




# Impact of Autonomic Computing on Process Industry

Walter Quadrini <sup>1,\*</sup>, Simone Arena <sup>2,\*</sup>, Sofia Teocchi <sup>1</sup>, Francesco Alessandro Cuzzola <sup>1</sup> and Marco Taisch <sup>1</sup>

<sup>1</sup> Department of Management, Economics and Industrial Engineering, Politecnico di Milano, 20156 Milano, Italy; sofia.teocchi@polimi.it (S.T.); francescoalejandro.cuzzola@polimi.it (F.A.C.); marco.taisch@polimi.it (M.T.)

<sup>2</sup> Department of Mechanical, Chemical and Industrial Engineering, University of Cagliari, Via Marengo 2, 09123 Cagliari, Italy

\* Correspondence: walter.quadrini@polimi.it (W.Q.); simone.arena@unica.it (S.A.)

## Abstract

Traditional sustainability frameworks in large scale production systems, such as Process Industry (PI) ones, often overlook operational resilience, creating a “resiliency gap” where systems optimized for efficiency remain vulnerable to disruptions. This study addresses this gap by proposing and empirically validating a Quadruple Bottom Line (4BL) framework that integrates resilience as the fourth pillar alongside economic, environmental, and social goals. The purpose is to evaluate the impact that Autonomic Computing (AC) can imply in this perspective. A Procedural Action Research (PAR) methodology was conducted across four distinct PI industrial cases (asphalt, steel, pharma, and aluminum). This involved the ECOGRAI framework to qualitatively link strategic companies’ objectives to shop-floor Key Performance Indicators (KPIs), guiding the assessment of AC systems. The results show benefits at a business level observed following the introduction of AC systems, which were implemented for enhancing resilience by managing ML model drift. Key findings include reduction in plant downtimes, decreases in waste (steel), reductions in gas consumption, and improved operator trust. This research provides empirical evidence that AC can make resilience an actionable component of industrial strategy, leading to measurable improvements across all four pillars of the 4BL framework. Its contribution is methodological and operational, aiming to demonstrate feasibility and causal plausibility.

**Keywords:** autonomic computing; self-X; MAPE-K; process industry; ECOGRAI; sustainability; KPIs; resilience; Quadruple Bottom Line

## 1. Introduction

The process industry (PI, including metals, pharma and chemicals, pulp and paper, cement, etc.) represents a challenging field where various innovative mechatronic technologies are continuously subject to experimentation in various directions to achieve improvements in terms of resources and energy consumption, operative costs, environmental emissions, work safety, etc. [1]. At the same time, this industrial field is considered “mature” because it is based on consolidated processes that can hardly be substituted by alternative ones guaranteeing the same production volumes and performance.

Historically, efforts to modernize these mature processes have focused on siloed optimization—targeting isolated metrics such as energy input per ton of product or reductions in specific raw material usage. While valuable, this approach is deemed insufficient, as the global push towards comprehensive sustainability, crystallized in frameworks like the Triple Bottom Line (3BL) [2] and the UN’s Sustainable Development Goals (SDGs), demands a more holistic perspective. The reasons for this perspective change have been discussed in



Academic Editor: Hao-Chiang Koong Lin

Received: 18 November 2025

Revised: 22 December 2025

Accepted: 4 January 2026

Published: 14 January 2026

**Copyright:** © 2026 by the authors.

Licensee MDPI, Basel, Switzerland.

This article is an open access article distributed under the terms and

conditions of the [Creative Commons](https://creativecommons.org/licenses/by/4.0/)

[Attribution \(CC BY\)](https://creativecommons.org/licenses/by/4.0/) license.

sociology, being referred to as the so-called “macro-trends” or “game-changers” [3], which, as noticed by Worthington and Patton [4], led the manufacturing sector to move from profit-oriented businesses [5] to policies oriented towards sustainability [6]. This perspective, also after slow phenomena such as the so-called “climate change” [7], or abrupt ones such as the COVID-19 pandemic [8], the 2020–2023 global chip shortage [9], and the current geopolitical situation [10], also affected PI, which is characterized by energy-intensive processes, high volumes, and values of raw materials stored as WIP and in stockpiles. This led the sustainability goals to be seen not as objectives hierarchically organized after the economic efficiency, but as interconnected components of a single and complex system, requiring new methods able to simultaneously face these goals.

However, most of the traditional sustainability frameworks do not consider the resiliency dimension, defined as the capacity of a system to anticipate, face, and adapt to unexpected disturbances [11]. Indeed, a process that is efficient and green under ideal conditions but prone to costly downtime when faced with raw material variability or equipment failure is not truly sustainable. This “resiliency gap” is a major challenge for PIs, where process disturbances can lead to cascading failures, significant material waste, and safety hazards. Addressing this gap requires moving beyond static optimization to enable dynamic and adaptive systems to self-correct, eventually through intelligent decision support systems [12]. This paper argues that to achieve true industrial sustainability, resiliency must be elevated to a fourth, co-equal pillar alongside the traditional TBL.

A primary barrier to implement such a multi-pillar framework is methodological. High-level business objectives (BOs) like “improve productivity” or “reduce CO<sub>2</sub> emissions” are often disconnected from the real-time, operational key performance indicators (KPIs) measured on the factory floor. Without a clear, quantitative link between these two levels, strategic goals remain abstract and operational improvements cannot be reliably mapped against their business impact. A structured methodology is needed to systematically cascade strategic objectives down to functional and, ultimately, measurable technical KPIs [13]. This ensures that actions taken at the process level are verifiably aligned with the company’s overarching sustainability and business targets.

To address this challenge, this paper adopts expanded view of industrial performance, moving beyond the traditional TBL to explicitly include resiliency as a fourth pillar alongside economic, environmental, and social goals. We present a procedural action research [14] study conducted within four PI enterprises that developed and deployed novel autonomic techniques to distinct PIs: asphalt, steel, pharmaceuticals, and aluminum production. The core contribution of this paper is threefold: (i) it proposes and empirically grounds a sustainability framework (under economic, environmental, social, and resiliency dimensions) as a more complete model for assessing industrial sustainability; (ii) it details a structured methodology that qualitatively links high-level business objectives to shop-floor KPIs, making strategic goals actionable and measurable; and (iii) it provides empirical evidence from real-world industrial use cases demonstrating how an AI-driven, human-in-the-loop system can deliver simultaneous improvements across all four dimensions of this expanded sustainability framework.

## 2. Materials and Methods

### 2.1. Sustainability and Resilience in the Process Industry

PIs transform raw materials at a scale usually via energy intensive and continuous operations. Their notable high footprint in terms of greenhouse gas emission, water consumption, and waste production positions this sector as both a challenge and a lever for meeting sustainability goals [15]. Indeed, the concept of sustainability has been significantly influenced by the so-called “Brundtland report” [16], which focused on inter-generational

equity, in particular after the recent macro-trend of the so-called “climate change” which increased the focus of consumers and public opinion on the environmental aspects of industry (e.g., the effects of manufacturing on resources such as air, water, land, and mined materials [17]). Therefore, traditional approaches of sustainability in the process industry have focused on optimization problems such as reducing resource usage [18], exploiting resources with a minor environmental footprint [19], and increasing energy efficiency, which matched well-known indicators such as life-cycle assessments [20].

However, other macro-trends such as the increased sensitivity about workers’ conditions and demographic phenomena (e.g., the workforce ageing which affects heavily industrialized areas such as Europe, Japan and, lastly, China [21]) highlighted the opportunity to not consider sustainability as a sole environment-related optimization objective, but to elevate it to an integrated property of the entire production system, including other policy targets. This approach materialized in the well-known 3BL [2], which defined three complementary pillars to address the economic, environmental, and social dimension of sustainability, also framed in the 3P (profit, people, planet) framework introduced by Kaptein and Wempe, which insists more on the interconnection and complementarity of these three aspects [22]. Additionally, recent global disruptions have exposed some limitations of this view too; the increasing frequency of supply chain shocks, geopolitical instability, and the undeniable impacts of climate change have forced industries to recognize that operational efficiency is fragile without operational resilience [23].

This has led to a paradigm shift in how sustainability is framed in advanced manufacturing—the focus is indeed moving from a static, efficiency-based model to a dynamic, resilience-based one. This emerging view frames sustainability as not just being about optimizing for a known present but also about preparing for an uncertain future. It requires a wider approach that balances economic, environmental, and social performance (the traditional 3BL) while simultaneously enhancing the system’s adaptive capacity, which has been formalized (with some lack of imagination) in a so-called “Quadruple Bottom Line” (4BL) for the last decade, where in the original 3BL a fourth line has been added (namely “Governance”), which is in charge of forecasting changes and applying mitigation actions through organizational policies [24].

Recent academic work has highlighted that there is no unique canonical formulation of a “fourth” bottom line: different elements of the literature add governance, institutional capacity, or resilience-related constructs as the fourth pillar depending on the disciplinary perspective and policy intent. For example, governance-oriented versions of the Quadruple Bottom Line emphasize institutional arrangements, stakeholder representation and corporate governance as the additional dimension needed to translate sustainability commitments into accountable action [25].

At the same time, a growing body of research in sustainability transitions and organizational studies treats resilience as either a complement to, or a reframing of traditional sustainability concerns. This literature argues that resilience and sustainability are related but distinct: traditional sustainability (economic, environmental, social) tends to emphasize long-term viability and normative targets, while resilience emphasizes the capacity to cope with and adapt to shocks and variability [26].

A parallel stream of work has focused on methods to operationalize resilience within organizations and production systems, explicitly recommending the use of measurable performance indicators and decision frameworks to integrate resilience with sustainability planning. Several recent contributions propose KPI-driven approaches, balanced-scorecard adaptations, or decision frameworks that explicitly combine resilience and sustainability objectives, thus enabling practical monitoring and managerial action rather than leaving resilience at a conceptual level [27].

These studies underline two practical implications particularly relevant for process industries: (i) resilience must be assessed with metrics that capture variability, recovery times, and adaptive capacity (not only mean performance), and (ii) resilience assessment benefits from being linked to concrete decision centres and control loops so that mitigation actions can be triggered and evaluated in operational time horizons [28].

Building on this literature, the originality of the 4BL framing [24] adopted in this study is primarily methodological and operational rather than conceptual. Instead of introducing an additional normative sustainability pillar, the fourth dimension is interpreted as an enabling layer that structures how sustainability objectives are monitored, governed, and acted upon in practice. Specifically, it provides a coherent mechanism to embed resilience-oriented considerations within existing sustainability assessment and management processes, supporting adaptive decision-making in process industry operations.

## 2.2. Autonomic Computing

The challenge of embedding resilience into complex industrial systems, however, requires tools capable of dynamic adaptation. In particular, this states the concerns in modern manufacturing scenarios where, after the spread of initiatives such as the so-called “Industry 4.0”, the production environments increased their data richness and interconnection, overwhelming the human operators’ and supervisors’ ability to intervene on the system’s controllers. A similar issue appeared in the early 2000s, when computer scientists noticed how the increase of complexity of computer networks would have made it hard for system administrators to cope with software errors and bugs.

To cope with this, in 2003, Kephart and Chess proposed a new paradigm named Autonomic Computing (AC) [29], a model inspired from biology that was aimed to allow multi-agent software systems to self-manage, relieving system administrators from most of the tasks related to problem identification, bug fixing, and restoring of the system’s performance after disruptive events. According to the same foundational work, an autonomic system is composed by an interconnected network of autonomic modules, each of which implement the self-management behavior by embodying four properties, defined as self-X or self-CHOP (after their first letters) [30]:

- self-Configuration: the ability of a module or a system to adapt its configuration automatically under normal operating conditions, based on high-level objectives or expected external changes (e.g., new requirements, workload variations, updates);
- self-Healing: the ability of a module to restore availability after a malfunction, by detecting, diagnosing, and correcting errors or failures whenever possible;
- self-Optimization: the ability of a module to improve its performance in a reactive and proactive way;
- self-Protection: the ability of a module to detect in advance external factors potentially leading to failures.

While these properties were originally conceived for software systems, they map directly onto general challenges of resilience and efficiency in industry. In this physical context, for example, self-Healing can be seen as the property of a controlled system to autonomously react after process disturbances, such as unexpected variations in raw material quality or equipment degradation, preventing production loss or quality issues. self-Optimization can be inflected by translating the continuous improvement of production recipes or the automatic retraining of AI models to maintain peak predictive performance, directly impacting resource efficiency. Finally, self-Configuration and self-Protection can relate to the system’s ability to adapt to changing production schedules and to prevent minor process deviations from cascading into major system failures.

Apart from the well-documented recent literature in discrete manufacturing where AC has been extensively used to cope with flexibility issues derived by mass-customization and usually translate in the need of an adaptive orchestration of the production assets [31–33], AC also has a history of successful cases in mitigating issues related to the adaptability of production systems, since retrieving relevant use cases from the literature has demonstrated it as a promising tool to influence PI scenarios by increasing the processes' robustness against change of conditions. For example, it has been employed in food production to estimate the control of the properties of meat after analysis of fodder and weather conditions [34]; in metal forming processes, it has been successfully deployed to make measurement systems able to react against sudden misplacements of sensors, which is a common issue of several PIs subdomains given the high energy released, in various forms, during the transformation from raw materials to finished products [35].

A standard control pattern for implementing such systems is the MAPE-K feedback loop. This model provides a structured and real-time, looping control strategy based on the following sequential phases:

- Monitor: the phase of the control or supervising system to collect data from sensors and other sources.
- Analyze: the phase of detecting anomalies, performance deviations, or opportunity for optimization.
- Plan: the phase of setting up an action plan to address the issue detected by the Analyze phase (e.g., retraining a machine learning model, modifying some control parameters, or restoring a software reinstalling an older patch version).
- Execute: the phase corresponding to the actuation of the aforementioned one.
- Knowledge: not a phase, but a common layer of instruction, policies, process models, and historical data provided to support the control/monitoring phases.

While AC is positioned as a key enabler for achieving resilience objectives, its successful adoption needs some tangible results across some real or real-like implementation scenarios, eventually considering all the four pillars of the 4BL. The challenge, therefore, shifts from siloed optimization to integrated performance measurement. This context frames the primary research inquiry: what is AC's impact on the economic, environmental, social, and resilience performance of PI?

### 2.3. KPI Identification

The aforementioned research question requires a further investigation about how to effectively measure the impact of a specific technology. To evaluate the impact of introducing new tools into an existing or new production process, and to effectively measure its impact, businesses usually rely on key performance indicators (KPIs). The selection of a proper methodology for defining these KPIs is therefore a critical step, as many conventional frameworks present significant drawbacks when applied to complex industrial environments. A careful examination of these limitations reveals the need for a more systemic and process-oriented approach. In this perspective, the identified KPIs can be regarded as the conceptual formalization of the "Monitor" phase within the MAPE-K loop, as they translate high-level objectives into observable and measurable parameters that guide data collection and subsequent analysis. Properly defined KPIs thus provide the foundation for an effective monitoring system, enabling both reactive and proactive management of the production process.

For instance, the widely used SMART (Specific, Measurable, Attainable, Relevant, Time-bound) pyramids methodology [36], while mainly used for ensuring that individual KPIs are well-ranked, falls short in providing a systemic structure. It offers conceptual rating scales for prioritizing KPIs but does not explicitly address the underlying actions

that must be manipulated to improve system performance. This leaves a gap between measuring an outcome and understanding how to influence it, making it difficult to create a truly actionable performance management system [36], since the framework focuses more on KPIs themselves, than on their integration within a cause-and-effect operational context [37]. Similarly, the Supply Chain Operations Reference (SCOR) model presents limitations when dealing with the roles of various actors within a process [38]. The SCOR methodology indeed tends to simplify the function of these actors, often treating them as mere inputs or outputs within the system, since it is mainly focused on supply-chain dynamics [39]. Furthermore, its fixed topology forces companies to frame their objectives into a rigid structure, limiting its applications for attributes of low-level functionalities (e.g., underestimating technology-related contribution, even if in the last releases digital technologies have been introduced [40]). The Goal-Question-Metric (GQM) approach [41], whose origin sits in software engineering, offers a logical, top-down path from high-level goals to specific metrics. While its structured procedure is an advantage for engineering analyses, its primary limitation lies in its abstraction from the underlying business or manufacturing processes: GQM is indeed mainly used to ensure that every metric serves a purpose, but it does not inherently provide a mechanism to model the process itself, not leading to the identification of bottlenecks and inefficiencies. As a result, while the resulting KPIs are logically sound, they may lack the operational context needed to drive process improvement on the factory floor. Other strategic frameworks, such as the Balanced Scorecard (BSC) [42], are powerful for aligning high-level business goals with different organizational perspectives (e.g., financial, supply chain). However, they often operate at a company level of abstraction; while widely used for strategic alignment, they typically lack the process-oriented granularity required to model the workflows and to focus on their decision nodes where new technologies, like autonomous computing systems, or humans—as the main decision makers—operate [43].


To address this gap, a more systemic and process-centric approach has been deemed needed. The ECOGRAI methodology [44], a framework recognized for its systematic approach to performance management and indicator design, offers a suitable solution. It is indeed a structured framework used to identify KPIs that can support an organization's strategic objectives. The multi-phase process begins with Phase 0, which involves a detailed analysis of the production system to identify key "Decision Centers" (DCs, the specific units where decisions are made and actions are taken). Previous research accomplished using process modeling tools like BPMN to visualize and understand the interactions between different agents and functions [45].

Following this, Phase 1 focuses on defining the objectives through a dual process: high-level business objectives (BOs) related to the organization's strategy are identified, often through analysis of business models. Concurrently, functional objectives (FOs) are identified for each DC, representing the specific goals within each unit. The critical step is then to explicitly link these two sets of objectives, therefore, a "split up" diagram or matrix is used to visualize and formalize how each FO contributes to the achievement of one or more BOs, ensuring a coherent alignment between operational functions and strategic goals. This structure acknowledges that while FOs support BOs, they are identified in a process that considers both top-down strategic needs and bottom-up functional realities. Subsequently, Phase 2 identifies the "drivers", defined as physical or logical variables that a DC can act on to achieve its assigned FOs. With this foundation, Phase 3 proceeds to the selection and definition of the KPIs themselves, reaching the distinctive triple set of objectives, drivers, and KPIs. The final stages, Phase 4 and Phase 5, involve an intervention on the information system to collect KPI data and integrating this system into the

broader production management information system, thereby embedding the performance measurement framework into the organization's daily operations [44].

#### 2.4. Research Framework: Procedural Action Research

To understand the impact of autonomic computing on sustainability and resilience, the work was framed into the form of procedural action research (PAR) [14], a structured form of action research [46] suited for technical and organizational improvement. This approach was chosen to systematically guide the collaborative process of problem definition and solution design. While embodying the participatory mindset of involving end-users directly, the research exploited the ECOGRAI framework and business process model and notation (BPMN) for process visualization. Introduced in 2016 by the Institute for Manufacturing, PAR is indeed a derivation of the better-known action research [46], which is characterized by an iterative process of inquiry designed to develop solutions through a participatory and collaborative activity. Procedural action research, on the other hand, tries to mitigate traditional action research limitations through a rigorous procedure, formally staged in five steps: definition, design, development, testing, and deployment; if the testing does not provide satisfactory results, it triggers a fifth step (namely refinement) which leads to a re-definition of the first step, as depicted in Figure 1.



**Figure 1.** Procedural action research [14].

This research framework benefitted from the opportunity to cooperate with four PI companies (involved in asphalt, steel, pharmaceuticals, and aluminum production) engaged in a research project (namely “self-X Artificial Intelligence for European Process Industry digital transformation”, funded by the European Commission through the Framework Programme Horizon Europe, through Grant Agreement number 101058715). This granted the authors the opportunity to continuously engage industrial counterparts committed in the adoption of AC technologies in their production sites, throughout an extended period of time, evaluating the AC lifecycle from its requirements' definition to its effective deployment.

The study's PAR approach was deliberately divided into two complementary stages to ensure a rigorous connection between process-level actions and business-level outcomes. The first stage allowed to map the operational landscape, using BPMN to model the production processes and formally identify the MAPE-K loops that the autonomic computing (AC) system is required to manage. The validated process models and control loops from this first stage became the inputs for the second stage. This subsequent stage was devoted to performance measurement, where the ECOGRAI methodology was applied to derive relevant KPIs. This structured approach allowed to ensure that the link between the AC-driven control action and their quantifiable business impact was transparent and grounded on solid bases.

More precisely, for the first stage, the steps were organized as follows:

- **Definition:** it materialized in the analysis, performed by the research institution leading this work, of the documentation related to the use cases' description and the related requirements reported by the industrial companies. This initial step was performed offline.

- **Design/Development:** these steps were carried on in conjunction with the industrial partners, and BPMNs based on the previous steps were designed. This step was performed through an initial workshop (one for each use case), which involved the leading research institution, the use case companies, and their system integrators. The initial workshop was followed by a series of interviews, conducted via email exchanges (the loop between Design and Development in Figure 1) where all the stakeholders concurred in validating or correcting the assumptions depicted in the drafted BPMN and in a final workshop where the BPMN was conceptually validated.
- **Testing:** this step consisted in a coherence validation of the BPMN via the Camunda software and was led by the leading research institution.
- **Refinement/Re-definition:** after a refinement phase which consisted of an asynchronous assessment, this step saw the identification of the MAPE-K loops in each of the pilots, as well as the eventual subdivision in the validated BPMN in more sub-models for the sake of clarity. This step has been performed by the leading research institute and was directly followed by the testing step to validate the coherency of the BPMNs and by an asynchronous assessment with the stakeholders.
- **Deployment:** after the last assessment, the BPMN was deemed validated for the second PAR stage.

The second stage, aimed at applying the ECOGRAI methodology to derive the KPIs, happened as follows:

- **Definition:** the leading research institution analyzed the documentation of the research project funding the activities. The strategic objectives targeted by the PI companies and agreed with the funding body (European Commission) were assumed as business objectives (BO), were listed case by case (ECOGRAI, phase 1), and were labeled after the type of objective they implemented according to the 4BL.
- **Design:** this step contained several inner loops of continuous improvement. A number between two and four collaborative workshops per use case, interspersed with corrective and validating email exchanges, were conducted by the lead research institution. This step aimed to:
  - a. Identify the decision centers (DCs) starting from the BPMNs outcome of the first stage (ECOGRAI, phase 0).
  - b. Identify functional objectives (FOs) for each of the DCs (ECOGRAI, phase 1).
  - c. Identify drivers for each of the FOs (ECOGRAI, phase 2).
  - d. Derive KPIs for each of the FOs/drivers (ECOGRAI, phase 3).
  - e. Couple the triplet FOs/drivers/KPIs with BOs via split up diagrams (ECOGRAI, phase 1).

This step was participated by the leading research institution, which led the activities and drafted the different ECOGRAI phases, by the end users, which validated or requested corrections to the drafts, and by the system integrators, which assessed the feasibility of retrieving the KPIs-related data.

- **Development:** this step (ECOGRAI, phase 4) was conducted by the end users and their system integrators, who formalized the rules to derive KPIs from the data available on the production shop and office floors. It materialized in telcos and email exchanges in a one-to-one format. The time-series KPIs reported in the results section are derived from production data collected by the industrial partners through their data infrastructures, including manufacturing execution systems (MES), laboratory information management systems, process sensors, and human-validated inputs generated during routine operations. In accordance with the ECOGRAI methodology, the study assumes the availability of systematic data logging for both automated process variables

and human-in-the-loop actions (e.g., alarm validation, recipe adjustments, retraining decisions), which are treated as notable events in the performance assessment.

- No additional instrumentation or experimental data acquisition campaigns were introduced for this study; instead, the analysis relies on industrial-grade, operational data streams that reflect real production conditions and constraints.
- Testing: this step was restricted to technical testing of the data pipelines available on the end users' sites and was attended by the end users' themselves and their system integrators. Therefore, refinements were exclusively of technical nature and did not materialize in any re-definition of the BOs.
- Deployment: this step (ECOGRAI, phase 5) involved the system integrators related to each of the end-users, they integrated the rules defined in the aforementioned one in the software architecture managing the AC functionalities. For sake of reference, this software architecture has already been presented to the scientific and practitioners' community [30] together with a reference implementation toolset that has been integrated in the use cases, therefore it will not be debated in this article.

### 2.5. Scope, Assumptions, and Methodological Boundaries

This study is grounded in the assumption that sustainability improvements in process industries can be meaningfully assessed through structured links between strategic objectives and operational KPIs. The adopted methodology assumes that (i) relevant production data are available and sufficiently reliable to populate KPIs, (ii) autonomic computing functionalities are deployed as decision-support systems within a human-in-the-loop setting, and (iii) organizational structures allow the identification of stable decision centers as required by the ECOGRAI framework.

The scope of the analysis is deliberately limited to operational and organizational impacts directly associated with AC-enabled functionalities. The study does not model market dynamics, macroeconomic conditions, or long-term investment effects, nor does it attempt to isolate AC impacts through counterfactual statistical designs. Instead, the research focuses on establishing transparent causal pathways between AC interventions, process-level performance indicators, and business objectives within the observed industrial contexts.

Accordingly, the results should be interpreted as context-dependent and exploratory, aimed at validating the applicability of AC as an enabler of resilience within a Quadruple Bottom Line framework, rather than as universally generalizable performance benchmarks.

### 2.6. Use Cases

The four use cases involved in the research are briefly summarized, explicating their AC requirements as follows [47]:

1. Asphalt industry use-case: "asphalt mix properties prediction, asphalt production process optimization and fault detection". The asphalt industry is quite a distributed process industry including a production site where several aggregates are mixed with bitumen, a laboratory quality system able to test the mechanical and volumetric properties of the material, and a complex network of dump trucks in charge of logistics and paving.

The AC-related tasks applied in the considered use case are as follows:

- a. Asphalt mix design and re-design by prediction of the theoretical produced material properties (both mechanical, volumetric, and composition properties). In other words, AC is supposed to predict in advance the compliance of the produced material with respect to the product specifications.

- b. Implementation of real-time fault-detection systems in the production-site PLC architecture (including motor failures, possible anomalous vibrations, burner efficiency deviations with consequent definition of an optimized burner setting).
    - c. Prediction of some of the material properties (e.g., the temperature of the mixture) upon reaching the paving zone. This tool allows to optimize the costs associated to material heating, transportation, and logistics for paving.
  2. Steel industry use-case: “scrap usage for EAF steel production”. The EAF (Electric Arc Furnace) is a conventional process used to melt recycled steel and operate an electric arc to produce recycled material. This type of process needs to face by means of AC technology the following tasks:
    - a. Ability to predict in advance the chemical composition of the produced steel-grade based on a “mix recipe” defining the amounts of raw material types used for each heat and optimizing consequently the EAF parameters. It is worth pointing out that the raw material utilized in the EAF process are recycled metal scraps classified in terms of chemical contents (in the following the term scrap-type will be used to refer to different categories of recycled scraps).
    - b. Planning the usage of the scraps by a moving-horizon cost optimization exploiting the available information concerning the production book, the current levels in the scrap yard, and the nominal chemical composition of every scrap. The planning tool can use the predictive tool defined at point “a” here.
    - c. Design and re-design of a “mix recipe” with respect to the usage of different scrap types. Indeed, for every heat it is necessary to satisfy the international standard requirements that specify for every chemical specie the min/max constraints according to the target steel-grade of the final product to be realized.
  3. Pharma industry use-case: “real time control for production of adrenosterone”. The pharma process is a pilot-size installation allowing, by means of an electro-chemical actuation, the conversion of cortisone to adrenosterone. The pilot installation aims at testing different possible production setups to verify experimentally what the most favorable ones are, allowing in the long term a reduction of operative costs, and particularly those related to electrical energy absorption and usage of consumable components like the electrodes that are subject to corrosion. This process is equipped with the following:
    - A power supply controlled by current intensity.
    - A FTIR (Fourier-Transform Infrared Spectroscopy) that measures vibrational modes of chemical bonds within molecules and records and infrared spectrum from which the concentration of the substances can be inferred.
    - An OCT (Optical Coherence Tomography) sensor for monitoring the electrodes where it is able to automatically evaluate the corrosion status from the OCT images.The tasks to be faced with this process are as follows:
    - a. Given the desired throughput of the chemical reaction to be realized it is necessary to find out automatically the most successful process settings in terms of maximization of the chemical reaction rate under acceptability constraints on the electric power consumption and on the electrode lifespan.
    - b. Regulating in real time the reaction rate by optimizing the settling time and the robust stability of the closed loop. This task will be faced by incorporating a predictive model of the process in the controller according to the so-called Internal Model Control (IMC) principle.
    - c. Automatically discover ineffective positioning of the OCT sensor being human actuated.

It is worth pointing out that the goal described at point (a) represents a multi-objective optimization problem with contrasting requirements since the reaction rate can be deterministically increased by increasing the electric current setpoint, but at the same time, an increment of the current setpoint unavoidably increases the electrodes corrosion rate and a consequent reduction of the electrode lifespan.

- Aluminum industry use-case: “scrap usage for aluminum production”. The aluminum use case is similar to the steel use case since the scrap yard and EAF processes are very similar. However, there are some differences concerning both the operative conditions (for instance the scrap raw material is systematically tested) but also in terms of challenges to be faced since it is limited to the “a” and “c” points already presented for the steel use case whereas the moving-horizon planning is not considered part of the use case (point “b”) if not as a future possible extension.

All the aforementioned requirements (summarized in Table 1) are current industrial needs of the selected use cases. In this case, they have been implemented by AC requirements, as AC is perceived as a tool to operationally mitigate the issues the requirements refer to. Additionally, the leitmotiv of the requirements is apparently connected (more or less strongly, requirement by requirement) with a need to make the production system robust against the flexibility of the processes. This highlights the opportunity of debating how AC can effectively imply operational and how these benefits can influence the strategic objectives of the cases.

**Table 1.** Use cases involved in the proposed research.

Use Case	Asphalt	Steel	Pharma	Aluminum
Core process	Aggregate mixing, heating, and paving	EAF steel recycling	Electrochemical reaction	Gas furnace aluminum recycling
Scale	Distributed production site with complex logistics	Large scale industrial production	Laboratory installation	Medium-scale industrial production
Key challenge	Material property prediction, process optimization, and fault detection	Scrap mix optimization, resource planning, and chemical composition prediction	Real-time control and multi-objective optimization	Scrap mix optimization and chemical composition prediction
AC application	Recipe design/re-design, fault detection, and logistics optimization	Predictive recipe design and scrap usage planning	Real-time process control and sensor monitoring	Predictive recipe design and diagnostics

All the AC-related tasks can be framed as MAPE-K loops. For instance, referring to the asphalt mix recipe design/re-design control task (the “a” task of asphalt UC), it allows not only to design a new product never produced before but also to re-design a recipe by defining what are the modifications to compensate a detected anomaly in previous production. In this case the MAPE-K modeling of this control task is represented by the following steps:

- M(onitoring) corresponds to the collection of the measured produced material quality characteristics.
- A(nalysis) corresponds to the verification if the quality characteristics correspond to the product requirements for the desired asphalt type.
- P(lanning) consists in designing a recipe or deciding what are the necessary trims to be applied to an existing recipe in case an adjustment is in place (re-designing).
- E(xecute) consists in actuating a possible modification of a recipe and transmitting it to the production site.
- K(nowledge) is represented by the international norms defining the standardized characteristics of a product, the recorded production data logs about the past pro-

duction and all possible past modifications applied to a product mix recipe with the corresponding quality results.

### 3. Results

This section analyses how AC functionalities are mapped to operational and organizational objectives within the proposed 4BL framework. In particular, the analysis focuses on the identification of DCs, FFOs, drivers, and KPIs using the ECOGRAI methodology, and on how these elements enable the translation of AC interventions into measurable sustainability and resilience outcomes. The purpose of this section is therefore to establish the analytical backbone linking AC mechanisms to BOs, which is subsequently validated through the industrial use cases.

The deployment of the first PAR stage for the asphalt use case led to the identification of the following actors in the BPMN derived from the processes under investigation: a centralized server (hosting the company's database related to production and operations, as well as its ERP), a data scientist (in charge of data analytics and AI training), a laboratory (in charge of quality control on produced asphalt), an automation system (over reigning the automated production plant), the production plant (in charge of transforming raw materials into asphalt mixture), the logistics and paving (in charge of moving the asphalt on the paving site and to, literally, pave the road), and the autonomic manager (a software system hosting the AC modules [30]).

The refinement of the BPMN allowed to identify the following interactions by the actors identified in the BPMN depiction and the MAPE-K loops designed to match the requirements presented in Section 2.5, where Table 2 depicts the matching and where the MAPE-K loops are identified according to Section 2.5.

**Table 2.** Identification of actors involved in the MAPE-K loops for asphalt use case.

MAPE-K Loop	Centralized Server	Data Scientist	Laboratory	Automation Plant	Production Plant	Logistics And Paving	Autonomic Manager
1a	✓	✓	✓	✓	✓		✓
1b	✓	✓		✓	✓		✓
1c	✓	✓				✓	✓

Consequently, the stage related to the steel use case defined via the BPMN analysis the following actors involved in the process: the plant's manufacturing execution system (MES), a laboratory (in charge of the quality control of melted steel), the EAF, the production manager, the scrap yard (where raw material is siloed), a data analyst (in charge of training the ML models to design production recipes), the laboratory (the actual production reactor, in its equipment and data acquisition system), and the autonomic manager. The matching with their involvement in the MAPE-K loops is depicted in Table 3.

**Table 3.** Identification of actors involved in the MAPE-K loops for steel use case.

MAPE-K Loop	MES	Laboratory	EAF	Production Manager	Scrap Yard	Data Analyst	Autonomic Manager
2a	✓	✓	✓				✓
2b	✓				✓	✓	✓
2c	✓			✓		✓	✓

The same approach was followed for the description of the Pharma use case, where the process modeling identified the following actors participating in the MAPE-K loops: a database (connected to an HMI and storing historical and live sensors' data), a data analyst (in charge of data analysis and ML models training), a scientist (supervising the chemical reaction outcome), and the autonomic manager. Since all the actors are involved in all the MAPE-K loops, no table is depicted.

The aluminum use case was again depicted as a BPMN, and the following actors have been identified: a centralized server (containing historical datasets and production recipes), a data analyst (in charge of the data analysis and the ML training), the melting department (in charge of the scrap fusion and its analysis), the production manager, the scrap yard (where raw materials are siloed and chemical cold tests are performed), and the autonomic manager. These actors take part in the different MAPE-K loops as depicted in Table 4.

**Table 4.** Identification of actors involved in the MAPE-K loops for the aluminum use case.

MAPE-K Loop	Centralized Server	Data Analyst	Melting Dept.	Production Manager	Scrap Yard	Autonomic Manager
4a	✓		✓	✓	✓	✓
4c	✓	✓				✓

The second PAR stage started with the identification of business objectives (BOs) for the four use cases. The identification started with a set of broad BOs common to all the use cases (set as a requirement by the funding entity). These BOs are listed in Table 5, where they are labeled after the pillar of the 4BL they frame into.

**Table 5.** BOs of use cases.

BO	Title	Economic	Environmental	Social	Resilience
1	Improve productivity	✓			
2	Improve quality	✓			✓
3	Improve response time	✓			✓
4	Reduction of raw material consumption	✓	✓		
5	Reduction of energy consumption	✓	✓		✓
6	Reduction of generated CO <sub>2</sub>		✓		
7	Increase of competences			✓	

The BOs have then actualized in each of the four use cases according to the specific businesses and processes of each of the four considered companies. For each use case, for sake of readability, in the body of the paper the BOs, the FOs, the drivers, and the KPIs connected to the autonomic manager, are only reported as deemed relevant for the current analysis. The Appendix A depicts the punctual identified items, where it can be noticed that some of the various BOs, FOs, drivers, and KPIs have been duplicated or merged, because of complete correlation, or because of unicity of the measurement source.

### 3.1. Asphalt Use Case

As depicted in Figure 2, the split up diagram lists all the BOs, FOs, drivers, and KPIs identified through the research. In order to properly identify the impact of AC, it is needed to start from the FO allocated to the autonomic manager, which is the DC where AC functionalities are embodied. A complete association list between DCs and FOs is

annexed in Table A5. In the asphalt case, the autonomic manager is supposed to minimize the detection time of the drift of ML models (FO14). This FO contributes to BOs related to the improvement of quality and reduction of response time. The specific declination of the cross-cases of BOs depicted in Table 5 is detailed in Table 6.

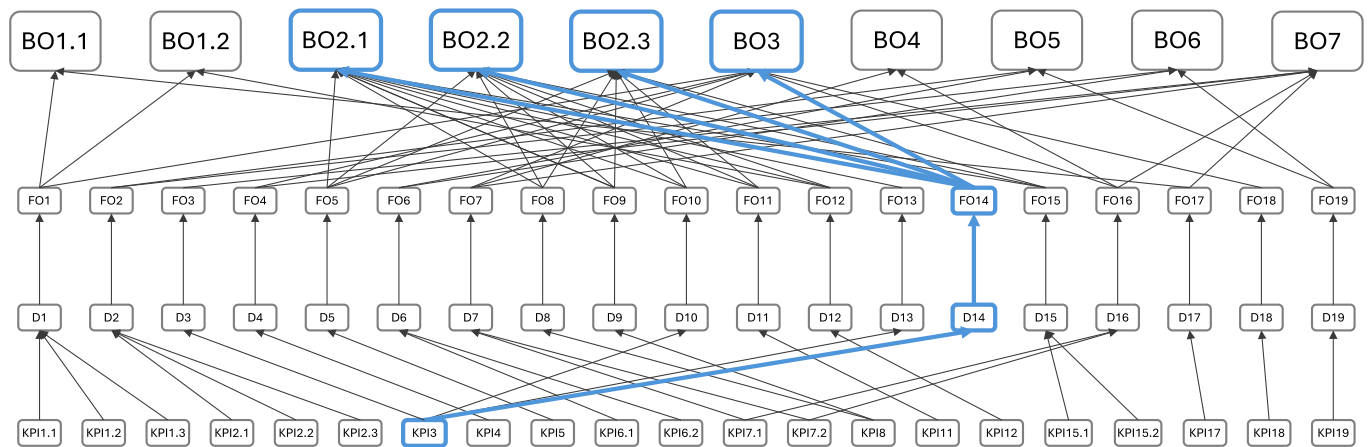


Figure 2. Asphalt split-up diagram.

Table 6. Autonomic manager-related BOs, FOs, drivers, and KPIs for asphalt use case.

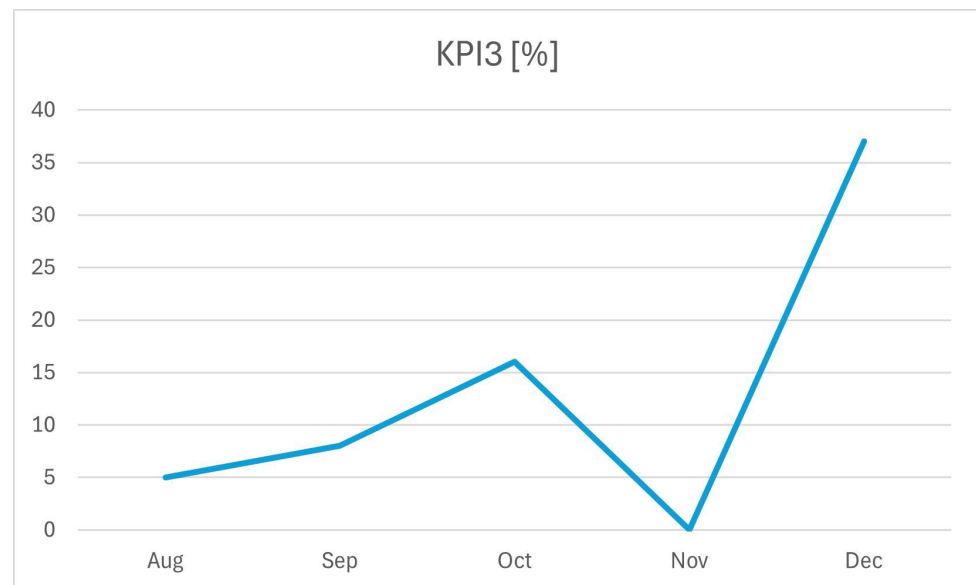
Item Class	ID	Description
BO	BO2.1	Improve adjustment of filler/bitumen rate
	BO2.2	Decrease deviations in dosing process
	BO2.3	Improve control of hot aggregates separation
	BO3	Improve optimization of maintenance schedule
FO	FO14	Optimize the detection time for a ML model drift vs. the process
Driver	D14	Definition of tolerances over the detection of possible model drift so as to reach a fair trade-off between the minimization of the detection time and the need of avoiding false alarms
KPI	KPI3	Number of datapoints collected after the last model training [% over the total number of available datapoints] outside a confidence interval of $\pm 5\%$ , $\pm 10\%$ , $\pm 15\%$ concerning the deviation between measurement and prediction

Moving down in the split up diagram, as a driver to control FO14, the tolerances over the detection of a model drift have been defined in order to achieve a trade-off between the minimization of the detection time and the avoidance of false alarms. The reader may note that, while the same FO can contribute to different BOs, the relationship between drivers and FOs are generally one-to-ones (unless duplicates).

At the operational level, KPIs are used to measure the performance of the FOs and, indirectly, the achievement of the BOs. For FO14 (and D14), the identified performance metrics have been found in the number of datapoints collected after the last model training (KPI13), since, if maintained in a reasonable interval at a data analyst’s discretion, it represents a trade-off between an excessive number of re-trainings (which could also imply overfitting issues) and a prompt detection of models’ drifts.

Table 6 depicts the punctual list of BOs, FOs, drivers, and KPIs connected with the autonomic manager.

KPI3, monitored for a period of six months and depicted in Figure 3, displays the monthly monitoring performed by the company, with the notable behavior noticeable in the month of December, where the weather conditions, which typically affect the asphalt production, may have led to some significant drift in the ML performance.



**Figure 3.** Asphalt KPI3.

On the other hand, and for the same period, concerning the BOs influenced by the FOs of the autonomous computing, the company reported the following achievements:

- BO2.1 (Improve adjustment of filler/bitumen rate): for mixes with low recycled content, the company registered in 2023 a percentage of mixes outside the tolerance range of 43.14% (195 of 452 mixes) which required reworking or corrections in the production plant; in 2024, this value decreased to 39.20% (138 out of 352 mixtures). For the mixes with high recycled content, the percentage of mixes outside the tolerance range was 21.01% in 2023 (25 of 119 mixes), 39.29% in 2024 (22 of 56 mixes), and decreased to 10.00% in the first months of 2025 (two of twenty mixes).
- BO2.2 (Decrease deviations in dosing process): the 2024 evolution of average ratios between actual and theoretical quantities used in the production of asphalt mix for aggregates components decreased by 21.8%, while the filler one decreased by 23.8%, the bitumen one increased by 8.6% and the reclaimed asphalt pavement (RAP) one increased by 7.4%. The temporary worsening observed in the RAP-related indicator between 2024 and early 2025 reflects a transition phase in which higher RAP percentages were intentionally tested following AC-supported recommendations (in the attempt to decrease the environmental impact of the production). This exploratory adjustment initially increased variability and reject rates. Such non-monotonic behavior is deemed consistent with adaptive process optimization under real production constraints.
- BO2.3 (improve control of hot aggregates separation) was not measured, since, because of the inherent nature of the asphalt manufacturing process, no traceability tools measuring the flow of each type of cold aggregates and relating them to the hot aggregates into the hot ones already separated by size, is currently available in the market.
- BO3 (Improve optimization of maintenance schedule): a reduction of 39% in the intermittent down-times of the plant was recorded.

### 3.2. Steel Use Case

The same approach was applied to steel use case, whose split-up diagram is depicted in Figure 4.

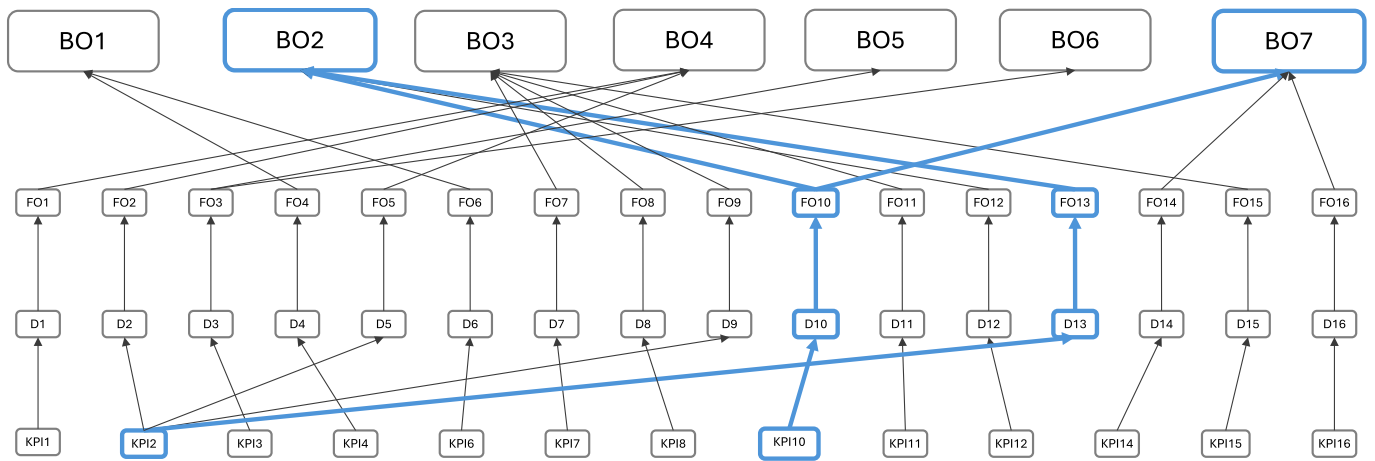


Figure 4. Steel split up diagram.

The steel use case saw two FOs assigned to the Autonomic Manager, namely FO10 (minimize number of occurrences before warning, warnings corresponding to false drift, maximize reliability of prediction of scrapyard criticalities) and FO13 (improve recipe accuracy). These FOs were influenced by the respective drivers (i.e., the definition of tolerances over detection of non-conformities, and the creation for production order for new recipe validation), which were monitored through the following KPIs: KPI2 (cost of additive material (chromium, nickel, and molybdenum) used per HeatID), KPI10 (number of occurrences before raising warning). Table 7 lists the aforementioned items.

Table 7. Autonomic manager-related BOs, FOs, drivers, and KPIs for steel use case.

Item Class	ID	Description
BO	BO2	Reduce chemical waste
	BO7	Introduction of a decision support system
FO	FO10	Minimize number of occurrences before warning, warnings corresponding to false drift, maximize reliability of prediction of scrapyard criticalities
	FO13	Improve recipe accuracy
Driver	D10	Definition of tolerances over detection of non-conformities
	D13	Creation for production order for new recipe validation
KPI	KPI2	Cost of additive material (chromium, nickel, and molybdenum) used per HeatID
	KPI10	Improvement of Mean Square Error (MSE) over the data points related to non-conformities

In the first month of 2025, after the introduction of AC in the production site, Figure 5 shows a decreasing trend in the amount of additions of chemical additives in the EAF process. On the other hand, Figure 6 depicts the alarms raised after laboratory testing: orange dots represent model deviations, while red dots indicate alarms, and purple dots indicate the alarms confirmed by the data analyst. More precisely, a single deviation exceeding the threshold for one heat does not necessarily imply a systematic issue; instead, the deviation must occur repeatedly over multiple heats before triggering an alarm.

To ensure a meaningful alarm generation, a threshold is set for the Mean Absolute Error (MAE) of Cu over a short period. If the MAE exceeds this threshold for a defined number of occurrences, an alarm is triggered, acknowledging a potential systematic anomaly in the scrap properties. It is worth noting that initially these alarms were raised after three deviations, which always triggered validations and correction actions by the data analyst, who retrained the ML model and set the trigger to five breaches. The consequent

improvements in retraining allowed the production process more affordable recipe designs, consistent with the decrease in the alarm points in the last weeks of production.

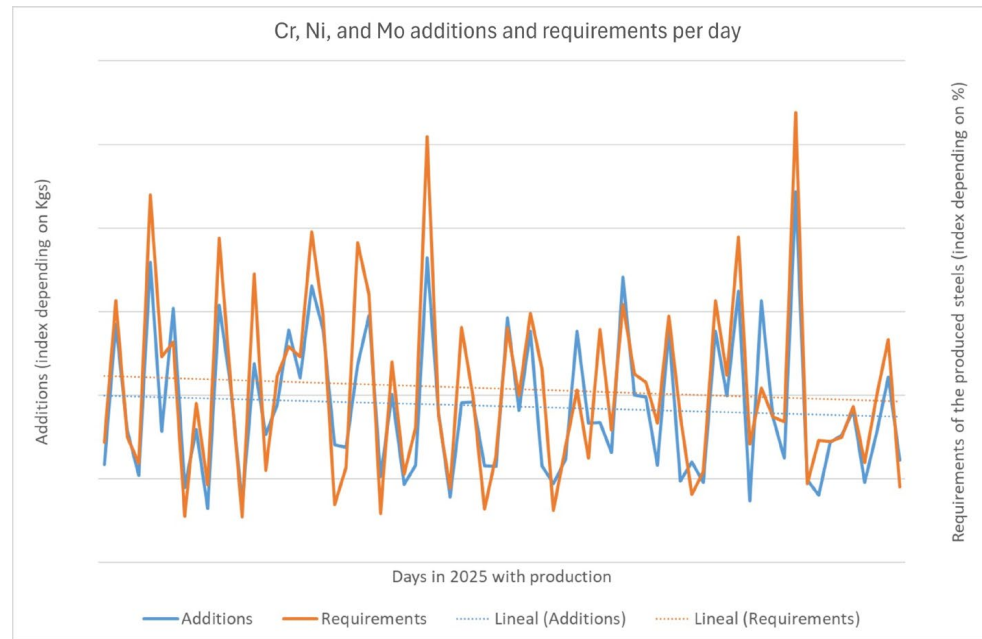


Figure 5. Steel KPI2.



Figure 6. Steel KPI10.

The BOs affected by the AC in the steel use case were as follows:

- BO2 (Reduce chemical waste): for this KPI, the company reported a decrease of waste of 6.1%, as a consequence of better adjustment of residual elements with respect to customers' specifications, fewer internal rejections due to quality defects, and a better controlled EAF, which induced a minor usage of O<sub>2</sub>, implying lower impurities generated in secondary metallurgy process.
- BO7 (introduction of a decision support system): after an internal survey conducted by the company, four interviewed operators over five approved the introduction of the decision support system, enabled with a scrap supervision tool to ease the understanding of the estimated chemical composition of the alloy.

### 3.3. Pharma Use Case

Similarly, the split up diagram of the pharma use case is depicted in Figure 7.

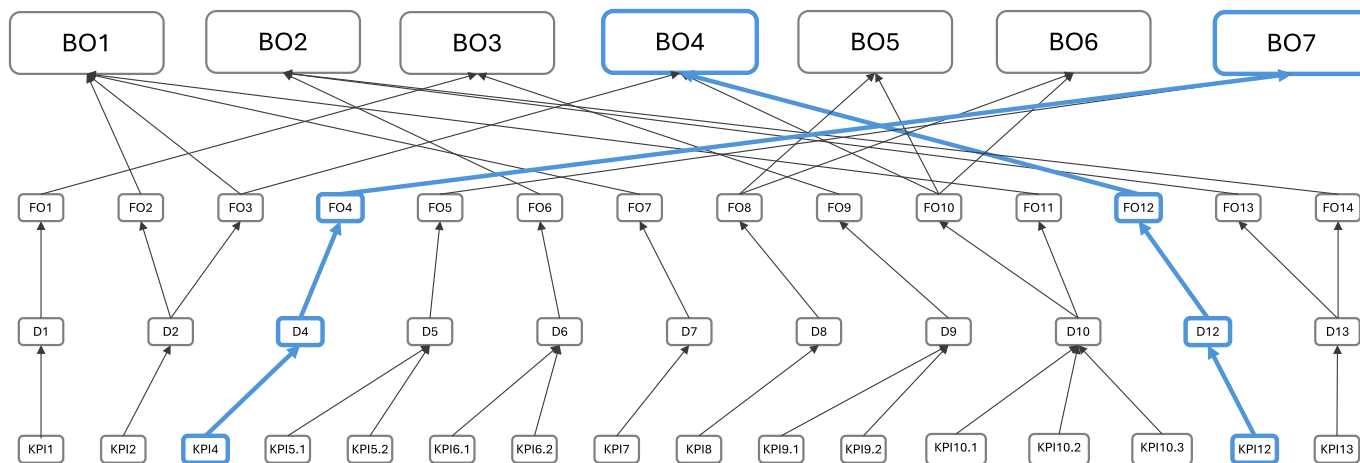


Figure 7. Pharma split up diagram.

The pharma use case allocated two FOs to the Autonomic Manager (namely FO4, minimizing the detection time for the possible wrong positioning of the OCT, and FO12, optimizing the detection time for a ML model drift), where the latter, symptom of an abrupt deviation in the experimental setup, can be influenced by the celerity in the possible manual intervention, while the latter one, symptom of a domain shift, can be controlled by the precise definition of tolerances over the detection of possible model drift. As depicted by Table 8, to monitor these variables KPI4 (time between experiment start and first actual correction of a wrong positioning of the OCT sensor) and KPI12 (number of times a deviation between measurements and predictions in adrenosterone percentage outside the confidence intervals defined by KPI9.2) are introduced. The iteration of experiments in the laboratory reactor (where one session is depicted in Figure 8) reported an average number of alarms of 4.3 per hour for what concerns the positioning of the OCT sensors, while the regression model supervising the reaction turned out to be robust enough to not provide any alarm in the tolerance range.

Table 8. Autonomic manager-related BOs, FOs, drivers, and KPIs for pharma use case.

Item Class	ID	Description
BO	BO4	Reduction of raw material consumption
	BO7	Embodiment of applications adapting to heterogeneous skill levels
FO	FO4	Minimize the detection time for the possible wrong positioning of the OCT
	FO12	Optimize the detection time for a ML model drift vs. the process
Driver	D14	Speed of the possible manual intervention during the execution of the chemical reaction to correct the possible wrong positioning of the OCT
	D12	Definition of tolerances over the detection of possible model drift
KPI	KPI4	Time between experiment start and first actual correction of a wrong positioning of the OCT sensor
	KPI12	Months for which KPI9.1 (namely “deviation between measurements and predictions in adrenosterone percentage”) was exceeded

Concerning the BOs, BO4 (decrease quality rejections) resulted in an improvement over several experiments [48] between 16.4% and 48.3% depending on different experimental parameters, while an improvement of 56% of image analysis accuracy diminished the image processing alarms and the OCT-related ones, increasing the operator’s trust in the system and his/her operational efficiency, correlated to BO7 (embodiment of applications adapting to heterogeneous skill levels).




Figure 8. Adrenosterone theoretical rate (green marks) against measured one (blue marks) under confidence interval (green area) over time [min].

### 3.4. Aluminum Use Case

Finally, Figure 9 depicts the split up diagram of aluminum use case.

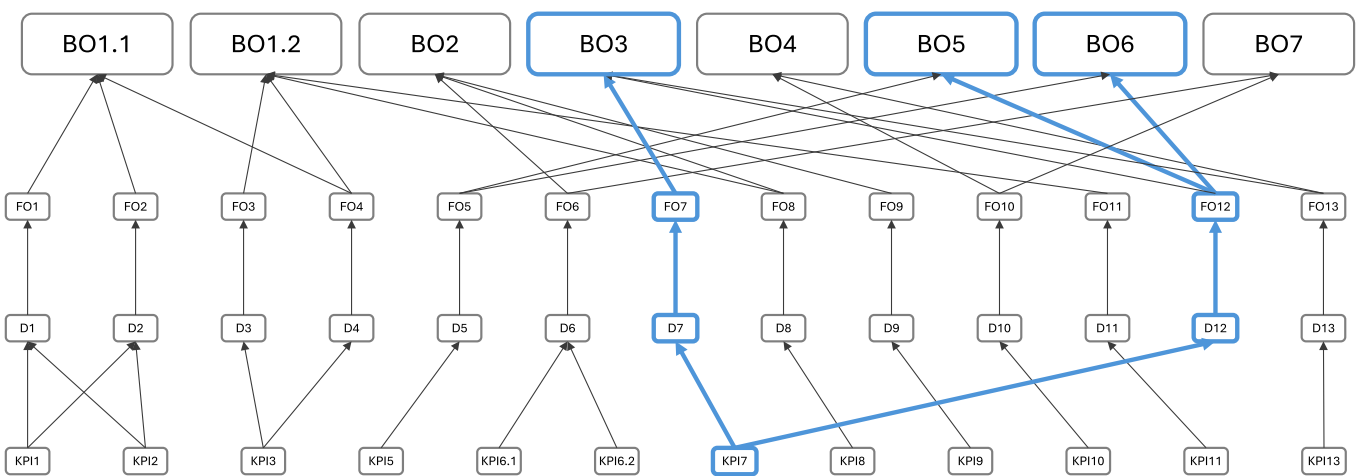


Figure 9. Aluminum split up diagram.

As depicted by Table 9, the FOs allocated to the autonomous manager for this use case are FO7 (data logging securing for possible model drift detection) and FO12 (promptness and effectiveness of diagnostics for a possible model drift), which are both driven by the control parameters of optimization of thresholds for a possible model drift detection. To monitor the effectiveness of the approach, KPI7 (ratio between number of ML anomalies validated and rejected by the data analyst) has been recorded over a period of six months of production and reported in Figure 10.

Table 9. Autonomic manager-related BOs, FOs, drivers, and KPIs for aluminum use case.

Item Class	ID	Description
BO	BO3	Increase aluminum feed
	BO5	Reduce overall gas consumption
	BO6	Reduce overall gas consumption
FO	FO7	Ensure data logging on recipe adjustments/chemical analysis for possible model drift detection
	FO12	Ensure prompt and effective diagnostics for a possible model drift
Driver	D7	Optimization of thresholds for decision making concerning a possible model drift detection
	D12	Optimization of thresholds for decision making concerning a possible model drift detection (same as D7)
KPI	KPI7	Ratio between number of ML anomalies validated and rejected by the data analyst




Figure 10. Aluminum KPI7.

In this case, since the ally melting is realized in a methane-powered furnace, BO5 and BO6 are both proportional to the amount of gas used to melt the metal scrap and maintain it in liquid phase.

Regarding the BOs, the company reported an increase in the aluminum feed of 6% (BO3, deemed proportional to the decreased time batches that aluminum had to wait in ladle furnaces because of chemical corrections needed after eventual drifts of the ML model), while the ratio between gas consumptions and amount of material produced (BO5 and BO6, expression of an increased efficiency of the recipes) was reported to have decreased by 5.49% in the six months of experimentation.

Across the four industrial cases, reported in Table 10, AC-enabled interventions apparently supported resilience-related outcomes by reducing variability, improving recovery after disturbances, or enabling controlled exploration of new operating conditions. While absolute KPI values are not directly comparable across contexts, normalized trend analysis reveals a common pattern in which short-term instability is followed by improved alignment with BOs once adaptation stabilizes.

Table 10. Comparative synthesis of BOs outcomes across the four industrial use cases.

Use Case	BO	Direction of Change	4BL Pillar	Interpretation
Asphalt	BO1.2	Mixed	Economic	Improved for low-recycled-content mixes. Improved after 2024 worsening for high-recycled content mixes.
	BO2.3	Not Available	Economic, Resilience	Not instrumented during observation period.
	BO3	Improved	Economic, Resilience	Improvement observed via reduced downtime episodes.
Steel	BO2	Improved	Economic, Resilience	Reduction observed following stabilization of ML predictions and fewer corrective alloying actions.
	BO7	Improved	Social	Evaluated through internal surveys.
Pharma	BO4	Improved	Economic, Environmental	Evaluated through experimental activities separately reported [48].
	BO7	Improved	Social	Cognitive load reduction through earlier anomaly detection.
Aluminum	BO3	Improved	Economic, Resilience	Throughput increase attributed to fewer recipe corrections.
	BO5/BO6	Improved	Economic, Environmental, Resilience	Gas reduction proportional to improved recipe stability.

In summary, this section inflected the ECOGRAI framework for assessing the impact of AC on sustainability in process industries according to the 4BL framework. By combining the 4BL perspective with the ECOGRAI methodology, the basic idea is to define indeed how DCs, FOs, drivers, and KPIs can be derived and linked to resilience, economic performance, environmental efficiency, and social factors. These elements constitute the core analytical keywords of the proposed approach and provide the conceptual foundation for the empirical analysis presented in the following sections.

## 4. Discussion

The empirical results from the four industrial pilots offer a detailed view of AC can be applied within PI production environments. A closer look at the findings reveals several interconnected themes that deserve further exploration.

One of the most consistent observations across the asphalt, steel, pharma, and aluminum use cases is the central role of the autonomic manager in addressing “model drift” (that in the specific considered cases can be identified as an effect of the so-called “domain shift”, widely explored by the computer science community [49]). This suggests that in these dynamic production settings, the imminence of the initial accuracy of an ML model is less important than the system’s ability to continuously monitor its performance and adapt to changing conditions. As a matter of example, the asphalt case showed a notable performance deviation in December 2024, which may be linked to weather conditions affecting production, highlighting the specific need for industrial resilience in PI: the prerogative to manage variability, whether from raw materials, environmental factors, or equipment degradation. The AC systems implemented appear indeed to act mainly as adaptive mechanisms, fulfilling the self-Configuration and self-Optimization properties of the AC paradigm to keep the performances of predictive models robustly constant over time. Across all use cases, the causal mechanism linking AC to sustainability outcomes operates through variance reduction rather than direct optimization. By detecting and correcting ML model drift, the autonomic manager stabilized production behavior, reducing unplanned deviations. This improved stability causally precedes reductions in waste, energy consumption, downtime, and operator cognitive load.

The application of the 4BL framework, which adds resiliency to the traditional economic, environmental, and social pillars, provides a structured lens that the reported BOs’ improvements can be interpreted through. Rather than considering these four pillars as separated targets, the ECOGRAI application shows how a bottom-up approach based on a technology that explicitly targets the resilience pillar can create positive cascade effects on the other three, suggesting that resilience can be considered a foundational element for industrial sustainability.

Indeed, three of the four considered use cases see the integration of AC directly contributing to BOs related to resilience targets, with the notable exception of the pharma case, where, coherently with the industrial sector, the high value of raw materials prioritizes the achievement of BOs aimed at material savings (therefore associated with environmental objectives).

Additionally, the KPIs analyzed in this work intentionally operate at a macroscale level, such as batch-level, heat-level, experiment-level, or monthly aggregated indicators, rather than at the level of individual sensor readings or control-loop dynamics. This choice reflects the study’s objective of linking AC interventions to BOs within an industrial sustainability framework, where decision-making happens at operational and managerial timescales.

Uncertainty is addressed implicitly through the use of confidence intervals, threshold-based alarms, repeated-occurrence validation, and human confirmation within the MAPE-K loops, without relying on explicit propagation of sensor-level measurement errors. In this sense, the reported KPIs should be interpreted as robust indicators of operational trends and system behavior, suitable for resilience and sustainability assessment, rather than as high-resolution physical measurements.

Regardless, this enhanced resilience translates into tangible economic and environmental benefits. For example, by improving the reliability of “production recipes” for material design and reducing process deviations (a resilience function), the steel and aluminum use cases achieved a more efficient use of resources. This produced, as a side effect, a reduction of chemical waste and lower gas consumption, respectively.

These outcomes are not achieved through a separate optimization effort but are a direct consequence of a more stable and predictable process, where the system's ability to prevent slow dynamics shifts or quickly corrects abrupt deviations (e.g., reducing plant downtimes in the asphalt use case) also has clear economic advantages by improving productivity and throughput. Parallely, the analysis highlights that improvements in certain targets may temporarily conflict with others. For instance, increasing RAP content in the asphalt use case supports environmental objectives but may initially worsen quality or efficiency indicators due to higher process variability. These trade-offs are intrinsic to real industrial decision-making and justify the adoption of resilience-oriented AC mechanisms capable of managing transient performance degradation during adaptation phases [50].

However, the social and resilience dimensions present a more nuanced picture. The "social" impact, for example, was measured through operator approval surveys and increased trust in the system. This points to the AC system's role as a collaborative tool rather than a fully autonomous agent, framing its role as a system intended for collaboration between the ML tools supposed to control the process and the operators and data analysts asked to supervise it. The system appears to enhance human capabilities by providing reliable data and diagnostics, which is a critical factor for successful technology adoption on the factory floor [47,51,52].

This "human-in-the-loop" aspect is a recurring pattern: in the steel, pharma, and aluminum cases, a "data analyst" plays a central role in validating alarms and retraining models. This interaction between the autonomic system and a human expert seems to be a key part of the MAPE-K loop in practice. For example, in the steel use case, the data analyst adjusted the alarm triggers after validating the system's warnings, which led to more affordable recipe designs, which frames the current implementation of AC in these contexts as more of an advanced decision support system than a monolithic software stack. Indeed, traditionally, monitoring complex industrial processes can lead to high cognitive load, as operators are constantly vigilant for anomalies and potential deviations. The AC system, by automating the continuous monitoring for model drifts and process violations, effectively reduces this cognitive burden, bringing the human expert's role to higher-level analytical and decision-making tasks. For instance, in the steel use case, the data analyst's responsibility shifted from constant manual data scrutiny to validating system-generated alarms and making decisions such as when to retrain a model or adjust alarm triggers. This transformation represents a tangible improvement in work quality: it fosters a less stressful work environment, where operators can leverage their expertise for decision making without the cognitive stress implied by mentally demanding tasks related to detection, and with the support of aggregated information. The pharma case further supports this, noting increased operator trust and efficiency due to the AC system. This collaborative partnership between human and machine, where technology empowers the expert rather than replacing them, directly contributes to the social pillar by improving job satisfaction, reducing stress, and potentially enhancing skill development within the workforce.

While these results are promising, it is important to acknowledge that other factors not controlled in this study, such as fluctuations in market demand, raw material costs, or scheduled maintenance, could have also influenced these business-level outcomes.

## 5. Conclusions

This paper tries to address the "resiliency gap" in traditional industrial sustainability frameworks by proposing and empirically testing a 4BL model that depicts resilience as a co-equal pillar alongside economic, environmental, and social ones [23].

Through a PAR study across four distinct PI settings (asphalt, steel, pharma, and aluminum) it shows how autonomic computing can act as an enabling technology to achieve these integrated objectives.

The findings confirm that AC, particularly through its self-Configuration, self-Healing, and self-Optimization properties, effectively enhances operational resilience by continuously monitoring and correcting for “model drift”. This was consistently observed across all use cases, where the autonomic manager was asked to ensure the robustness of ML models against process variability. Improvements in resilience, such as reduced plant downtimes and faster correction of process deviations, directly translated into economic gains (e.g., increased aluminum feed), environmental progress (e.g., reduction of chemical waste and energy consumption), and social benefits (e.g., increased operator trust and approval). In doing so, this work provides empirical evidence that closes the identified gap, showing that high-level sustainability objectives can follow from shop floor improvements through a structured approach.

Despite these positive outcomes, the study has several limitations that are supposed to be debated in future research. A primary limitation lies in the methodological framework used. Indeed, while the ECOGRAI methodology is instrumental in systematically linking strategic BOs to shop-floor KPIs, this connection remains fundamentally qualitative. “Split up” diagrams show that a KPI influences a BO, but they do not provide a quantitative model to predict how much a change in a KPI will impact the BO (e.g., the energy consumption decrease could also be due to planned shutdowns or to a contraction of the market). Furthermore, the variations observed in the BOs could have been experienced due to external factors such as market demand fluctuations, planned shutdowns, seasonal effects (in particular for asphalt use case whose logistics and paving activities are sensitive to weather), organizational learning, and increased operator familiarity with the installed solutions. While these factors may have contributed to the observed improvements, the ECOGRAI framework ensures that only KPIs directly linked to AC-managed decision centers were considered, reducing—but not eliminating—attribution ambiguity.

Building upon the empirical insights gained from this study, a critical direction for future research should involve a more rigorous quantitative assessment of the AC system’s impact. While our current work demonstrates the observed improvements in KPIs and business objectives, the proprietary nature of the industrial data precluded access to the raw process and BOs data needed for statistical analyses, such as regression modeling, confidence interval estimation, or statistical tests for significance. Future studies should aim to overcome these data access challenges to enable a reliable quantification of the causal links between AC interventions and sustainability outcomes. Such quantitative validation would significantly strengthen the evidence base for the proposed 4BL framework and the effects of AC in process industries, enabling more precise strategic planning and investment justification.

Furthermore, the PAR approach, while beneficial for ensuring industrial relevance and stakeholder buy-in, has its own inherent limitations: the highly collaborative and iterative nature of PAR means that the research process is deeply embedded within the specific context and constraints of the partner companies: the definition of BOs, for instance, was influenced by the funding entity’s requirements, and the implementation of KPIs was dependent on the technical capabilities of the system integrators. This makes it challenging to generalize the findings with statistical certainty, as the outcomes are contingent on the unique organizational, technical, and human factors of each case.

Another limitation of the study can be found in the quantity of the use cases (which do not exhaustively represent the entire domain of PI) and in their intrinsic nature: since they have been engaged through a funded research mechanism, their selection may contain

an intrinsic bias, limiting the pool to the companies engaged in publicly funded research activities. More precisely, the characteristics of such companies, which could limit the generalizability to the entire PI domain, can be found in their digitalization level, which is supposed to be mature enough to provide interfaces to collect raw data for analysis and control, their size or business volume, which should meet the financial robustness criteria to make them eligible for funding. Additionally, their internal organization is usually supposed to have a dedicated research and development unit.

A last limitation concerns the varying observation period across the different use cases: data collection periods ranged from one month for the steel use case to six months for the aluminum use case, reflecting the practical constraints of industrial collaborations and data availability. While the scope of this study was focused on demonstrating the applicability and benefits of AC in different PI settings, longer and standardized observation periods and seasonality analyses, particularly for the asphalt use case where climatic conditions are relevant, could provide some additional insights into the effectiveness of the deployed solutions.

Building on these limitations, several directions for future research emerge. First, the consistent reliance on a “human-in-the-loop”, with a data analyst playing a critical role in validating alarms and retraining models, suggests that the current implementation is more of an advanced decision support system: future work should explore pathways to greater autonomy, potentially by automating the analysis and planning phases of the MAPE-K loop to move closer to the original vision of self-managing systems.

Then, future studies could explore the development of more sophisticated KPIs that capture the dynamics of human–computer interaction, such as the ratio of validated alarms to corrective actions (Asphalt KPI 7.1) or the ratio of suggested actions to operator selections (Asphalt KPI 6.2), which were identified but not reported on in the results. This would provide deeper insights into the social pillar and the effectiveness of the collaborative socio-technical system.

Additionally, while the AC ability to contribute to PI resilience was qualitatively demonstrated, the long-term sustainability of the results, particularly concerning the impact of prolonged seasonal variations, could not be exhaustively assessed. Future research, ideally conducted with extended data access and dedicated long-term monitoring, will be crucial to investigate seasonal effects and validate the AC’s robustness over more extended operational cycles. This would provide a more complete understanding of its applicability and benefits after environmental and operational seasonality.

However, by grounding the abstract principles of AC in the real industrial production cases of PI, this work offers a qualitatively validated approach for integrating resilience as an actionable and measurable component of modern manufacturing strategy.

**Author Contributions:** Conceptualization, W.Q., S.A., and F.A.C.; methodology, W.Q.; formal analysis, W.Q., and F.A.C.; investigation, W.Q., S.T., and F.A.C.; resources, W.Q.; data curation, S.T.; writing—original draft preparation, W.Q.; writing—review and editing, S.A., F.A.C., and M.T.; supervision, S.A., F.A.C., and M.T.; project administration, F.A.C.; funding acquisition, M.T. All authors have read and agreed to the published version of the manuscript.

**Funding:** This work has been supported by the projects self-X Artificial Intelligence for European Process Industry digital transformation (s-X-AIPI) Autonomous, scalable, trustworthy, intelligent European meta Operating System for the IoT edge-cloud continuum (aerOS), and European Light-house to Manifest Trustworthy and Green AI (ENFIELD) which have respectively received funding from the European Union’s Horizon Europe research and innovation program under grant agreements 101058715, 101069732, and 101120657.

**Institutional Review Board Statement:** Not applicable.

**Informed Consent Statement:** Not applicable.

**Data Availability Statement:** The original contributions presented in this study are included in the article. Further inquiries can be directed to the corresponding authors.

**Conflicts of Interest:** The authors declare no conflicts of interest and the funders had no role in the design of the study; in the collection, analyses, or interpretation of data; in the writing of the manuscript; or in the decision to publish the results.

## Abbreviations

The following abbreviations are used in this manuscript:

3BL	Triple Bottom Line
4BL	Quadruple Bottom Line
AC	Autonomic Computing
AI	Artificial Intelligence
BO	Business Objective
DC	Decision Centre
EAF	Electric Arc Furnace
FO	Functional Objective
FTIR	Fourier-Transform Infrared Spectroscopy
HITL	Human In The Loop
KPI	Key Performance Indicator
MAPE-K	Monitor, Analyze, Plan, Control-Knowledge
MES	Manufacturing Execution System
ML	Machine Learning
OCT	Optical Coherence Tomography
PAR	Procedural Action Research
PI	Process Industry
PLC	Programmable Logic Controller
RAP	Reclaimed Asphalt Pavement
SMART	Specific, Measurable, Attainable, Relevant, Time-bound

## Appendix A

This appendix lists the complete set of BOs, FOs, drivers, and KPIs for the four use cases.

**Table A1.** Asphalt-actualized BOs.

	BO	Actualized BO
1	Improve productivity	BO1.1 Improvement of logistics factory-jobsite
		BO1.2 Decrease quality rejections
2	Improve quality	BO2.1 Improve adjustment of filler/bitumen rate
		BO2.2 Decrease deviations in dosing process
		BO2.3 Improve control of hot aggregates separation
3	Improve response time	BO3 Improve optimization of maintenance schedule
4	Reduction of raw material consumption	BO4 Increase recycling
5	Reduction of energy consumption	BO5 Reduction of generated CO <sub>2</sub>
6	Reduction of generated CO <sub>2</sub>	BO6 Decrease consumption of diesel and heavy fuel
7	Increase competences	BO7 Embodiment of applications adapting to and enriching heterogeneous skill levels

**Table A2.** Steel-actualized BOs.

BO		Actualized BO	
1	Improve productivity	BO1.1 BO1.2	Increase metal yield optimization Decrease downstream quality rejections
2	Improve quality	BO2.1 BO2.2	Improve adjustment of residual elements to customer specifications Decrease O <sub>2</sub> content in EAF
3	Improve response time	BO3.1 BO3.2	Improve prediction of EAF process parameter (T, O <sub>2</sub> , and P) Decrease frequency of recipe update
4	Reduction of raw material consumption	BO4	Decrease additive usage
5	Reduction of energy consumption	BO5	Reduction of EAF and ladle furnaces energy consumption
6	Reduction of generated CO <sub>2</sub>	BO6	Reduction of EAF and ladle furnaces energy consumption
7	Increase competences	BO7	Introduction of a decision support system

**Table A3.** Pharma-actualized BOs.

BO		Actualized BO	
1	Improve productivity	BO1	Increase reaction efficiency
2	Improve quality	BO2	Reduce chemical waste
3	Improve response time	BO3	Decrease operators' reaction time
4	Reduction of raw material consumption	BO4	Decrease quality rejections
5	Reduction of energy consumption	BO5	Improve ratio between energy consumption and Active Pharmaceutical Ingredient (API) produced
6	Reduction of generated CO <sub>2</sub>	BO6	Improve ratio between energy consumption and Active Pharmaceutical Ingredient (API) produced
7	Increase competences	BO7	Embodiment of applications adapting to heterogeneous skill levels

**Table A4.** Aluminum-actualized BOs.

BO		Actualized BO	
1	Improve productivity	BO1.1 BO1.2	Improve metal yield Cost reduction
2	Improve quality	BO2	Decrease downstream quality rejections
3	Improve response time	BO3	Increase aluminum feed
4	Reduction of raw material consumption	BO4	Decrease NaCl and KCl fluxes
5	Reduction of energy consumption	BO5	Reduce overall gas consumption
6	Reduction of generated CO <sub>2</sub>	BO6	Reduce overall gas consumption
7	Increase competences	BO7	Improve process transparency

**Table A5.** DCs, FOs, and drivers for asphalt use case.

DC	FO	Driver
Centralized Server	FO1: ensure real-time production monitoring and raising of fault-detection alarms	D1: tuning knobs for real-time plant monitoring and respecting real-time constraints in reporting alarms concerning diagnostics
	FO2: set-up of the optimal burner configuration	D2: systematic execution of the ML model for optimal selection of the burners at every time the decision can be subject to verification
	FO6: ensure human-in-the-loop efficiency for alarm reporting	D6: systematic acknowledging by the operator of an alarm after analysis of the situation
	FO8: ensure availability of labeled data for anomalies	D8: dashboard/data-exploration features for guaranteeing data reliability for subsequent use
	FO13: efficient massive filtering of data from database for ML retrain	D13: tuning knobs for filtering out data useless for ML training
	FO17: ensure data availability for measurements at paving site	D17: systematically classify data concerning measurements at paving site according to production lots for traceability and correlatability to production lots

Table A5. Cont.

DC	FO	Driver
Data Scientist	FO3: ensure the selection of the appropriate data by filtering for ML (re)training FO9: ensure correct labeling for anomalies FO11: ensure supervision of quality parameters estimated by ML model FO12: ensure by ML model that validated production recipes do not correspond to a risk of incurring an anomaly on material properties FO16: ensure human-in-the-loop efficiency for data-scientist action after an alarm (same of FO7)	D3: tuning knobs for filtering out data useless for ML training D9: dashboard and data-exploration features for accurate labeling of data into anomaly and normal scenarios D11: systematic execution of the ML model to optimally estimate quality parameters D12: systematic execution of the ML model to verify if a production recipe might generate a quality problem in terms of volumetric/mechanical properties D16: all the HMI features assisting the data-scientist for the verification of an alarm and an appropriate intervention
Laboratory	FO15: ensure availability of data concerning the analysis of material sampling collected at plant	D15: systematically classify samples according to production lots for traceability and correlatability between analysis-data
Automation Plant	FO4: ensure the data collection and synchronization for plant ML models (re) training FO7: ensure human-in-the-loop efficiency for alarm reaction FO10: ensure efficient anomaly detection and data exportation (by comparing measured and theoretical data and possible data aggregation logics) FO19: ensure the required target parameters for paving	D4: tuning knobs of the algorithms for signal processing and data synchronization versus production lots tracking D7: all the HMI features assisting the operator for the verification of an alarm and an appropriate intervention D10: tuning knobs for comparing measured data with theoretical data so as to limit the number of false alarms in anomaly detection D19: systematically apply the target temperature at production plant suggested by the ML model so as to guarantee the temperature at paving site
Logistics and Paving	FO18: ensure availability of the measurements and repeatability of the process FO20: ensure the transportation of the material at the paving site under the predicted conditions	D18: execution of maintenance activities on the equipment and on the sensors/remote connection to paving site D20: systematically apply the target temperature at production plant suggested by the ML model so as to guarantee the temperature at paving site (same as D19)
Autonomic Manager	FO14: optimize the detection time for a ML model drift vs. the process	D14: definition of tolerances over the detection of possible model drift so as to reach a fair trade-off between the minimization of the detection time and the need of avoiding false alarms

Table A6. DCs, FOs, and drivers for steel use case.

DC	FO	Driver
MES	FO1: maximize output value under order production constraints FO6: allocate recipe to different customer FO7: selection of input data related to other recipes FO9: definition of new recipe where assign product to FO11: selection of input data related to anomaly FO14: provide reliable data about raw materials availability	D1: recipe selection D6: selection of customer (OTIF, forecasted) D7: granularity of selected recipes vs. Stock Keeping Unit (SKU) D9: selection of parameters for new recipe D11: quality of input data D14: inventory check frequency
Laboratory	FO4: minimize time and ensure reliability	D4: number of analyses derived from HeatID
EAF	FO3: follow process parameters FO5: reach requirements with respect to additives cost	D3: current/temperature per HeatID D5: quantity of additive material used for HeatID
Production Manager	FO1: maximize output value under order production constraints FO5: reach requirements with respect to additives cost FO6: allocate recipe to different customers	D1: recipe selection D5: quantity of additive material used for HeatID D6: selection of customer (OTIF, forecasted)
Scrap Yard	FO2: ensure required materials	D2: selection of scrap silos
Data Analyst	FO8: assign product to existing recipe FO12: check recipe reliability FO13: improve recipe accuracy FO15: optimize time horizon FO16: optimize usage of "good scrap"	D8: selection of recipe (group) to assign the product to D12: definition of tolerances over non-conformities characteristics D13: creation for production order for new recipe validation D15: decision of time parameters (frequency of updates and horizon) D16: Usage of scrap silos
Autonomic Manager	FO10: minimize number of occurrences before warning, warnings corresponding to false drift, maximize reliability of prediction of scrapyard criticalities FO13: improve recipe accuracy	D10: definition of tolerances over detection of non-conformities D13: creation for production order for new recipe validation

**Table A7.** DCs, FOs, and drivers for pharma use case.

DC	FO	Driver
Database	FO3: ensure effectiveness of experiment setpoints	D3: possible manual intervention during the execution of the chemical reaction to reduce effects of possible exceptions (same of D2)
	FO5: ensure HITL efficiency for alarm reporting	D5: snoozing of an alarm after analysis of the situation
	FO11: efficient saving in the database for subsequent elaboration including ML model retraining for metadata + process data	D11: tuning knobs of the methodology to classify the experiments (e.g., in categories like “effective working point” or “inefficient working points”) to be used subsequently in the cost optimization activity trading off between different contrasting objectives (same of D10)
	FO14: efficient massive filtering of data from database for ML retrain	D14: tuning knobs for filtering out data useless for ML training (same of D13)
Data Analyst	FO2: ensure precision in following the defined experiment targets	D2: possible manual intervention during the execution of the chemical reaction to reduce effects of possible exceptions
	FO10: ensure correct data-labeling of the experimental results	D10: tuning knobs of the methodology to classify the experiments (e.g., in categories like “effective working point” or “inefficient working points”) to be used subsequently in the cost optimization activity trading off between different contrasting objectives
	FO13: ensure the selection of the appropriate data by filtering for ML retrain	D13: tuning knobs for filtering out data useless for ML training
Laboratory	FO6: ensure accuracy for attaining the defined experimental targets	D6: tuning knobs of the controller influencing its robustness and the regulation settling time
	FO7: ensure monitoring of the process including reaction progress and electrode wear	D7: tuning knobs of the algorithms for signal processing and image recognition including denoising (e.g., for electrode wear estimation from OCT) and for data synchronization
	FO8: ensure availability of the equipment and repeatability of the process	D8: execution of maintenance activities on the equipment and on the sensors
	FO9: ensure efficiency for the analysis of the experiment results (evaluate the prediction reliability vs. measured data and compare the prediction on the electrode wear estimation)	D9: tuning knobs of the methodology to compare measurement with prediction (e.g., at present, the division of the experiment duration in 20 ranges with independent evaluation of the corresponding RMSE)
Scientist	FO1: maximize the knowledge on the process and evaluate possible needs of new experiments	D1: granularity of the datapoints collected versus the working space for the process
Autonomic Manager	FO4: minimize the detection time for the possible wrong positioning of the OCT	D4: fastness in the possible manual intervention during the execution of the chemical reaction to correct the possible wrong positioning of the OCT
	FO12: optimize the detection time for a ML model drift vs. the process	D12: definition of tolerances over the detection of possible model drift

**Table A8.** DCs, FOs, and drivers for aluminum use case.

DC	FO	Driver
Centralized Server	FO3: ensure satisfaction of the order selling requirements (chemical limits, yield stress limits, tensile stress limits, etc.)	D3: safety stock levels of each silo
	FO9: allocate heat to different customer	D9: selection of customer
Data Analyst	FO10: ensure the automatic suggestion of feasible recipes for the target production	D10: input of ML model
	FO13: ensure correction of a possible process drift and consequent improvement in recipe accuracy	D13: offline redefinition of the recipes used in the past

Table A8. Cont.

DC	FO	Driver
Melting Department	FO1: ensure right analysis chemical composition of scraps	D1: under the assumption that for every reception of raw material a chemical analysis is always executed, tolerances/tuning knobs on raw material classification (to dispatch the material on the silos)
	FO5: follow melting process parameters	D5: amount of gas and target temperature
	FO6: minimize time and ensure chemical feasibility	D6: maximal number of alloy adjustments for the heat
	FO8: reach requirement with respect to additives cost	D8: quantity of additive material used per heat
Production Manager	FO3: ensure satisfaction of the order selling requirements (chemical limits, yield stress limits, tensile stress limits, etc.)	D3: safety stock levels of each silo
	FO8: reach requirement with respect to additives cost	D8: quantity of additive material used per heat
	FO9: allocate heat to different customer	D9: selection of customer
	FO11: ensure selection of the optimal recipe in terms of feasibility and costs	D11: recipe selection
Scrap Yard	FO2: ensure right size and right positioning of scraps in the silos	D2: optimize the criterion for material storing
	FO4: ensure required materials (pick the material from the programmed silos defined by the production manager)	D4 selection of scrap silos
Autonomic Manager	FO7: ensure data logging on recipe adjustments/chemical analysis for possible model drift detection	D7: optimization of thresholds for decision making concerning a possible model drift detection
	FO12: ensure prompt and effective diagnostics for a possible model drift	D12: optimization of thresholds for decision making concerning a possible model drift detection (same as D7)

Table A9. KPIs for asphalt use case.

KPI	Description
1.1	Accuracy (defined as sum of true positives and true negatives divided by total events) and precision (defined as true positives divided by false positives) of the fault-diagnostic ML model
1.2	Ratio between theoretical production and actual one [tons]
1.3	Accuracy and precision of the AM (computed as per KPI1)
2.1	Ratio between the number of times the ML model for optimal selection of the burners is executed and the number of production lots executed at plant
2.2	Ratio between fuel consumptions [m <sup>3</sup> ] and amount of material produced [tons]
2.3	Ratio between electrical power consumption [kWh] and amount of material produced [tons]
3	Number of datapoints collected after the last model training [% over the total number of available datapoints] outside a confidence interval of $\pm 5\%$ , $\pm 10\%$ , $\pm 15\%$ concerning the deviation between measurement and prediction
4	Reliability index of the measurement archived from all the sensors
5	Number of maintenance interventions
6.1	Number of alarms validated by the plant operators
6.2	Ratio between number of times the operator selected the burner configuration suggested by the ML model and number of times the ML model for burner selection has been executed
7.1	Number of alarms validated by the plant operators and followed by a corrective action
7.2	Ratio between downtime for faults and production time excluding programmed maintenance periods [%]
8	Ratio between the number of anomalies labeled and total recorded number of anomalies [%]
9	Same as KPI8
10	Same as KPI3
11	Ratio between the number of measurements collected from the laboratory about properties of the material and the number of production lots in the same period
12	Percentage of recipes inside the asphalt mix design recipe “technology book” (related to granulometric properties) whose simulation by the ML models corresponds to a risk of generating an anomaly during production.
13	Same of KPI3
14	Number of produced batches before a warning for model drift and after the last model deployment
15.1	Percentage of lots whose laboratory check matches with a production recipe
15.2	Percentage of recycled asphalt in the final OTIF batches
16	Same as KPI7
17	Percentage of used recipes for the weekly production that are compliant with paving site-activity
18	Number of maintenance interventions
19	Number of datapoints collected (in [%] over the total number of available datapoints) outside a confidence interval of $\pm 5\%$ , $\pm 10\%$ , $\pm 15\%$ concerning the deviation between measurement and prediction (for the temperature at paving site)
20	Same of KPI19

**Table A10.** KPIs for steel use case.

KPI	Description
1	Estimation of cost of raw materials per HeatID
2	Cost of additive material (chromium, nickel, and molybdenum) used per HeatID
3	O <sub>2</sub> and energy consumption per HeatID
4	Ratio between number of analyses and of HeatIDs per week
5	Same of KPI2
6	OTIF (forecasted) on EAF production orders
7	Number of recipes maintained in the database for every steel grade (variability in function of steel grade)
8	Number of chemical composition non-conformities per week of production
9	Same of KPI2
10	Number of occurrences before raising warning
11	Improvement of Mean Square Error (MSE) over the data points related to non-conformities
12	Number of HeatIDs (for the same recipe) without non-conformities before next detection of non-conformity
13	Same as KPI2
14	Number of errors notified by the scrapyard operator
15	Number of non-conformities detected in the next time window
16	Deviations of each scrap silos with reference to target levels

**Table A11.** KPIs for pharma use case.

KPI	Description
1	Minimal density of saved experiment datapoints without meaningful repetition
2	For each experiment, ratio between number of exceptions requiring a manual intervention and number of actual manual interventions executed
3	Same of KPI2
4	Time between experiment start and first actual correction of a wrong positioning of the OCT sensor
5.1	Ratio between number of acknowledged alarms concerning the wrong positioning of the OCT sensors and number of alarms concerning the wrong positioning of the OCT sensors
5.2	Ratio between number of correct alarms concerning the wrong positioning of the OCT sensors and number of alarms concerning the wrong positioning of the OCT sensors
6.1	Correction amplitude executed by the controller for every experiment to reach 95% of the steady state, for every experiment
6.2	Settling time of the controller for every experiment to reach 95% of the steady state
7	Percentage of images out of focus
8	Number of maintenance interventions
9.1	Deviation between measurements and predictions in adrenosterone percentage
9.2	Number of datapoints collected outside a confidence interval of $\pm 5\%$ , $\pm 10\%$ , $\pm 15\%$ concerning the deviation between measurement and prediction
10.1	Ratio between the amount of produced adrenosterone [mg] in multiple and production time [min] from the last change of electrodes and the effective production time
10.2	Ratio between the amount of produced adrenosterone [mg] and the consumed electrical power [Wh]
10.3	Ratio between the amount of produced adrenosterone and the input raw material
11	Same of KPI10
12	Months for which KPI9.1 was exceeded
13	Recalculation of the model loss after a ML retraining (at least the last weeks) verifying lack of overfitting.
14	Same of KPI13

**Table A12.** KPIs for aluminium use case.

KPI	Description
1	Ratio of "robust" silos classification of the material vs. chemical composition/number of silos deliveries
2	Ratio between standard deviation for each silo and limit tolerance on the standard deviation for every chemical composition
3	Number of withdrawals from silos under the safety stock level
4	Same as KPI3
5	Ratio between gas consumptions and amount of material produced
6.1	Number of heats requiring an order reassignment
6.2	Ratio between number of heatIDs analyzed by laboratory triggering alloy adjustment and total number of laboratory analyses
7	Ratio between number of ML anomalies validated and rejected by Data Analyst
8	Ratio between real quantity produced and theoretical quantity planned
9	Ratio between OTIF production orders and total programmed production orders
10	Ratio between number of feasible recipes obtained by the ML model and number of suggested recipes per heatID
11	Ratio between the number of times the optimal recipe per cost is selected and the total number of recipe suggestions
12	Same as KPI7
13	Number of MAE, MSE, R <sup>2</sup> per retraining

## References

1. Dwivedi, A.; Sassanelli, C.; Agrawal, D.; Moktadir, M.A.; D'Adamo, I. Drivers to Mitigate Climate Change in Context of Manufacturing Industry: An Emerging Economy Study. *Bus. Strategy Environ.* **2023**, *32*, 4467–4484. [[CrossRef](#)]
2. Elkington, J. Towards the Sustainable Corporation: Win-Win-Win Business Strategies for Sustainable Development. *Calif. Manag. Rev.* **1994**, *36*, 90–100. [[CrossRef](#)]
3. Avelino, F.; Wittmayer, J.; Kemp, R.; Haxeltine, A. Game-Changers and Transformative Social Innovation. *Ecol. Soc.* **2017**, *22*, 41. [[CrossRef](#)]
4. Worthington, I.; Patton, D. Strategic Intent in the Management of the Green Environment Within SMEs: An Analysis of the UK Screen-Printing Sector. *Long Range Plan.* **2005**, *38*, 197–212. [[CrossRef](#)]
5. Friedman, M. A Friedman Doctrine—The Social Responsibility of Business Is to Increase Its Profits. *The New York Times*, 13 September 1970. Available online: <https://www.nytimes.com/1970/09/13/archives/a-friedman-doctrine-the-social-responsibility-of-business-is-to.html> (accessed on 3 January 2026).
6. Holliday, C.O.J.; Schmidheiny, S.; Watts, P. *Walking the Talk: The Business Case for Sustainable Development*, 1st ed.; Routledge: London, UK, 2017.
7. Imbabi, M.S.; Carrigan, C.; McKenna, S. Trends and Developments in Green Cement and Concrete Technology. *Int. J. Sustain. Built Environ.* **2012**, *1*, 194–216. [[CrossRef](#)]
8. Obradović, T.; Vlačić, B.; Dabić, M. Open Innovation in the Manufacturing Industry: A Review and Research Agenda. *Technovation* **2021**, *102*, 102221. [[CrossRef](#)]
9. Mohammad, W.; Elomri, A.; Kerbache, L. The Global Semiconductor Chip Shortage: Causes, Implications, and Potential Remedies. *IFAC-Pap.* **2022**, *55*, 476–483. [[CrossRef](#)]
10. Gryparis, E.; Papadopoulos, P.; Vokas, G.A.; Psomopoulos, C.S. The Impact of Covid and Ukrainian War in EU Electricity Demand and Carbon Emissions in Power Generation under High Penetration of Electric Vehicles. *Appl. Energy* **2025**, *393*, 126022. [[CrossRef](#)]
11. Hollnagel, E.; Woods, D.D.; Leveson, N. *Resilience Engineering: Concepts and Precepts*; CRC Press: London, UK, 2006.
12. Monostori, L. Cyber-Physical Production Systems: Roots, Expectations and R&D Challenges. *Procedia CIRP* **2014**, *17*, 9–13. [[CrossRef](#)]
13. Gladysz, B.; Krystosiak, K.; Buczacki, A.; Quadrini, W.; Ejsmont, K.; Kluczek, A.; Park, J.; Fumagalli, L. Sustainability and Industry 4.0 in the Packaging and Printing Industry: A Diagnostic Survey in Poland. *Eng. Manag. Prod. Serv.* **2024**, *16*, 51–67. [[CrossRef](#)]
14. Ilevbare, I.; Dusch, B.; Templeton, P. *A Framework and Methodology for Creating Business Tools and Processes*; Institute for Manufacturing: Cambridge, UK, 2016.
15. Yu, W.; Patros, P.; Young, B.; Klinac, E.; Walmsley, T.G. Energy Digital Twin Technology for Industrial Energy Management: Classification, Challenges and Future. *Renew. Sustain. Energy Rev.* **2022**, *161*, 112407. [[CrossRef](#)]
16. World Commission on Environment and Development. *Our Common Future*; Oxford University Press: Oxford, UK, 1987; ISBN 978-0-19-282080-8.
17. Brent, A.C.; Rohwer, M.B.; Friedrich, E.; Blottnitz, H. von Status of Life Cycle Assessment and Engineering Research in South Africa. *Int. J. Life Cycle Assess.* **2002**, *7*, 167–172. [[CrossRef](#)]
18. Karkou, E.; Angelis-Dimakis, A.; Parlapiano, M.; Savvakis, N.; Siddique, O.; Vyrkou, A.; Sgroi, M.; Fatone, F.; Arampatzis, G. Process Innovations and Circular Strategies for Closing the Water Loop in a Process Industry. *J. Environ. Manag.* **2024**, *370*, 122748. [[CrossRef](#)] [[PubMed](#)]
19. Walsh, C.; Thornley, P. The Environmental Impact and Economic Feasibility of Introducing an Organic Rankine Cycle to Recover Low Grade Heat during the Production of Metallurgical Coke. *J. Clean. Prod.* **2012**, *34*, 29–37. [[CrossRef](#)]
20. Jacquemin, L.; Pontalier, P.-Y.; Sablayrolles, C. Life Cycle Assessment (LCA) Applied to the Process Industry: A Review. *Int. J. Life Cycle Assess.* **2012**, *17*, 1028–1041. [[CrossRef](#)]
21. Chen, R.; Xu, P.; Song, P.; Wang, M.; He, J. China Has Faster Pace than Japan in Population Aging in next 25 Years. *Biosci. Trends* **2019**, *13*, 287–291. [[CrossRef](#)]
22. Kaptein, M.; Wempe, J. *The Balanced Company: A Theory of Corporate Integrity*; Oxford University Press: Oxford, UK, 2002; ISBN 978-0-19-925550-4.
23. Alibašić, H. (Ed.) Examining the Intersection of Sustainability and Resilience. In *Sustainability and Resilience Planning for Local Governments: The Quadruple Bottom Line Strategy*; Springer International Publishing: Cham, Switzerland, 2018; pp. 85–98, ISBN 978-3-319-72568-0.
24. Alibašić, H. (Ed.) Defining, Initiating, and Reviewing Sustainability and Resilience Planning. In *Sustainability and Resilience Planning for Local Governments: The Quadruple Bottom Line Strategy*; Springer International Publishing: Cham, Switzerland, 2018; pp. 1–16, ISBN 978-3-319-72568-0.
25. Manta, F.; Tarulli, A.; Morrone, D.; Toma, P. Toward a Quadruple Bottom Line: Social Disclosure and Financial Performance in the Banking Sector. *Sustainability* **2020**, *12*, 4038. [[CrossRef](#)]

26. Scordato, L.; Gulbrandsen, M. Resilience Perspectives in Sustainability Transitions Research: A Systematic Literature Review. *Environ. Innov. Soc. Transit.* **2024**, *52*, 100887. [[CrossRef](#)]
27. Molenda, P.; Groneberg, H.; Schötz, S.; Döpfer, F. Resilience Balanced Scorecard: Measuring Resilience of Manufacturing Companies at Multiple Levels. *Procedia CIRP* **2023**, *120*, 189–194. [[CrossRef](#)]
28. Sesana, M.M.; Dell’Oro, P. Sustainability and Resilience Assessment Methods: A Literature Review to Support the Decarbonization Target for the Construction Sector. *Energies* **2024**, *17*, 1440. [[CrossRef](#)]
29. Kephart, J.O.; Chess, D.M. The Vision of Autonomic Computing. *Computer* **2003**, *36*, 41–50. [[CrossRef](#)]
30. Quadrini, W.; Cuzzola, F.A.; Fumagalli, L.; Taisch, M.; De Luca, G.; Calderaro, M.; Marzano, M.G.; Marguglio, A. A Reference Architecture to Implement Self-X Capability in an Industrial Software Architecture. *Procedia Comput. Sci.* **2024**, *232*, 446–455. [[CrossRef](#)]
31. Müller, T.; Kamm, S.; Löcklin, A.; White, D.; Mellinger, M.; Jazdi, N.; Weyrich, M. Architecture and Knowledge Modelling for Self-Organized Reconfiguration Management of Cyber-Physical Production Systems. *Int. J. Comput. Integr. Manuf.* **2022**, *36*, 1842–1863. [[CrossRef](#)]
32. Van, H.M.; Da Silva, A.S.; Knissel, T.; Weiss, G. Self-Adaptive Service Deployment for Resilience of Smart Manufacturing Architectures. In Proceedings of the 2023 IEEE International Conference on Autonomic Computing and Self-Organizing Systems Companion (ACSOS-C), Toronto, ON, Canada, 25–29 September 2023; IEEE: Toronto, ON, Canada, 2023; pp. 146–151.
33. Zheng, P.; Xia, L.; Li, C.; Li, X.; Liu, B. Towards Self-X Cognitive Manufacturing Network: An Industrial Knowledge Graph-Based Multi-Agent Reinforcement Learning Approach. *J. Manuf. Syst.* **2021**, *61*, 16–26. [[CrossRef](#)]
34. García, R.; Aguilar, J.; Toro, M.; Pérez, N.; Pinto, A.; Rodríguez, P. Autonomic Computing in a Beef-Production Process for Precision Livestock Farming. *J. Ind. Inf. Integr.* **2023**, *31*, 100425. [[CrossRef](#)]
35. Millara, Á.F.; Molleda, J.; Usamentiaga, R.; García, D.F. Profile Measurement of Rails in a Rolling Mill: Implementing and Evaluating Autonomic Computing Capabilities. *IEEE Trans. Ind. Appl.* **2019**, *55*, 5466–5475. [[CrossRef](#)]
36. Shahin, A.; Mahbod, M.A. Prioritization of Key Performance Indicators: An Integration of Analytical Hierarchy Process and Goal Setting. *Int. J. Product. Perform. Manag.* **2007**, *56*, 226–240. [[CrossRef](#)]
37. Ahn, H.; Clermont, M.; Hespings, J. The Unit Effect from a Performance Management Perspective. *Int. J. Product. Perform. Manag.* **2021**, *71*, 1895–1915. [[CrossRef](#)]
38. Stephens, S. Supply Chain Operations Reference Model Version 5.0: A New Tool to Improve Supply Chain Efficiency and Achieve Best Practice. *Inf. Syst. Front.* **2001**, *3*, 471–476. [[CrossRef](#)]
39. Maas, R.; Shevtshenko, E.; Karaulova, T. Supply Chain Quality Improvement Based on Customer Compliance. In *Proceedings of the Technological Innovation for Connected Cyber Physical Spaces*; Camarinha-Matos, L.M., Ferrada, F., Eds.; Springer Nature Switzerland: Cham, Switzerland, 2023; pp. 230–242.
40. Srhir, S.; Jaegler, A.; Montoya-Torres, J.R. Uncovering Industry 4.0 Technology Attributes in Sustainable Supply Chain 4.0: A Systematic Literature Review. *Bus. Strategy Environ.* **2023**, *32*, 4143–4166. [[CrossRef](#)]
41. Daskalantonakis, M.K. A Practical View of Software Measurement and Implementation Experiences Within Motorola. *IEEE Trans. Softw. Eng.* **1992**, *18*, 998–1010. [[CrossRef](#)]
42. Kaplan, R.S.; Norton, D.P. The Balanced Scorecard—Measures that Drive Performance. *Harv. Bus. Rev.* **1992**, *70*, 71–79. Available online: <https://hbr.org/1992/01/the-balanced-scorecard-measures-that-drive-performance-2> (accessed on 3 January 2026).
43. Johanson, U.; Skoog, M.; Backlund, A.; Almqvist, R. Balancing Dilemmas of the Balanced Scorecard. *Account. Audit. Account. J.* **2006**, *19*, 842–857. [[CrossRef](#)]
44. Doumeings, G.; Clave, F.; Ducq, Y. ECOGRAI—A Method to Design and to Implement Performance Measurement Systems for Industrial Organizations—Concepts and Application to the Maintenance Function. In *Benchmarking—Theory and Practice*. IFIP Advances in Information and Communication Technology; Springer: Boston, MA, USA, 1995.
45. Quadrini, W.; Mohammadianghovahloo, M.; Sassanelli, C.; Terzi, S.; Fumagalli, L. Defining Performance Indicators for a Multi-Stakeholder Manufacturing Digital Platform: An ECOGRAI Adaptation. In *The Future of Industry: Human-Centric Approaches in Digital Transformation*; Appolloni, A., Kumar, V., Kuzmin, E., Akberdina, V., Eds.; Springer Nature Switzerland: Cham, Switzerland, 2024; pp. 489–504, ISBN 978-3-031-66801-2.
46. Cunningham, B. Action Research: Toward a Procedural Model. *Hum. Relat.* **1976**, *29*, 215–238. [[CrossRef](#)]
47. Cuzzola, F.A.; Quadrini, W.; Teocchi, S.; Pinzone, M.; Fumagalli, L. Evaluating Humachine Dynamics in Process Industry AI Applications. *J. Manuf. Technol. Manag.* [[CrossRef](#)]
48. Fink, E.; Ibrahim, M.Y.S.; Wolfgang, M.; Kappe, C.O.; Cantillo, D. Real-Time Quantitative Analysis of Electrode Surfaces in Synthetic Organic Electrochemistry via Optical Coherence Tomography. *ACS Electrochem.* **2025**, *1*, 2217–2222. [[CrossRef](#)]
49. Luo, Y.; Zheng, L.; Guan, T.; Yu, J.; Yang, Y. Taking a Closer Look at Domain Shift: Category-Level Adversaries for Semantics Consistent Domain Adaptation. In Proceedings of the 2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), Long Beach, CA, USA, 15–20 June 2019; pp. 2502–2511.

50. Singh, A.K. Smart Cement Plants: Enhancing Productivity and Sustainability with AI. *Int. J. Res. Appl. Sci. Eng. Technol.* **2025**, *13*, 2801–2807. [[CrossRef](#)]
51. Wesche, J.S.; Langer, M.; Sonderegger, A.; Landers, R.N. Editorial to the Virtual Special Issue: Human-Automation Interaction in the Workplace: A Broadened Scope of Paradigms. *Comput. Hum. Behav.* **2022**, *134*, 107335. [[CrossRef](#)]
52. Rasmussen, J. Outlines of a Hybrid Model of the Process Plant Operator. In *Monitoring Behavior and Supervisory Control*; Sheridan, T.B., Johannsen, G., Eds.; Springer US: Boston, MA, USA, 1976; pp. 371–383, ISBN 978-1-4684-2523-9.

**Disclaimer/Publisher’s Note:** The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.