students' attention toward the world, so that they may encounter what the world demands of them. To exist "in an adult way," he adds, means making the question "whether our desires are truly worthy of being desired" a

living question. In an age when technologies are designed to capture and fragment attention, defining teaching as "education in attention to the world" sounds almost like a full-fledged media education program: teaching students to pause, to look deeply, and to cultivate a long-term perspective, whereas the digital environment pushes toward quick, reactive responses.

Educating Through Non-Neutral Technologies?

Teachers are ethically called upon to make students aware of *the non-neutrality* of platforms, their commercial interests, and the hidden mechanisms of social engineering. Indeed, we must distinguish between media as mere transmission channels (fiber optics, radio waves)—where the purposes and uses are governed by the user—and media as complex environments. Take social networks, for example: the architecture of the environment does not merely facilitate interaction but shapes choices, beliefs, and behaviors according to logics invisible to those who are subject to them. This gives rise to the paradox highlighted by Sapienza University researcher Marco Fasoli: the more digital technologies evolve, the more sophisticated they become in stimulating the most ancestral circuits of our brain.

On a parallel track to Fasoli's, but with a more directly socio-economic perspective, David Buckingham argues that the major platforms that now mediate our daily lives should not be treated as neutral technological infrastructures, because "they are not simply technology companies: they are also media companies," which increasingly provide us with "the means of representation and communication indispensable to modern life." Drawing on Melvin Kranzberg's famous maxim, Buckingham emphasizes that technology "is neither good nor bad, nor even neutral," and this also applies to the algorithms that organize our searches, populate our feeds, and select the content we consume: far from being automatically "objective and true," they are the result of concrete choices about what data to collect, how to correlate it, and whom it should serve. It follows that schools cannot limit themselves to equipping students to "use the tools well"; they must cultivate in them the critical ability to recognize how these tools mediate "the relationship between individuals, knowledge, and power."

It is on this very ground of non-neutrality that international research is now driving media education, going beyond the mere—albeit fundamental—competence in the use of tools. The concepts of *datafication* and *platformization*, developed by Felicitas Macgilchrist and Andreas Weich, encourage a critical examination of the infrastructures (algorithms, business models, data exploitation) within which digital practices are embedded, rather than shifting the responsibility for structural problems onto the individual. The work of Hintz, Dencik, and Wahl-Jorgensen on digital citizenship in a datafied society points in the same direction. The stakes have risen even higher with generative artificial intelligence, which lies at the heart of the debate on guidelines for computer science and AI in schools: *critical GenAI literacy* requires not only knowing how to use *chatbots* but also understanding their democratic implications and their intertwining with power and market dynamics.

The 2026 Guidelines: An Opportunity Not to Be Wasted

This discussion comes at a decisive moment. In Italy, the new National Guidelines for preschool and elementary school, adopted by Ministerial Decree No. 221 of December 9, 2025, and published in the Official Gazette on January 27, 2026, entered into force **on February 11, 2026**, and will be fully implemented starting in the **2026–2027** school year. For the first time, they introduce computer science as early as elementary school and frame it within an approach that, in the Ministry's own words, goes beyond the purely instrumental use of technology to foster critical digital literacy, including an initial introduction to artificial intelligence. These same guidelines will take effect progressively starting with the 2026–2027 school year, followed shortly thereafter by guidelines for high schools (published in draft form in April 2026, with implementation scheduled for September 2027) and by guidelines for technical schools. At the European level, the convergence is clear: on **March 5**, **2026,** the European Commission published four guidance packages for teachers on the ethical use of AI and data in education, on combating disinformation, on selecting high-quality digital content, and **on teaching computer science**—emphasizing that computer science "is more than just coding" and must enable students to understand and shape technology. At the same time, as part of the Digital Education Action Plan and DigComp 2.2/3.0, the EU has consolidated the European framework for digital competencies, for which *AICA* has overseen certification in Italy for decades, facilitating its incorporation into key ministerial documents.

These are all developments that, on paper, support the priority advocated here: the cultivation of a critical perspective is recognized as an end in itself, not as an afterthought. But this is precisely where the crux of the matter lies: the very terms "digital competencies," "computer science," and "artificial intelligence" can be interpreted either in terms of purely measurable qualifications or, in the Biestaian tradition, in terms of subjectivation.

In this regard, Marco Gui, in his latest book *Un patto sociale per l'educazione digitale (A Social Pact for Digital Education)*, strongly emphasizes the need for an educational alliance capable of freeing the relationship with technology from the logic of passive adaptation and reorienting it toward a shared project of human, cultural, and civic growth. Gui proposes a "social pact" structured around three pillars: a gradual approach to digital autonomy, the development of critical thinking skills, and collective action among schools, families, and institutions (including pediatricians), thereby freeing digital education from the current confusion and the dangerous distortions of digital capitalism.

Therefore, to prevent the 2026 guidelines from being reduced to yet another wave of instrumental training, media education must become their backbone rather than merely a rhetorical framework: the text of the new documents points in the right direction, but it will be concrete practice that determines whether computer science in schools will produce citizens capable of critically engaging with the digital world or merely more efficient users.

Conclusions

The true pedagogical priority of contemporary schools does not lie in the simple integration of technologies into teaching, but rather in the development of a critical and informed media education. For years, the race to build digital infrastructure has absorbed vast economic and planning resources, without, however, producing systemic improvements. Critical media analysis—the only approach capable of transforming a passive user into an autonomous individual, aware of the hidden logic behind screens—has instead been relegated to the margins.

Sherry Turkle, a distinguished sociologist and psychologist at *MIT*, in *Reclaiming Conversation* (2015, republished in 2025 with a new chapter on generative AI), documents how connected devices have made us and our students "forever elsewhere"—constantly projected elsewhere—and how schools that succumb to their allure risk undermining the very fundamental skills (deep attention, listening, empathy, slow thinking) they are called upon to cultivate. Another authoritative perspective confirms the same order of priorities: first, we educate people *about* media; then—and only then, effectively—we educate people *through* media.

It is worth reiterating, however, that this awareness in no way implies a rejection of digitalization, nor of the opportunities it offers to the educational context. Being critical of technology, as Neil Selwyn—one of the leading contemporary scholars of educational technology—reminds us, does not mean being hostile to it, but rather engaging in constructive, realistic reflection that is open to imagining alternatives. Above all, it means shifting the question from *"What works?"* to *"What is the purpose of education?"*, which is the fundamental question that Biesta poses in opposition to the age of measurement, and which restores to media education its logical and moral priority over the teaching of tools.

It is no coincidence, then, that Illich's thought is currently experiencing a new wave of interest: from the pages of L.M. Sacasas's *The Convivial Society* to the anthology *Thinking With Ivan Illich* edited by Carl Mitcham (2023), all the way to *the Convivial Machine Learning* course at the 2026 Digital Humanities Summer Institute, the most recent scholars are calling for Illich's criteria of openness, accessibility, non-dependence, and user control to become the compass by which every new educational tool—from social media to generative artificial intelligence—is evaluated. A truly *convivial* relationship between education and technology, therefore, does not arise from chasing after youth culture nor from its sterile rejection, but rather from restoring to schools their highest mission: to redirect students' attention toward the world, and to enable them to ask themselves—even when facing a screen—whether what they desire is truly worthy of their desire.

Bibliography

Australian Government, *Online Safety Amendment (Social Media Minimum Age) Act 2024* (in force since December 10, 2025), Parliament of Australia, Canberra 2024.

Biesta G., *Il mondo al centro dell'educazione. Una visione per il presente*, tab edizioni, Rome 2023 (orig. ed. *World-Centred Education*, Routledge, 2022).

Biesta G., *The Rediscovery of Teaching*, Routledge, London–New York 2017.

Buckingham D., *The Media Education Manifesto*, Polity, Cambridge 2019 (Italian trans. *Un manifesto per la media education*, ed. by G. Cappello, Mondadori Università, Milan 2020).

Calvani A., *Manuale di tecnologie dell'educazione*, new expanded and updated ed., ETS, Pisa 2001.

Council of Europe, *Draft Recommendation on Artificial Intelligence Literacy* (adopted by the CDEDU, 9th plenary session, Strasbourg, March 2026).

European Commission, *Guidelines for teachers on the ethical use of AI and data; on digital literacy; on digital education content; on teaching informatics*, as part of the Digital Education Action Plan, 2025.

Decree of the Minister of Education and Merit No. 221 of December 9, 2025, *Indicazioni nazionali per il curricolo della scuola dell'infanzia e del primo ciclo d'istruzione* (Gazzetta Ufficiale, January 27, 2026).

Department for Science, Innovation and Technology (UK), "Social media to be banned for under-16s in landmark government move to give kids their childhood back," gov.uk, press release of June 15, 2026.

Farina M., Lavazza A. (eds.), *Philosophy, Expertise, and the Myth of Neutrality*, Routledge, London–New York 2024.

Fasoli M., "Il paradosso delle 'nuove' tecnologie e i suoi pericoli cognitivi," in *Atti del convegno AISC*, Naples 2013.

Fasoli M., Gui M., "Stress digitale. Gli effetti collaterali della sovrabbondanza comunicativa," in *Aggiornamenti Sociali*, March 2015, pp. 218–227.

Gui M., *Il digitale a scuola. Rivoluzione o abbaglio?*, Il Mulino, Bologna 2019.

Gui M., Gerosa T., "Strumenti per apprendere o oggetti di apprendimento?" in *Scuola Democratica*, Milan 2019.

Gui M., *Un patto sociale per l'educazione digitale. Attori, evidenze e prospettive dopo l'era degli entusiasmi*, Mondadori Università, Milan 2026.

Hintz A., Dencik L., Wahl-Jorgensen K., *Digital Citizenship in a Datafied Society*, Polity Press, Cambridge 2019.

Illich I., *Tools for Conviviality*, Harper & Row, New York 1973 (Italian trans. *La convivialità*, Mondadori, Milan 1974).

Livingstone S., Sefton-Green J., *The Class: Living and Learning in the Digital Age*, NYU Press, New York 2016.

Macgilchrist F., Weich A. (eds.), *Postdigital Participation in Education*, Springer, Cham 2023.

Mitcham C., “Tools for Conviviality: Argument, Insight, Influence,” in *Thinking With Ivan Illich*, Conspiratio, 2023.

Respi C., Gui M. et al., *EYES UP (EarlY Exposure to Screens and Unequal Performance)*: final project report, Milan 2025 (www.benesseredigitale.eu).

Sacasas L.M., *The Convivial Society*, newsletter since 2020 (theconvivialsociety.substack.com).

Selwyn N., *Education and Technology: Key Issues and Debates*, 3rd ed., Bloomsbury Academic, London 2022.

Tech Oversight Project, “Unsealed Court Documents Show Teen Addiction Was Big Tech’s ‘Top Priority’,” report of January 25, 2026 (techoversight.org).

Turkle S., *Reclaiming Conversation: The Power of Talk in a Digital Age*, Penguin Press, New York 2015; new ed. with a chapter on generative AI, 2025 (Italian trans. *La conversazione necessaria*, Einaudi, Turin 2016).



Christian Fiori

christian3.fiori@gmail.com

Christian Fiori is a digital educator and researcher in the field of critical studies on emerging technologies, specializing in the ethical and social implications of their development and use, as well as their impacts on the education and health of young people. In addition to his technical training in Electronic and Telecommunications Engineering at the University of Bologna and his experience in quantum computing with CINECA, he studied the Philosophy of Information at the University of Urbino, earned a Master’s degree in the Philosophy of Digital Technology and Artificial Intelligence at the University of Udine, and obtained a diploma in Advanced Specialization in Ethics and AI at the Pontifical Antonianum University in Rome. He is currently *a digital educator and ethicist* at AICA, serves on the editorial board of AICA’s two journals, BRICKS and MONDO DIGITALE, and is a member of the Organizing Committee for the DIDAMATICA event.

A Few Sentences, One Discipline

Computer Science in the New National Guidelines for the First Cycle of Education

Michael Lodi, University of Bologna;
Violetta Lonati, University of Milan

Abstract

The new National Guidelines for preschool and the first cycle of education, published in the Official Gazette in January 2026, introduce for the first time competencies, learning objectives, and knowledge in computer science, integrated into the subjects of mathematics and technology. We present an analysis of the passages dedicated to computer science, organized around three fundamental scientific core areas—data and information, algorithms, and programming—and the two domains of digital awareness and creativity. Finally, we discuss the conditions necessary for the curriculum to be implemented in the classroom: in-service teacher training, reform of university programs for initial teacher education, high-quality materials, and, looking ahead, the recognition of computer science as an independent discipline.

Keywords: computer science; National Guidelines; first cycle of education; computer science education; teacher training.

1 Computer Science in the Curriculum

With the ministerial decree of December 9, 2025, published in the Official Gazette on January 27, 2026, the new National Guidelines for preschool and the first cycle of education[1] introduce, for the first time, competencies, learning objectives, and knowledge[2] explicitly related to computer science (MIM – Ministry of Education and Merit 2025). Computer science elements are integrated into

[1] Which includes elementary school and lower secondary school.

[2] Knowledge that, unlike competencies and objectives, is not prescriptive but helps to better understand the reasoning of the Guidelines' authors.

mathematics and technology, and will be introduced in classrooms starting in the 2026–2027 school year, beginning with the first grades of elementary and lower secondary school. This is a development that risks going unnoticed, as it is summarized in just a few very concise sentences: no new school subject is being created, yet those sentences contain a precise plan, and it is worth reading them carefully.

Why teach computer science in school, and starting in elementary school? Automated data processing is the invisible engine behind apps, platforms, markets, and public services: anyone who lacks the concepts to understand how data represents reality—what it means for a decision to be delegated to an algorithm, or how a program does what it does—views this world as a mere spectator. While the tools may change, those concepts remain constant, and they can be taught from an early age, just like the fundamentals of mathematics or science: the much-discussed "computational thinking" (Wing 2006) cannot develop without the study and practice of computer science (Lodi et al. 2017; Lodi and Martini 2021; Lodi 2020).

In this article, we offer an analytical reading of the passages dedicated to computer science in the new Guidelines: where they come from, what they say, and, above all, what kind of computer science lies behind them.

2 A New Development with Deep Roots

The inclusion of computer science in the National Curriculum Guidelines comes after decades of a troubled relationship between computer science and the Italian school system, which can be summarized by identifying several "historical" phases: the 1980s, with experiments using LOGO and the interesting proposals put forward by Giovanni Prodi (Prodi 1990; Barozzi 2010), which later merged into the National Computer Science Plan (Barozzi and Ciarrapico 2003), focused on the constructive aspects of mathematics and education in algorithmic thinking; the long era of the "computer license," during which studying computer science meant using software applications; the years of "coding" and "computational thinking," which were incorporated into the curriculum without a clear connection to the discipline (Lodi and Martini 2024). The new Guidelines finally adopt the term *"computer science"* and refer to the European Union Council Recommendation of November 23, 2023 (Council of the European Union 2023), which calls for the introduction of computer science education from the start of compulsory education, as confirmed by subsequent guidelines (European Commission and Directorate-General for Education, Youth, Sport and Culture 2026).

Despite the prevailing institutional orientation in the past, however, there have always been initiatives aimed at promoting a different vision of computer science in schools—one understood as a scientific discipline with its own concepts, methods, and cultural foundations, not reducible to the mere use of technologies or the acquisition of programming skills (Nardelli 2020). In December 2017, CINI (the National Interuniversity Consortium for Computer Science), together with university professors of computer science (GRIN) and computer engineering (GII), published a "Proposal for National Guidelines for the Teaching of Computer Science in Schools" (CINI, Laboratorio Nazionale Informatica e Scuola 2017), which we helped draft. The proposal

drew on international experiences, particularly the British *Computing* curriculum—mandatory in elementary school since 2014 (Department for Education 2013)—and adapted its spirit to the format of Italian curriculum documents, with goals and objectives organized into five areas: **data and information, algorithms, programming, digital awareness, and digital creativity**.

The draft of the new Guidelines, released in March 2025 for public debate [3], reproduced the objectives and learning outcomes of the CINI proposal almost verbatim; the version published in the final text reworks that content, making it more accessible to non-computer scientists, but also consolidating and, in part, reducing it. The result is short sentences that condense—sometimes into a single line—complex conceptual frameworks: let's take a closer look at them, following the five areas of the CINI proposal, with the first three focused primarily on mathematics and the last two on technology.

3 Digital Competencies or Computer Skills?

Before addressing the objectives, it is worth starting with the cultural premises, which state that "there is a difference between 'digital competencies' (understood as mastery of the effective, safe, and informed use of digital devices, tools, and technologies) and 'computer skills,' the latter referring to an understanding of how data (represented digitally, that is, in numerical/symbolic form) can be processed automatically by 'agents' that mechanically execute procedures devised by humans" (MIM – Ministry of Education and Merit 2025, 25).

This is a crucial distinction. It establishes, in a national curriculum document, that knowing how to use technologies and understanding computer science are two different things: confusing them has been the prevailing misconception of the last thirty years, much as if biology were the science of microscopes. The novelty lies entirely in automatic processing, and the mathematical premise reaffirms this: "What sets computer science apart from the digital representation of data is that it allows data to be processed in a fully automated manner through the use of a (computer) device, which acts as a mere mechanical executor of a computational procedure conceived and designed by humans" (MIM – Ministry of Education and Merit 2025, 120). Humankind has been designing algorithms for centuries, and the word itself derives from the name of al-Khwārizmī, a 9th-century mathematician; what computer science adds is that these procedures are executed by agents lacking intelligence, common sense, and free will (Nardelli 2022, chap. 3). This changes the way these procedures must be conceived and described, as well as the effects their execution—which can be replicated on a massive scale—has on society.

The goal of computer science in the first cycle is for students to acquire the cultural perspective of the discipline, achieved by "modeling—using data that is appropriately organized and represented—the information relevant to the objectives to be achieved; and by designing algorithms, that is, precise and non-

[3] https://www.mim.gov.it/documents/20182/0/Nuove+indicazioni+2025.pdf/cebce5de-1e1d-12de-8252-79758c00a50b?version=1.0&t=1741684578272

ambiguous (in reference to a specific executor) of procedures (to achieve these objectives) that lend themselves to automation; using artificial, directly executable languages to express algorithms as computer programs, so that they can be executed automatically by "mechanical agents"» (MIM – Ministry of Education and Merit 2025, 120–21). In a single sentence, three major ideas that structure the scientific fields: information is expressed through data, digitally encoded; algorithms operate on data to perform tasks and solve problems; programs express algorithms in a form that a machine can execute (Bell et al. 2018).

4 Data and Information

By the end of elementary school, students should be able to "express information through various types of data and encode such data, including digitally"; as early as third grade, "select and use objects or symbols to represent information"; and by fifth grade, "select combinations of symbols to represent structured data." Among the required knowledge is the "representation of data and information and the difference between them."

That distinction, which may sound new to many, is a profound concept (Lonati and Monga 2025) and can be taught very early on. A piece of data, on its own, means nothing: "10" can be an age, a price, a day, a speed (see Figure 1). It becomes meaningful when the recipient knows how to interpret it within a context.

Fig. 1. A piece of data, on its own, says nothing: "10" can be an age, a price, a speed, etc.

Conversely, the transition from information to data always involves a choice. Which piece of data we choose to convey a piece of information depends on the purpose and context: to identify one of the two polygons on the left side of Figure 2, I can use color (or the number of sides, the position above/below, …); but when there are four polygons, color alone is not enough (nor is the number of sides, nor the position above/below).

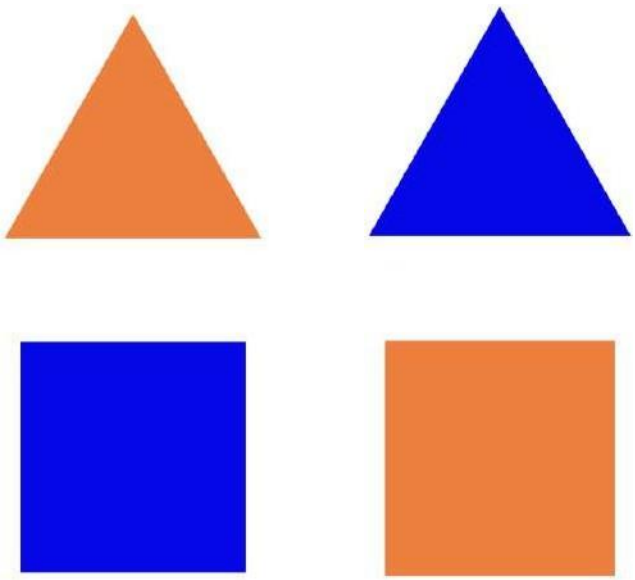


Fig. 2. Four polygons in which a single piece of data—such as color, shape, or position—is not enough to uniquely identify an object.

Choosing data means modeling: every model retains certain aspects of reality and excludes others, depending on the purpose. It is no coincidence that the competency standard for secondary education calls for "representing information in relation to the task at hand": it is the task that determines which data are needed, which are superfluous, and when two representations are interchangeable.

A further step is encoding: any data "of various kinds" (numbers, but also text, images, and sounds) can be represented as a sequence of symbols chosen from a finite alphabet—even as few as two symbols—provided that the rules of the system allow the data to be reconstructed without loss or ambiguity.

For this reason, the learning objectives emphasize making the interpretation of symbols explicit—that is, the legend. Much of what is needed, however, can be introduced by drawing on contexts already familiar to students in the classroom and reinterpreting them through a computational lens: map legends, writing systems, and playing cards covered with numbers and symbols (Boldi et al. 2025).

It is here that the "binary system," listed among the topics covered in secondary school, finds its computational significance: as a paradigmatic example of a coding system, to be explored through discovery (what will this sequence encode? by what rules?), and not as yet another conversion procedure to be performed mechanically. The ability to convert data of various types into digital representations is one of the reasons for the extraordinary versatility of computing devices: numbers, text, images, sounds, videos, and even programs can be stored, copied, transmitted, and manipulated by the same type of machine because they are converted into a common form (Bell et al. 2021). This is why a computer is not a machine specialized for a single task—unlike many traditional analog devices—but a universal processor capable of manipulating symbols through programs that perform specific tasks that vary each time.

5 Algorithms

By the end of elementary school, students are expected to be able to "describe procedures for performing practical tasks using algorithms." Specifically, by the end of third grade, the objective is to "describe everyday activities in words using sequences of precise and unambiguous steps and know how to carry them out"; among the required knowledge, the

"the concept of an algorithm and its rigorous execution."

School, and mathematics in particular, has always been full of algorithms: columnar addition, calculating the greatest common divisor, and constructions using a straightedge and compass. The computer science perspective introduces two shifts in meaning. The first concerns unambiguity, which is not a generic call for precision: a description, according to the National Guidelines, is unambiguous only in reference to a specific *performer*, whose basic operations—that is, their "catalog of actions"—must be explicitly stated. This is the weak point of the recipe metaphor, which is ubiquitous in textbooks (Monga et al. 2025): "wait a few seconds and try tasting it again" works for a cook with good judgment, whereas an automatic performer cannot decide what "a few" means. Describing the preparation of a sandwich in writing and seeing it

carried out to the letter by someone who “acts like an automaton” shows children exactly this: the designer and the executor are separate; the execution will take place later, with no opportunity for clarification; and any ambiguity will take its toll (Bissoli et al. 2026; Lodi et al. 2026). Even the concluding part of the third-grade objective—that “knowing how to carry them out”—thus takes on greater depth: interpreting the role of the executor, carrying out to the letter—physically or mentally—the steps described by others, is the most direct way to discover where ambiguities lie.

The second shift concerns the relationship between algorithms and problems. In school settings, an algorithm is often a tool to be learned and applied, so much so that in mathematics education, “algorithmic learning” refers precisely to the ability to carry out problem-solving methods (Crippa, Branchetti, et al. 2025). In computer science, the relationship is reversed: the algorithm is the end goal, the solution to a general problem, where a “general problem” refers to an entire class of cases that share the same structure but differ in their data (Crippa, Lonati, et al. 2025). Three questions asking whether three different chocolate bars cost more or less than the average price per kilogram are three instances of the same problem: they can be solved one by one, or a procedure can be designed once and for all to solve them all, and then its execution can be delegated. Conceptualization often involves breaking things down: identifying simple subtasks and combining them—as when, for example, to sort a collection of objects, one starts each time by figuring out how to find the smallest one. Designing, analyzing, and comparing algorithms, and understanding why a method works rather than merely applying it, are creative and reflective activities—the opposite of mechanical execution. This is the direction of the secondary-level objective: “to examine an algorithm or a program to understand its behavior, identify any defects, and correct them (debug).” The algorithm becomes an object to be studied, tested, and improved.

6 Programming

By the end of elementary school, students should be able to “write and understand simple programs, expressed in basic programming languages for educational purposes, and evaluate their suitability for the task to be automated”; as early as third grade, “write simple programs and verify, by running them, whether they perform the intended task and correct them if necessary”; in fifth grade, instructions “for reacting to events, selection (with elementary conditions), and repetition” are introduced, along with the goal of “translating an algorithm into a program and examining its behavior, including for the purpose of correcting it.” In middle school, the following are added: the

“concept of a variable,” programs “structured into modular components,” “functions and procedures,” and the ability to evaluate a program’s outcome “by following its steps and keeping track of the values of the variables.” Note that the Guidelines always refer to “programming”: never to “coding,” a term that has expanded over the years to encompass activities lacking in computational content (Francalanci 2020; Lodi and Martini 2024).

To be effectively executable, an algorithm must be written in the form of a program (Lodi and Martini 2025a): its expression in an artificial language with precise rules, understandable to humans and, at

the same time, executable4 by a machine (Lonati et al. 2022). The languages designed "for educational purposes" cited in the text exist and are well-established, ranging from visual block-based programming (Resnick et al. 2009) to simple text-based languages such as LOGO (Papert 1980); environments in which one can observe the automatic executor performing actions from its "catalog" based on instructions given by the programmer. The constructs listed are the building blocks of every program: sequential actions, repetition ("repeat while/repeat until," "repeat N times"), selection upon the occurrence of a condition ("if... then..."), and reaction to an event ("when you click on…"). Repetition and selection also apply to entire blocks of instructions, such as when drawing a square by repeating the pair "draw a side, turn 90 degrees" four times.

Two details deserve attention. The first: students are asked to write and understand, because reading someone else's program, predicting its behavior, and modifying it to see what happens is even more powerful practice than writing from scratch—especially today, when programs can be written with the help of artificial intelligence (Lodi and Martini 2025b; Izu et al. 2019). Second: correctness is defined not in absolute terms but in relation to "the task one wishes to automate," and verification through execution introduces an almost experimental practice consisting of predictions, tests, and revisions, in which the error—the bug—is part of ordinary work rather than a failure.

7 Awareness and Creativity: The Technological Aspect

The Guidelines entrust the remaining two areas to technology, with a clear premise: information technologies "must be presented to students as tools that function on the basis of the theoretical and methodological findings of computer science—that is, as a means and not as an end"—thereby avoiding "methodological errors based on digital competencies defined merely as the use of tools" (MIM – Ministry of Education and Merit 2025, 140–41).

Digital literacy, as outlined in the objectives, is grounded in understanding: "recognizing the presence of computers in daily life and distinguishing the Internet as a communication network from the services and content it provides" in elementary school, progressing to an understanding in secondary school of "the basic architecture (physical and functional) of a computer system" and the "fundamental mechanisms of the Internet." Distinguishing the network from the services that travel through it (for example, web content, messaging, streaming), knowing what "secret codes" are for protecting privacy, understanding "the value of personal and sensitive data," and the "social risks associated with the systematic collection of data": these are the conceptual foundations without which digital citizenship education is reduced to a list of prohibitions. Secondary school curricula also include

"critical evaluation of data found online and of information provided by generative artificial intelligence systems," with a fundamental principle that the text clearly states: "that it is human beings who

[4] After a multi-layered translation process.

maintain control over decisions based on computer systems that may have a significant impact on people" (MIM – Ministry of Education and Merit 2025, 141).

Finally, digital creativity shifts students from consumption to production: "developing a positive attitude toward computer applications by recognizing their potential as tools for personal expression," "creating simple digital content and applications for expressive purposes using appropriate environments and modifying or combining what is already available." That "modifying or combining" legitimizes a valuable—and deeply computational—practice: learning by building upon existing artifacts, exploring them, reusing them, and recombining them. This is the logic of *creative computing* and *creative learning*, following in the tradition of Papert's constructionism (the creator of LOGO) and forming the basis of Scratch, where programming means designing interactive media for creative and personal expression (Papert 1980; Brennan and Resnick 2012; Brennan et al. 2014; Resnick 2017).

8 With Mathematics, but with Its Own Identity

Placing computer science within the realm of mathematics and technology is a pragmatic choice: with unchanged schedules and staffing levels, it was difficult to imagine anything else. The connection to mathematics is, however, real and should not be wasted: rigor, abstraction, *problem-solving*, and reasoning are common ground, and the algorithm is a true bridge between the two disciplines (Crippa, Branchetti, et al. 2025). The computer science perspective can enrich the mathematical one, moving from the solution of a single exercise toward the general problem and the automation of its solution; conversely, the mathematical sensibility for justification and proof resonates with the attention that computer science places on the correctness of algorithms and programs.

But for this encounter to be fruitful, the identities of each discipline must remain recognizable: research on interdisciplinarity teaches that coordination and mutual transformation require a preliminary phase in which each discipline clarifies its own specificity (Sbaraglia 2023; Lodi et al. 2025). Their coexistence within the curriculum is an opportunity that must be cultivated. Computer science can certainly provide conceptual tools to enhance the teaching and learning of mathematics; however, the objectives outlined in the Guidelines suggest a broader perspective: computer science emerges as the discipline that offers a specific interpretive lens on the now-ubiquitous phenomena in which automated information processing plays a role. This is the cultural dimension that the Guidelines themselves refer to in the section "Why We Study Mathematics" (MIM – Ministry of Education and Merit 2025, 120).

9 Turning These Sentences into School Practice

The crucial question remains: who will teach all of this? The Guidelines introduce computer science without dedicated class hours or specialized teachers, entrusting it to teachers who, for the most part, have never studied it. This is not their fault: until now, computer science has never been a school subject, and the degree program in Primary

Education Sciences (LM-85 bis) covers educational technologies but includes no instruction in computer science or its pedagogy (MIUR 2010). Furthermore, the rushed update of textbooks runs the risk of repeating the mistakes made in the texts of the past decade: misleading metaphors, "coding" as a catch-all term, and algorithms without an executor (Monga et al. 2025; Capecchi et al. 2025).

The experience in the United Kingdom is instructive. When Computing became mandatory in 2014, the greatest challenge was not writing the curriculum but implementing it: a few years later, a report by the Royal Society highlighted a school struggling precisely with teacher preparation (The Royal Society 2017), and the government responded in 2018 with a public investment of 84 million pounds in the National Center for Computing Education (NCCE), a national network to provide training and resources to computer science teachers in England (The Royal Society 2019). The lesson is that a curriculum only thrives when accompanied by a structured training plan—both in-service and, above all, at the initial training level: we need reform of the LM-85 bis law and certification programs that include computer science and its pedagogy; we need high-quality materials, guides, annotated examples, and communities of practice where teachers and researchers work together. Pilot projects—some designed by interdisciplinary research groups—are already getting underway, but a structured plan is needed: success cannot be left to the goodwill of individuals.

And, looking ahead, we need the step that these Guidelines prepare for without actually taking it: the recognition of computer science as an independent school subject, with its own class hours and teachers, throughout the entire compulsory education period, in line with the European Recommendation and a path already taken by many countries. For those who have been working on these issues for years, the few sentences included in the Guidelines represent an almost unexpected milestone; for the Italian school system, they mark a starting point. Behind each of them, as we have tried to show, lies an entire discipline: now we must create the conditions for it to truly enter the classroom.

Bibliographic References

Barozzi, Giulio Cesare. 2010. *Giovanni Prodi (1925–2010): Un Ricordo Personale*. https://matematica.unibocconi.eu/articoli/giovanni-prodi-1925-2010-un-ricordo-personale.

Barozzi, Giulio Cesare, and Lucia Ciarrapico. 2003. "Il Piano Nazionale Per l'informatica." *Bollettino Dell'Unione Matematica Italiana. La Matematica Nella Società e Nella Cultura*, Serie VIII, 6-A (3): 441–61.

Bell, Tim, Paul Tymann, and Amiram Yehudai. 2018. "The Big Ideas in Computer Science for K-12 Curricula." *Bulletin of EATCS* 124: 36–46. http://bulletin.eatcs.org/index.php/beatcs/article/view/521.

Bell, Tim, Paul Tymann, and Amiram Yehudai. 2021. *The Big Ideas of K-12 Computer Science Education*. https://www.csse.canterbury.ac.nz/big-ideas/BigIdeas-webdocument-15-Feb-2021.pdf.

Bissoli, Giorgia, Agnese Del Zozzo, Michael Lodi, and Alberto Montresor. 2026. *Sulle orme di Milù – Guida di Informatica per l'insegnante*. Teacher's guide, Mondadori Education, Milan, ISBN 9791220419833.

Boldi, Arianna, Sara Capecchi, Violetta Lonati, and Mattia Monga. 2025. “Playcards, Pasta and Rectangles: Teaching Information Representation and Data Encoding in 4th Grade.” *Proceedings of the 30th ACM Conference on Innovation and Technology in Computer Science Education V. 1* (New York, NY, USA), ITiCSE 2025, 354–60. https://doi.org/10.1145/3724363.3729091.

Brennan, Karen, Christan Balch, and Michelle Chung. 2014. *Creative Computing Curriculum Guide*. https://scratched.gse.harvard.edu/guide/curriculum.html.

Brennan, Karen, and Mitchel Resnick. 2012. “New Frameworks for Studying and Assessing the Development of Computational Thinking.” *Proceedings of the 2012 Annual Meeting of the American Educational Research Association* (Vancouver, BC, Canada). https://scratched.gse.harvard.edu/ct/files/AERA2012.pdf.

Capecchi, Sara, Michael Lodi, Violetta Lonati, and Anna Morpurgo. 2025. “L’informatica negli attuali libri di testo della scuola secondaria di primo grado, alla luce delle Indicazioni Nazionali 2025.” In *La Didattica della Matematica con e per la scuola*, edited by Giorgio Bolondi, Bruno D’Amore, and Federica Ferretti. Bonomo Editore.

CINI, Laboratorio Nazionale Informatica e Scuola. 2017. *Proposta Di Indicazioni Nazionali Per l’insegnamento Dell’informatica Nella Scuola*. Consorzio Interuniversitario Nazionale per l’Informatica. https://www.consorzio-cini.it/images/Proposta-Indicazioni-Nazionali-Informatica-Scuola-numerata.pdf.

European Commission, and Directorate-General for Education, Youth, Sport and Culture. 2026. *Linee Guida Per l’insegnamento Dell’informatica: Strategie Pratiche Per Le Classi Europee*. Publications Office of the European Union. https://doi.org/10.2766/0709808.

Council of the European Union. 2023. *Raccomandazione del Consiglio del 23 novembre 2023 sul miglioramento dell’offerta di abilità e competenze digitali nell’istruzione e nella formazione*. https://eur-lex.europa.eu/legal-content/IT/TXT/?uri=OJ:C_202401030.

Crippa, Marta, Laura Branchetti, Luca Lamanna, and Violetta Lonati. 2025. “Conoscenza procedurale, algoritmi e problemi: spunti di riflessione da una prospettiva interdisciplinare tra matematica e informatica.” In *La Didattica della Matematica con e per la scuola*, edited by Giorgio Bolondi, Bruno D’Amore, and Federica Ferretti. Bonomo Editore.

Crippa, Marta, Violetta Lonati, Luca Lamanna, and Laura Branchetti. 2025. “Educare al pensiero algoritmico: riflessioni e proposte per il primo ciclo.” In *ITADINFO 2025. Terzo Convegno Italiano sulla Didattica dell’Informatica*, edited by Maria Angela Pellegrino and Veronica Rossano. Ledizioni.

Department for Education. 2013. *National Curriculum in England: Computing Programmes of Study*. https://www.gov.uk/government/publications/national-curriculum-in-england-computing-programmes-of-study.

Francalanci, Lucia. 2020. “Coding.” *Italiano Digitale* 12 (1): 95–101. https://doi.org/10.35948/2532-9006/2020.3311.

Izu, Cruz, Carsten Schulte, Ashish Aggarwal, et al. 2019. “Fostering Program Comprehension in Novice Programmers: Learning Activities and Learning Trajectories.” *Proceedings of the Working Group Reports on Innovation and Technology in Computer Science Education* (New York, NY, USA), ITiCSE-WGR ’19, 27–52. https://doi.org/10.1145/3344429.3372501.

Lodi, Michael. 2020. “Informatical Thinking.” *Olympiads in Informatics* 14: 113–32. https://doi.org/10.15388/ioi.2020.09.

Lodi, Michael, Agnese Del Zozzo, Alberto Montresor, and Giorgia Bissoli. 2026. *Insegnare Informatica - Una guida per docenti del primo ciclo*. Vol. 1.

Lodi, Michael, Violetta Lonati, and Simone Martini. 2025. “Informatica nelle Indicazioni di Matematica: un’opportunità interdisciplinare.” In *La Didattica della Matematica con e per la scuola*, edited by Giorgio Bolondi, Bruno D’Amore, and Federica Ferretti. Bonomo Editore.

Lodi, Michael, and Simone Martini. 2021. “Computational Thinking, Between Papert and Wing.” *Science & Education* 30 (4): 883–908. https://doi.org/10.1007/s11191-021-00202-5.

Lodi, Michael, and Simone Martini. 2024. “Coding, pensiero computazionale e informatica: quale evoluzione possibile?” *ITADINFO 2024. Secondo Convegno Italiano sulla Didattica dell’Informatica* (Genoa), 20–24.

Lodi, Michael, and Simone Martini. 2025a. “Algoritmi, programmi e linguaggi.” *Nuova Secondaria* 43 (3): 80–83.

Lodi, Michael, and Simone Martini. 2025b. “Le Notional machine per l’apprendimento della programmazione, anche nell’era dell’AI generativa.” In *ITADINFO 2025. Terzo Convegno Italiano sulla Didattica dell’Informatica*, edited by Maria Angela Pellegrino and Veronica Rossano. Ledizioni.

Lodi, Michael, Simone Martini, and Enrico Nardelli. 2017. “Abbiamo davvero bisogno del pensiero computazionale?” *Mondo Digitale*, no. 72: 1–15. https://mondodigitale.aicanet.net/2017-5/articoli/MD72_02_abbiamo_davvero_bisogno_del_pensiero_computazionale.pdf.

Lonati, Violetta, Claudio Mirolo, and Mattia Monga. 2022. “Di cosa parliamo quando parliamo di programmi?” *Mondo Digitale* 20 (97): 1–12.

Lonati, Violetta, and Mattia Monga. 2025. “Rendere concreta l’informazione per elaborarla.” *Nuova Secondaria* 43 (3): 74–79.

MIM – Ministry of Education and Merit. 2025. *Indicazioni nazionali per il curricolo della scuola dell’infanzia e del primo ciclo d’istruzione*. https://www.mim.gov.it/documents/20182/10554370/curricolo_web.pdf/.

MIUR. 2010. *Decreto 10 settembre 2010, n. 249 – Regolamento concernente la definizione della disciplina dei requisiti e delle modalità della formazione iniziale degli insegnanti*.

Monga, Mattia, Renzo Davoli, Luca Lamanna, and Luca Forlizzi. 2025. “Il concetto di algoritmo nei libri di testo: spunti per una riflessione.” In *La Didattica della Matematica con e per la scuola*, edited by Giorgio Bolondi, Bruno D’Amore, and Federica Ferretti. Bonomo Editore.

Nardelli, Enrico, ed. 2020. *Coding e Oltre: L’informatica Nella Scuola*. Lisciani Scuola.

Nardelli, Enrico. 2022. *La Rivoluzione Informatica: Conoscenza, Consapevolezza e Potere Nella Società Digitale*. Themis.

Papert, Seymour. 1980. *Mindstorms: Children, Computers, and Powerful Ideas*. Basic Books.

Prodi, Giovanni. 1990. "Quale informatica per la scuola media?" *L'Insegnamento Della Matematica e Delle Scienze Integrate* 13 (4): 355–58. https://www.centromorin.it/home/pubblicazioni/riviste/speciale.asp?cs=2.

Resnick, Mitchel. 2017. *Lifelong Kindergarten: Cultivating Creativity through Projects, Passion, Peers, and Play*. MIT Press. https://doi.org/10.7551/mitpress/11017.001.0001.

Resnick, Mitchel, John Maloney, Andrés Monroy-Hernández, et al. 2009. "Scratch: Programming for All." *Communications of the ACM* 52 (11): 60–67. https://doi.org/10.1145/1592761.1592779.

Sbaraglia, Marco. 2023. "Teaching Informatics to Novices: Big Ideas and the Necessity of Optimal Guidance." PhD thesis, Alma Mater Studiorum – University of Bologna. https://doi.org/10.48676/unibo/amsdottorato/10914.

The Royal Society. 2017. *After the Reboot: Computing Education in UK Schools*. Royal Society. https://royalsociety.org/topics-policy/projects/computing-education/.

The Royal Society. 2019. *Policy briefing on teachers of computing*. https://royalsociety.org/-/media/policy/Publications/2019/21-08-19-policy-briefing-on-teachers-of-computing.pdf.

Wing, Jeannette M. 2006. "Computational Thinking." *Communications of the ACM* 49 (3): 33–35. https://doi.org/10.1145/1118178.1118215.



Michael Lodi

michael.lodi@unibo.it

Michael Lodi is an adjunct professor and research fellow in the Department of Computer Science – Science and Engineering at Alma Mater Studiorum – University of Bologna and a member of the CINI Laboratory "Informatics and School." He studies the challenges and methodologies involved in introducing students to programming, the philosophical aspects

of computer science as a scientific discipline (e.g., computational thinking, big ideas), and the teaching of cryptography. He teaches students of computer science and other disciplines, as well as teachers in training. He was part of the editorial team for the CINI proposal on Guidelines for Teaching Computer Science in Schools. He conducts training courses for teachers and has led workshops for children and teenagers. He is a co-author of teaching materials and guides for computer science in elementary school.



Violetta Lonati

lonati@di.unimi.it

Violetta Lonati is a researcher in the Department of Computer Science at the University of Milan, where she is one of the founders of the ALaDDIn laboratory and the Bebras dell'Informatica project, a team-based computer science competition for schools. She is a member of the CINI "Computer Science and Schools" Laboratory and was part of the editorial team for the CINI proposal on Guidelines for Teaching Computer Science in Schools. She conducts research on computer science education and is involved in teacher training—topics on which she teaches courses and collaborates with the research community in mathematics education.

Calibra: An AI-Based Framework for Personalizing Learning Pathways

Nicola Illuzzi, Claudio Vergini, and Giorgio Tullio De Negri, Calibra Project

Abstract

Calibra arose from the need to move beyond the one-size-fits-all approach to education, proposing a model capable of analyzing a learner's profile and translating it into more targeted instructional choices. The relationship between learner characteristics and teaching methods, after all, has long been a central theme in research on the adaptation of educational pathways (Cronbach and Snow, 1977). Calibra begins with an assessment of life skills, soft skills, and learning styles; it constructs a "Basic Functioning Snapshot" and translates it into a "Preferred Learning Style," intended as a roadmap for accessing training. This paper describes the model's architecture, its white-box logic, and the limitations to be verified through future applied studies.

1. Introduction

In contemporary education, the idea still persists that making content available means, almost automatically, making it learnable; however, a course can be well-written, well-organized, even elegant, yet remain distant from those who are supposed to benefit from it: this happens in schools, at universities, in companies, and in vocational training. The same content, presented in the same way, never produces the same effect on everyone, because every learner has different prior knowledge, expectations, motivations, cognitive paces, emotional experiences, and ways of accessing information. The issue, therefore, is not merely the quality of the content, but the way in which that content connects (or fails to connect) with the person who is supposed to learn it.

The most widespread educational models often follow a specific logic: standardize, make the content replicable, and distribute it efficiently; this is an organizational necessity, not a flaw in itself. The problem arises when this standardization is mistaken for neutrality: certain teaching methods have long been criticized precisely because they risk treating the learner as a passive recipient—no different from any other user—rather than as an active participant, with their own background, in the construction of knowledge (Freire, 1970).

Even when the content and the teacher's expertise are the same, the learning experience can vary greatly. Some learners need concrete examples, others a guided content map, and still others images, simulations, discussion, or practical application. It is not always a matter of individual ability; rather, the compatibility between how the content is delivered and the person's learning style often comes into play. The medium, sequence, language, and format are not obvious choices; they shape how the message is received and alter the way it is understood (McLuhan, 1964).

This gap between available content and content that is actually learned becomes even more evident in digital, modular, or short courses. In microlearning, for example, the short duration does not in itself guarantee effectiveness: brevity can be useful, of course, but the content can also become a "pill" that is consumed quickly and forgotten just as quickly. The literature shows that what primarily influences training outcomes are the quality of design, the alignment between objectives and format, the potential for practical application, and the course's ability to support learning over time (Taylor and Hung, 2022). Personalization means creating conditions that are better aligned with the learner's profile in terms of content format, order of presentation, level of guidance, examples, feedback, interaction, and assessment. Some learners may need more guided steps and step-by-step explanations, while others may require greater autonomy, hands-on activities, or more concise materials: personalization does not change the learning objective, but rather the way in which the learner is guided to achieve it.

Artificial intelligence can make an important contribution, especially when it helps identify training needs, learning patterns, and engagement indicators. However, AI-based adaptive systems still suffer from well-known issues, such as decision-making mechanisms that lack transparency, personalization criteria that are not fully verifiable, and evaluations that often focus solely on satisfaction or completion (Kabudi et al., 2021). For this reason, an *AI-based* training model should not merely adapt content but should also explain the how and why of that adaptation. Another risk inherent in the use of AI concerns the rigid application of "learning styles": the idea that a person learns better simply because the content is aligned with a predefined category has been widely discussed and criticized (Pashler et al., 2008).

The *AI-based* training framework Calibra interprets individual differences not as fixed labels, but as functional and modifiable configurations; in other words, Calibra identifies learner types but is also able to recommend alternative conditions that might facilitate access to content at a certain stage, for example. Calibra personalizes learning paths through the integrated analysis of three broad areas: *life skills*, soft skills, and *learning styles*. It was not designed as a diagnostic tool for classifying users; rather,

it is a methodological support that seeks to understand how the learner organizes their experience, relates to the context, and accesses content, proposing solutions to make the relationship between the individual profile and instructional design more coherent.

The initial assessment allows for the creation of the *"Basic Functioning Snapshot*," a functional and dynamic snapshot of the learner within the learning context. From this, *the "Preferred Learning* Style" is then derived—understood not as a fixed modality, but as a guide to the most suitable conditions through which content can be made more accessible and applicable. The crucial step is the transition from profiling to design. As previously noted, Calibra aims not only to describe the learner but to translate that description into learning choices: content format, level of structuring, guidance, emphasis on practical experience, use of visual, narrative, analytical, or reflective elements, and individual or group approaches. This approach is consistent with the principles of cognitive load theory, according to which the effectiveness of learning also depends on how information is organized and presented in relation to available cognitive resources (Sweller et al., 2011).

A central element of the model is the *white-box* logic: in data- and AI-driven instruction, collecting information on learner behavior is useful only if that information is linked to clear, consistent, and traceable design decisions; *learning analytics*, in fact, should not be limited to recording what the user does, but should contribute to improving the learning environment on multiple fronts (Ferguson, 2012). Transparency is not a technical detail; it is a methodological requirement: a system that offers a personalized learning path without making its rationale understandable risks producing recommendations that are personalized only in appearance. Conversely, a *white-box* model allows for human oversight, the reconstruction of decisions, the correction of any distortions, and the progressive improvement of the quality of personalization (Guidotti et al., 2018).

The Calibra framework was developed by Live Connection Srl, in partnership with Top Spin Consulting, as part of applied research and methodological design activities related to the People Method; it thus emerged within a broader process of experimentation and formalization of training models that place the learner at the center, are oriented toward personalization, and translate assessment data into operational instructional decisions. Calibra represents a specific evolution of People: it does not merely organize training content or modules, but proposes a method for interpreting the learner's functional profile and transforming it into traceable design guidelines.

This paper represents the first public formalization of Calibra, developed based on the experience gained by professionals at Live Connection Srl and Top Spin Consulting in the design and implementation of professional and university training programs supported by digital technologies; however, the framework has not yet been the subject of previous scientific publications. For this reason, the article does not intend to present Calibra as a definitively validated system, but rather as a methodological proposal under development, designed to integrate instructional design, artificial intelligence, and human oversight within an explainable framework.

The following section describes Calibra as a framework for *AI-based* personalization of training. We will analyze how the assessment works, the role of the three macro-areas—which define the type of information to be collected on users—the transition from *the "Basic Functioning Snapshot"* to *the "Preferred Learning* Style" (the two key moments in Calibra's operation)—and the function of the PLS as an indicator of "educational accessibility" in instructional design. The following sections focus on the traceability of the Calibra system's decisions, on the methods of training personalization—which place the learner at the center and distinguish Calibra from other models—and on the cross-curricular nature of the approach (through two very different examples), as well as on the role of AI, concluding with the framework's original features, its limitations, and potential future developments of the work.

The goal is not to present a definitively validated model, but to propose an evolving system designed to bridge digital resources and learner-centeredness.

2. Method: Criteria for Constructing the Model

The approach of this paper is essentially descriptive; it does not present an experimental validation of the model or a psychometric validation of the assessment. Instead, it proposes the formalization of an *AI-based* framework for the personalization of learning pathways: before measuring the effectiveness of an adaptive system, it is necessary to clarify what it adapts, what information it processes, and what logic it uses to translate the learner's profile into instructional choices. Calibra is therefore described as an evolving application model, based on assessment, adaptive instructional design, decision traceability, and human oversight; within the debate on artificial intelligence in education, this approach does not reduce AI to mere automation technology but seeks to make it a tool capable of rethinking the relationship between the instructor, data, instructional design, the learner, and learning (Azevedo et al., 2024).

Furthermore, the literature on adaptive systems emphasizes the need to make personalization mechanisms explicit, because an adaptive technology that cannot be explained raises methodological concerns (Xie et al., 2019). Traceability is therefore part of the model; every step must be reconstructible: the *white-box* approach addresses this need—that is, preventing personalization from becoming an opaque output, perhaps elegant, but not open to discussion. In AI-based systems, explainability has not only technical value but also educational value, as it allows for the correction of distortions, maintains human control, and makes learners more aware of their own learning styles (Adadi and Berrada, 2018).

The model's development stems from the integration of contributions from pedagogy, instructional design, adaptive learning, artificial intelligence applied to education, and affective neuroscience. The sources were selected based on three criteria: relevance to the concept of personalized learning and usefulness in supporting a non-diagnostic interpretation of the user's profile.

In terms of accessible design, Calibra aligns with Universal Design for Learning, particularly in its focus on diverse modes of engagement and content presentation (CAST,

2018). This latter function does not involve offering endless variations of the same material, but rather making the content accessible through different pathways, without losing its value.

In constructing the macro-areas, the model uses *life* skills as a general educational framework but reinterprets them from a functional perspective. The focus is not on measuring "how much" a person possesses a particular competency; rather, it is on understanding how thought, emotion, and responsiveness contribute to the way a person approaches the learning experience (World Health Organization, 1994). Research on the relationship between emotion and cognition shows that these dimensions work together in the selection and organization of information (Pessoa, 2008). Soft skills, on the other hand, are considered in terms of their relational nuances: not as static qualities, but as dispositions that influence how the learner relates to the task, the group, the instructor, and the context (Robles, 2012). The third area concerns the various modes of accessing and processing experience. Here, the model engages with experiential learning and learning channels, but without transforming them into rigid categories: the use of these references is orientational, not deterministic; it serves to describe conditions for accessing content, not to pigeonhole the learner into a static profile (Kolb, 1984; Fleming, 2001).

This article does not yet measure the empirical impact of Calibra; however, it identifies the dimensions that will need to be observed in subsequent developments: *engagement*, persistence, cognitive load, perceived clarity, applicability, and transfer to real-world situations. Future evaluation, moreover, should not be limited to satisfaction.

3. Results

Calibra's process is divided into three main steps: learner assessment, construction of *the Basic Functioning Snapshot* (BFS), and definition of *the Preferred Learning Style* (PLS)—a sequence that distinguishes between data collection, profile interpretation, and training selection. The assessment gathers information about the learner across the three macro-areas; the BFS organizes this information into a dynamic and functional framework; and the PLS translates this analysis into operational guidelines for designing the training program.

3.1 The Three Broad Areas of the Assessment

As previously mentioned, the Calibra framework has identified three broad areas as the data sources for learner assessment: *life skills*, soft skills, and *learning styles*. *Life skills* are understood as functional dimensions of one's relationship with experience and are categorized into three axes: thought, emotion, and reactivity. Thinking concerns the cognitive organization of information; emotion indicates the weight of affective resonance in the construction of meaning; reactivity describes the way in which the individual responds to stimuli, demands, and complexity. In this framework, the emotional component is not ignored or separated from cognition, because learning is never merely the rational processing of information: emotion,

motivation, perceived security, and the relationship with the context all influence understanding and the willingness to continue the learning journey (Immordino-Yang and Damasio, 2007).

Calibra draws on soft skills, above all, in their relational dimension; they are not considered static competencies, but rather configurations that influence how the learner relates to the educational context: collaboration or competition, individual or group work, empathy or distance, flexibility or rigidity. It is not a matter of choosing which approach is better; rather, it is a matter of understanding which setting can best support learning at that moment.

The third broad area concerns the ways in which the experience is accessed and processed. In Calibra, *learning styles* are interpreted as dynamic frameworks, not as definitive categories; they integrate channels of engagement, experiential modalities, and levels of awareness. The distinction between felt, thought, and practical experience allows us to view learning as a multilayered process in which affective, cognitive, and operational components intertwine (Panksepp, 1998).

3.2 *Basic Functioning Snapshot* and *Preferred Learning Style*

The key step is translating the profile into instructional decisions. Based on the *Basic Functioning Snapshot* and the *Preferred Learning Style*, Calibra guides the choice of method, format, sequence, level of guidance, degree of interaction, activities, and assessment methods. Personalization concerns the way in which the learner is guided toward learning. The content is not simplified or diluted; rather, the format in which it is presented changes. The model thus seeks to go a step further than more standardized training, viewing individual differences not as obstacles to design but as valuable material for building more accessible and relevant learning paths. One profile may suggest a more structured learning path, another may call for more opportunities for practical application, and yet another may indicate a need for visual, narrative, analytical, or collaborative mediation. The system's function is also to make these choices more intentional and traceable, not to replace the judgment of the instructor or instructional designer.

The *Basic Functioning Snapshot (*BFS) represents the first point of synthesis in the process. It is a functional and dynamic snapshot of how the learner approaches the learning experience—not a diagnosis, not a label. The BFS describes how the person tends to relate to content, tasks, context, and perceived complexity; its usefulness lies in interpreting combinations—not individual preferences—and in identifying which conditions seem to facilitate or hinder learning. A person may prefer individual work but benefit from guided group discussions; they may appear self-directed yet become disoriented when faced with poorly organized content: the BFS serves precisely to identify these patterns, avoiding oversimplifications.

The BFS yields the *"Preferred Learning Style*," which is not a definitive answer; it is not intended as a fixed "learning style," but rather as a guiding framework for the conditions that can make content more accessible at a certain stage of learning. The model does not, therefore, establish that a person can

learn in only one way; rather, it posits that certain forms of access may facilitate entry into the learning path, and then guide the learner toward more varied and less conventional methods.

The PLS also serves a metacognitive function: making the learner more aware of the conditions that facilitate their own learning can foster a less passive and more self-regulated role for the learner within their educational journey. Personalization, therefore, affects not only the form of the content but also the individual's ability to observe their own learning style (Flavell, 1979).

3.3 The PLS: "Educational Accessibility" and Instructional Design

One of the model's most significant outcomes is the definition of the PLS as an indicator of "educational accessibility": this term refers to the set of conditions through which content can be not only presented to the learner but also received, understood, processed, and transformed into knowledge. Accessibility, therefore, is not simply the availability of the material: content may be online, open, and downloadable, yet remain difficult to navigate for those who are supposed to learn it.

Access to education is not a closed space; it is the entry point to an evolving journey. Calibra does not argue that the learner must remain in the mode that is immediately most functional; on the contrary, the model can start from more favorable conditions to guide the learner toward more complex, less familiar, or more challenging forms of learning. The PLS therefore serves a dual function: it is both facilitative and, in and of itself, (meta)formative. This idea ties into the possibility of tailoring the learning path to different levels of experience—more immediate and embodied, more symbolic, more reflective; learning can proceed through images, diagrams, narratives, reinterpretations, and applications, not always in the same order.

The practical outcome of the model is the transition from the profile to instructional design. As previously mentioned, the PLS guides decisions regarding format, sequence, guidance, interaction, application, and assessment. The content retains its objectives but is presented in different forms based on the identified accessibility conditions. In practice, the model can steer the learning path toward more structured materials, practical examples, simulations, maps, discussion sessions, reflection exercises, or group activities. Instructional design is not left to intuition but is grounded in information collected and interpreted through an explicit logic. In this way, Calibra distinguishes itself from a generic training system: it does not suggest content based on superficial preferences but builds a relationship between the functional profile, learning objectives, and conditions of use. Personalization becomes an integral part of instructional design, not a technical afterthought added at the end.

3.4 Traceability of Outputs

The formalization of the framework has also led to the need for traceable outputs. Every step must be reconstructible: information gathered, factors considered, the BFS summary, the PLS definition, and the criteria used to generate the training proposal. This traceability allows the instructor, designer, and

those responsible for training to understand why a course was recommended and, if necessary, to modify it.

The outputs are also understandable to the learner; receiving feedback on one's learning style can increase awareness and engagement, because traceability supports continuous improvement from various perspectives: if a learning path results in low participation, low engagement, or difficulty in understanding, the training choices can be revised, since the relationship between the individual, the context, and the training is dynamic and must be observed in their interaction (Bandura, 1986). The value of Calibra lies not only in its use of AI, but in its ability to make the relationship between assessment, its interpretation, and instructional design transparent.

4. Discussion

4.1 Personalization and Cross-Curricularity: Two Examples

Calibra addresses a common problem in digital professional training: many courses claim to be personalized, but often this personalization is limited to content selected by the instructor for everyone, flexibility in timing, or minor content variations for individual learners. In digital contexts, a great deal of data is collected on user behavior; this information is not always used to personalize the learner's educational path but, rather, for marketing purposes. *Learning analytics* gains value when it goes beyond mere monitoring to become a tool for understanding and improving the learning experience (Siemens, 2013). In this regard, Calibra occupies the space between assessment, design, and review; it does not merely observe whether a learner completes or abandons a course, but attempts to formulate hypotheses about the conditions that make that course more or less accessible. Personalization thus becomes a design process, not a mere technological function.

One of Calibra's most innovative and significant aspects, moreover, lies in its cross-cutting nature across education as a whole. Although the framework was initially conceived as a tool to support learning processes, its operating principle is applicable to any context in which one or more people need to convey information, skills, knowledge, or guidance to others. In fact, the idea behind Calibra is that there is no universally or ontologically effective form of communication, but that every learning process can benefit from a greater understanding of the recipient's individual characteristics—their uniqueness. For this reason, Calibra lends itself to applications ranging from university and vocational education to healthcare, psychology, personnel selection, career guidance, and counseling.

During the assessment phase, students are given questionnaires designed to explore their *life skills*, soft skills, and *learning styles*. *Life skills* help identify aspects related to thinking patterns (*decision-making*, etc.), emotional-affective patterns, and reactive patterns (problem-solving, etc.). Soft skills, on the other hand, help bring to light competencies that cut across professional fields of application,

such as communication, empathy, collaboration, leadership, and autonomy. Finally, preferred learning styles—for example, a prevalence of perceptual-experiential or more reflective and conceptual approaches—are investigated through models such as VARK, along with additional dimensions related to how experiences and information are processed. Based on the questionnaire responses and feedback on how the individual interacts, the BFS is determined, from which a PLS is derived—the learning style most consistent with one's BFS.

Take, for example, a professor who needs to design and launch a university-level medical course for 120 students:

Phase 1 - Assessment

Before the course begins, enrolled students receive a series of online questions via Calibra, accessible through a QR code or link. The questionnaires explore three main dimensions and collect the fundamental data needed to understand individual and collective characteristics.

Phase 2 – Profiling

Calibra processes the collected data and provides the instructor with an overview of the class. The system highlights the students' predominant profiles, their main strengths, potential challenges, and emerging learning needs. For example, the teacher may find that 60% of students prefer visual learning and that certain communication and collaboration skills are underdeveloped. In addition to the overall snapshot of the class, Calibra also provides individual profiles that the teacher can access.

Phase 3 - Adapting the Content

Based on the information gathered, the instructor will find it easier to plan and design the course. The professor of medicine will decide to supplement lectures with concept maps, infographics, and diagrams to align with the prevailing learning style profile (LSP) in the classroom. He will also introduce practical exercises, simulations, and guided discussions to foster interaction among learners and facilitate the transfer of theoretical knowledge to professional practice. To further support interpersonal skills, he will incorporate group activities and structured discussion sessions. Instructional design is thus tailored to the students' actual characteristics rather than to generic assumptions: the future implementation of Calibra may include the option to directly generate instructional materials in the format best suited to the student's PLS, while keeping the content as unchanged as possible.

Phase 4 - Adaptive Learning Experience

Throughout the course, Calibra will be able to collect feedback and learning data through additional questionnaires, enabling the instructor to monitor the effectiveness of the strategies adopted and adjust content formats, methods, and tools should new needs arise. Training thus becomes a dynamic and adaptive process in which instructors and students participate in a journey that is continuously calibrated to the characteristics and evolution of both the group and the individual, through a personalized, flexible, and learner-centered instructional model.

In the final analysis, it is worth noting Calibra's applications beyond explicit educational contexts, such as the doctor-patient relationship: in this case, Calibra could support both the healthcare professional and the patient in establishing a more personalized therapeutic communication. By gathering information regarding the patient's relationship with reality and their emotional and communicative patterns, the system could help the physician understand how to present a diagnosis, explain a treatment, or encourage adherence to prescriptions. Consider a patient who struggles to follow treatment instructions correctly—whether due to medical difficulties or simply because the instructions are complex. Calibra would serve as a supportive bridge of knowledge to foster compliance and help navigate even difficult moments, such as the treatment process. Calibra's analysis might reveal, for example, that the patient learns best through auditory cues and is more responsive to concrete examples than to abstract explanations. The physician could then supplement the consultation with an audio recording of the prescriptions and dosages, or explain what to look for to best monitor the patient's health. For healthcare professionals, Calibra can serve as a tool for reflecting on their communication and interpersonal skills—their soft skills, strengths, and weaknesses—to chart paths for further learning and development, either independently or with the guidance of Calibra's professionals.

From this perspective, Calibra serves as a tool capable of adapting the way information is conveyed to the characteristics of the recipient, thereby enhancing the effectiveness of learning processes and skill development across a variety of professional contexts.

4.2 Beyond the Data

The greatest risk is turning the profile into a category; Calibra avoids this oversimplification by treating BFS and PLS as dynamic configurations: they do not define who the learner is, but rather indicate which conditions can facilitate access to content at a given stage of the learning journey. Talking about *"Preferred Learning Style*," therefore, does not mean claiming that a person can learn effectively in only one way. It means recognizing that the relationship between the learner and instructional content can be enhanced by certain initial conditions. This approach is consistent with instructional design that takes into account diverse conditions and modalities. Learning requires structure, but also mediation; clear objectives, but also forms of access compatible with the learner (Gagné, 1985). Calibra updates this need with an *AI-based* approach, seeking to make the link between profile and instructional choice more explicit.

4.3 The Role of AI

In the Calibra model, artificial intelligence does not replace the teacher, trainer, or instructional designer. Technology alone guarantees nothing—neither personalization nor effectiveness; what matters is how it is integrated into an instructional model grounded in explicit rules, learning objectives, and human oversight. AI becomes useful when it helps manage a level of complexity that would be difficult to sustain in traditional contexts. An

instructor may intuit that different learners have different needs; often, however, they lack the tools and time to transform this intuition into differentiated learning paths. Calibra addresses this issue by gathering information, organizing it, and guiding decision-making—but without disregarding human supervision. The *"white box*" approach is also crucial: in adaptive systems, the risk is not only that the algorithm will make a mistake, but that no one will be able to understand why it produced a certain solution or to what extent that solution is valid. In education, every decision is delicate and can influence how learners perceive the content, their own abilities, and their relationship with learning; for this reason, a *human-centered* approach requires systems that are reliable, explainable, and controllable (Shneiderman, 2020).

In the current version of Calibra, an algorithm has been implemented to define *learning styles* and perform the corresponding clustering; this algorithm combines an *AI-driven* process with a deterministic process for final verification and control. Currently, the system uses calls to the Gemini 2.5 Flash LLM, Google's artificial intelligence model, designed to provide rapid responses, low costs, and strong performance even for complex or large-scale tasks. It is engineered for low latency and efficient token usage, making it well-suited for contexts requiring speed, scalability, and resource optimization. Among its strengths are advanced reasoning—which allows it to handle complex requests in a more structured way—and multimodal capabilities, enabling it to process various types of input, including text, images, audio, and video. Looking ahead, the use of additional LLMs or *AI* expert systems is already being explored to develop new and optimized *features* for Calibra.

4.4 Microlearning and Professional Training

Calibra appears particularly well-suited for microlearning and professional training contexts, where time is limited and learners quickly assess the usefulness of the course. In corporate settings, training is offered not only to enrich knowledge but also to solve problems, perform better, and apply what is learned to the workplace. In these contexts, the format of the content is crucial. If poorly designed, microlearning risks becoming a sterile framework of short, disconnected content: brevity may facilitate access, but it does not guarantee continuity or the transfer of knowledge. Furthermore, in professional settings, personalization can influence whether learners continue: a course perceived as distant, too abstract, or poorly connected to one's own context may be abandoned even if the content is valid. Conversely, a course that quickly connects with how the learner understands and applies concepts can foster engagement and continuity; in asynchronous learning, this matters even more.

4.5 Educational Accessibility and Operational Implications

The concept of educational accessibility is important to the model. Traditionally, accessibility in education is associated with the physical ability to access content or the removal of technical barriers. Calibra broadens this perspective: content is truly accessible when it can be understood,

processed, transformed into knowledge, and applied by the learner. This approach shifts the focus from the availability of knowledge to its actual usability. A resource may be online, open-access, and well-designed, yet still prove of little use if its format does not meet the learner's needs. The PLS thus becomes a tool for identifying a more suitable entry point, without reducing complexity or compromising the objectives of the learning path.

The operational implications of the accessibility goal affect teachers, instructional designers, and organizations. For instructors, Calibra can support the differentiation of teaching methods; for instructional designers, it provides a foundation for building modular and adaptive learning paths; and for organizations, it serves as a tool to make training more relevant and assessable. Calibra can, therefore, make these choices more intentional: not merely distributing content, but suggesting which format can make that content more accessible for a specific learner profile and in a specific context.

5. Limitations and Future Developments

This paper considers Calibra as a proposed method, not as a definitively validated model. The first limitation, therefore, concerns the descriptive nature of the work. The structure of the model—the transition from *the "Basic Functioning Snapshot" to the "Preferred Learning Style"* and, consequently, the translation of the profile into instructional guidelines—has been formalized theoretically but requires empirical validation to assess its stability, reliability, and impact on learning outcomes.

A second limitation concerns the three macro-areas of assessment: it will be necessary to verify to what extent *life skills*, soft skills, and *learning styles* are capable of describing configurations that are truly functional for instructional design. It will be necessary to understand, for example, whether the identified patterns remain stable over time, whether they change based on content, and whether they can be effectively updated through interaction with the learning pathway. For the PLS—defined as a guideline for educational accessibility rather than a rigid category—clear interpretive criteria will be needed to prevent misuse of the profile.

The model's predictive power must also be verified: Calibra hypothesizes that greater consistency between the profile, content format, and course structure can foster *engagement*, comprehension, adherence, and practical application, but this hypothesis must be tested through practical studies, larger samples, different contexts, and measures that go beyond mere satisfaction.

Future evaluation should integrate behavioral indicators, subjective feedback, and practical application tests: access to the training program, completion of the program, return after interruptions, perceived clarity, cognitive load, and knowledge transfer in real-world situations must be analyzed in conjunction with the initial design of the learner profile. *Learning analytics*, in fact, becomes useful only when it links the collected data to information about the learner's learning and to meaningful instructional decisions (Long and Siemens, 2011).

Then there is the issue of artificial intelligence: even though Calibra is designed according to a *white-box* logic, the AI component must be governed by criteria of transparency, data protection, human oversight, and responsible updating. Every automated or semi-automated decision must be traceable, especially when it affects the training proposal offered to the learner.

Future developments must therefore proceed in four directions: validation of the assessment, ongoing verification of the PLS, development of blended metrics, and refinement of the *white-box* logic.

6. Conclusions

Calibra proposes a model of personalized learning based on a simple yet profound idea: learning depends not only on the quality of the content, but also on the relationship between content, form, context, and individual ways of accessing knowledge. From this perspective, learning is not merely the transmission of information, but a process in which the learner interprets, organizes, and transforms the experience.

Calibra's original contribution becomes clearer when the model is compared with certain families of systems already analyzed in the literature on personalized learning. *Adaptive learning* and *personalized learning* systems have helped overcome the rigidity of one-size-fits-all learning paths by using data on preferences, outcomes, characteristics, and usage behaviors to adapt content, activities, or instructional sequences (Xie et al., 2019; Kabudi et al., 2021). *Intelligent Tutoring Systems*, on the other hand, are designed to provide individualized instruction, feedback, and support based on structured representations of the learner's knowledge, characteristics, and difficulties, which are updated in light of progress made throughout the learning journey (Ma et al., 2014; Mousavinasab et al., 2021). Selection and personalized recommendation systems applied to *Technology-Enhanced Learning* help select resources, materials, or activities deemed most relevant to the user, while Universal Design for Learning provides an important framework for designing digital environments and content accessible to a diverse range of learners (Manouselis et al., 2013; CAST, 2018).

Calibra aligns with these approaches but does not fully coincide with any of them. Its originality lies not only in content adaptation, nor in the production of training resources, nor solely in the analysis of behaviors recorded on the platform. The distinctive feature is the construction of a methodological model that links the initial assessment of the user's profile, *the "Basic Functioning Snapshot*," *"Preferred Learning Style*," and traceable instructional choices. In other words, Calibra does not merely propose a certain piece of content to a specific learner, but seeks to answer a broader question: under what design conditions can this content become more accessible, understandable, and applicable for this profile at this stage of the learning journey? The model's main contribution, therefore, lies in linking assessment and design. Through the BFS and the PLS, Calibra seeks to identify conditions that can make the content more accessible, understandable, and applicable in relation

to the identification of the learner's profile. The profile is not used as a label, but as a guiding tool to tailor format, sequence, guidance, interaction, etc.

A second element of originality concerns the way the learner's profile is interpreted. In adaptive systems, there is a risk that personalization will turn into classification: the user is associated with a cluster, a preference, or a behavioral pattern and receives content consistent with that category. Calibra, on the other hand, interprets the BFS and PLS as functional and dynamic configurations, not as fixed labels. The PLS does not define what the learner "is," but rather identifies conditions for educational accessibility that can facilitate entry into the program and that can be updated, corrected, or expanded over time.

An important aspect of Calibra is that AI does not function merely as an automation engine, but rather as a tool to assist in interpreting the profile and generating design recommendations, with the teacher and instructional designer remaining primarily responsible for supervising, validating, and, if necessary, modifying the choices suggested by the system. Calibra cannot replace teachers, trainers, or instructional designers; rather, it offers a design support tool in which AI serves an enabling function: the value of the model lies not in the technology itself, but in the ability to use it to make the relationship between the learner and the learning path more conscious and adaptive.

Finally, Calibra adopts the *white-box* approach as an integral part of the model: traceability becomes a methodological requirement, not a technical option. Educational personalization, especially when supported by AI, must be explainable, and its steps must be reconstructible; for this reason, the framework places value not only on the effectiveness of the output but also on the transparency of the process that produces it. The model's original contribution, therefore, lies in integrating personalization, educational accessibility, instructional design, and the traceability of choices into a single architecture. This integration further distinguishes Calibra from other adaptive systems for the personalized delivery of educational content.

In conclusion, Calibra represents a methodological approach to moving beyond a "one-size-fits-all" educational model without falling into opaque or reductive personalization. Its goal is not to adapt knowledge to the learner's immediate preferences, but to create more favorable conditions so that each person can access the content, understand it, apply it, and progressively expand their own way of learning.

Bibliography

Adadi, A., & Berrada, M. (2018). Peeking inside the black-box: A survey on explainable artificial intelligence (XAI). IEEE Access, 6, 52138–52160. https://doi.org/10.1109/ACCESS.2018.2870052

Azevedo, R., Città, G., Gentile, M., & Ifenthaler, D. (2024). Editorial. Artificial Intelligence and the paradigm shift of teachers' role. Italian Journal of Educational Technology, 32 (1), 5–7. https://doi.org/10.17471/2499-4324/1358

Bandura, A. (1986). Social foundations of thought and action: A social cognitive theory. Prentice-Hall.

CAST. (2018). Universal Design for Learning Guidelines version 2.2. CAST.

Cronbach, L. J., & Snow, R. E. (1977). Aptitudes and instructional methods: A handbook for research on interactions. Irvington Publishers.

Ferguson, R. (2012). Learning analytics: Drivers, developments and challenges. International Journal of Technology Enhanced Learning, 4 (5/6), 304–317. https://doi.org/10.1504/IJTEL.2012.051816

Flavell, J. H. (1979). Metacognition and cognitive monitoring: A new area of cognitive-developmental inquiry. American Psychologist, 34 (10), 906–911. https://doi.org/10.1037/0003-066X.34.10.906

Fleming, N. D. (2001). Teaching and learning styles: VARK strategies. N.D. Fleming.

Freire, P. (1970). Pedagogy of the oppressed. Continuum.

Gagné, R. M. (1985). The conditions of learning and theory of instruction (4th ed.). Holt, Rinehart and Winston.

Guidotti, R., Monreale, A., Ruggieri, S., Turini, F., Giannotti, F., & Pedreschi, D. (2018). A survey of methods for explaining black box models. ACM Computing Surveys, 51 (5), 93. https://doi.org/10.1145/3236009

Hwang, G. J., Xie, H., Wah, B. W., and Gašević, D. (2020). Vision, challenges, roles and research issues of Artificial Intelligence in Education. Computers and Education: Artificial Intelligence, 1, 100001. doi:10.1016/j.caeai.2020.100001

Immordino-Yang, M. H., & Damasio, A. (2007). We feel, therefore we learn: The relevance of affective and social neuroscience to education. Mind, Brain, and Education, 1 (1), 3–10. https://doi.org/10.1111/j.1751-228X.2007.00004.x

Kabudi, T., Pappas, I., & Olsen, D. H. (2021). AI-enabled adaptive learning systems: A systematic mapping of the literature. Computers and Education: Artificial Intelligence, 2, 100017. https://doi.org/10.1016/j.caeai.2021.100017

Kapp, K. M., & DeFelice, R. A. (2019). Microlearning: Short and sweet. Association for Talent Development.

Kolb, D. A. (1984). Experiential learning: Experience as the source of learning and development. Prentice Hall.

Long, P., & Siemens, G. (2011). Penetrating the fog: Analytics in learning and education. EDUCAUSE Review, 46 (5), 30–40.

Ma, W., Adesope, O. O., Nesbit, J. C., and Liu, Q. (2014). Intelligent tutoring systems and learning outcomes: A meta-analysis. Journal of Educational Psychology, 106(4), 901–918. doi:10.1037/a0037123

Manouselis, N., Drachsler, H., Verbert, K., and Duval, E. (2013). Recommender Systems for Learning. Springer. doi:10.1007/978-1-4614-4361-2

Mayer, R. E. (2009). Multimedia learning (2nd ed.). Cambridge University Press. https://doi.org/10.1017/CBO9780511811678

McLuhan, M. (1964). Understanding media: The extensions of man. McGraw-Hill.

Mousavinasab, E., Zarifsanaiey, N., Niakan Kalhori, S. R., Rakhshan, M., Keikha, L., and Ghazi Saeedi, M. (2021). Intelligent tutoring systems: A systematic review of characteristics, applications, and evaluation methods. Interactive Learning Environments, 29(1), 142–163. doi:10.1080/10494820.2018.1558257

Panksepp, J. (1998). Affective neuroscience: The foundations of human and animal emotions. Oxford University Press.

Pashler, H., McDaniel, M., Rohrer, D., & Bjork, R. (2008). Learning styles: Concepts and evidence. Psychological Science in the Public Interest, 9 (3), 105–119. https://doi.org/10.1111/j.1539-6053.2009.01038.x

Pessoa, L. (2008). On the relationship between emotion and cognition. Nature Reviews Neuroscience, 9 (2), 148–158. https://doi.org/10.1038/nrn2317

Robles, M. M. (2012). Executive perceptions of the top 10 soft skills needed in today's workplace. Business Communication Quarterly, 75 (4), 453–465. https://doi.org/10.1177/1080569912460400

Shneiderman, B. (2020). Human-centered artificial intelligence: Reliable, safe & trustworthy. International Journal of Human-Computer Interaction, 36 (6), 495–504. https://doi.org/10.1080/10447318.2020.1741118

Siemens, G. (2013). Learning analytics: The emergence of a discipline. American Behavioral Scientist, 57 (10), 1380–1400. https://doi.org/10.1177/0002764213498851

Sweller, J., Ayres, P., & Kalyuga, S. (2011). Cognitive load theory. Springer. https://doi.org/10.1007/978-1-4419-8126-4

Taylor, A. D., & Hung, W. (2022). The effects of microlearning: A scoping review. Educational Technology Research and Development, 70 (2), 363–395. https://doi.org/10.1007/s11423-022-10084-1

World Health Organization. (1994). Life skills education for children and adolescents in schools. World Health Organization.

Xie, H., Chu, H.-C., Hwang, G.-J., & Wang, C.-C. (2019). Trends and development in technology-enhanced adaptive/personalized learning: A systematic review of journal publications from 2007 to 2017. Computers & Education, 140, 103599. https://doi.org/10.1016/j.compedu.2019.103599



Nicola Illuzzi

nicolailluzziroma@gmail.com

Nicola Illuzzi is a dentist and independent practitioner in Rome. He serves as a council member of the Dentists' Registry Commission of the Order of Physicians, Surgeons, and Dentists of Rome and as a member of the Order's Telemedicine and Artificial Intelligence Commission. He is involved in teaching, administration, and scientific coordination within undergraduate and graduate programs. He is the Director of the Master's Program in Instructional Design at the University of San Raffaele in Rome and focuses on healthcare innovation, professional communication, and the integration of education, digital technologies, and artificial intelligence.



Claudio Vergini

Claudio Vergini is a digital entrepreneur with over thirty years of experience in educational and technological innovation. He has built his career at the intersection of e-learning, educational technologies, and the digital transformation of learning systems. He has held roles in educational leadership and

strategic coordination, leading complex projects for businesses, public administrations, and national organizations, with a growing focus on AI-driven solutions.



Giorgio Tullio De Negri

Giorgio Tullio De Negri is a psychologist and psychotherapist, registered with the Order of Psychologists of Lazio (No. 22660). He has gained clinical experience working with children, adolescents, and adults, providing psychological support and psychotherapy using a psychodynamic approach. In his work, he has integrated digital, technological, and video game tools as therapeutic and communicative resources, using them to foster emotional expression, rapport, and patient engagement. He treats anxiety, depression, relationship difficulties, and supports personal growth, conducting sessions both in person and online.

ICDL Artificial Intelligence: Italy's First Certification in AI Literacy

Antonella Fancello, AICA

Abstract

Since 1961, AICA has been supporting Italian citizens, students, and professionals on their journey toward digital literacy. With the ICDL Artificial Intelligence 2.0 Syllabus, the Association is launching the first Italian certification dedicated to AI literacy, in response to the requirements introduced by Article 4 of the European Regulation on Artificial Intelligence (AI Act), Italy's first law on AI (Law 132/2025), and the guidelines issued by UNESCO and the Ministry of Labor. This article outlines the regulatory and historical context that made this certification necessary and analyzes the five sections of the 2.0 Syllabus, which cover the fundamentals of AI, its practical applications, risks and ethics, generative tools, and basic prompt engineering techniques. AI Literacy is understood not as a technical skill but as a democratic civic competency: a universal prerequisite for consciously exercising one's rights in the age of algorithms.

Keywords: *AI Literacy, AICA, ICDL, digital certification, AI Act, digital skills, artificial intelligence, digital citizenship, prompt engineering, generative AI*

1. The Legacy of a Sixty-Year Mission

From automatic computation to autonomous algorithms, on a long journey spanning the Third and Fourth Industrial Revolutions, AICA—the Italian Association for Information Technology and Automatic Computation—has been supporting schools and the professional world for over sixty years in fostering the growth and development of informed digital citizenship. Founded in 1961, the Association has contributed to every technological turning point while keeping its mission unchanged: to make digital skills accessible to the entire community, regardless of age, profession, or cultural background.

In the midst of the third industrial revolution—which, beginning in the 1960s, saw the gradual spread of personal computers and later the internet—through the European Computer Driving License (ECDL,

now ICDL), AICA's mission has been to certify basic digital skills and raise awareness about the safe use of the internet and office automation applications. Millions of Italians have earned ECDL/ICDL certification: students seeking a credential valuable in the job market, workers needing to retrain, public employees tasked with modernizing their administrative procedures, and seniors who wanted to actively participate in the digital life of their communities.

That journey has left a profound mark on the Italian education system. It has demonstrated that digital skills certification is not merely a tool for individual assessment, but a vehicle for public policy capable of raising a country's average level of technological literacy. It has built a widespread network of test centers distributed throughout the country—schools, universities, and vocational training centers—making certification accessible even to geographically remote areas and segments of the population farthest from innovation hubs.

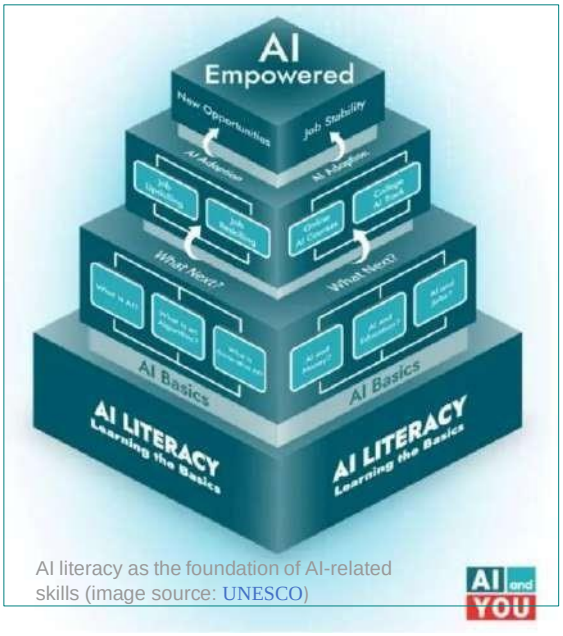
Today we are in the midst of the fourth industrial revolution—the data revolution: never before has there been such a vast amount of data, generated by all objects connected to the network (IoT, Internet of Things), processed by Big Data and artificial intelligence systems based on Machine Learning and Deep Learning, which extract "knowledge" from this massive volume of data to "learn" how to generate new value, thanks to unprecedented computing power. In this radically transformed landscape, AICA has chosen to channel its expertise into a new mission it considers non-negotiable: promoting and certifying AI Literacy—literacy in artificial intelligence.

This is not a change in direction, but rather a continuation of our mission adapted to the times. Thirty years ago, AICA guided citizens and professionals through their first steps with personal computers. Twenty years ago, it helped them understand the internet. Today, the challenge is of a different magnitude: it is no longer about learning to use a tool, but about understanding what happens when we entrust a tool to make decisions on our behalf. It is a challenge that concerns not only programmers and technology experts, but every person who interacts with a digital system—whether it's the robot that autonomously cleans the house, the algorithm that decides whether or not you should get a bank loan, or the recommendation system that selects the news stories that will appear in your morning feed.

1.1 The Regulatory Context: When Law Meets Education

AICA's decision to invest in AI literacy stems not only from an internal strategic vision but also responds to a specific regulatory mandate that has taken on increasingly defined contours at the European and national levels in recent years.

On August 1, 2024, the European Regulation on Artificial Intelligence, known as the AI Act



AI literacy as the foundation of AI-related skills (image source: UNESCO)

(EU Regulation 2024/1689), entered into force. This is the world's first comprehensive regulatory framework dedicated specifically to artificial intelligence, and its Article 4 introduces an obligation that is unprecedented in the technology regulatory landscape: all those who develop, distribute, or use artificial intelligence systems must possess an adequate level of skills, knowledge, and understanding of the systems they use. This is neither a recommendation nor a policy objective: it is a legally binding obligation for all twenty-seven member states of the European Union. The AI literacy requirement established by Article 4 of the AI Act is part of a broader approach: that of holding accountable not only the producers of AI systems, but also those who deploy them and their end users. The Regulation recognizes that an artificial intelligence system, however technically sophisticated, is not neutral: its decisions have real impacts on people's lives, and those involved in its use—at any level of the supply chain—have a responsibility to understand how it works, what data it uses, what errors it may make, and what rights users have with respect to it.

At the national level, on October 10, 2025, Italy's first law on artificial intelligence (Law No. 132 of September 23, 2025) entered into force, transposing and supplementing the European framework while reaffirming the central role of training as a tool for implementing obligations related to AI literacy.

At the same time, the UNESCO framework on the ethics of artificial intelligence (Recommendation on the Ethics of Artificial Intelligence, 2021) identifies AI literacy as a universal prerequisite for the exercise of human rights in the age of algorithms, emphasizing that unequal access to these skills risks creating a new "digital divide" that goes far beyond a simple lack of connectivity. The Italian Ministry of Labor's guidelines include AI literacy among the skills to be promoted in workers' continuing education and in active employment policies, recognizing that the transition to a labor market permeated by artificial intelligence requires systematic, rather than occasional, training support.

Table 1 – Regulatory and Institutional Framework for AI Literacy

Regulatory / Institutional Source	Main Reference	Relevance to AI Literacy
EU AI Act (Reg. 2024/1689)	Art. 4 – AI Literacy Requirement	Requires providers, deployers, and users of AI systems to have an adequate level of

Regulatory / Institutional Source	Main Reference	Relevance to AI Literacy
		skills and awareness regarding the systems they use
Italian AI Law (Law No. 132/2025)	Articles 4, 5, 9 – Promotion of AI Literacy	Italy's first AI law, in effect as of October 10, 2025; reaffirms the educational requirement
UNESCO – Recommendation on AI Ethics (2021)	Global Framework for AI Literacy	Identifies AI literacy as a universal prerequisite for the exercise of human rights in the age of algorithms
Ministry of Labor	Guidelines on Digital Skills in National Collective Bargaining Agreements	Include AI literacy among the skills to be promoted in continuing education for workers and in active labor market policies
DigComp 3.0 (European Commission, 2022)	Cross-cutting integration of AI across the framework's areas	Moves beyond the concept of AI competency as an isolated module: artificial intelligence becomes an enabling technology that cuts across all aspects of digital citizenship

Source: Author's analysis

1.2 AICA's Response: Italy's First AI Literacy Certification

It is within this regulatory and cultural framework that AICA, in collaboration with the ICDL Foundation (the international body that manages ICDL certifications worldwide), has developed and launched the ICDL Artificial Intelligence Syllabus, Version 2.0: the first Italian certification explicitly dedicated to AI literacy.

This is not the first time AICA has certified skills related to artificial intelligence. It had already developed, together with the ICDL Foundation, an initial version of the "AI – Artificial Intelligence" Syllabus, which covered the fundamentals of traditional artificial intelligence. It had also integrated AI topics into the Digital Citizenship Syllabus. It had also incorporated artificial intelligence into the most recent DigComp 2.2 and DigCompEdu certifications. But the new Syllabus 2.0 represents a qualitative leap forward compared to everything that came before it.

The reason is historical: November 2022 marked a profound turning point in the landscape of artificial intelligence. The launch of ChatGPT and the subsequent proliferation of systems based on Large Language Models (LLMs) and generative AI brought AI out of laboratories and tech companies, placing it in the hands of tens of millions of ordinary people within a matter of months. What experts call a "disruptive" transformation—such was the significance of its impact on the professional world—has rendered obsolete any AI training that fails to account for this universe of generative tools, their mechanisms, their potential, and, above all, their risks.

By the time (June 2026) when machine-generated content has already surpassed human-generated content online, Stanford tells us that in just three years, generative AI is used by 53% of the

of the world's population—a pace faster than that of the spread of personal computers and the internet (Stanford University, AI Index Report 2026).

The ICDL Artificial Intelligence 2.0 Syllabus was therefore developed as a structured response to this transformation: a training program divided into five sections that covers the entire journey of understanding AI, from theoretical foundations to the practical use of generative tools, with a strong emphasis on ethical issues, security, and individual and collective responsibility.

2. AI Literacy as a Civic Competency: Why It Is Not a Technical Matter

Before delving into the analysis of the five sections of the Syllabus, it is worth pausing to consider a conceptual point that serves as the central theme of the entire certification: AI Literacy is not a technical skill. It is a civic competency.

This distinction is fundamental and must be clearly explained, as it risks being misunderstood. To say that AI Literacy is a civic competency does not mean that those who possess it know how to program algorithms, train machine learning models, or optimize neural network architectures. It means something deeper: that a person can recognize when they are interacting with an artificial intelligence system, knows their rights regarding that system, is capable of critically evaluating its responses rather than passively accepting them, and understands the implications of those interactions on their own life, on the lives of others, and on society as a whole.

The most accurate analogy is with the ability to read. Knowing how to read does not mean knowing how to write novels or understanding the deep syntax of language: it means having access to the meaning of the texts that surround us, being able to participate in the social and cultural life of one's community, and exercising critical thinking when faced with the information we receive. Functional illiteracy—the condition of those who can technically decipher letters but do not understand the meaning of what they read—excludes people from full citizenship. Similarly, a lack of AI literacy produces a form of "digital illiteracy" that excludes people from understanding the world in which we already live.

Every interaction we have with "intelligent" systems is also a choice of responsibility. We can accept, correct, reformulate, or halt the output of an AI system, always keeping in mind that these systems can make mistakes silently and often in far-reaching ways. This awareness is not a luxury for experts: it is a democratic necessity. The more people understand how these systems work, the more the collective decisions concerning them—such as how to regulate them, how to use them in schools and hospitals, and how to protect ourselves from their misuse—will be truly informed and participatory.

AI Literacy, as defined by the ICDL Artificial Intelligence 2.0 certification, is therefore a cultural response to a structural transformation. It is the recognition that living and working in the age of artificial intelligence requires new forms of awareness: knowing when you are interacting with

an automated system, understanding what data it is based on, knowing how to evaluate and, if necessary, reformulate its conclusions—and doing so without blindly delegating, while always giving due weight to the human responsibility for judgment, choice, and decision-making.

3. The ICDL Artificial Intelligence 2.0 Syllabus: Structure and Content

The ICDL Artificial Intelligence 2.0 Syllabus is structured into five thematic sections, each of which explores a fundamental aspect of the relationship between people and artificial intelligence systems. The course follows a progressive approach: it begins with an understanding of the fundamentals and progresses to the practical use of tools, always guided by an ethical thread that runs through all sections. The certification is designed for the "ICDL Core module audience": suitable for all citizens, including students, it does not require prior technical skills.

Table 2 – Structure of the ICDL Artificial Intelligence 2.0 Syllabus

Section	Title	Main Thematic Areas
1	Understanding AI Systems	Definition of AI, ML, and Deep Learning; foundation models; generative AI; autonomous agents; AI literacy as a civic competency
2	Practical Applications of AI	Home automation, transportation, public administration, education, healthcare, finance, creative industries, scientific research
3	Risks, Best Practices, and Ethics	Privacy, security, algorithmic bias, deepfakes, misinformation, accountability, social and environmental impact
4	Generative AI Tools	Text, image, audio, video, and code generators; RAG systems; selection criteria; account setup
5	Basics of prompt engineering and using the results	Prompting techniques, layering, zero-shot, few-shot, Chain of Thought, prompt chaining; refining and sharing outputs

Source: ICDL Artificial Intelligence 2.0 Syllabus, AICA 2026

3.1 Section 1 – Understanding Artificial Intelligence Systems

The first section of the syllabus lays the conceptual foundation for the entire course. Its goal is not to train IT professionals, but to provide anyone—regardless of their educational background—with the basic conceptual tools needed to understand what artificial intelligence is, how



it works in its most basic principles, and how it differs from other types of software.

The section introduces the distinction between the three evolutionary levels of artificial intelligence: narrow AI, which is what we interact with every day and excels at specific tasks such as facial recognition or machine translation; artificial general intelligence (AGI), a hypothetical system not yet realized but often touted as a reality by major players in the AI sector, capable of reasoning across any domain like a human being; and artificial superintelligence, a scenario that remains theoretical but is the subject of intense ethical and scientific debate. This distinction is important because it helps people manage their expectations and fears: much of what is presented in the media as "artificial general intelligence" is, in reality, narrow AI, with extraordinary capabilities in specific areas but without any form of general understanding or consciousness.

A particularly relevant aspect of this section concerns the basic technical functioning of AI systems. The Syllabus describes the three fundamental ingredients of any intelligent system—data (the raw material), algorithms (the processing instructions), and models (the internal representation that the system constructs to generalize)—and introduces the concepts of machine learning, artificial neural networks, and deep learning in language accessible to non-experts. Of particular relevance is the introduction of the so-called "black box" model problem, which is not

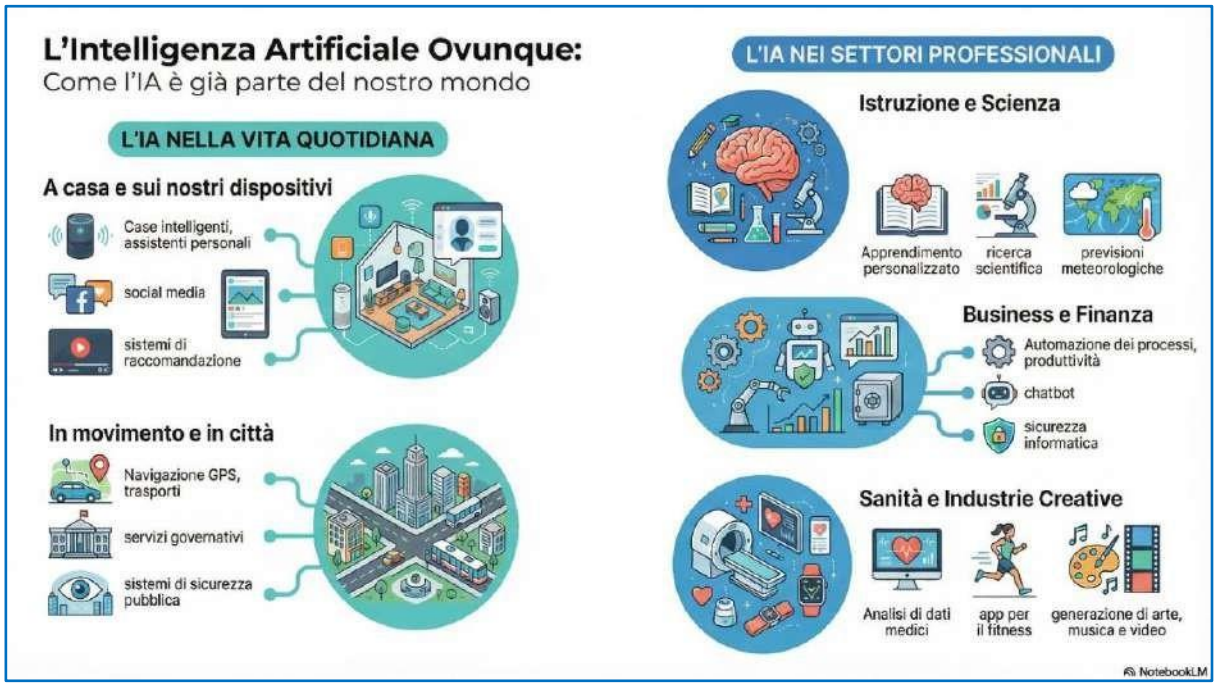
absolutely prohibited by the AI Act: the intrinsic difficulty of many deep learning systems, in which the internal decision-making process is not interpretable even by their own creators, with profound implications for transparency and accountability.

The section then addresses the phenomenon of generative artificial intelligence—the most visible and discussed frontier of contemporary AI—explaining the concept of prompts, how Large Language Models (LLMs) work, Natural Language Processing (NLP), and the evolution toward multimodal systems capable of working simultaneously with text, images, audio, and video. Finally, special attention is given to new AI agents: systems capable not only of responding to instructions but also of planning and carrying out sequences of autonomous actions to achieve a goal, opening up unprecedented scenarios of autonomy that require new reflections on human control.

The most important concept in the entire first section—and throughout the entire training program—is AI literacy. The syllabus defines it as the ability to understand what AI is, how it functions in its essential principles, what its limitations and risks are, and how to interact with it in a critical and informed manner. This definition is deliberately broad: it includes the cognitive dimension (understanding), the critical dimension (evaluation), the ethical dimension (responsibility), and the civic dimension (informed participation). The section strongly argues that this competency is urgently needed, especially for young people and ordinary citizens—not to become experts, but to ensure they are never completely alienated from issues that affect them.

3.2 Section 2 – Practical Applications of Artificial Intelligence Systems

The second section invites us to look at the world around us with fresh eyes, recognizing where and how artificial intelligence is already present in daily life and in professional fields. It is a section that is deliberately concrete and rich in examples, because one of the main obstacles to AI literacy is the perception—



often fueled by a certain type of media coverage—that artificial intelligence is a phenomenon of the future or a technology exclusive to laboratories and large tech companies. In reality, AI is already here, already active, already intertwined with every aspect of human life.

In the home, the Syllabus explores the role of voice assistants (Amazon Alexa, Google Assistant, Apple Siri), smart thermostats that learn the habits of a household's residents and optimize energy consumption, and recommendation algorithms that curate content on streaming platforms and social media. This last example is particularly instructive: it shows how AI is not merely a tool at the user's service, but also a system that actively shapes their information experience, with implications ranging from filter bubbles to the manipulation of public opinion.

In the professional sphere, this section explores artificial intelligence in education (adaptive learning systems, educational chatbots, advanced grading tools, automatic plagiarism detection), in healthcare (medical image analysis, epidemic prediction, acceleration of drug discovery), in finance (predictive risk analysis, fraud detection), in the workplace (automation of repetitive processes, digital assistants for customer service), and in the creative industries (generation of images, music, videos, and text).

A particularly insightful section is the one dedicated to scientific and environmental applications, such as systems that solve one of the most complex problems in molecular biology—predicting the three-dimensional structure of proteins—a prime example of how AI can revolutionize scientific research. Similarly, environmental monitoring through AI-processed satellite imagery, advanced weather forecasting, and the management of smart power grids demonstrate how these technologies can serve as essential tools in addressing global climate challenges.

The educational goal of this section is not merely informative. Understanding concrete examples of AI in the real world is the essential starting point for authentic AI literacy: what we do not know, we cannot evaluate, and what we do not evaluate, we cannot govern. Recognizing AI in everyday contexts is the first step toward being able to ask the right questions: Who designed this system? On what data was it trained? What are its limitations? Who is responsible if it makes a mistake?

3.3 Section 3 – Risks, Best Practices, and Ethical Considerations

The third section is the critical core of the Syllabus. It is not merely about knowing what AI does, but about understanding the risks it entails, how to protect oneself, and how to use it responsibly. This section most directly addresses the regulatory mandate of Article 4 of the AI Act, which requires not only



technical knowledge but also an awareness of the risks and ethical implications of artificial intelligence systems.

The section begins with an in-depth analysis of the issue of privacy. Every artificial intelligence system requires data to function, and that data often comes from the users themselves, through active input (information we enter directly) and passive tracking (monitoring of browsing behavior, geographic location, and purchases). This section explains how this data is used for personalization, advertising profiling, and training new models, and what concrete risks this poses to individual privacy: from a lack of transparency regarding data management, to the re-identification of individuals from datasets presented as "anonymized," to systematic tracking that progressively erodes freedom of thought and behavior.

Closely linked to privacy is cybersecurity. The Syllabus highlights how AI is both a target and a tool for cyberattacks: AI systems contain large amounts of sensitive data and, if compromised, can cause enormous damage; but AI is also used by cybercriminals to make their attacks more sophisticated—from automatically generated phishing via LLMs to voice cloning for telephone fraud, to the creation of credible fake identities.

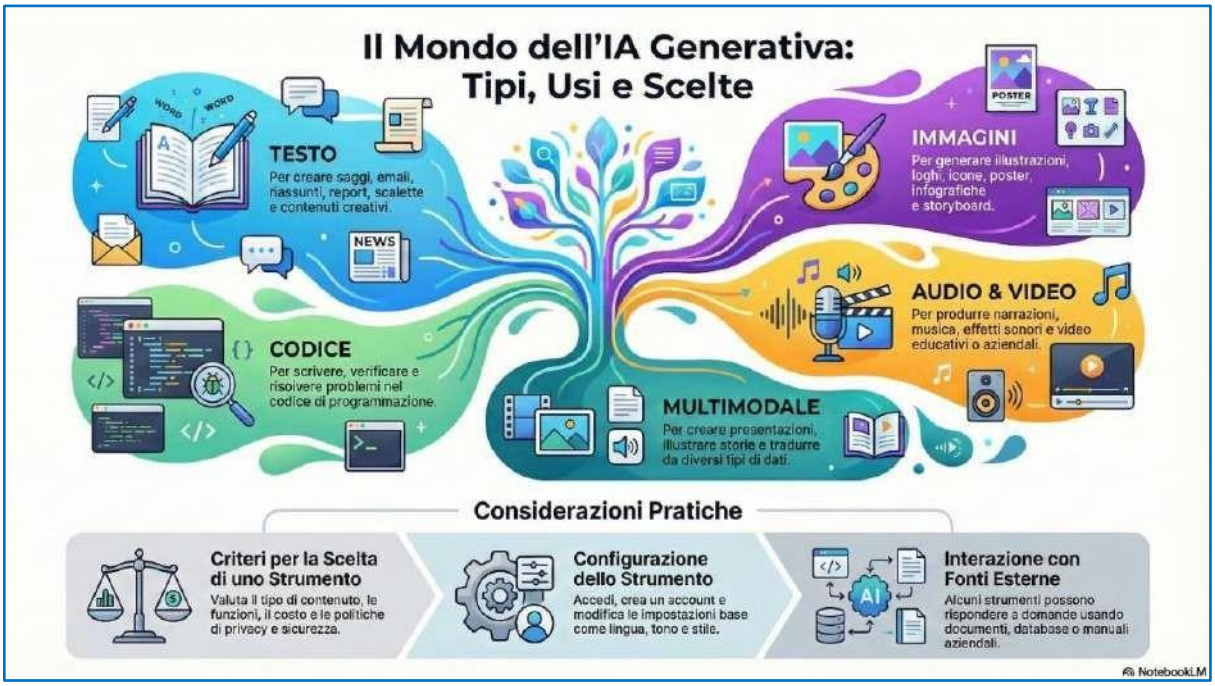
A particularly relevant issue is that of algorithmic bias: the unconscious prejudices that become embedded in AI models when they are trained on data that reflects stereotypes and inequalities present in human society. There are concrete examples of how facial recognition systems have shown significantly higher error rates for people with darker skin tones, how algorithms used in the judicial system to assess the risk of recidivism have disproportionately penalized specific communities, and how recruitment tools have reproduced the gender discrimination of past decades. In all these cases, the problem does not stem from intentional malice on the part of the programmer, but from a design that failed to take into account the quality and representativeness of the training data.

This section devotes considerable space to the topic of disinformation and deepfakes: images, audio, and video in which a real person's face or voice is replaced or manipulated with a level of realism that is increasingly difficult to distinguish from the authentic. The implications are enormous: from personal defamation to political manipulation, from the creation of false evidence to the erosion of collective epistemic trust—our shared ability to distinguish truth from falsehood—with potentially devastating consequences for democracy and public debate. The Syllabus provides methodological tools for the critical evaluation of digital content: questions to ask when encountering any information found online, visual and stylistic red flags, and strategies for verification through multiple, reliable sources.

The section concludes with an analysis of the social and environmental impact of AI. On the social front, the Syllabus calls for addressing the transformation of the labor market and the risk of growing economic inequalities between those who possess the skills to leverage AI and those who are excluded from it; the danger of excessive reliance on AI, which delegates to machines decisions that should remain human; and the issue of mass surveillance. On the environmental front, it is also important to address the major issue of the hidden energy costs of artificial intelligence: the enormous consumption of electricity and water required, respectively, to power and cool the data centers that host AI systems, whose energy consumption for some training sessions is comparable to—or even exceeds—that of a transatlantic flight.

3.4 Section 4 – Understanding Generative AI Systems

The fourth section delves into practical matters: What generative AI tools are actually available today, how do they work, how do you choose them, and how do you configure them? From an educational standpoint, developing this section can serve as an up-to-date guide to the landscape of generative tools, which has expanded extraordinarily since November 2022 and continues to evolve at



an impressive pace.

Text generators—currently OpenAI's ChatGPT, Anthropic's Claude, Google's Gemini, and Microsoft's Copilot—are likely the most widespread and versatile tools. The Syllabus provides a general overview of the most relevant use cases for students and citizens: drafting various types of texts, summarizing complex documents, explaining concepts in accessible language, supporting creative writing, answering questions, and generating structures and outlines. At the same time, it clearly highlights a structural limitation of these systems that users must always keep in mind: language models do not "know things" in the human sense of the term, but rather generate statistically plausible text sequences based on patterns learned during training. When they do not know the correct answer, they tend to invent one with the same confidence with which they would provide correct information—a phenomenon that we humans have termed "hallucination." For this reason, any output produced by a generative system must always be critically reviewed, with particular attention to factual data, dates, citations, and specific claims.

Image generators use different architectures (primarily diffusion models) and open up possibilities that were once unthinkable for anyone without a background in graphic design. From illustrations for

school presentations to concept art for creative products, these tools drastically lower the skill threshold required to engage in visual production. This section illustrates their practical applications while also highlighting the ethical implications: the ability to reproduce the styles of real artists without consent, the presence of visual biases in training data, and the risk of their use to create false or manipulative content.

The section then explores the generation of video, audio, music, and source code. Regarding this last category, the section emphasizes how AI coding tools have already radically transformed software development practices: not only for experienced programmers, for whom they act as a productivity booster, but also for those approaching programming for the first time, by drastically lowering the barrier to entry and transforming debugging from an often-frustrating task into a guided and educational process.

One category of particular interest is that of RAG (Retrieval-Augmented Generation) systems: tools capable of answering questions by drawing on specific documents or databases provided by the user. These systems overcome one of the main limitations of generic language models—the lack of up-to-date or domain-specific knowledge—and find very practical applications within organizations: a company can build an internal Q&A system that answers employees' questions by drawing on company policies; a legal department can query an archive of court rulings; a public institution can make its regulations accessible through a conversational assistant.

The section concludes with a practical guide to selecting and configuring these tools: selection criteria (type of content required, intended use, cost, privacy policies), specific features to evaluate for each category, and instructions for creating an account and customizing settings. Particular attention is given to the privacy and data security policies governing the information entered into these tools: a crucial issue, especially in professional and organizational contexts, where the information processed may be confidential or sensitive.

3.5 Section 5 – Basics of Prompting and Using Results

The fifth and final section is the most practical of the entire course: how to communicate effectively with a generative AI system to get the desired results, how to refine those results, and how to use them responsibly. Prompt engineering—the discipline that studies how to write effective instructions for AI systems—is presented not as a technical skill reserved for specialists, but as a practical ability that is learned and honed through practice, accessible to any informed user: the art of asking the right question.



The foundation of prompt engineering is simple yet often overlooked: the quality of the result depends to a very large extent on the quality of the instruction provided to the system. A vague request produces a vague response; a precise, contextualized, and well-structured request produces a precise, useful, and often surprisingly effective response. The Syllabus identifies five fundamental elements of an effective prompt: the clarity and specificity of the instruction (precise constraints that guide the system without limiting its creativity); context (background information that helps the system calibrate the tone, register, and content of the response); role (asking the system to adopt a specific perspective to activate a relevant knowledge base); the purpose of the desired text; and constraints regarding format, length, and style.

A particularly useful technique for complex tasks is "layering": constructing the prompt in successive layers of information, gradually adding details and constraints rather than trying to include everything in a single long and complex instruction. Layering allows you to refine the result iteratively, fine-tuning the output with additional instructions without having to rewrite everything from scratch each time.

The section also introduces some more advanced techniques, presented in a way that is accessible to non-experts. The zero-shot technique involves asking the system to perform a task without providing examples: the system uses its prior knowledge to carry it out. The few-shot technique, on the other hand, involves providing one or a few examples of the desired result before making the actual request: this is particularly useful for specific or non-standard tasks, where demonstration is more effective than abstract description. The Chain of Thought technique involves asking the system to show its reasoning step by step before providing the final answer: this produces more accurate and verifiable answers to complex problems, because making the process explicit reduces the likelihood of logical errors. Finally, prompt chaining is the technique of linking multiple prompts in sequence, using the output

from one as input for the next: this allows complex tasks to be broken down into manageable steps, yielding higher-quality results than a single all-encompassing request.

The section concludes with a guide on using and sharing results. Several tools offer direct export functionality to standard formats (Word, PDF, PowerPoint, Excel), and the Syllabus emphasizes the importance of knowing how to choose the most appropriate format based on the intended use.

Finally, the most important aspect highlighted in the last section is once again the ethical one: transparency regarding the use of AI (in many academic and professional contexts, it is proper to disclose when content has been produced or significantly assisted by AI) and the responsibility to verify (before sharing any information generated by AI, it is the user's duty to verify its accuracy, especially for factual data and citations). Sharing erroneous AI-generated content without verification is not merely a technical problem: it is an ethical responsibility and, in certain contexts, a legal one as well.

4. Conclusions: A Cultural Choice, Even Before a Technical One

With the ICDL Artificial Intelligence 2.0 Syllabus, AICA proposes something more ambitious than mere technical training: it aims to help educate people capable of asking the right questions, without blindly delegating, while ensuring that human responsibility for judgment, choice, and decision-making remains where it belongs. This goal is at once educational, civic, and, ultimately, political in the highest sense of the term: that of a community that takes care of itself in the face of the most rapid and profound technological transformations in modern history—as articulated, not least, in an extraordinary encyclical by Pope Leo XIV.

The journey into artificial intelligence does not end with a certification. On the contrary, this is precisely where it begins. The field of AI is evolving at an unprecedented pace: tools that seem cutting-edge today will become obsolete within a few years, and new capabilities and ethical challenges will continually emerge. The most valuable skill this program can offer is not knowledge of a fixed set of tools or techniques, but the ability to learn, adapt, and think critically in the face of change.

A citizen who understands the fundamental principles of AI, who can critically evaluate digital content, who knows their privacy rights, and who can use generative tools responsibly and competently: this is the person capable of navigating the present with intelligence and helping to shape the future with awareness. And this is precisely the goal of AI Literacy according to AICA: not a niche interest reserved for experts, but a democratic necessity for all citizens in the age of algorithms.

References

- AI Act – Regolamento (UE) 2024/1689 del Parlamento europeo e del Consiglio del 13 giugno 2024. Official Journal of the European Union, July 12, 2024.
- Stanford University, AI Index Report 2026.
- ICDL Foundation (2026). ICDL Artificial Intelligence Syllabus 2.0. Dublin: ICDL Foundation.
- Legge 23 settembre 2025, n. 132, "Disposizioni in materia di intelligenza artificiale." Gazzetta Ufficiale della Repubblica Italiana, No. 236, October 9, 2025.
- Ministry of Labour and Social Policies (2025). Decreto ministeriale n. 180 del 17 dicembre 2025 – Linee guida per l'implementazione dell'Intelligenza Artificiale nel mondo del lavoro.
- Encyclical Letter "Magnifica Humanitas" of the Holy Father Leo XIV, "Sulla Custodia Della Persona Umana Nel Tempo Dell'intelligenza Artificiale" [St. Peter's, May 15, 2026].
- UNESCO (2021). Recommendation on the Ethics of Artificial Intelligence. Paris: UNESCO.
- European Union (2022). DigComp 2.2: The Digital Competence Framework for Citizens. Luxembourg: Publications Office of the European Union.
- Fancello, A. (2026). ICDL – Intelligenza Artificiale. E-Book Syllabus 2.0. AICA – Associazione Italiana per l'Informatica ed il Calcolo Automatico. Milan.



Antonella Fancello

antonella.fancello@aica.it

AICA board member and trainer, she has produced training support materials on the content of the "ICDL Artificial Intelligence" Syllabus. She has been involved in digital transformation and public informatics for nearly 30 years, having studied under Prof. G. Duni, the first Italian theorist to propose that a file could have legal validity in an administrative proceeding. Formerly a professor of "E-Government – Digital Administration" in the Department of Law at the University of Sassari, and formerly a professor at the same university of "Digital Practices for Citizen Empowerment," she also collaborates with the University of Cagliari and numerous institutions, companies, and schools. She is a senior consultant for government agencies under the Presidency of the Council of Ministers, including FormezPA and the Agency for Digital Italy.

She is an expert on digital transformation at the ANCI Academy of Digital Municipalities.

A digital activist, blogger, and keynote speaker at numerous events of national significance, she has always been engaged in the legal and philosophical debate on rights and technologies, the digital divide, the digital agenda, the process of digitizing the Italian public administration, and the more recent impact of artificial intelligence systems on people's lives.

Building a “Better Rifle”

Palantir, Artificial Intelligence, and the Ethics of the “Technological Republic”

Christian Fiori, AICA

Abstract

In “The Technological Republic,” Alexander C. Karp, CEO of Palantir, and Nicholas W. Zamiska accuse Silicon Valley of betraying its original mission and call for a new alliance between the software industry, the state, and the defense sector. This article critically examines the book in light of Palantir’s current practices—from the ImmigrationOS platform for ICE to its partnership with the Israeli Defense Ministry, and its handling of UK NHS health data—as well as the rapid convergence between tech giants and security apparatuses, a trend whose origins can be traced to the “capitalism of fragmentation” described by Quinn Slobodian. Five ethical issues are identified—ranging from dual-use to algorithmic opacity to the blurring of corporate and national interests—leading to the conclusion that the true ethical challenge lies in determining which systems, with what safeguards, and under what democratic oversight we entrust to software.

Keywords: artificial intelligence, Palantir, technology ethics, dual-use, mass surveillance, algorithmic accountability, national security.

An Uncomfortable Manifesto

It is rare for the CEO of a technology company to write a political treatise. This is precisely what Alexander C. Karp, CEO of Palantir Technologies, and his “right-hand man” Nicholas W. Zamiska have done with *The Technological Republic* (2025; Italian ed. Mondadori): not a corporate memoir nor a popular essay on AI, but a declaration of intent regarding the relationship between the software industry, the state, and military power in the West.

The thesis is clear: Silicon Valley, born from the intertwining of scientific research and the war effort (from military electronics developed in the Bay Area to the Internet, which also came about thanks to DARPA funding), has betrayed its original calling, falling back on “superficial idealism” and unambitious consumer products, to the point of abdicating its duty to contribute to national defense. Hence the call to strengthen ties with governments and to put engineering talent back at the service of a “national project.” The argument deserves attention because Palantir does not merely theorize it but puts it into practice: an explicit philosophy of technology translates into software actually deployed on the front lines, on battlefields, and in healthcare systems. The ethical question is not abstract: what happens when the “engineering mindset” described in the book encounters mass surveillance, algorithmic target identification, and the sensitive data of millions of citizens?

The Authors’ Positions

The duty to defend. Karp and Zamiska argue that those who build technologies intended to permeate every aspect of collective life have a responsibility to take a stand and, in particular, to contribute to national security. Neutrality—understood as the refusal, common among many engineers, to work for military purposes—is not a virtue in their view, but a form of disguised disengagement.

The West and organized violence. The book takes up and makes its own a famous assessment by political scientist Samuel Huntington: the rise of the West was not determined by the superiority of its ideas or values, but by its superiority in the use of organized violence. For the authors, this is not a cause for shame but for awareness: the logic is that of deterrence—that is, maintaining a technological advantage sufficient to discourage adversaries and, thus, avoid war.

Against Nihilism. A substantial portion of the book is an indictment of what the authors describe as moral relativism and the “superficial nihilism” of contemporary cultural and technological elites. Slogans like *“Don’t be evil”* are interpreted as a symptom of a generation of brilliant technologists trained to identify evil and distance themselves from it, rather than to take on the more uncomfortable task of operating in an imperfect world. The proposed remedy is the rediscovery of a shared culture, a common identity, and shared civic rituals.

The engineer’s ethics. Finally, the authors champion a results-oriented culture: what matters is what works, the direct connection between those who create the software and those who use it, without the need for grandstanding. It is, in their view, an ethics of doing as opposed to an ethics of appearing.

It must be acknowledged that some of these arguments address real problems: the responsibility of those who design systems that affect the lives of millions of people is a serious ethical issue, and the claim that technology is “neutral” technology has long been disproved. The critical point, as we shall see, lies not in the questions the book poses, but in the answers it proposes and what these ultimately legitimize in practice.

From Theory to Practice: Palantir Today

Palantir, founded in 2003 with initial funding that included an investment from *the CIA*, builds data analytics platforms and AI systems for defense, intelligence, law enforcement, and the private sector. Its products—*Gotham* (military and intelligence), *Foundry* (corporate and public data), and the most recent *AIP*—integrate diverse sources into a single operational view. Three recent deployments illustrate the ethical stakes better than any argument.

Immigration: ImmigrationOS. In April 2025, the U.S. agency ICE awarded Palantir approximately $30 million for ImmigrationOS, a platform that provides “near real-time” visibility into the entire
“lifecycle” of immigration, from identification to tracking to deportation. It is based on an infrastructure that has linked Palantir to ICE since 2013 and is now valued at over $145 million. Civil rights advocates and lawmakers have denounced the aggregation of public and private data—including, according to some reports, databases from other federal departments—to profile targets. The problem is twofold: there is no independent audit or transparency regarding error rates, and “accuracy” depends on data that is notoriously flawed. When the data is incorrect, “precise” targeting hits the wrong person.

War: The Partnership with Israel. In January 2024, just a few months after the start of the war in Gaza, Palantir announced a “strategic partnership” with the Israeli Defense Ministry for “war-related” missions. Several investigations link its technologies—in particular *AIP*, designed to accelerate target identification—to the targeting systems used in the conflict; UN Special Rapporteur Francesca Albanese cited
“reasonable grounds” to believe that the platform was used in unlawful military operations, causing disproportionate civilian casualties. Palantir co-founder Peter Thiel’s response to those questioning him about the military use of AI was revealing: he wanted to “defer” the decision to Israel, rather than “judge” it from the outside. This is an explicit renunciation of independent moral judgment on the part of the tool’s provider.

Healthcare: The British NHS’s data platform. Not all cases are military. In 2023, NHS England awarded Palantir a 330 million-pound contract for the *Federated Data Platform*, designed to unify fragmented clinical data. The contract award, criticized for its timing and lack of transparency, ended up in litigation. In the spring of 2026, it emerged that external personnel, including Palantir staff, would have “unlimited” access to identifiable patient data prior to pseudonymization, and that *opt-out* rights would not apply under a directive that had never been published; an entire region (Greater Manchester) refused to participate. There are no targets or drones here, but the question is just as crucial: Who really controls citizens’ most intimate data once a private operator takes root at the heart of a public infrastructure?

After the Book: The Prophecy Accelerates

What makes *The Technological Republic* more than just a pamphlet is that, in the months that followed, the merger between the tech industry and the state that the book envisions came to pass at a surprising speed—and went far beyond Palantir.

In June 2025, the U.S. Army established Detachment 201: Executive Innovation Corps, enlisting four executives—including Shyam Sankar, Palantir's CTO, along with the chief technical officers of Meta and OpenAI—as reserve lieutenant colonels, without any conventional military background. It is the book's thesis brought to life: no longer companies selling software to the government, but managers in uniform who remain in their corporate positions. Sankar welcomed the initiative with words that could have come straight from Karp's book, calling the union between American innovation and military missions "the key to our 20th-century triumphs." Yet, as many have noted, these new officers would not shy away from their companies' business dealings with the Department of Defense: the intermingling celebrated in the book becomes a textbook case of conflict of interest.

The phenomenon is systemic. In July 2025, the Pentagon awarded contracts worth up to $200 million each to Google, OpenAI, Anthropic, and xAI: large AI models are entering the heart of military infrastructure. But it was in early 2026 that a case put to the test the idea—dear to the book—that "integrity comes at a price." In a $200 million deal, Anthropic set two non-negotiable conditions: no mass surveillance of citizens, and no autonomous weapons that act without human accountability for the use of force. The deal fell through. The administration ordered federal agencies to abandon their products, and the Pentagon labeled them a risk to national security. A few hours later, OpenAI announces its own agreement—with guarantees similar to those that had just cost its competitor dearly.

It is the perfect reversal of the book's thesis. Where Karp and Zamiska see moral cowardice in the refusal to serve the state, here a company sets clear ethical limits and discovers that the "price of integrity" can be retaliation by the government. The question is no longer whether companies should collaborate, but whether they can do so under conditions they deem acceptable: and it is precisely this moral judgment—the very thing the book would like to set aside—that returns to center stage.

Ethical Issues

1. The logic of the "better rifle" and dual-use. The title of a chapter, *"Building a Better Rifle*," sums up an instrumental ethic: the technician makes the tool more effective, while its use is left to others. This is exactly what Langdon Winner, in *Do Artifacts Have Politics?* (1980), called "technological somnambulism": building the object while pretending that its political consequences are someone else's problem. But the technologies at stake are *dual-use* and, above all, *enabling*: they do not merely carry out an order; they shape the space of

possible decisions. A system that generates lists to identify people to be targeted or expelled is not a more accurate rifle, but rather an apparatus that redefines who becomes the target.

2. The dilution of responsibility. "Deferring" moral judgment to the armed forces is the contemporary version of an old objection: "I merely provided the tool." Twentieth-century philosophy of technology (from Günther Anders to Hans Jonas) thoroughly dismantled this defense, showing how untenable it is, in the age of high-powered technologies, to separate the responsibility of the designer from that of the user. Paradoxically, it is the very same book that argues engineers *must* take a stand: but the stance it calls for is to side *with* the state and its objectives, not to assume responsibility for the effects of the tool.

3. Opacity, surveillance, and the rule of law. In both the migration and military spheres, the systems described operate without independent audits, without transparency regarding criteria and error rates, often functioning as veritable "black boxes." Here lies the most sensitive issue: the accountability of algorithms, the right to an explanation, and the ability to verify automated decisions are now recognized and, in part, codified principles, as demonstrated by the framework of the European AI Act and its classification of "high-risk" systems. An AI that affects personal freedom or the lives of others without adequate mechanisms for oversight and challenge stands in direct conflict with the rule of law, no matter how noble the stated purpose may be.

4. The blurring of corporate interests and national interests. The book celebrates the "blending" of commercial and national objectives; but when a private actor becomes the irreplaceable infrastructure for immigration, defense, or healthcare—to the point of being contracted as a "sole provider"—the issue becomes one of sovereignty and power. This is the tension that Lewis Mumford identified when he distinguished between a "democratic technique"—distributed and on a human scale—and an "authoritarian technique"—centralized and control-oriented—whose ultimate form is the megamachine: the union of organization, technical knowledge, and coercion. The convergence between large software companies and security apparatuses—even down to the symbolic gesture of executives in uniform—evokes a possible 21st-century megamachine. The authors themselves describe tech companies as small
"Nations," true city-states: a keen observation, but one that should be read as a warning, not as an ideal. It is precisely this warning that historian Quinn Slobodian addresses in his book *Crack-Up Capitalism*, in which he traces the history of market fundamentalists who, for decades, have been eroding the nation-state with free zones and enclaves where capital operates beyond the reach of democratic control. It is no coincidence that Slobodian's account opens with Peter Thiel, who wrote in 2009: "I no longer believe that freedom and democracy are compatible." Having set aside the idea of escaping the state, Thiel discovered that it is far more convenient to operate within it: contracts with ICE, the military, and the British National Health Service show how the utopia of a thousand nations ultimately gives way to the reality of large public contracts. Because in a democratic state

, one can dissent, vote, or leave; from the Palantir republic, which no one has elected, one cannot resign.

5. *Anti-nihilism* as a shortcut. The "West versus nihilism" dichotomy is rhetorically powerful, but it risks shutting down debate rather than fostering it. Turning every ethical doubt—from refusing to work on a weapons system to demanding data safeguards—into a sign of moral weakness or "woke culture" means delegitimizing the very debate the book claims to want to open. Ethics does not consist in the absence of convictions, but in the often difficult ability to subject them to criticism.

Why This Matters to Us

It would be intellectually dishonest to dismiss the book as mere advocacy: Karp's arguments have their own merit. Deterrence is real: if democracies cede their technological advantage, far worse regimes will seize it; and a blanket refusal to work with the government does not resolve the moral problem—it merely shifts it to less scrupulous suppliers, as the Maven parable illustrates. But this is precisely where the book's rhetorical deception lies: presenting us with a stark choice between neutral techno-utopianism and unconditional enlistment, while leaving the real ethical work in the shadows—namely, *which* systems, with *what* safeguards, under *what* democratic oversight, and for *what* verifiable purposes. On that front—the only one that matters—the book is reticent; and this reticence, on the part of those who already own the platforms, is not a distraction but a strategy.

The underlying unease that remains is that the question is no longer abstract; when a single private actor becomes the irreplaceable infrastructure for immigration, war, and healthcare across entire nations—and its executives wear uniforms without relinquishing their shares—the distance between "building the tool" and "governing people" has already vanished, while we are still debating whether it even exists. Codes of ethics (ACM, IEEE) are a compass, not an autopilot, and Luciano Floridi's *"soft ethics"* reminds us that the decisive question is not what the law permits, but what is right to do when everything is permitted. The "engineering mindset" celebrated in the book is valuable as long as we ask *what purpose* something that works serves, and *who* pays when it works too well; it ceases to be valuable the very moment we stop asking that question.

The Technological Republic thus has the unintended merit of voicing aloud what many prefer to keep silent: this is why it should be read—not to agree with its answers, but so we can no longer pretend not to hear the questions. The answer will not be yet another slogan, but the patient, stubborn, and unphotogenic construction of transparency, verifiability, and democratic oversight commensurate with the power we entrust to software. This is indeed a task for engineers—provided we do not discover too late that we have left it to someone else.

Bibliography:

- A. C. Karp, N. W. Zamiska, *La repubblica tecnologica. Come l’alleanza con la Silicon Valley plasmerà il futuro dell’Occidente*, Mondadori, Milan 2025 (orig. ed. *The Technological Republic*, 2025); foreword by Federico Rampini. *Quotations and arguments are paraphrased from the Italian edition.*
- ImmigrationOS / ICE: American Immigration Council, “ICE to Use ImmigrationOS by Palantir” (Aug. 2025); ACLU, “All the Ways Palantir Is Assisting Trump’s Removal Campaign” (2026); *Wired*, “ICE-Palantir ImmigrationOS” (2025). Federal ID Contract 70CTD022FR0000170, ~30 M.
- Partnership with Israel: Palantir press release, January 2024 (“strategic partnership”); Business & Human Rights Resource Centre; Drop Site News; report by UN Special Rapporteur Francesca Albanese, *From economy of occupation to economy of genocide* (2025). Statements by P. Thiel (July 2024) on “deferring” to Israel.
- NHS Federated Data Platform: contract worth ~330 M£ (2023); *Financial Times* and *The Register* (May 2026) on “unlimited” access to identifiable data; debate in the House of Commons (Hansard, April 2026); rejection by the Greater Manchester ICB.
- *Detachment 201 — Executive Innovation Corps*, U.S. Army (press release of June 13, 2025; among those appointed was Palantir CTO Shyam Sankar); Pentagon (CDAO) contracts worth up to $200 million each awarded to Google, OpenAI, Anthropic, and xAI (July 2025); the Anthropic/OpenAI–Pentagon controversy (February 2026): Anthropic’s “red lines” on mass surveillance and autonomous weapons, its designation as a security risk, and OpenAI’s agreement (NPR, *Wall Street Journal*, *Built In*).
- The earlier Project Maven/Google case (2018) on employees’ objections to military work.
- L. Winner, *“Do Artifacts Have Politics?*” (in *Daedalus*, 1980); *The Whale and the Reactor,* University of Chicago Press, 1986;
- G. Anders, *Die Antiquiertheit des Menschen* (C.H. Beck, 1956)*; Italian trans. L’uomo è antiquato,* Bollati Boringhieri, 2007;
- H. Jonas, *Das Prinzip Verantwortung* (Insel Verlag, 1979)*; Italian trans. Il principio responsabilità,* Einaudi, 1990;
- L. Mumford, *Technics and Civilization* (Harcourt, 1934)*, Italian trans. Tecnica e cultura,* Donzelli, 2013; *The Myth of the Machine* (Harcourt, 1967–70), *Italian trans. Il mito della macchina,* Donzelli, 2019;
- L. Floridi, *Etica dell’intelligenza artificiale*, Raffaello Cortina Editore; English ed. *The Ethics of Artificial Intelligence*, Oxford University Press, 2023.
- Q. Slobodian, *Il capitalismo della frammentazione. Gli integralisti del mercato e il sogno di un mondo senza democrazia*, trans. by G. Garbellini, Einaudi, Turin 2023 (orig. ed. *Crack-Up Capitalism: Market Radicals and the Dream of a World Without Democracy*, 2023).

- Regulatory Framework: EU Regulation on Artificial Intelligence (*AI Act*), risk categories, and transparency obligations. Codes of Ethics: ACM Code of Ethics; IEEE Code of Ethics.

Editor's note: Sensitive topic. Statements regarding the war context include allegations and assessments by third parties (UN, NGOs, media outlets) and do not constitute definitive judicial findings.



Christian Fiori

christian3.fiori@gmail.com

Christian Fiori is a digital educator and researcher in the field of critical studies on emerging technologies, specializing in the ethical and social implications of their development and use, as well as their impacts on the education and health of young people. In addition to his technical training in Electronic and Telecommunications Engineering at the University of Bologna and his experience in quantum computing with CINECA, he has studied the Philosophy of Information at the University of Urbino, earned a Master's degree in the Philosophy of Digital Technology and Artificial Intelligence at the University of Udine, and obtained a diploma in Advanced Specialization in Ethics and AI at the Pontifical Antonianum University in Rome. He is currently *a digital educator and ethicist* at AICA, serves on the editorial board of AICA's two journals, BRICKS and MONDO DIGITALE, and is a member of the Organizing Committee for the DIDAMATICA event.

AICA's International Activities

Marina Cabrini, CEPIS - Member of the Board

Abstract

On May 27, 28, and 29, the annual in-person meeting of the CEPIS Council took place in Warsaw. As has been the case for several years, meetings of the various working groups were held on this occasion to take stock of the situation regarding various issues. This article outlines the structure of CEPIS and its working groups to provide an overview of its activities.

Keywords: AICA, CEPIS, LSI, Women in ICT, DiversIT, ICODSIP, Skill[set]

1. Introduction

This column aims to provide insight into the activities that AICA carries out at the international level, primarily through the Council of European Professional Informatics Societies (CEPIS), of which AICA was a founding member in 1989 along with eight other associations, and which now has 28 active members. This first article will describe CEPIS's structure and how it operates within Europe, while future issues of Mondo Digitale will delve deeper into the activities of the various working groups operating within CEPIS.

2. CEPIS

CEPIS stands for the Council of European Professional Informatics Societies and was founded back in 1989 with the aim of promoting and enhancing high standards among IT professionals in Europe. It is a European network representing more than 200,000 ICT professionals and providing a coordinated voice for these professionals at the European level. It also participates in European projects in partnership with companies and universities.
Its governing body consists of a Council of representatives from all member associations and a Board of six directors plus the President, who are elected by the Council. The Council meets twice

year to deliberate on CEPIS's strategies and financial management. The Board, on the other hand, serves as a guiding body for CEPIS and all its activities.

The plenary meetings of the CEPIS Council are held every year in May and November. Generally, only one of these meetings takes place in person, and it is usually the May meeting, to which both representatives of CEPIS member associations and all members of the working groups are invited.

This year, the May meeting was held in Warsaw, coinciding with the 45th anniversary of the founding of the Polish Informatics Association (Polskie Towarzystwo Informatyczne, PTI). It was three days of intense activity, featuring numerous meetings—both of individual working groups and in plenary sessions—as well as an official welcome reception for all attendees. A total of more than 50 experts from across Europe attended these sessions. They had the opportunity to meet formally during the meetings and informally during the various breaks, establishing new collaborations among CEPIS member associations on issues of inclusion, digital skills, ICT professionals, and European policies.

(https://cepis.org/cepis-meetings-warsaw/)



3. CEPIS Working Groups

Several working groups are active within CEPIS, addressing various areas of ICT. Each group includes experts from different sectors, selected by CEPIS member associations from among their own members, who work together to organize various activities and initiatives. AICA itself is represented in all these working groups by some of its members, who were nominated and selected as experts in a wide range of digital fields—from IT training to legal and security issues, to the

presence of women in the ICT sector. Each group meets at regular intervals, in person or online depending on its needs and preferences, to address issues that arise in the various fields and decide on a course of action.

The groups active within CEPIS are as follows:

Informatics Education

This group is dedicated to promoting computer science education in schools, with the aim of providing students with the tools they need to navigate today's increasingly digitized society. It is a founding member of *the Informatics for All Coalition*.
(https://cepis.org/informatics-education-group/)

Women in ICT

This group aims to promote women's participation in ICT by advancing gender equality and encouraging the next generation to pursue studies and careers in the ICT field. It is one of CEPIS's most active groups and has for several years been offering *DiversIT*, a *roadmap* that organizations and companies can use to assess the level of integration of women working in ICT professions within their organizations.
(https://cepis.org/women-in-ict/ and https://cepis.org/diversit-charter/)

Legal and Security Issues

The Legal and Security Issues (LSI) group is composed of cybersecurity experts who monitor legislation and proposed laws related to ICT and cybersecurity at the European level. This group's activities also include formulating CEPIS's official positions vis-à-vis various European institutions through position papers and direct testimony before the European Parliament. It participates directly in the work of related associations such as ECSO and ENISA. (https://cepis.org/legal-and-security-issues/)

Digital Skills Policy

This group analyzes documents related to digital skills policies issued by European institutions and promotes digital skills and certifications for workers. (https://cepis.org/digital-skills-policy-group/)

Information Society

This group was established to keep CEPIS member associations informed about current ICT topics, such as the Digital Single Market, accessibility, and the impact of artificial intelligence on society, both through a series of questionnaires sent to CEPIS member associations and through the preparation of a series of booklets. (https://cepis.org/information-society/)

Research, Development, and Innovation

This group aims to promote and analyze research, development, and innovation in ICT. Furthermore, it serves as a platform to connect CEPIS member associations with potential partners so they can participate in European-level projects.

If an Italian professional wishes to join one of these groups, they may do so, provided they are an AICA member, as working group members are nominated to CEPIS by their respective IT associations.

4. Upcoming CEPIS Activities

CEPIS organizes events related to the world of ICT in collaboration with its national associations upon their request. In particular, for the third consecutive year, CEPIS is organizing the Skill[set] conference (formerly known as ICODSIP), dedicated to the exchange of ideas and information on digital skills. This year, the conference—organized in collaboration with the German Informatics Association (Gesellschaft für Informatik, GI)—will take place in Berlin on December 3 and 4, 2026, and will focus, as always, on digital skills, ICT professionalism, and computer science education. The call for papers is now open: you can find all the relevant details here, and AICA members are invited to submit their work so they can present the activities carried out in Italy at a conference with a European scope. (https://skillset.gi.de/ and https://cepis.org/save-the-date-skillset-conference-2026/)

Bibliography

CEPIS Council in Warsaw: https://cepis.org/cepis-meetings-warsaw/
Informatics Education Group: https://cepis.org/informatics-education-group/
Women in ICT Group: https://cepis.org/women-in-ict/
DiversIT Charter: https://cepis.org/diversit-charter/

LSI Group: https://cepis.org/legal-and-security-issues/
Digital Skills Policy Group: https://cepis.org/digital-skills-policy-group/
Information Society Group: https://cepis.org/information-society/
CEPIS and GI Skill[set] Conference: https://cepis.org/save-the-date-skillset-conference-2026/
CEPIS and GI Skill[set] Conference: https://skillset.gi.de/
CEPIS and GI Skill[set] Conference – Call for Papers: https://skillset.gi.de/call-for-papers



Marina Cabrini

mcabrini@mgeng.com

Marina Cabrini is an IT consultant who has long been involved in the development of numerous IT certifications, both in Italy and internationally, and is a member of several working groups. She also works for AICA on the Italian localization of materials related to these certifications when they are not originally developed in Italian, and is responsible for ensuring the quality of the training materials and tests used for certifications in Italy. She is convinced of the importance of women's participation in STEM fields and currently serves as vice chair of the CEPIS Women in ICT group, in addition to being a member of the CEPIS Board of Directors.

Digital Ecology

Giovanni Franza, AICA

Abstract

Digital Ecology is a column that aims to explore the "dark side" of computer science—its impact on our ecosystem—starting with what it consumes (land, water, energy) and extending to what it produces (GHG, e-waste)—not from a Luddite perspective of rejection, but with the aim of raising awareness that, even in ICT, there are no "free lunches" and that we must also consider the costs, not just the—albeit splendid—opportunities offered to us.

Keywords: ICT, Green Software, CO2 equivalent, e-waste, GHG, DNSH

1. Introduction

This issue of Mondo Digitale marks the launch of a new column, provocatively titled "Digital Ecology," to illustrate the connections between ICT and the environment. But, let's be clear: the aim is not to explore what is already being done (well) using ICT to optimize resource use, but rather to shed light on the "dark side" of ICT—the costs and energy consumption associated with building and maintaining ICT infrastructure.

2. The Problem

Although ICT is enormously pervasive today (consider, on the one hand, the impact of AI, and on the other, that of the IoT—starting with home automation), there is little awareness of the costs associated with this pervasiveness. The purpose of this column is precisely to shed light on these aspects, providing insights for interpreting and understanding the phenomenon.

There are several reasons why this phenomenon is little known:

- The pervasiveness of ICT means that IT tools are increasingly used by non-technical staff, who (understandably) are unaware of the technical aspects of how they work but simply use them according to the manufacturers' instructions;

- The proliferation of devices perceived as simple, for general use, (relatively) inexpensive, and small in size (primarily smartphones, but also laptops) obscures the extent to which these devices depend, for their operation, on resources that are invisible because they are distant yet crucial—such as telecommunications networks and the large data centers that power the cloud;
- The prevalence of flat-rate plans, not tied to usage, leads to a failure to assess the amount of resources consumed and to compulsive use of these devices—a situation, incidentally, that is encouraged by the service providers themselves;
- the difficulty of tracking the environmental impact of ICT "from cradle to grave": the excavation of staggering amounts of material to extract essential components such as rare earth elements; the amount of energy required to manufacture and ship devices; the challenge of accurately assessing electricity consumption during use; and the problem of end-of-life disposal of the devices themselves (e-waste).

3. A Perspective

It is particularly important to understand these issues, given that we live in a finite world with limited resources. The consumption of these resources depends not only on the major players (known as MAMAA—Meta, Apple, Microsoft, Amazon, Alphabet) or on governments, but, to a large extent, on users' habits—that is, on all of us.

It is not easy to determine what constitutes good habits if we do not understand the effects of our actions. And to understand these effects, it is important to have models of how ICT systems function. This is precisely the purpose of this column: to provide guidance and models related to the ecological footprint of ICT and its use.

This footprint is not limited to electricity consumption during the usage phase; it begins with the consumption of materials for manufacturing and ends with the disposal of devices: understanding these aspects is not easy, nor is quantifying them (this is referred to as LCA—life cycle assessment).

To simplify, we can say there are three main categories:

- the consumption of resources for the production and use of ICT devices and supporting infrastructure;
- *the* production of greenhouse gases (*GHG*);
- disposal (e-waste).

3.1 Resource Consumption

This manifests itself in several specific ways:

- mineral extraction to obtain the materials needed for ICT, notably rare earth elements;

- land use for the siting of data centers;
- water consumption associated with the cooling needs of these data centers;
- electricity consumption throughout all stages of an ICT device's life cycle.

3.2 Emissions of Climate-Altering Gases

Also known as GHGs (greenhouse gases); they are measured in metric tons of CO2 equivalents:

- emissions occur at every stage of an ICT device's life cycle because they are linked to energy production—any form of energy, regardless of its intended use—whether it involves mineral extraction or powering servers;
- Greenhouse gases have a cumulative effect: their presence traps solar radiation and thus leads to a continuous rise in the average surface temperature; therefore, an increase in these gases is not necessary—even simply maintaining them at current levels results in ongoing global warming.

3.3 Disposal

At this stage, we're talking about e-waste, or electronic waste:

- Disposing of end-of-life components still results in a waste of the materials they contain—materials whose extraction and refining were costly and which are now mostly thrown away, when we should be finding ways to extract them and put them back into circulation (*secondary raw materials*);
- land is consumed due to the size of the landfills where these materials end up—land that, in many cases, is not located in wealthy countries but in the Global South;
- greenhouse gas emissions result from the energy consumed in transporting and treating this waste;
- there is pollution caused by both improper treatment and storage;
- there are all the social issues related to *eco-mafias* and the shipment of this waste to countries in *the Global South* for cheap processing that is harmful to the environment and the health of the people involved.

These are all crucial aspects of the problem, regarding which awareness should be raised to encourage the adoption of best practices in the use of information technology.

4. Raising Awareness

Precisely because of its position, AICA—which has been committed to the development of IT skills since its founding—has decided to address these issues in *Mondo Digitale* through a column designed to stimulate interest and raise awareness of the problem.

Future installments of this column will explore various aspects of this issue:

- an overview of the problem, illustrating the scale of the issue;
- an introduction to the LCA (Life Cycle Assessment) methodology;
- the impacts of smartphone use;
- software-induced energy consumption and green software;
- the concept of DNSH (Do No Significant Harm).

4.1 Disposal

In the next issue, we will explore the issues involved in the various phases to provide a rough estimate of the problem; we will offer an overview of how the various phases are interconnected and, above all, the magnitudes involved, with the aim of giving all readers a basic understanding of the problem.

4.2 LCA (Life Cycle Assessment)

We will provide an introduction to LCA (Life Cycle Assessment), explaining the techniques used to estimate the environmental impact of manufacturing and using an IT device. The goal is to help readers understand the complexity of these products and the process of assessing their environmental impact, while illustrating the various techniques used to make these estimates.

4.3 Impacts of Smartphone Use

The impacts of using devices such as smartphones will be illustrated, starting with where energy consumption actually occurs (telecommunications networks and data centers that provide cloud services) and concluding with the impact of how and for how long the devices are used; In this section, we will discuss the impact of “flat-rate” plans, with the aim of providing readers with a framework for understanding the impact of their own habits.

4.4 Software-Induced Energy Consumption and Green Software

For a more technical audience—including developers as well as those simply curious—we will explain how to address the typically non-functional requirement of energy consumption related to application use (software-induced energy consumption), introducing the concept of *green software* and the techniques for developing it.

4.5 The Concept of DNSH (Do No Significant Harm)

To also address the regulatory aspect, we will introduce the concept of DNSH (Do No Significant Harm) in the creation of new facilities and services, and discuss how one approach might involve the reuse of existing devices.

Bibliography

ISO (2006a). Gestione ambientale - Valutazione del ciclo di vita - Principi e quadro di riferimento (ISO 14040:2006). International Organization for Standardization, Geneva.

ISO (2006b). Gestione ambientale - Valutazione del ciclo di vita - Requisiti e linee guida (ISO 14044:2006). International Organization for Standardization, Geneva.

ITU-T (2015). Metodologia per la valutazione dell'impatto ambientale delle tecnologie dell'informazione e della comunicazione (Recommendation L.1410). International Telecommunication Union, Geneva.

Kern, E., Dick, M., Naumann, S., Santarius, T. (2013). Green Software and Green Software Engineering – Definitions, Measurements, and Quality Aspects. In: Proceedings of the First International Conference on Information and Communication Technologies for Sustainability (ICT4S), ETH Zurich, Zurich.

Sissa, G. (2024). Le emissioni segrete. L'impatto ambientale dell'universo digitale. Il Mulino, Bologna.

European Commission (2021). Orientamenti tecnici sull'applicazione del principio «non arrecare un danno significativo» (DNSH) a norma del regolamento sul dispositivo per la ripresa e la resilienza (2021/C 58/01). Official Journal of the European Union, Brussels.

European Parliament and Council of the European Union (2020). Regolamento (UE) 2020/852 relativo all'istituzione di un quadro che favorisce gli investimenti sostenibili e che modifica il regolamento (UE) 2019/2088 (Regolamento Tassonomia). Official Journal of the European Union, L 198/13, Article 17, Brussels.



Giovanni Franza

gfranza@mgeng.com

Giovanni Franza, PhD, is a computer scientist with over forty years of experience in communication systems and information technology. His main areas of interest include Linux systems, networks, cybersecurity, and the sustainability of ICT infrastructure. For years, he has been analyzing the environmental impact of digital technologies, promoting energy-efficient practices, equipment reuse, and the use of open-source software to reduce the IT sector's ecological footprint.

LIFO avant la lettre

Pierangelo Indolfi, Polytechnic University of Bari

Abstract

The abstract data structure known as the stack is traditionally introduced in programming courses through its formal definition, followed by application examples. This article proposes a reverse approach: to show that the LIFO principle emerges spontaneously in fields of knowledge far removed from computer science, long before anyone thought to formalize it. Archaeology, neurobiology, psychoanalysis, network protocols, layered operating systems, and decontamination protocols in high-biological-risk areas: each of these contexts exhibits the same underlying structure, one in which the order of placement determines the reverse order of access. The Tower of Hanoi occupies a special place in this gallery: its case reveals that the stack is not a tool chosen to solve the problem, but the deep structure of the problem itself. The article discusses the educational implications of this approach, with particular reference to courses in Medical Systems Engineering, where the application domain provides motivating contexts of extraordinary concreteness. The underlying thesis is epistemological: formalizing a data structure is not reinventing the wheel, but recognizing a face that was already there, and returning it to the world in the form of a tool capable of producing new things.

Keywords: data structures, stack, LIFO, computer science education, cognitive metaphors, medical systems engineering, recursion, Tower of Hanoi

1. Introduction

Presenting an abstract data structure to engineering students in the standard order—formal definition followed by application examples—produces a predictable effect: the student grasps the mechanism without understanding its meaning. They know what a stack is, they know it operates according to the LIFO principle, and they know how to use the push and pop commands. What they don't understand is why that structure exists, or what real-world need prompted someone to isolate it and give it a name.

Data structures originate from the observation of recurring phenomena waiting to be formalized, while textbooks only mention them much later. The stack, in particular, is a model that emerges spontaneously in fields very distant from one another, long before computer science existed as

a discipline. Archaeology encounters it every time an excavation brings to light overlapping layers of civilizations. Neurobiology recognizes it in the architecture of the central nervous system (MacLean, 1990). Psychoanalysis uses it as a foundational metaphor for the o (Freud, 1899). Network protocol theory formalizes it as the encapsulation principle (Tanenbaum, 2003).

Showing students all of this before they even write a single line of code is an act of intellectual honesty regarding the nature of ideas, as well as a more effective pedagogical choice (Papert, 1980).

2. The Layering Model

The defining property of a stack—namely, that the last element added is the first to be retrieved—is an inevitable consequence of a fundamental physical process: layering. Whenever objects, events, or structures accumulate over time without mixing, the result is a sequence in which the order of addition determines the order of access, and the most recent element occupies the most exposed position. This process is so pervasive that it spans scales and fields of knowledge that are difficult to compare. The model is nevertheless recognizable everywhere, with the same geometry: something is added on top, or outside, or in front, and to reach what lies deep within, one must work backward through everything that had been deposited layer by layer.

When a student first encounters the formal definition of a stack, they are given a name and a notation for something they have already encountered countless times, in contexts they would never have thought to connect. The teacher's task is to make visible a structure already present in the student's experience, one that is simply waiting to be recognized (Polya, 1945).

The following sections analyze some of the most significant manifestations of this model, chosen for the diversity of the fields they span and for the light each sheds on a different aspect of the structure we are describing.

3. Metaphors, One by One

3.1 Archaeology

When an archaeologist begins an excavation, time unfolds in reverse order. The artifacts closest to the surface belong to the most recent civilizations; those deeper down, to the most ancient ones. Each layer is sealed by the one above it, and access to any of them requires the removal of all those deposited subsequently. It is not possible to extract a Roman artifact without first passing through the medieval, Renaissance, and modern layers that cover it.

This sequence is not a methodological choice made by the archaeologist: it is a physical property of stratified soil. The order of deposition, determined by time, becomes the reverse order of access. It is the very definition of LIFO, written into the soil millennia before anyone thought to formalize it.

Archaeology also offers a further insight that textbooks rarely make explicit: stratification does not destroy information, but preserves it in the order in which it was produced. Each layer is a snapshot of the moment it was deposited. The sequence of excavation, when conducted correctly, is a time machine that travels backward through history.

3.2 The Central Nervous System

In the 1960s, American neurologist Paul MacLean proposed the triune brain model: a structure in which three overlapping systems—the reptilian brain, the limbic system, and the neocortex—reflect three distinct stages in the evolution of vertebrates (MacLean, 1990). The oldest structures, which regulate basic vital functions, are located at the center and deep within. The more recent ones, which house higher cognitive functions, envelop them from the outside.

MacLean's model has been partially revised by contemporary neurobiology, which has questioned its rigidity, but it retains its descriptive validity as a metaphor for evolutionary accumulation. What came first is at the bottom; what came later is at the top. To reach the archaic structures, one must pass through—at least anatomically—the more recent ones. Evolutionary stratification produces a geometry that any computer scientist would immediately recognize.

3.3 Freudian Stratification

Sigmund Freud built his model of the psyche on an explicit spatial metaphor: repressed contents sink into the unconscious, buried beneath successive layers of psychic life (Freud, 1899). Analytical work is a process of excavation, and Freud himself repeatedly used the archaeological analogy to describe it. The therapist removes, layer by layer, the most recent psychic material to reach the deep nuclei, which are also the oldest and most resistant.

The direction is the same as that found in archaeology and in the triune brain: time flows downward, and moving back toward the origin means retracing everything that had been deposited layer upon layer above it. Freud was not describing a stack, but he used the same geometry.

A Neapolitan proverb, "'A capa è 'na sfoglia 'e cipolla"—the head is a layer of onion—captures the same insight with popular conciseness: the mind consists of concentric layers, each of which conceals the one beneath it, and one cannot reach the center without passing through the outer layers. The proverb does not need Freud to be true, and it is probably older than he is. It is also, as we shall see, a surprisingly accurate description of certain computer systems.

3.4 The ISO-OSI Model and Encapsulation

The ISO-OSI model describes the architecture of communication networks as a sequence of seven layers, each of which adds a wrapper to the data as it moves down the stack and removes it as it moves back up (Tanenbaum, 2003). During transmission, each layer adds its own header; during reception, it removes them in reverse order, until the original data is delivered to the application.

In this case, the LIFO principle does not refer to the temporal order of storage, but rather to the logical order of encapsulation. The symmetry between transmission and reception is perfect, and it is not coincidental: it is a structural property of any layered system in which each layer is unaware of what lies above and below it.

3.5 Virtual Machines and the Onion Structure

The onion metaphor, evoked by the Neapolitan proverb, proves surprisingly apt for describing the architecture of layered systems. Andrew Tanenbaum, in describing the structure of operating systems, uses the term “virtual machine” in a precise and more general sense than is commonly understood: each level of the system exposes an interface to the level above it that abstracts everything below it, creating the illusion of a machine with specific capabilities (Tanenbaum, 2013). The user at level N does not know—and does not need to know—how level N-1 actually implements what is being offered to them. In this sense, each level is a virtual machine for the level above it.

Just as with the layers of an onion and those of the Freudian psyche, it is not possible to reach the center without passing through the outer layers. The difference from a biological onion, however, is significant: the layers of an onion grow all at once, while those of a layered system are defined sequentially, and this sequence determines the order in which they can be traversed. The concentric geometry of the onion combines with the temporality of the stack.

It is worth noting that the term “virtual machine” is also used, in a narrower and more recent technological sense, to refer to software such as VMware or VirtualBox, which emulate an entire hardware environment. This is a case of homonymy: the concept is different, even though the word is the same.

3.6 The Narrow Elevator

Among all the examples collected, the elevator is the only one in which LIFO behavior does not stem from an inevitable physical or biological property, but from a rational and deliberate choice. In a small elevator, where it is not possible to change positions once you have entered, it makes sense for those who need to get off on the upper floors to enter first: this way, when the time comes, those who need to get off first are already near the door.

No physical law dictates this rule. If people enter in the wrong order, the elevator still works, but the person who needs to get off first ends up trapped at the bottom. In this case, LIFO is a norm of collective behavior that optimizes the use of a shared resource. It is worth noting that this is also the situation most closely resembling certain scheduling problems in operating systems, where the order of process execution is deliberately chosen to maximize efficiency (Tanenbaum, 2015).

3.7 The Decontamination Protocol in a Cleanroom

In the medical field, personnel working in areas with high biological risk wear a series of personal protective equipment (PPE) items in a precise, predetermined order: first, items that come into contact with the body; then, those that protect the most exposed entry points; and finally, those that protect the mucous membranes. Each item adds a layer of protection to the previous one.

The removal of PPE at the end of a procedure must strictly follow the reverse order. The reason is biological: the outer surface of each piece of PPE is potentially contaminated, and removing an inner layer before an outer one exposes the hands or face to contamination from the outer layer during subsequent maneuvers. Documented cases of Ebola infection among healthcare workers have been attributed precisely to errors in the removal sequence (WHO, 2014).

Among all the examples collected in this section, the decontamination protocol is the one in which the LIFO property ceases to be a convenience or an emergent property and becomes a biosafety constraint. Violating the removal order does not result in a logical error detected by a compiler—it poses a risk to the operator's life.

4. The Tower of Hanoi as a Special Case

Among all the examples collected, the Tower of Hanoi occupies a special position. In the other cases, the LIFO model emerges in phenomena that exist independently of computer science, and computer science recognizes and formalizes them. In the Tower of Hanoi, the relationship is reversed: the problem was deliberately constructed as a mathematical puzzle, but its solution reveals that the stack is the underlying structure of the problem itself, not a tool chosen for the convenience of implementation.

The puzzle, devised by the French mathematician Édouard Lucas in 1883 (Lucas, 1883), consists of transferring a tower of disks of decreasing diameter from a starting peg to a destination peg, using an auxiliary peg, with the constraint that a disk of larger diameter can never be placed on top of one of smaller diameter. The minimum number of moves required to transfer a tower of n disks is $2n - 1$.

The recursive solution is elegant and well-known: to move n disks from peg A to peg C using B as an auxiliary peg, first move the top n-1 disks from A to B, then move the largest disk from A to C, and finally move the remaining n-1 disks from B to C. Recursion is applied to each of the two subproblems until the base case is reached—that is, a single disk, which is moved directly (Cormen et al., 2009).

What this description does not make immediately apparent is that each recursive call adds a frame to the call stack and that the execution sequence strictly follows the LIFO principle: the last call added to the stack is the first to be resolved. A disk cannot be moved until all those stacked on top of it have been removed—this is the very definition of a stack, applied to physical objects made of wood or ivory.

The iterative solution to the Tower of Hanoi problem requires the use of an explicit stack to simulate what recursion does implicitly. Eliminating recursion does not eliminate the stack: it simply makes it visible. This is a concrete demonstration of a principle that holds in general: recursion and the stack are two sides of the same structure, and every recursive algorithm contains a stack, regardless of whether or not it appears in the code (Knuth, 1968).

5. Educational Implications

5.1 Order of Presentation

Established practice dictates that the stack be introduced after the student has already gained some familiarity with programming: the structure is defined, its fundamental operations are described, and some applications are illustrated. This order has its own logic, but students who follow this sequence risk learning the mechanism without grasping its meaning: they know how to use `push` and `pop` without having a solid mental image to hold onto, and they know the structure without understanding why it exists.

An alternative approach is to reverse the sequence. The metaphors discussed in the previous sections make this approach feasible even for students who have not yet been introduced to the formal concept of a stack, since they draw on the experiences and knowledge these students already possess.

5.2 Classroom Discussion as a Starting Point

Before launching a development environment, students are asked to list situations from daily life or their academic studies in which something accumulates in layers and must be retrieved in reverse order. The ensuing discussion almost invariably yields examples that cover the entire spectrum explored in this article: some mention archaeology, others the network layers studied in a previous course, and still others their own backpacks. The instructor gathers and organizes the insights the students already possess. Only at this point is the stack defined, the LIFO property formalized, and the code written. The formalism is presented as the culmination of a journey the student has already undertaken, not as a starting point requiring an effort of abstraction from scratch.

This sequence offers an additional advantage when it comes to the Tower of Hanoi. When the puzzle is presented as an exercise after the student has already internalized the LIFO pattern through metaphors,

the recursive solution does not appear as a mathematical trick to be memorized. Instead, it appears as the natural transposition of something the student recognizes: to move the largest disk, one must first remove all those above it—exactly as in archaeology, exactly as when removing PPE in a sterile room. Recursion becomes understandable because the pattern it implements is already familiar (Polya, 1945).

5.3 The Application Domain as a Motivating Context: The Case of Medical Systems

Everyday metaphors have immediate evocative power, but they risk remaining mere illustrations if they are not anchored to the professional field the course is aimed at. The decontamination protocol described in Section 3.7 offers this opportunity in a particularly direct way: for a student of medical systems engineering, this is not simply another example, but a context in which they might one day find themselves designing or certifying an actual computer system (WHO, 2014).

Proposing the modeling of that protocol as a programming exercise—with a PPE class, a VestitionProtocol class that manages a LinkedList used as a stack, and a verifyRemoval method that flags sequence violations—means asking the student to translate into code something they have already understood at the clinical level. The data structure is recognized as the natural tool for modeling a real biological constraint, and the stack ceases to be a convention to be accepted and instead becomes the answer to a question posed by the domain itself.

6. Conclusion

William Butler Yeats, in a short and dense poem, imagines a woman applying makeup in front of a mirror not out of vanity, but for recognition: she is searching, beyond the conventions that time has deposited, for the face she had before the world was created (Yeats, 1933). The gesture is the opposite of what it seems: it does not add, but digs.

Anyone who formalizes a data structure performs a similar act. The stack was not invented by the person who wrote the first compiler: it was already present in the soil that the archaeologist excavates, in the architecture of the nervous system, in the geometry of the psyche, in the rigorous sequence with which a healthcare worker removes protective gear in a sterile room. Computer science discovered it, recognized it for what it was, and gave it a name and an interface. It found the face that was already there before the world was created.

This does not reduce computer science to the role of reinventing the wheel. Formalization is not a passive discovery: it is a creative act with consequences that the original model did not contain. Stratified soil has existed for billions of years, but no geologist had ever thought to use its logic to manage the recursive calls of an algorithm (Knuth, 1968). The Freudian psyche had the geometry of a stack, but Freud could not

have been able to write an operating system. The added value of computer science lies precisely in this transfer: taking an ancient model, transferring it to a new context, and discovering that there it produces results unthinkable in its original context.

Yeats's woman, after rediscovering her original face, does not keep it to herself: she wants her beloved to love that thing—the one that existed before the world was created. It is an act of restitution. Computer science does something similar: it takes a model that reality had already produced, formalizes it, and returns it to the world in the form of a tool. With that tool, we build compilers, operating systems, network protocols, and software for managing cleanrooms. The model was ancient; the things it makes possible are new. The question of how many other data structures have a natural history similar to that of the stack remains open—and it's worth leaving it open. The queue, the list, the tree, the graph: each of them likely has its own gallery of pre-computational manifestations, its own set of domains in which the model emerges even before it is named. Exploring them using the same method applied in this article constitutes an educational research program that this text does not exhaust but, hopefully, helps to launch.

Bibliography

Cormen, T.H., Leiserson, C.E., Rivest, R.L., Stein, C. (2009). Introduction to Algorithms, 3rd ed. MIT Press, Cambridge (MA).

Freud, S. (1899). Die Traumdeutung. Franz Deuticke, Leipzig–Vienna. [Engl. transl.: The Interpretation of Dreams, Basic Books, New York, 1955.]

Knuth, D.E. (1968). The Art of Computer Programming, Vol. 1: Fundamental Algorithms. Addison-Wesley, Reading (MA).

Lucas, É. (1883). Récréations Mathématiques. Gauthier-Villars, Paris.

MacLean, P.D. (1990). The Triune Brain in Evolution: Role in Paleocerebral Functions. Plenum Press, New York.

Papert, S. (1980). Mindstorms: Children, Computers, and Powerful Ideas. Basic Books, New York.

Polya, G. (1945). How to Solve It. Princeton University Press, Princeton (NJ).

Tanenbaum, A.S. (2003). Computer Networks, 4th ed. Prentice Hall, Upper Saddle River (NJ).

Tanenbaum, A.S. (2013). Structured Computer Organization, 6th ed. Pearson, Upper Saddle River (NJ).

Tanenbaum, A.S., Bos, H. (2015). Modern Operating Systems, 4th ed. Pearson, Upper Saddle River (NJ).

WHO — World Health Organization (2014). Guidance on Personal Protective Equipment in the Context of Filovirus Disease Outbreak Response. WHO Press, Geneva.

Yeats, W.B. (1933). The Winding Stair and Other Poems. Macmillan, London.



Pierangelo Indolfi

indolfi@gmail.com

Pierangelo Indolfi teaches the Software Design Lab course for Medical Systems Engineering students at the Polytechnic University of Bari. With over thirty years of experience in ICT education at technical high schools, he has served as a Cisco Networking Academy instructor and an ICDL supervisor, and holds certifications in Arduino, Raspberry Pi, and digital competencies (DigCompEdu C2). He is the recipient of the 2023 MIT App Inventor Foundation Teacher Excellence Award. His research interests include computer science education, abstract data structures, and the epistemological foundations of computer science education.