



*Fondazione Centesimus Annus
Pro Pontifice*

CAPP France

Artificial Intelligence

in the Light of the Social Teaching of the Church

INTRODUCTION

For most people, life is understood, at least in part, in terms of necessity – that is, satisfying needs and filling gaps: lack of food, lack of money, lack of health, and other such deficiencies. This shared state of necessity among people can lead to the development of a social life that fosters solidarity and concern for the vulnerable, but it can also give rise to violence in all its forms, often fueled by innovation, in the name of a Darwinian “struggle for life.”

Human creativity, at its best, strives to make this necessity less burdensome for the greatest number of people, through the use of technology. The Industrial Revolution of the 19th century marked a significant milestone in this regard. Since then, the harnessing of energy and physical force, coupled with high-quality education provided on a mass scale, has led to a significant improvement in the living conditions of a large portion of the population.

This rapid evolution, however, has revealed major social and spiritual flaws, stemming from attempts by a small number to dominate and enslave entire peoples in the name of force and concentrated power, in various forms and for various utopian motives. This is the very *raison d’être* of the encyclicals *Rerum Novarum*, written in 1891, and *Centesimus Annus*, which revisited it in 1991, both of which emphasize the importance of preserving the dignity of the worker in the face of big capital and the indispensable role of civil society in promoting the common good. They emphasize that progress has meaning only if it remains at the service of the whole human person and does not sacrifice freedom and solidarity to the idols of production and omnipotence.

The 21st century, in turn, will be marked by a major transformation driven by the spread of artificial intelligence (AI), which enables knowledge to be gathered, connected, and processed on a scale beyond any human capacity.

While AI is not a new technology (its theoretical foundations date back to the 1950s), it is the rise in computing power and the availability of massive volumes of training data that have very recently enabled the emergence of large language models (LLMs) comparable to immense networks of

billions of neurons, with a speed and scope that far exceed human capabilities. These LLMs are at the heart of artificial assistants such as Copilot, Gemini, Claude, Llama, and of course ChatGPT, whose Version 4, released in December 2022, was the first to be accessible to the public, with the success we all know.

This artificial “intelligence” immediately reveals the ambition behind its development. For some, it is not merely a matter of creating assistants but entities deliberately designed in anthropomorphic terms, endowed with specific qualities that place them, at the very least, on equal footing with humankind. The terms are already flooding into everyday language: co-pilots, tutors, agents... To avoid any anthropomorphic temptation that might lead us to believe this AI is a living being, we should strictly refer to it as an “Artificial Intelligence System” to clearly emphasize that it is a machine. Furthermore, this risk of misplaced anthropomorphism will only grow with the gradual emergence of humanoid robots, equipped with this AI, which will take on an increasingly prominent role in workplaces, healthcare settings, and homes.

Beyond the anthropomorphic risk – the potential gravity of which should not be underestimated – this contemporary technological development is both fascinating and concerning for at least three reasons.

First, people today are more aware than ever of the necessities that weigh upon them, not only on an individual level but also on a large scale – whether in terms of poverty, which drives migration on a dramatic scale, or in terms of the fragility of their natural environment, which has been affected by a drastic reduction in biodiversity. But at the same time, humans are the creators of new tools such as AI, which, by their very nature, lack this compass of necessity. There is therefore no built-in alignment between what drives human beings and what drives these new “humanized machines”.

Furthermore, the human mind has known, since time immemorial, that there is no absolute determinism, whether in the social or personal sphere, and this is what makes the quest for freedom a defining feature of modern times. Humans can choose to break free from their past decisions, and even from their reason, to follow their hearts. The creation of artificial intelligence, by contrast, belongs to the realm of strict mathematical processing, resulting from complex lines of code processing information across a broad spectrum. Between freedom, the heart, and mathematical determinism, there is therefore a delicate balance to maintain.

Finally, the inherent lack of consciousness in an AI entity means that, by nature, it is not easy to assign it a responsibility that would lead to autonomy. During the 19th-century Industrial Revolution, the 1855 "Limited Liability Act" in the United Kingdom established the foundational concept of "limited liability" associated with a legal entity whose corporate purpose is identified and known. Moreover, it never occurred to any legal scholar to treat these legal entities as mere simulacra of human beings. When there is no longer a corporate purpose, no explicit governance associated with an entity whose very contours are evolving, what can be done when responsibility can no longer clearly be attributed to either a human person or a legal entity? Imagine a medical misdiagnosis made by a fully autonomous AI entity. Against whom would the patient turn?

Thus, when we consider the situation from different angles, we see that there is no simple alignment between humans and their AI creations. Necessity, freedom, and responsibility are therefore essential points of consideration if we wish to make this postmodern AI revolution a factor in improving the human condition. The unfortunate experiences of the 20th century have sufficiently shown us that the absence of realistic critical inquiry can easily give way to serious utopian excesses. It is precisely because we firmly believe that the AI revolution has the potential to transform personal and social life for the better that we see the urgent need for a thorough reflection on the subject, leading to clear standards. This post-industrial revolution at our doorstep must therefore be accompanied by this genuine process of reflection to foster harmony.

This situation makes it necessary to view this new chapter of humanity—still to be written—both through a human lens and through a Christian one. Indeed, we know that in God, all things are harmoniously united, and this is the heart of the Kingdom of Heaven. On our Earth, it takes all of humanity’s effort to create an environment that fosters the common good.

Drawing on the experience gained during the industrial revolution of the 19th and 20th centuries, Christians have developed a unique, consensual, and universal framework centered on four key concepts: the common good, human dignity, solidarity, and subsidiarity. At the heart of this humanism, the person of Jesus Christ offers a unique and essential point of reference for three reasons. First, he leads us to focus on caring for the weak before considering the benefits for the strong. Furthermore, he promotes the alignment of the aspirations of the weak and the strong, reminding them that their status as weak or strong is only temporary during the fleeting span of their lives. And finally, by revealing to every person that, beyond reason alone, the human heart has the capacity to welcome God himself, which makes them superior to any artificial creation. This perspective, which is naturally skeptical of any form of fetishistic, scientific, or ideological idolatry, thus makes the contribution of Christians essential to the emergence of a society that is ever more favorable to human life freed from all enslavement of body and mind. This point is fundamental, for the adoption of this new framework of life within which AI must find its proper place requires the emergence of individual and collective trust that is grounded in respect for the preeminence of the human person.

Since we are still at the very beginning of the widespread use of AI, our goal is not primarily to pass judgment after the fact on established situations, but rather to contribute concretely and positively to guiding the field’s evolution toward greater humanity. This constructive work is necessary and legitimate because it addresses a growing concern among most decision-makers, particularly white-collar workers, who are on the front lines. Their fears include the following: are we heading toward a world without human oversight, where situations will be analyzed by algorithms and decisions made and communicated by AI-powered tools? In corporate life, in hospitals, in agricultural activities, on the battlefield, will we see the disappearance of humans, who have become a source of potential disorder due to their unpredictable moods and their inconsistent professionalism? Will perfection become the exclusive domain of increasingly sophisticated machines, while humans will need to be managed with ever-greater precision in the name of social harmony?

Since *Rerum Novarum*, the Church has not waited for AI to ask the right questions about humanity’s place in successive technological revolutions, and what stance decision-makers and workers should adopt to ensure that technology serves without enslaving. It is striking to note that virtually all of the recommendations on AI issued by the Holy See in February 2020 following the “The Good Algorithm” conference at the Vatican, in collaboration with IBM and Microsoft, and published in the “Rome Call for AI Ethics,” have been largely incorporated into the European AI Act of August 2025.

This strikingly confirms the relevance and wisdom that the four dimensions of the CST can bring to enlighten not only consciences but also professional practices regarding the adoption and use of AI.

We will therefore explore these four dimensions in the following chapters, both from the perspective of doctrine and its concrete application:

1. *Absolute and uncompromising respect for human dignity;*
2. *The pursuit of the common good: that is, working together to achieve a positive impact on the world;*
3. *The systematic application of the principle of subsidiarity;*
4. *Solidarity, both within and outside the company.*

As Christians, we are particularly challenged. What does Jesus, the Incarnation of the Word, tell us, and more generally, what does all biblical wisdom inspire in us regarding this AI, which is itself a mathematized imitation of the word?

To answer this question, we will draw from the Scriptures and, above all, the Social Teaching of the Church the insights necessary to understand this technological revolution that has taken us by surprise.

ARTIFICIAL INTELLIGENCE AND HUMAN DIGNITY

What does the Catholic Social Teaching tell us?

The concept of human dignity plays a major role in the Social Teaching of the contemporary Church. It finds its origin, it must be emphasized from the outset, in the gift of God, who created man in his image. As the *Compendium* states in no. 108, citing *Gaudium et Spes*: “being in the image of God the human individual possesses the dignity of a person, who is not just something, but someone. He is capable of self-knowledge, of self-possession and of freely giving himself and entering into communion with other persons. And he is called by grace to a covenant with his Creator, to offer him a response of faith and love that no other creature can give in his stead.”

Specifically, in 271: “This subjectivity gives to work its particular dignity, which does not allow that it be considered a simple commodity or an impersonal element of the apparatus for productivity. Cut off from its lesser or greater objective value, work is an essential expression of the person, it is an ‘actus personae’. Any form of materialism or economic tenet that tries to reduce the worker to being a mere instrument of production, a simple labor force with an exclusively material value, would end up hopelessly distorting the essence of work and stripping it of its most noble and basic human finality. The human person is the measure of the dignity of work: ‘In fact there is no doubt that human work has an ethical value of its own, which clearly and directly remains linked to the fact that the one who carries it out is a person’.”

What conclusion can we draw from this?

From the outset, therefore, man and machine occupy two radically and essentially different levels. They cannot, therefore, under any circumstances be placed on the same level, and this applies to every human being. No form of anthropomorphism is conceivable here. The human being, because he is created by God and can welcome God, will always remain infinitely superior to the machine, even if it is “intelligent.”

As noted, the term “intelligence” is understood in many ways. In circles dealing with AI, and in line with the pervasive materialism of our societies, it is understood as the ability to solve problems or, at the very least, to act in response to situations. Hence the term “artificial intelligence.” Understood in this way, human intelligence can be surpassed. But from the perspective of the Christian faith, the mind, consciousness, and especially intelligence cannot be reduced to the technical resolution of problems.

From this perspective, one might even imagine that machines have a processing capacity far superior to that of humans, without necessarily concluding that humans are inferior in dignity. Is the dignity of a human being any less worthy than that of a jet airplane simply because it does not possess the same capabilities? The fact is that human beings possess a radical ontological superiority over machines: they are conscious, endowed with moral conscience, capable of authentic relationships, and have a heart—things a machine will never have. Humans are relational beings, possessing sensitivity. In short, a person. And such a person, created in the image of God, is the only one open to the infinity of divine transcendence. As for intelligence, we believe rather that it is a capacity to perceive the true, the beautiful, and the good, a key element of which is contemplation, which of course also serves to find solutions when needed—that is, to organize reality—but this is only one of its applications.

In summary, while humans may in some cases be technically inferior, they are always ontologically superior, for they possess within themselves a purpose and free will that machines cannot have.

While machines may sometimes, through an anthropomorphic sleight of hand, resemble humans, only humans are called to eternal life in God.

This radical difference between man and machine will prevail, particularly about issues such as respect for human rights, in light of principles such as truth, freedom, justice, and love, and more specifically, the dignity of work.

What are the implications for business leaders?

We will therefore consider the development of AI in general and any specific application considering its impact on the full development of human beings, whether for better or for worse.

This will shed light on several areas discussed elsewhere. The most obvious is, of course, that of work understood as employment. The full dignity of the human person presupposes that they play an active role in contributing to society, that is, through remunerative and dignified employment. However, AI will destroy many jobs and, notably, many jobs in fields traditionally considered more intellectually demanding and thus more valued. It would, however, be counterproductive to adopt a principle of systematically defending such jobs by resisting AI, at the risk of losing all competitiveness relative to competitors. We must instead embrace this opportunity to delegate to machines activities which, precisely because they can be delegated to a machine, do not in themselves possess as humanizing a character as we might have thought.

But by reshaping the contours of individual responsibilities, AI can also give rise to a new hierarchy of subordination. The emergence of machines within organizations that were previously essentially constituted of humans creates a need for a human-machine relationship that ensures everyone has a role that makes sense from a human perspective. For example, a new hierarchy may emerge between the role deemed more noble—that of the person controlling the machine—and that of the workers who feed it. It will therefore be essential to ensure that people displaced from their previous jobs are supported in their transition to new roles that are sufficiently rewarding, including for the most disadvantaged or those most affected by digital illiteracy.

More specifically, even though predictions in this area should be viewed with caution, AI is likely to change not only the nature of work in a vast number of jobs—time will tell whether for better or for worse—but above all to shape new individual relationships with work and perhaps with the meaning of work itself. If this is the case, it could directly affect the conception of the dignity of that work, and thus potentially the dignity of the person. In practical terms, where a person might have taken pride in a particular task they mastered well, AI is capable, in many fields, of performing all or part of that task better than they can. People may therefore see being replaced as a personal affront, a threat to their dignity. Yet what underpinned this dignity was not the work itself, but the aspect of the work that stems from human distinctiveness—whether it be human intelligence and creativity, dedication, relationships with colleagues and the company, or the pursuit of a job well done, a concept that is exclusively human. He will therefore have to do two things: reassess what he needs to do by redistributing tasks differently between himself and the machine, and above all, clarify what truly gives meaning to his work and where he finds his dignity as a human being at work. This will not necessarily be found in the part that can be automated and thus transferred to a machine. In doing so, they will need to restore authentic human relationships—the key to the dignity of human work—which have often already been damaged by digitalization or the bureaucratization of the postmodern era.

The rise of “Shadow AI” in companies—that is, the covert use of ChatGPT by employees—speaks volumes about the unease that has set in since the arrival of AI in the workplace (or in schools). Is an employee (or student) who gets help from ChatGPT without saying so just as shameful as a chef who

secretly uses frozen foods? The preservation of our dignity will also depend on the attitude we adopt toward these tools and how we allow ourselves to grow (or not) through the use of these new tools.

More fundamentally, human dignity is also the dignity of responsibility. This responsibility cannot be delegated to machines and must therefore be fully and exclusively upheld by people. Dignity here is not a demand, but rather a task that falls to us, without shifting it onto machines.

As required by the European AI Regulation, it is particularly essential that any AI system be continuously audited (validation of its sources or databases, verification of its operation) to assess its actual effects on people, not only in terms of unfair discrimination, but also in terms of simple unintended side effects or unsuspected algorithmic drift. AI is not necessarily created with malicious intent, but harmful effects can arise without a clear explanation as to why. Humans must always have the final say.

The question of human dignity in the face of AI is not limited to the place of humans in the workplace of tomorrow or to the issue of the hierarchy between person and machine. It also concerns the way in which humans inhabit their time and their inner lives—these two realms where their individual freedom is at stake. Yet AI, by enabling ever-more-precise automation of actions, decisions, and even thoughts, is profoundly transforming humanity's relationship to these two dimensions. The digital economy has already turned attention into a resource to be captured and monetized. There is a significant risk that AI will further intensify this commodification, reducing attention—this expression of our inner life directed toward ourselves, toward others, and toward God—to a mere measurable and exploitable flow.

In the face of this, Christian tradition recalls, in addition to the value of work (“negotium”), the value of “otium”, that time freed from productivity, devoted to contemplation, prayer, and gratuitousness. The Sabbath, in Revelation, is not merely rest for the body but rest for the gaze: a voluntary suspension of dominion over the world in order to welcome anew the gift of Creation. Thus, human dignity does not lie in hyperactivity but in the capacity to direct time toward what is meaningful.

If we succeed in preventing AI from enslaving humanity to a frenetic spiral toward ever-greater productivity—and thus, ultimately, toward an acceleration of work—AI could, conversely, free us from a number of tasks and become a paradoxical instrument of liberation from time. This will only happen on the condition that this freed-up time is not immediately reinvested in other utilitarian activities or swallowed up by futile screen time, but rather devoted to relationships, culture, and contemplation. Humans remain responsible for how they use the time that technology gives them back. AI must therefore also be judged by what it makes possible: not to produce more, but to love more. The real question for the business leader will therefore be how the introduction of AI could serve as a catalyst for the humanity of the men and women working in the company, rather than their alienation. This will surely involve reinventing the way we work, for work shapes the person, who cannot be content merely to contemplate, pray, or interact.

A second fundamental point concerns the person themselves. Only a human being is a person, that is, a being endowed with reason, freedom, and conscience, capable of relationship and transcendence. A legal entity or a digital entity is merely a legal or technical fiction, useful certainly, but incapable of interiority, conversion, forgiveness, or love. To equate AI with persons would be to confuse the instrument with the creature, to introduce a form of technological idolatry.

The Christian faith teaches us that no artificial intelligence, however powerful, will ever be able to participate in grace or in the sacraments. It cannot, of course, pray, receive communion, or go to confession, for these acts imply an embodied freedom and an interiority open to God. The danger, therefore, is not that the machine will become human, but that man will become a machine by forgetting the spiritual dimension that grounds his dignity. Where AI simulates emotions or religious discourse, we must tirelessly recall the difference between simulacrum and mystery, between the appearance of love and true love.

This demand for discernment also calls for a specific pedagogy of AI within educational institutions. It is not merely a matter of learning to use the tools, but of understanding the virtues and limits of an algorithmic world. Schools and universities (especially Catholic ones) play a prophetic role here: forming minds capable of handling technology without submitting to it. This requires curricula that integrate both technical competence and moral training: courses in ethics applied to AI, case studies on concrete ethical dilemmas, and workshops on collective discernment where technological choices are examined in the light of the Word of God, so as to ultimately develop a genuine anthropological depth in each person that goes beyond technical mastery and ethics.

Finally, human dignity in the digital age requires asceticism, in the noble sense of the term: an exercise of freedom in the face of the fascination of power and the continuous flow of information. Human beings, called to self-mastery, must learn to detach themselves from the constant demands of the connected world: this is what we might call digital temperance. This asceticism can take concrete form in “algorithm fasts”: moments of voluntary disconnection, measured use of predictive tools, and vigilance regarding the veracity of sources. Likewise, a digital examination of conscience can become a contemporary spiritual exercise: what have I entrusted to the machine today? Have I let AI think or choose in my place? Have I once again fallen silent to listen to God and my brothers?

This digital sobriety is not nostalgia for a vanished past; on the contrary, it is a condition of inner freedom. It puts technology at its proper place: that of a tool in the service of personal growth, and not a subtle master who shapes our desires. A dignified person is one who knows how to say no; and this ability to resist temptation, to pause, to contemplate, is perhaps the highest sign of their likeness to God.

ARTIFICIAL INTELLIGENCE AND THE COMMON GOOD

What does the Social Teaching tell us?

The Church's Social Teaching emphasizes the primacy of the common good, affirming that "far from being the object or passive element of social life" the human person "is rather, and must always remain, its subject, foundation and goal" (Compendium of the Social Teaching of the Church, no. 106). Thus, the development and use of artificial intelligence must be directed toward the service of humanity and respect for the dignity of every individual. It reminds us that technology must be at the service of the human person and not the other way around (Compendium, no. 458), and that no technological advance can justify the marginalization of the human person, nor the creation of new forms of exclusion or domination.

AI must be used in such a way as to promote social justice, solidarity, and the integral development of human communities. As Social Teaching emphasizes, "it should be stressed that progress of a merely economic and technological kind is insufficient. Development needs above all to be true and integral." (Caritas in Veritate, no. 23). This implies ethical regulation, systematic oversight, and constant vigilance to ensure that technological innovation remains at the service of the common good and does not become an instrument of increased inequality or a loss of meaning for people.

What lessons can we draw from this?

The emergence of Artificial Intelligence (AI) in our lives, and particularly in our businesses, represents a significant challenge regarding the creation and sharing of value. AI is not free, and we must therefore consider how we pay for it, who benefits from any increase in productivity, what the impact is on working hours and compensation, and so on. What future lies ahead for our business models and our long-standing mechanisms for sharing value, fair compensation for work, our social models, and thus the way in which each of us contributes to the common good?

Even more seriously, with tools such as ChatGPT, Claude, Gemini, or Grok becoming the benchmarks of daily life—true contemporary "household gods"—modern Lares—AI could potentially lead to a form of technological idolatry. Just as Laban, disoriented, searched for his household idols in Jacob's camp, modern man, drawn by the promise of simplicity and efficiency, risks quickly becoming entirely dependent on these digital life aids, blindly following the machine's instructions like a full-scale GPS for life and thus losing touch with reality—and with it, all concern for others and the common good. In China, AI is already replacing the traditional astrologer, becoming a new modern oracle. This clearly shows that AI is already being used to guide or influence human decisions, often at the expense of personal and authentic reflection.

Given such dependence, one might legitimately ask: who benefits—and will benefit—from AI? Today, 90% of AI investment takes place in the United States and China, and these countries are home to more than 90% of the sector's unicorns. We have entered an era of subtle digital colonialism where AI generally serves not the common good but primarily those who control it: big tech companies, certain governments, or powerful investors. Will Europe manage to avoid becoming subservient to foreign AI, thereby allowing them to capture the bulk of the economic value created? What would then become of its economic and social model, the free will of its citizens, and its cultural wealth?

What are the implications for business leaders?

Given the exponential evolution of technology, the impact of AI on work is certainly unimaginable within our current frameworks. From a very practical standpoint, and without waiting for the emergence of a global consensus—which will take time to materialize due to regulatory differences between Europe, the United States, and China—it appears essential that public bodies and the governing boards of organizations—companies, NGOs, foundations, and others actively engage in ethical and responsible reflection on AI. Discussions on transparency, ethics, and social impact must become strategic priorities, with a sense of urgency comparable to that required in the face of a major systemic risk.

The probabilistic management of individuality by a small number of actors with access to vast databases could enslave us to a form of determinism. Our human intelligence would thus risk being reduced to mere problem-solving capabilities, neglecting the profound richness of human experience, which includes creativity, contemplation, and freedom.

The common good requires clear sector-specific ethical boundaries, delineating what humans can and cannot delegate to machines. AI, due to its analytical and decision-making power, already impacts vital domains: health, education, security, and justice. In these spheres, the CST urges us to reiterate forcefully that certain decision-making thresholds cannot be delegated. We can assist and support, but we cannot judge or condemn in the name of an algorithm. The principle of ethical precaution must here align with that of subsidiarity: AI can inform decision-making but never replace it.

Thus, in healthcare, automated diagnosis cannot supplant the physician's responsibility toward the patient. In education, algorithmic personalization must remain a tool at the teacher's disposal, not a substitute for the educational relationship. In law enforcement and the justice system, the use of facial recognition or profiling must be strictly regulated, excluding any recourse to social scoring, predictive discrimination, or digital reputation systems. These are non-negotiable "red lines," for to tamper with judgment regarding a person is to tamper with human dignity itself.

In all these areas, the CST urges us to prioritize prudence over the thrill of efficiency. It reminds us that "not all that is technically possible is morally acceptable" (Directory for Catechesis, June 25, 2020), and that true wisdom consists in rejecting certain powers in order to better serve humanity. The prohibition, far from being a limitation, becomes here a higher form of freedom.

A use of AI that promotes the common good would involve freeing up time to allow people—in addition to gaining dignity at work—to devote time to pursuing their vocation and what makes them fully human: caring for the most vulnerable, education, artistic and cultural creation, as well as spiritual growth and a relationship with God. Thus, the time saved thanks to AI could be reinvested in a deeper and more fulfilling human dimension. The widespread use of AI will therefore require a major educational effort to develop applications that respect human dignity, ethics, and the common good. In short, the goal is to prioritize a community-based and solidarity-driven approach rather than extreme individualism, consciously choosing to "grow together" rather than "win alone."

But building the common good in the digital age first requires a genuine information ecology, which is inseparable from the virtue of truth. Technological progress, by multiplying the possibilities for creating and disseminating content, has fundamentally altered our very perception of reality. Generative AI systems can instantly produce texts, images, voices, and videos indistinguishable from reality. This unprecedented power poses a direct threat to social trust, for without trust in the truth, no human community can survive.

Faced with the proliferation of “deepfakes” and deceptive artifacts, a new responsibility arises: that of ensuring the traceability and verifiability of AI-generated content. Humans have the right, but also the duty, to seek the truth and not allow themselves to be trapped in an algorithmic relativism where truth and falsehood are treated as equivalent. This is why, under the European AI Act, any AI that disseminates or mediates information must explicitly identify the origin and method of content generation.

The common good also has a cultural and spiritual dimension. By standardizing language models and recommendation systems, AI tends to homogenize thought and impoverish cultural diversity. It is unlikely that there will be as many AI models (LLMs) as there are countries or cultures, leading to a form of cultural homogenization dominated by the most powerful, creating a new form of digital invasion or colonization—a new law of the strongest (and the richest). Protecting linguistic and cultural diversity in the age of AI—as well as copyright, already weakened by digital technologies—means protecting humanity from being reduced to a single statistical mindset.

Finally, technological omnipotence comes at a price: the considerable depletion of essential natural resources. The digital universe is not immaterial; it relies on a heavy infrastructure that is voracious in its consumption of resources: energy, rare metals, water, and natural spaces. Some data centers already consume amounts of water comparable to those of entire cities, and the production of semiconductors involves extraction chains that are often incompatible with environmental protection. From the perspective of integral ecology, business leaders must therefore adopt a responsible attitude toward the growing use of AI, taking into account the natural resources consumed to power it. Measuring these data, coupled with measurable goals for resource efficiency, would make the hidden costs of digital technology visible, enable us to take responsibility for them, and, of course, encourage greater resource efficiency.

Beyond the necessary vigilance, the CST places solidarity at the heart of the vision of shared growth, specifically in the service of all and by all. AI offers rich functionalities by facilitating access to diverse knowledge—accessible quickly, at a lower cost, and with a higher level of quality—which, if made available to as many people as possible, can help find better solutions to the challenges facing humanity and its stakeholders—companies and individuals—in fields as varied as health, education, and nutrition. Open source, as a common good to be shared (and as an illustration of the Universal Destination of Goods), offers, from this perspective, opportunities for the development of innovative solutions available to all.

ARTIFICIAL INTELLIGENCE AND SOLIDARITY

What does the Social Teaching tell us?

Solidarity is a foundational principle of social life, which Saint John Paul II defines as “the firm and persevering determination to commit oneself to the common good; that is to say to the good of all and of each individual, because we are all really responsible for all.” (Sollicitudo Rei Socialis, no. 38). It is inseparable from justice and charity, and requires special attention to the poorest, the excluded, and the vulnerable.

This principle aligns with the fundamental conviction that authentic progress must benefit everyone, not just a privileged fraction of humanity. As Fratelli Tutti reminds us, fraternity and social friendship are indispensable to avoid the logic of selection and abandonment, which seriously threatens our societies when an innovation—even a beneficial one—is not equitably accessible.

In other words: there is no true progress if the most vulnerable are left behind, or if the system’s winners irreparably widen the gap with the losers.

What lessons can we draw from this?

Solidarity is defined in the Church’s Social Teaching as “the need to recognize in the composite ties that unite men and social groups among themselves, the space given to human freedom for common growth in which all share and in which they participate”. The Social Teaching of the Church thus emphasizes three key dimensions of solidarity with which AI comes into direct conflict: the bonds that unite humans, human freedom, and shared growth. It is against the backdrop of these three pillars that our ability to build AI in the service of a renewed solidarity—in society and within our companies—can be judged.

AI in businesses and organizations is a source of both fascination and dread: fascination with the power of this technological tool to optimize processes and perform complex tasks in record time; fear of visiting factories managed by AI and devoid of any human presence, or virtual call centers, where humans have been replaced by bots trained to sometimes respond to customers better than humans; yet without truly “understanding us,” and thus lacking genuine empathy. But can we conceive of solidarity without empathy?

AI offers an opportunity to question solidarity as we experience it in a modern and postmodern society. The “mechanical” solidarity characteristic of industrial societies was no longer based on strong, homogeneous social bonds, but on the division of labor and the interdependence of individuals (interdependence that does not necessarily imply solidarity). Each person has a specialized and complementary role, which creates social cohesion based on differentiation and cooperation. This mechanical solidarity, particularly in European societies, tends toward centralization at the state level and through centralized structures such as Social Security, via mechanisms for the redistribution of taxes and social contributions, which constitute the main levers of material solidarity.

Raymond Aron (French philosopher) had anticipated that this centralization, this quest for equality, could come at the expense of individual freedom and responsibility, and identified this aspect as a point of concern. One of the antonyms of solidarity is individualism, a characteristic of our modern and postmodern societies.

However, modern society also relies on less visible, deep-rooted, and pervasive forms of solidarity—which vary in strength depending on the society and economic system—that, as noted above, are closer to the workings of organic solidarity and have persisted in industrial society: within families,

in associations, and among members of religious communities. These forms of solidarity should be maintained, or even become central again, in an “artificialized” society.

In the corporate world, this tension between equality and freedom crystallizes around redistribution, via taxes and payroll and employer contributions. Organic solidarity, on the other hand, takes the form, for example, of spontaneous fundraising efforts when a colleague’s child is sick, or “impact days” – solidarity initiatives – organized by certain companies to support the broader community in which they operate.

Another obstacle stems from the very nature of AI’s operation, particularly its energy intensity: these factors, which create or significantly amplify existing inequalities, call for a new form of solidarity akin to that established for vaccine access during the COVID-19 pandemic, lest we widen gaps that will be even more difficult to bridge: inequalities in access to water and energy, essential resources for data processing and storage centers; but also inequalities in the speed of AI adoption, which faces differing conceptions of time across cultures and is therefore likely to rapidly amplify initial gaps (in Africa or India). Finally, solidarity may also decline when access to proprietary AI, and thus to content enrichment, is limited or restricted; this leads to the formation of “self-reinforcing” communities that amplify opinions, severing ties between individuals within the same society or between societies.

What are the implications for business leaders?

In theory, AI could enable greater solidarity by performing a vast amount of work – and doing so exponentially – without human intervention, with the value generated by AI’s increased productivity distributed equitably to the greatest number of people (based on predefined algorithms). Humans could thus have even more free time, particularly to care for others, their loved ones, and their families. Beyond the potential loss of meaning linked to humans abandoning the workplace, humans – supposedly freed to manage their own time – become increasingly subject to AI (and those who define the algorithms), with AI thus serving as both the generator and distributor of value and wealth.

This vision of AI-driven solidarity may seem appealing because AI is more consistent, generates fewer errors, and requires neither forgiveness, gratitude, nor love; it could manage the interdependencies between humans and corporate reorganizations more quickly and with purported neutrality, instructing everyone to do this, then that, following appropriate and personalized training. It therefore has the potential to propose or impose a form of solidarity in its own image, consistent with its algorithm and its language.

But these logics of extreme automation have rarely had favorable consequences for human solidarity. The space for human freedom and responsibility would be irrevocably reduced – except, and this is no small paradox, for those who write the algorithms; often the very people who promote this convenient vision of an AI purportedly acting as an agent of this “mechanical” solidarity. Solidarity managed by AI ultimately carries the risk of reinforcing a purely materialistic approach to solidarity, with a deterioration in the fair distribution of value, to the detriment of spontaneous solidarity – arising from a willingness, a freer and more enduring initiative originating from the individual.

In reality, AI offers the opportunity for new forms of solidarity, succeeding both the mechanical and the organic forms. The three pillars of solidarity as proposed by the CST – namely, connections, freedom, and shared growth – seem, from this perspective, legitimate for refocusing solidarity on what makes humanity unique. This redefined solidarity must view AI as capable of positioning humans as subjects of work, rather than objects – thereby responding to the fundamental changes

affecting work more generally, and particularly since the onset of COVID-19. Humans liberate themselves by becoming the subjects of their work, using AI to free themselves from low-value tasks and focus on other tasks with higher added value for themselves, their organization, and society. The role of the socially responsible company is to empower its employees to become active participants in their work and in this adaptation—training in AI tools is therefore the first essential lever of solidarity. It is because AI will be accessible to all that it will become not a center of power or an instrument of exploitation, but an instrument of empowerment and freedom.

The CST also places the bonds that unite people and social groups at the heart of its definition. Paul Valéry (French poet) wrote: “Holy language, the honor of mankind.” AI, built on languages, must be deployed with this issue of human connections in mind. A dehumanized world cannot be one of solidarity, because the ultimate decision may be made by the machine to accomplish its task, without taking into account the factors that define essentially human solidarity. For example, in a company, certain AI systems will be more efficient than an intern; yet an intern is learning and needs someone with more experience to train them. Solidarity in the age of AI therefore calls for—or even demands—greater accountability, for example on the part of decision-makers, who must choose to hire and train an intern rather than opt for the apparent convenience of AI for simple and introductory tasks, even if this requires more effort on the manager’s part. The same applies to intergenerational solidarity within companies.

This responsible approach is also illustrated by prioritizing the relationship with a salesperson and the solidarity-driven choice of human interaction, rather than the delivery and recommendation—presented as personalized—offered by an AI. The commitment to solidarity is thus a conscious, resolute preference for human and creative uncertainty. These examples constitute new islands of solidarity (connections, freedom, shared development), and the company can be one of these islands thanks to the multitude of precise, targeted, thoughtful, conscious, and free choices (and struggles) made by its managers and leaders, in its own interest of mastering its levers of differentiation—namely culture, creativity, and innovation—and beyond mere operational and financial performance; choices that can be communicated to the company’s stakeholders, including customers and shareholders, with solidarity thus becoming a hallmark of an employer brand.

AI is redefining the role of human beings, particularly in the workplace—the CST can define a new, augmented form of solidarity. The word “augmented” is etymologically linked to authority. The all-powerful man is succeeded by a man whom solidarity—marked by freedom, connection, and shared development—can enhance, by leading him to accept authority and responsibility: that of his conscience and of love. True solidarity cannot be automated. AI can help, but it will never know how to love. Human support, patient listening, and care for the most vulnerable can never be reduced to lines of code. In other words, AI can enhance our capabilities, but only charity enhances our humanity.

Reflecting on the risks posed by AI thus offers us the opportunity to redefine the bond between people—true love, which has nothing to do with emotional materialism. It is this horizon of love that nourishes the hope for a world and individuals transformed by AI and united in solidarity. The real challenge of tomorrow is therefore not the coexistence of humans and AI, but the way in which human beings can harness these new technological capabilities to serve the most vulnerable.

ARTIFICIAL INTELLIGENCE AND SUBSIDIARITY

What does the Social Teaching tell us?

Subsidiarity, a key principle of Social Teaching, asserts that what can be adequately accomplished by the individual, the family, or intermediate bodies must not be taken away or absorbed by a higher authority; on the contrary, the higher level must support, coordinate, and complement without substituting itself. It aims to promote responsible subjectivity and initiative, for the sake of the common good and in conjunction with solidarity.

Saint John Paul II specifies that “a community of a higher order should not interfere in the internal life of a community of a lower order, depriving the latter of its functions, but rather should support it in case of need and help to coordinate its activity with the activities of the rest of society, always with a view to the common good.” (Centesimus Annus, no. 48).

This principle condemns paternalism and technocracy, which deprive actors of their decision-making power, while recognizing that a higher authority must intervene when lower levels are truly powerless or failing. Applied to the technological context, subsidiarity requires that tools—including AI—remain ordered to the freedom and responsibility of individuals and communities, foster their capacity for action, and do not establish an opaque centralization of choices that would shape social life without oversight or informed participation.

What perspective can we draw from this?

Subsidiarity is the exercise of authority at the exact level needed to grant each employee the autonomy compatible with the organization’s required cohesion. It enables everyone to maximize their contribution to collective effectiveness. It thus reveals a powerful and ever-accessible potential for performance, as well as for meaning and job satisfaction.

Subsidiarity is even more powerful the greater the efficiency and competence of employees. AI, insofar as it can amplify employee performance and is designed to enable them to mobilize their skills at their own level, should therefore—all other things being equal—serve as a lever for developing the exercise of subsidiarity.

In the educational and professional spheres, subsidiarity also implies empowering individuals in their engagement with AI. The tool must not be a pretext for intellectual abdication. On the contrary, each individual must develop the skills, maturity, and moral sense necessary to use it judiciously. A comprehensive education—technical, ethical, and spiritual—becomes indispensable to enable everyone to act with full knowledge of the facts. The company, like the educational community, must thus ensure the growth in responsibility of its members, by giving them both the means and the freedom to act.

It is in this enlightened autonomy that the heart of subsidiarity lies: not a mechanical decentralization of procedures, but a moral liberation—that of the person capable of governing their own actions. But for this lever to be used wisely, we must ensure that the two essential issues of responsibility and trust are effectively addressed, since subsidiarity also involves maximizing individual responsibility (empowerment), balanced by the need to ensure accountability.

To reflect on how AI can help make companies more agile and effective in their implementation of subsidiarity, we must therefore examine whether and how AI increasingly provides the means for greater responsibility and greater trust. Subsidiarity indeed calls for institutional restraint in the face of the temptation of constant control. AI, through its ability to measure, predict, and classify everything, offers unprecedented surveillance capabilities. But a society that seeks to control everything loses its soul even before its freedom. The use of AI in administration, public policy management, or economic life must remain proportionate, reversible, and transparent. We must

accept that a degree of unpredictability, slowness, and error remains in human life: these are signs of freedom and creativity. Subsidiarity is thus also experienced as an asceticism of power, a voluntary renunciation of technical omniscience to better make room for human consciousness.

Ultimately, subsidiarity in the age of AI could be expressed as follows: AI empowers everyone to perform better without replacing what gives people's work its meaning; to delegate without shifting responsibility; to automate without dehumanizing. It is a pedagogy of freedom applied to technology, a learning process of cooperation and discernment, where humans remain the subjects and beneficiaries of progress.

What are the implications for business leaders?

The use of AI can only foster the "power to act" if it is implemented by people who can harness its potential – not just the desire or the feeling that they know how to do so. Training and learning that instills this ability in a manner tailored to each individual's situation is therefore an essential prerequisite. This learning must provide a realistic understanding of how to use AI effectively, as well as its limitations and the necessary safeguards.

We know that AI raises questions about the legal liability of the legal entities that use it – issues that will most likely require charters to be established at the national or international level. Among individuals, subsidiarity similarly raises questions of mutual responsibility that must be clarified to effectively apply the principle of subsidiarity. Within the company, it is unreasonable to wait for external solutions that will only be developed gradually; otherwise, the use of AI would spiral out of control. It is therefore the company's responsibility to progressively clarify how it is rethinking liability in a context transformed by AI.

There is no doubt that decision support through AI, or even decision-making by AI applications, constantly raises the question of the responsibility of the individuals who make up the company and the chain of delegation of responsibilities.

Subsidiarity, supported and implemented through this collective effort, must enable everyone to cultivate what is unique to them and what only they can master due to their position within the organization. But subsidiarity is only realistic because it is also accompanied by the requirement of accountability: we know how often this requirement – which is, incidentally, indisputable – is carried out through procedures that are overly bureaucratic and therefore ineffective: it is neither good nor effective to control everything. This is an area where the coordinated implementation of AI could allow us to be more responsive but also, paradoxically, more "human" due to AI's ability to facilitate dialogue between people.

On a more personal level, subsidiarity calls on each of us to renew our sense of service and trust. It is based on the conviction that every person, at their own level, is capable of sound judgment when acting in the light of their conscience. AI, when used properly, can support this growth in responsibility by offering everyone a better understanding of reality and the consequences of their actions. But trust cannot be automated: it is built on mutual recognition, patience, and listening. From this perspective, subsidiarity becomes a pedagogy of fraternity – a learning process regarding the proper use of machines by humans, in the service of humanity, and not the other way around. It is on this condition that technology, far from crushing freedom, will become an instrument of unity and moral growth, not a substitute for shared responsibility.

Ten Recommendations for Business Leaders

1. Affirm the principle of “people first.”

Specify in writing that the human person is the ultimate goal and that their dignity and the authenticity of human relationships must always be preserved. Never delegate to machines decisions that affect life, justice, or freedom.

2. Draw red lines and prioritize uses.

Prohibit autonomous weapons, mass surveillance, social ratings, or profiling. Promote uses that improve the quality of work and free up time for care, education, creativity, and relationships.

3. Establish subsidiary governance.

Create a multidisciplinary committee, close to operational realities, with the authority to raise alerts and suspend deployments when necessary. Decide, test, evaluate, and correct in short cycles, with transparency.

4. Ensure transparency, truthfulness, and traceability.

Make it mandatory to identify AI-generated content (watermarking), ensure transparency regarding its use and training data (and their potential biases).

5. Measure and reduce the material footprint.

Publish a comprehensive environmental report that includes the digital footprint of the company (energy, water, rare metals, location). Set sustainability goals: energy efficiency, resource sharing, renewable energy.

6. Provide training for mindful and responsible use.

Incorporate technical modules as well as modules on applied ethics and collective discernment into corporate education programs. Empower everyone to take responsibility for the impact of their digital decisions.

7. Promote digital moderation.

Ensure the right to disconnect, moments of digital silence, and practices of moderation: limit hyper-connectivity, verify sources, and cultivate mindfulness and genuine presence.

8. Prevent groupthink.

Ensure that AI systems do not standardize thinking. Diversify sources, encourage a plurality of approaches, and protect teams’ freedom of conscience and creativity.

9. Keep people in the loop and ultimately accountable.

Define, by job function, which decisions cannot be delegated. Document who makes decisions, based on what criteria, and how to challenge them. AI can advise, but it does not judge.

10. Share the gains and invest in people.

Link productivity gains to training, mobility, and fair compensation mechanisms. Reinvest in customer-facing, creative, and care-related roles: where people remain irreplaceable.

Conclusion

The emergence of artificial intelligence, particularly mass-market generative AI, marks not only a technological revolution but potentially an anthropological turning point. It compels us to redefine the foundations of civilization: the primacy of the person over the tool, of service over power, of meaning over efficiency. AI can elevate humanity, or impoverish it; it can free us from the burden of necessity, or plunge us into new dependencies, depending on how it is and will be used. The difference will depend on the moral compass we apply to these uses. When it comes to technology, history shows that we always overestimate short-term changes (1–2 years) but greatly underestimate long-term impacts (10–20 years). This is due to the exponential nature of technological evolution, which our brains mistakenly perceive as linear, leading us to initial denial and then to being caught off guard and inadequately prepared for each technological revolution.

Faced with the speed, but also the unpredictability, of these changes, the Catholic Social Teaching offers a valuable compass for this discernment—dignity, the common good, solidarity, and subsidiarity. These four guiding principles are not abstract concepts: they outline a concrete way to govern technology.

- Dignity reminds us that the human person is the measure of all things, and that no machine, however powerful, can claim the same value;
- The common good directs power toward justice and sharing, rejecting monopolies of knowledge and wealth;
- Solidarity broadens the perspective: all progress must benefit everyone, not just a few;
- Finally, subsidiarity restores trust in and for humanity: it affirms that it is necessary to grant each person the powers corresponding to their sphere of responsibility and not to take them away—even through AI—but also to assist as needed and, exceptionally, to step in.

To govern AI, and not be governed by it, is to balance two requirements: clear limits and fruitful freedoms. It is to affirm that there are domains—life, conscience, justice, truth—that cannot be entrusted autonomously to algorithms without degrading humanity itself. But it is also recognizing that, when properly guided, technology can become a lever for liberation: by relieving us of menial tasks, it can make us more available to appreciate a work well done, to relationships, to creation, to contemplation, to prayer.

The issue is therefore as much spiritual as it is ethical. To preserve truth is to preserve freedom itself. For freedom lives only in the light of truth: when it strays from it, it dissolves into relativism, and then into servitude. In a world saturated with artificially generated images and discourse, where everything can seem plausible, the responsibility of the business leader, the researcher, the educator, and the politician is to be a guardian of reality. Speaking the truth, even against the tide, becomes an act of moral resistance and enlightened governance.

But even more profoundly, the question that AI poses to each of us is one of meaning. Not “What can the machine do?” but “What does humanity want?” What does it wish to pass on, protect, love, and contemplate? AI offers a silent test of our spiritual maturity: do we still know how to choose what uplifts us, and reject what degrades us?

After decades of digital life spent—at times—performing repetitive, mechanical tasks ourselves behind a keyboard and a screen, AI could paradoxically set us free—allowing us to finally lift our heads and aspire to greater transcendence; that is to say, for us Christians, to actively set out in the footsteps of Christ.

Unfortunately, no one can predict whether AI will ultimately prove beneficial to humanity overall, but, freed from the repetitive motions that clutter our days—from these robotic tasks finally entrusted to machines—we might rediscover the meaning of our gaze and our breath. This regained time, if used well, could be devoted to what makes us greater: the satisfaction of a job well done, relationships, contemplation, and the quest for truth, beauty, and goodness. Perhaps this technological revolution, if harnessed properly, will invite us to aspire to greater transcendence, to rediscover the inner silence from which all wisdom springs, and to draw closer to God, the source of all true intelligence.

In response to intelligent machines, we must offer the intelligence of faith. AI allows us to accumulate data, connect signals, and infer decisions; faith, on the other hand, connects people, enlightens hearts, and leads to the living Truth. The intelligence of faith does not deny reason; it transcends it by opening it to the light of mystery. It discerns where AI calculates; it hopes where AI predicts; it loves where AI commands. In a world fascinated by efficiency, faith restores intelligence to its primary vocation: to understand in order to serve better. The Church does not fear artificial intelligence, but reminds us that the only true intelligence is that which leads to communion, not to mechanisms of domination. Thus, the believer is not invited to flee from technology, but to use it with faith, ceaselessly seeking how every advance can become a path of encounter, humility, and praise.

Machines will inexorably replace humans in all measurable, quantifiable tasks, in which algorithms will never cease to improve. But we Christians believe in the immeasurable. Human existence does not begin and end with what can be measured; it is far more than that. It is not about excess but about the immeasurable. God is, by definition, the one who cannot be measured; otherwise, he is not God. And man, created in his image, contains a dimension of the immeasurable. It is this humanity that will always elude the machine and makes the human being infinitely superior.

If Jesus were to speak to us today about artificial intelligence, he would undoubtedly speak to us neither of code nor of algorithms. He would speak to us of the human heart. He would remind us that “Nothing that enters one from outside can defile that person; but the things that come out from within are what defile.” (Mk 7:15). He would invite us not to fear the machine, but to guard against the pride of wanting to make ourselves gods through it. He would reiterate that true intelligence is not about knowing or predicting, but about loving; that true power lies in service, and that true light comes from grace. In this context, the parable of the talents takes on new resonance: the master entrusts to each according to his ability, not so that he may keep it, but so that he may make it bear fruit. He will return, and will ask an account not of profitability, but of faithfulness. Artificial intelligence multiplies the capacities put at the service of talents: it increases power tenfold, accelerates production, and automates creation. But humanity remains accountable for the meaning of what it multiplies. To use AI without a moral purpose is to bury one’s talent in the digital earth, to render it barren in the master’s forgetfulness. Conversely, directing this power toward the common good, dignity, justice, and peace is truly making the talents received from God bear fruit. The Last Judgment will not concern the power of the algorithms we have designed, but the fruitfulness of our hearts: what have we done with what we have received?

Understood in this way, the use of AI becomes a field of trial and hope: a trial of our freedom, a hope for our conversion. It can enhance our capabilities, but only charity enhances our humanity. The challenge is not to coexist with digital machines, but to remain fully human in their use. And this is possible only by keeping alive, at the heart of technology, the flame of the Spirit that enlightens, unites, and sanctifies.

For humanity’s use of machines to work, we must recall a crucial distinction: we speak—inappropriately—of artificial intelligence, but never of artificial wisdom. And this is no mere coincidence. Intelligence processes information; wisdom orders life. Intelligence calculates; wisdom contemplates. Intelligence coordinates means; wisdom discerns ends. It is precisely this wisdom—a gift of the Holy Spirit—that is uniquely human and that we must urgently cultivate today in a world of intelligent machines. For if we abandon wisdom, we will leave increasingly powerful Artificial Intelligence without direction, and that power would eventually turn against us. Wisdom calls

people to prudence, justice, and truth. It is the inner compass that no digital machine can replicate. More than ever, our age calls for wise men – not merely engineers of probability, but artisans of meaning and seekers of Truth, capable of uniting science and conscience, power and goodness, knowledge and love.

Paris, April 2026.