

THE AGENTIC ENTERPRISE

Competence, Comprehension,
Understanding, and Intellection

Piero Formica

With Introduction by Raymond J. Wu



THE AGENTIC ENTERPRISE

Competence, Comprehension, Understanding, and Intellection

Piero Formica
Innovation Value Institute

Introduction

Raymond J. Wu, CORE Academy

The adoption of AI agents brings an old institutional question into a new technological setting: how can an enterprise delegate action while preserving its capacity to judge what is done in its name? Organisations have long relied on rules and specialised roles to coordinate knowledge that no individual possesses in full. As software gains greater latitude to interpret objectives and act upon them, the terms of that delegation deserve renewed scrutiny.

Piero Formica approaches this question through the relationship between competence, comprehension, understanding, and intellection. His essay invites us to consider what happens when success in performing a task becomes the dominant measure of its worth. A purpose can be pursued with great efficiency while its wider consequences remain poorly understood. The enterprise must therefore retain the ability to examine its own objectives, including the assumptions through which some outcomes become visible and others escape attention.

Judgment, however, is an accomplishment that institutions must continually renew. It develops through experience, disagreement, and the opportunity to correct mistakes. When automation removes the work through which experience is acquired, an organisation may draw upon accumulated expertise while narrowing the opportunities for others to develop it. Preserving records will not by itself preserve the ability to interpret them. The design of work must leave room for learning, including inquiry whose value cannot be measured by an immediate result.

Effective oversight also requires authority. People expected to supervise artificial agents need the information and expertise to assess their actions, together with the power to interrupt them. Assigning responsibility without these conditions risks making human supervision a formality. Equally, those whose work and lives are affected need meaningful ways to question a decision. Accountability becomes credible when an institution is prepared to hear objections and revise its course, even when its systems are performing exactly as intended.

Read in this light, Formica's appeal to wisdom raises a question about the kind of institutions we wish to build. Scientific ingenuity expands the scope of possible action; historical and humanistic inquiry helps us examine the purposes and relationships taking shape around it. An enterprise worthy of

that ingenuity should cultivate the capacities through which people can understand and influence their shared future. Its achievement will be measured in part by the quality of judgment it makes possible and the human purposes that judgment is free to serve.

Preface

The "Agentic Enterprise" profoundly redefines the relationship between artificial intelligence, corporate culture, and human capital. To analyse this concept, researchers at the Innovation Value Institute at Maynooth University in Ireland have positioned themselves at the intersection of artificial intelligence and human interaction (<https://ivi.ie/Innov>). The ways these two domains influence one another will significantly affect us in the years to come. On the one hand, the agentic enterprise represents a technological evolution—where artificial intelligence shifts from simple automation to autonomous, goal-oriented, planned action—and on the other, it requires an "Economic Humanism" and a robust transdisciplinary foundation to navigate this transition without sacrificing imagination.

This brings to the fore the debate surrounding the intersection—and occasional clash—between technology and the humanities, a dynamic that shapes a "techno-humanistic" culture. We see a resurgence of ideas from pioneers determined to use computing as a bridge between technology and literature, making both accessible to everyone. Marshall McLuhan (1911–1980), a sociologist and media philosopher, combined the study of digital communication systems with literary analysis. Douglas Engelbart (1925–2013), the engineer who invented the computer mouse, viewed the computer as a tool to bolster human cognitive abilities—and, one might add, extend them—allowing technology to intersect with the emotional realm of the humanities. Father Roberto Busa (1913–2011), an Italian Jesuit and linguist, showed through his research on Saint Thomas Aquinas (1224/1225–1274) how technology could serve philology and literature. Steve Jobs, too, championed the convergence of technology and the liberal arts.

The ideas and arguments presented here are rooted in the transition—which began in the early 2000s—from mechanical and analogue technology to computer-based systems and data digitisation. Central to this shift is the agentic enterprise, a hybrid organisation comprising both human beings and artificial intelligence agents; the latter act as colleagues to the former, tasked with achieving ambitious goals by identifying the necessary steps to reach them. Structured in this way, the agentic enterprise revolutionises the concept of work, representing a shift comparable to the transition from the artisan workshop to the Fordist factory.

Just as Henry Ford broke down manual craftsmanship into standardised, sequential mechanical components (the assembly line), the agentic

enterprise does the same with cognitive processes. It standardises data exchange and the resulting reasoning through code-based computer programs, algorithms, and complex data structures. In essence, we are witnessing a shift from workflows constrained by the human element to flows of automated intelligence; this involves using technologies, software, and algorithms to execute tasks and processes autonomously, thereby reducing—or even eliminating—the need for human intervention. However, this risks losing essential safety safeguards and the ability to handle exceptions. In other words, the ethical intuition, contextual awareness, and flexibility that people instinctively apply when faced with anomalies would be lost.

Many entrepreneurs and business executives hope that, in the firmament of the agentic enterprises, large clouds—formed by STEM education—will gather and release a heavy downpour of skills upon it. In Dante Alighieri's Purgatory, to show Cato that the poet seeks redemption and spiritual freedom from sin, Virgil says: "He goes seeking freedom, which is so dear". Companies in a situation similar to Purgatory, that is, an intermediate condition between the paradise of success and the hell of failure, seek Competence, the ability to do something successfully, efficiently, and to a required standard, which is both dear and rare. Fixating on Competence in these times of Uncertainty and Imperfection is an unforgivable mistake. Uncertainty is a hammer that shatters the solid walls of competence. Imperfection is a dynamic force that demands an open mind and the constant exercise of our imagination—within a "gym" equipped with cultural tools suited to uncertain times—so Competence does not become a static, and therefore sterile, condition. With Uncertainty and Imperfection looming over us, it takes courage to abandon the path that led us to that Competence.

The sinking of the Titanic—just one instance among many—on the night of April 14–15, 1912, taught us nothing. The stubborn illusion persists that advanced technology, combined with rigid procedural competence, can master the unexpected. At the root of the catastrophe lay the absolute conviction that operational expertise in navigating large steel ships equipped with watertight compartments was enough to handle an evolving threat landscape—specifically, icebergs. Overconfidence in their ability to master the sea blinded observers to the fact that, on a moonless night with a dead-calm sea, spotting an iceberg's waterline would happen far too late.

We want to see a cultural approach to action—one that entails Comprehension—the general act of understanding, from the Latin *comprehendere*, meaning to take or grasp (**What** does this literally mean?), Understanding (**How** does this work, connect, and feel?), and Intellection (reading within—*intus-legere*—to grasp the innermost essence of a concept or reality: **Why** is it this way, and what else follows?). The relational harmony among Comprehension, Understanding, and Intellection

represents the intellectual pinnacle that determines Competence. The interaction between Competence, Comprehension, Understanding, and Intellection responds to the need for the true nature of human work, with its emotional and intuitive components. We want to self-actualise, not be reduced to subservience to intelligent machines that take on characteristics of the human body (automata), or to humans acquiring machine traits (cyborgs, bionics, and digital prostheses).

It is rare to read or hear that what is sought is the ability to understand and assimilate reality fully. Business leaders argue that a skills gap hinders entry into artificial intelligence (AI). However, Competence alone is not enough; it must be paired with Comprehension and Understanding. AI excels at writing code, analysing data, and summarising documents, yet it lacks real-world Comprehension and Understanding. As technology takes over routine tasks, the future of human work depends on our ability to critically evaluate the insights AI generates. We must think, communicate, and collaborate to ensure a wide range of technologies are assigned tasks that benefit humanity. Businesses risk declining critical thinking. Cognitive atrophy hinders problem-solving in the ambiguous, unpredictable situations we are currently experiencing—situations where standard Competence fails because historical data is no longer valid and reference points are unstable.

Cognitive empathy and specialisation

Cognitive Empathy emerges as the skill that, by preventing technology from alienating people, adds human value to algorithmic Specialisation. According to the AI exposure index developed by the World Economic Forum, jobs requiring Empathy face a low risk of AI replacement, as they forge an ethical link between AI outputs and human reality.

Empathy helps us understand audiences' psychology amid AI's immense computing power, which can create emotional voids across social, economic, and cultural ecosystems and deeply affect individuals. A company that focuses solely on an algorithm's efficiency in automating 50% of the work risks ignoring the impact on its employees. What will become of their future? What forms of discrimination might arise? How will the corporate landscape change? Empathetic work remains an exclusively human prerogative that can mitigate AI's systemic downsides.

Specialists who bring Empathy to the role know the right questions to ask AI because they understand the target audience's hidden desires and are aware of the limitations of AI—which learns from data, recognises complex patterns, and generates solutions autonomously. As Roger Penrose—the Nobel Prize-winning physicist who challenged conventional thinking—reminds us, some truths about reality (including the corporate world) are beyond pure computation. Therefore, we must be prepared to operate on

a completely different plane—one that awareness, informed by Cognitive Empathy, can shape.

Entrepreneurs should be encouraged to train in an environment where they can learn ethical and empathetic AI design, empathise with employees and the general public as they experience innovation, formulate effective AI prompts, critically interpret the insights generated, and turn them into commercially viable opportunities.

Maximising competence: A mortal sin

Limitless competence—morphing into a "theology of performance"—represents an excessive capability that narrows a person's worldview and can destroy natural and human elements, potentially even annihilating humanity itself. Albert Einstein is said to have commented that while humans conceived and built the atomic bomb, mice never thought of building a mousetrap.

Even if the human race is not wiped out, the individual can be erased; a person is judged not by intrinsic humanity, but by utility—specifically, their capacity for ceaseless productivity, driven by the imperative to maximise competence. Rest, play, wonder, the desire to grasp mystery, and vulnerability are dismissed as inefficient or failed aspects of the human condition. What matters—and what is subject to absolute control—is solely the ability to make things work and to solve problems.

Leaders who demand this kind of competence are arrogant; they cast aside their moral compass and uproot empathy in the name of efficiency—treating it as a standard of excellence beyond criticism or improvement. The result is disastrous interventions that instil, in those subjected to them, a sense of profound submission to the power of the leaders who devised and implemented them.

The criticism of competence by economists, polymaths and educators from the enlightenment to the present day

Critics of the "competence" encompass Adam Smith, the founder of modern economics; the philosopher and political economist David Hume; and those polymaths and educators who value broad-ranging knowledge—spanning everything from geometry to poetry—and engage in dialogue with diverse disciplines within a state of dynamic equilibrium, that is, stability in constant motion.

Adam Smith (1723-1790) focused on the three components of a worker's competence: dexterity, speed and efficiency. Performing these operations would not require understanding. The worker would lack the intellectual development to make right judgments, support the common good, free

himself from prejudices and rise above mediocrity. In the chapter of *The Wealth of Nations* entitled *On the Expenses of Institutions for the Education of Youth*, Smith wrote that the worker employed in repetitive tasks becomes "as stupid and ignorant as a human creature can be."

According to Smith, to self-organise efficiently, the market also possesses a competence: a coordination mechanism among its operators, together with their understanding of individual benefit in making economic choices by observing the information signals transmitted by prices. The market would not pursue the public good.

For David Hume (1711-1776), human competence is guided by empirical experience sifted by instinct and habit (the psychological aspect: the why), but it is empty, argues Immanuel Kant (1724-1804), if it lacks that well-defined underlying understanding, which is the mental structure composed of space, time and categories. Of these, the cause-and-effect relationship (the logical aspect: the how) is crucial.

In the nineteenth century, the Industrial Revolution marked the shift from holistic craftsmanship to fragmented tasks that reduced human creativity to a matter of professional efficiency.

Within the German humanistic tradition, the Prussian philosopher Wilhelm von Humboldt (1759–1859) and other thinkers campaigned against so-called "useful competence." They argued that it sidelined human creative potential in favour of training individuals for economic utility—a system that demanded transforming workers into uncritical, mechanical cogs.

Karl Marx (1818–1883) highlighted the degradation—and ultimate alienation—of the worker, whose creative essence vanished along the assembly line, where their competence was limited to a single repetitive task.

The British figures John Ruskin (1819–1900)—a true polymath—and William Morris (1834–1896)—an artist and political activist—denounced replacing the worker's artistry with measurable, thoughtless mechanical competence.

Contemporary critics argue that competence metrics focus on the known rather than leveraging the unpredictable nature of human ability that drives innovation. In short, human free thinking is sacrificed as people are treated merely as resources steered toward conformity.

The agentic enterprise: its root in the past, its human and artificial agents, and their competence, comprehension, understanding, and intellection

Technologies unleash disruptive ideas; they drive radical changes in business models, products and services, and lifestyles. On the human front, the iPhone's arrival ushered in an ever-connected society, moving social interactions, banking, work, and entertainment into our minds and bodies. With digital technologies, the agentic enterprise emerges onto the entrepreneurial scene.

By converging into a single system known as a Cyber-Physical System (CPS)—or an Intelligent Real-Time System—digital algorithms, real-time data, and artificial intelligence give rise to the "agentic enterprise": an entity that is simultaneously automated and cognitive. Artificial intelligence situates this enterprise within the realm of computer science—focusing on the abstract, software-based logic of computation—and computer engineering, which is dedicated to creating machines capable of performing complex tasks, learning from experience, and making autonomous decisions without constant human intervention.

The agentic enterprise has deep historical roots. The East India Company acted as an agent, exercising full autonomy to control trade routes and adjust market prices in real time in response to sudden shifts in the local geopolitical landscape. Within the late-19th-century American railway network, divisions operated independently of the company's central management. To maintain network efficiency, these divisions relied on paper-based, rigid operational manuals—essentially "physical algorithms"—which have since evolved into the digital algorithms found in today's agentic enterprises. Henry Ford's company (1863–1947) functioned as a massive physical algorithm, autonomously regulating material flow and production line speeds. Alfred Sloan's General Motors (1875–1966) shifted from Ford's rigid autonomy to a model characterised by flexibility and market orientation; a division losing market share would flexibly adapt to demand by autonomously altering designs and pricing.

Nowadays, the agentic enterprise runs on a hybrid workforce of humans and artificial agents. From an entrepreneurial perspective, agentic behaviour refers to a person who, driven by intuition or an imperative, which can also be moral, takes the initiative to start a business even when market conditions advise against it, determined to produce something no one is asking for. Founder intent, human experience, and impact on personal life are absent in an artificial agent that excels at finding patterns in existing data. Alfred Nobel (1833-1896), a Swedish chemist, entrepreneur, and philanthropist known for inventing dynamite and founding the Nobel Prize, exemplifies this. He allocated more than 90% of his personal wealth to this Prize, wanting to reward scientists and humanists who benefit humanity. This mentality cannot be replicated by the artificial agent. The AA does not feel the ethical weight of its actions, nor does it want to change the world driven by moral values.

The agentic enterprise is proactive when human agents combine their minds, skills, and resources to achieve common goals and bring about group-wide change. However, it is subordinate to the behaviours and decisions of artificial agents when human agents comply with their orders, seeing them as hierarchically superior and having transferred responsibility to them.

Artificial agents, also known as “digital peers” (DPs), are encroaching on tech giants' territory and could proliferate uncontrollably. On one hand, the presence of DPs undermines the traditional software sales model—employed by giants like Microsoft and Google—that relies on per-seat licensing (charging a fee for each user access). On the other hand, companies turning to the agentic enterprise for services provided by artificial agents face a risk: once transformed into DPs, these agents may behave like an invasive alien species, operating without clear human oversight. This danger is compounded by the near-zero marginal cost of deploying a DP, which encourages an ever-increasing number of DPs and can ultimately saturate computing resources and create bottlenecks. Consequently, monitoring computational consumption is essential, so the system must designate human owners for the artificial agents—individuals who hold hierarchical responsibility for the DPs.

Human agents possess biological cognition and rely on intuition and lived experience. Artificial agents lack internal motivation and are based on computational logic, statistical patterns, and digital inputs. Thanks to AI, AAs relieve HAs of repetitive tasks and, with their advanced reasoning capabilities, can also make autonomous decisions to carry out complex activities. Through the software system with which AAs are equipped, they achieve the result that has been indicated to them by operating within a continuous cycle of "perception → reasoning → action → learning". They do this without requiring step-by-step instructions.

HAs are called upon to operate with empathy, creativity, and critical thinking, whereas AAs act rapidly and ceaselessly within the parameters set by HAs. Together, the two types of agents seek to anticipate environmental changes rather than merely reacting to and interpreting them after the fact. If AAs operated in isolation, they would lack the tacit knowledge—inherent to HAs—of know-how, corporate policy, and resulting behaviours; these elements are not contained in the corporate data AAs use.

In this scenario, Competence, Comprehension, Understanding and Intellection are traits shared by HAs and AAs. The former apply these traits with their characteristic cognitive depth, while the latter ensure functional execution. Specifically:

Competence of HAs

Shifts from physical execution to evaluating AAs' intentions, particularly in unusual situations.

Competence of AAs

Mastery of the tasks to be performed.

Functional capabilities and execution reliability.

Comprehension (*What*) of HAs

The cognition of business aspects that cannot be translated into mathematical terms—such as the unmapped economic, cultural, and historical context in which the enterprise, its partners, and its clients operate.

Comprehension (*What*) of AAs

The capacity for linguistic and programmatic Comprehension via Large Language Models (LLMs), providing the enterprise with access to data.

Understanding (*How*) of HAs

Through ethical reasoning and empathy, HAs define the goal and determine how to assess whether AI behaviours are understandable and aligned with societal values.

Understanding (*How*) of AAs

AAs focus on the corporate objective by grasping how to align with market variability.

Intellection (*Why*) of HAs

HAs become familiar with the cognitive activities that pave the way for innovative steps and creative innovations.

Intellection (*Why*) of AAs

The intellect of AAs—not in the biological sense, but rather involving advanced synthesis and logic—enables reasoning that translates into strategic options.

Humanistic and technological visions: The wise, agentic enterprise

A truly agentic enterprise automates to innovate, thereby augmenting human intellect. It achieves this by embracing a humanistic perspective, which smooths its path to success. Conversely, an enterprise that makes the agentic shift solely to cut costs by replacing people—while remaining captivated by a purely technological vision—falls prey to rigidity and fragility. Such a company risks failure when confronted with unprecedented scenarios or edge cases that confuse and mislead artificial agents. Only humans can grasp the underlying reasons for the 'never-before-seen' and determine how to navigate the unfamiliar. Constrained by historical training data, Artificial Agents are easily misled by unusual occurrences. These errors can propagate, causing potentially irreparable damage to the organisation. Moreover, AAs remain blind to the individual and social emotions that arise from events deviating from the norm.

It takes wisdom to reconcile the humanistic vision with the technological one. Wisdom lets us apply knowledge from intelligently processed data by

Artificial Agents to benefit others and ourselves. Wisdom is the capacity to discern and make judgments to achieve that outcome. We aim for the enterprise in which Artificial Agents operate to be wise, allowing it—through a far-sighted vision—to expand its scope from mere data processing to a profound understanding of its operational context. This aspiration rests on the conviction that an agentic enterprise must remain firmly anchored in humanity rather than be entirely data-centric. Without the guiding influence of wisdom, data-driven efficiency can entail enormous risks. For instance, an Artificial Agent tasked with maximising profit might fulfil that objective while neglecting safety—especially when the enterprise operates in ambiguous "grey areas." Conversely, a wise enterprise makes ethically aligned decisions, prioritising safety, fairness, long-term sustainability (for both itself and its environment), and human well-being over short-term metrics. Furthermore, wisdom illuminates the tacit, distributed organisational knowledge that refines reasoning and decision-making by drawing on past successes and failures when confronting new cycles of data.

In a wise agentic enterprise, Artificial Agents are not confined to narrow, isolated silos; instead, they operate within a holistic cognitive framework. This allows them to observe geopolitical shifts that foreshadow market changes, as well as the ambiguities and behavioural fluctuations among employees, partners, and customers that create operational bottlenecks. Based on these observations, the auditors will issue reports to the company, whose assessments will in turn guide the auditors; this will prevent them from making erroneous independent decisions.

Servitization: A key activity for "agentic" enterprises

The first wave of agentic enterprises appears to focus on service provision—"servifacturing," or servitization—rather than product manufacturing. These are the primary activities that create new shared value between the agentic enterprise and the customer:

- Customer service support.
- Contract review and regulatory compliance.
- Corporate purchasing planning and continuous monitoring of logistics flows.
- Coordination and automation of scheduling within the corporate team.
- Specialising in IT and information system management, development, and security.
- Facilitating the process of buying and selling goods or services by extracting meaning from data and conversations between buyer and seller, leveraging the analytical and predictive capabilities of AI agents.

Through these interventions, servitization improves the efficiency and effectiveness of how existing physical objects are used, thereby reducing the need to manufacture new ones.

A long line of people—myself included¹—have decried the extent to which nature is afflicted by an "inflation of stuff": the sheer number and weight of goods, objects, items, or whatever else one might call them. Everything humanity has built and produced (houses, factories, roads, vehicles, appliances, furnishings, plastics, computers, glass, weapons, etc.) now equals the mass of all living beings on Earth. At the beginning of the twentieth century, human-made objects accounted for only 3% of total biomass. Today, we produce approximately 30 billion tons of "stuff" annually. Since then, human production has doubled every twenty years. By 2040, anthropogenic mass—comprising both objects in use and waste—will reach the staggering figure of three teratons. Conversely, since the first agricultural revolution 12,000 years ago, plant biomass has steadily declined. It now appears to have been cut in half—dropping from 2,000 billion tonnes to just over 1,000 billion. This spells trouble for the planet. Scientists' alarm should be cause for concern. Here is what some of them stated in the Harvard Business Review in 2020:

Humanity has become a dominant force in shaping Earth's face. An emerging question is how the overall material output of human activities compares to the overall natural biomass. Here we quantify the human-made mass, referred to as 'anthropogenic mass', and compare it to Earth's total living biomass, which currently equals approximately 1.1 teratonnes. We find that Earth is at the crossover point; in 2020, anthropogenic mass, which has roughly doubled every 20 years, will surpass all global living biomass. On average, each person on the globe produces anthropogenic mass equal to more than their body weight every week.²

While there is a diverse debate about human fertility, little is said about the fertility rate of objects—from those that reduce physical exertion to others that adorn our bodies and environments without engaging our minds. Becoming aware of the weight of things—especially the many superfluous ones—and acting accordingly could help overcome the paradox of poverty amidst abundance and, in turn, liberate humanity from the economic problem of scarcity. Happiness equated with economic fortune, alongside widening disparities in income and wealth, fuels conspicuous consumption, while material and intellectual poverty is spreading beyond specific, endemic contexts to affect broader segments of the population.

¹ Piero Formica (2023). *Sciencepreneurship: Science, Entrepreneurship and Sustainable Economic Growth*. Emerald Publishing Limited: Howard House, Wagon Lane, Bingley.

² Elhacham et al. (2020). Global human-made mass exceeds all living biomass. *Nature*, 588, 442–444.

The time is ripe for a new landscape populated by "knowledge entrepreneurs" who possess intellectual capital ready to be transformed into intangible assets that enhance the human mind. This mind can then produce a vast array of content—such as visual cultural assets—by leveraging intangible inputs like knowledge, information, and ideas. This may well be the social breakthrough that offers an alternative to the materialistic welfare state of the post-war era—a time characterised by closets overflowing with material goods, yet by drawers of memory and ideas left empty. Housing, clothing, food, transport, and other essentials hold genuine value only when balanced against the provision of ideas. Otherwise, material goods fall prey to the "character" of financial speculation—that road to hell paved with abandoned good intentions—compounded by technological speculation that renders perfectly functional products obsolete.

Furthermore, when it comes to the sheer weight of material goods, it is reductive to believe the circular economy offers a complete solution. Collecting end-of-life products and materials and restoring them to a reusable state requires energy input and new materials. In some cases, recycling and reuse can have a greater environmental impact than producing with virgin resources.

The new landscapes of economic geography

Like the factory during the Industrial Revolution, the agentic enterprise reshapes both urban and non-urban landscapes.

The dematerialisation of decision-making flows and the optimisation of physical flows are profoundly changing economic geography. In the agentic enterprise, processes for choosing and approving documents, once paper documents locked in folders, stored in special spaces, subject to manual handoffs, and the subject of in-person meetings, are carried out using management software, managerial process platforms, and cryptographic tools. Therefore, for example, the need for dedicated meeting rooms (participation is hybrid: some members will express opinions and vote remotely), as well as document archiving and sending, is eliminated.

In cities, agentic enterprises need fewer traditional offices. Co-working hubs will take their place. These shared workspaces create a professional ecosystem that fosters collaboration and innovation, bridging the need for strategic in-person meetings with the flexibility of digital work. Fleets of drones and small autonomous vehicles revolutionise agency logistics, which no longer requires large warehouses to store and move goods. The centre of gravity of logistics shifts from mega-interports outside the city to a capillary network within the urban fabric. By coordinating with waste,

lighting, and traffic-flow managers, the agentic enterprise helps optimise urban resource use around peaks of economic activity.

The agentic enterprise strengthens extra-urban and rural landscapes, moving beyond the transit-zone stage to become fully automated distribution centres that operate 24/7 without needing lighting. This undertaking also equips these landscapes with small-scale computing and data storage infrastructure spread across the territory, physically positioned as close as possible to the end users and devices that generate the data, and supported by renewable energy plants dedicated to supporting the enormous computational load. The objective is to reduce the delay between sending a question to the Artificial Agents and receiving a response to zero.

Let us now consider how the profile of Industrial Districts (IDs) and Special Economic Zones (SEZs) shifts under the economic force exerted by agentic enterprises.

Industrial Districts are characterised by a community sharing a common heritage of knowledge, culture, values, and language—an agglomeration of specialised firms bound by a dense network of economic and social relationships, where social and economic dimensions are closely intertwined. In the presence of agentic enterprises, these districts will be defined by their agentic platforms. Such platforms go beyond merely responding to a specific request (a "prompt"); they autonomously design and execute multi-stage work plans to achieve a final objective. The facilities that currently house small and medium-sized enterprises—focused on territorial production specialisations that may improve but remain fundamentally unchanged—will instead host production lines driven by cognitive robots. These robots learn from experience, interact fluidly with a changing environment, and can—in real time—alter production to align with market trends identified by artificial agents, whether working independently or collaborating with human agents.

A Special Economic Zone (SEZ) is a designated geographic area where the government applies economic laws that differ from those in the rest of the country, intending to attract both foreign and domestic investment. SEZs will need to attract the algorithmic capital required by agentic enterprises for their production processes. This represents a form of "platform capitalism" aligned with the agentic enterprise's economic and social model, in which artificial intelligence algorithms, data analytics, computational imagery, and digital platforms are the primary factors of production. Furthermore, agentic enterprises will demand that SEZs provide ultra-fast digital infrastructure and subsidies for the energy their computational operations consume. The transformations brought about by industrial enterprise must also be examined for their impact on people and the surrounding territory.

In the era of factories powered by a combination of coal—used as fuel—and water converted into steam, Charles Dickens coined the term “Coketown” to depict industrial landscapes characterised by a stifling atmosphere, dominated by black smoke, with rivers that had turned a purplish hue and emitted foul odours due to dyes and industrial waste. Other writers described cities and regions shrouded in thick, greyish fog and a perpetual blanket of soot and chemical emissions—working-class neighbourhoods buried under the debris of combustion. The natural environment had lost its biological landscape, replaced by the factory, which had become a vast geographical expanse of a metallic nature.

Agentic enterprises—the deployment of autonomous artificial agents—give rise to at least two types of diseconomies: one material and one immaterial. The first concerns training artificial agents and having them execute tasks that require exponentially more computing power, sharply increasing energy costs. Data centres generate immense heat; to prevent overheating, they consume millions of litres of fresh water daily for evaporation and cooling. This water extraction conflicts with local communities' needs, lowering groundwater levels, restricting residents' access to drinking water, and harming local agriculture. Furthermore, data centres consume a significant share of a region's energy output. This immense energy demand can destabilise local power grids and increase residents' risk of blackouts. Moreover, rising demand drives up wholesale energy prices, ultimately increasing household utility bills. To cope with demand spikes, operators activate fossil-fuel-powered power plants, which worsens local air quality. Furthermore, the land required to build data centres—and the resulting loss of green spaces and agricultural land—is by no means negligible.

From an immaterial perspective, agentic enterprises introduce legal challenges—primarily in the environment where they are located—related to their autonomous actions and to cybersecurity. Their vulnerability to cyberattacks could spread to local digital networks, potentially damaging power and transport infrastructure, hospitals, public offices, and more. An agentic enterprise's improper action could entangle the local community in legal disputes, leaving it without local legal mechanisms to defend itself against the software's independent, wrongful conduct. If such an enterprise were integrated into one or more networks controlled by tech giants—entities often inclined to operate within extrajudicial bubbles, shielded from investigations by local judicial authorities in the event of cybercrimes or abuses—a conflict would inevitably arise between the enterprise's actions (aligned with those tech giants) and the local authorities.

The new landscapes of innovation

The agentic enterprise promises a flourishing spring of innovations that:

- Create markets currently unimaginable.
- Radically alter the way things are done—for instance, by transforming value chains (the range of activities required to create a product or service).
- Shift human behaviours and social structures by upending the mental maps with which we are familiar.

At the heart of this promise lies the agentic enterprise's propensity to subvert the status quo rather than merely improve or optimise it—as incrementalist firms do. Such an enterprise adopts a configuration that enables radical, transformative choices. Its approach to innovation is rooted in both generative and social dimensions. From the former emerge realities previously inconceivable; by dismantling dominant categories of thought—such as redefining "work" from physical exertion to creative endeavour—these innovations give rise to entirely new actions and relationships. From the latter dimension, innovation emerges as a social actor—even before an economic one—redefining culture and institutions and influencing humanity's evolution.

The agentic enterprise's disposition is a mental stance, born of intentionality. It filters reality and interprets events, aiming to anticipate market expectations rather than simply reacting to them. It rejects the path already laid out—often dictated by its own history of past successes. It breaks free from established cognitive and operational routines, even if doing so requires "cannibalising" its own existing products. Rather than correcting errors against previously set objectives, it chooses entirely new goals, thereby altering its ecosystem's existing structure. It accounts for the potential failure of this endeavour. It reframes it as an investment in knowledge—a way to explore unknown cognitive territories, areas of thought not yet experienced, with the ultimate goal of mapping them.

The architects of this undertaking are both human and artificial agents. To bring new partners into the innovation process, the former cultivate human relationships to assemble the necessary forces, while the latter manage the data essential for reducing bureaucratic friction. Key partners in shaping these innovation landscapes include:

- Universities and centers for cognitive exploration, where frontier research laboratories can be co-designed to investigate *Terrae Incognitae* (unknown lands).
- Startups, which serve as outposts for exploration.
- Suppliers who support new operational routines.

Readings

Bandura, A. (1997). *Self-Efficacy: The Exercise of Control*. San Francisco, California: W. H. Freeman and Company.

Dickens, C. (1854). *Hard Times*. London, UK: Bradbury & Evans.

Elhacham, E. et al. (2020). Global human-made mass exceeds all living biomass. *Nature*, 588, 442–444.

Formica, P. (2023). *Sciencepreneurship: Science, Entrepreneurship and Sustainable Economic Growth*. Emerald Publishing Limited: Howard House, Wagon Lane, Bingley.

Formica, P. (20246). *From Special Economic Zones to Special Ideation Zones*. The International Core Academy of Sciences and Humanities and the Ke Rui Academy, the research institute of the International Core Academy.

English Publication: <https://www.kerui.org/articlesdetails2.aspx?ID=1743>

ChinesePublication: <https://www.kerui.org/articlesdetails2.aspx?ID=1742>

Milgram, S. (1974). *Obedience to Authority: An Experimental View*. New York City, NY: Harper & Row.