electronics

MDPI

Review

Impact of EU Laws and Regulations on the Adoption of Artificial Intelligence in Cyber–Physical Systems: A Review of Regulatory Barriers, Technological Challenges, and Cross-Sector Implications

Bo Nørregaard Jørgensen * and Zheng Grace Ma

SDU Center for Energy Informatics, Maersk Mc-Kinney Moller Institute, The Faculty of Engineering, University of Southern Denmark, 5230 Odense, Denmark; zma@mmmi.sdu.dk
* Correspondence: bnj@mmmi.sdu.dk; Tel.: +45-29125778

Abstract

Artificial intelligence is increasingly embedded in cyber–physical systems that coordinate sensing, computation, communication, and control across critical and semi-critical physical environments. Within the European Union, however, its adoption is shaped not only by technological maturity and economic value, but also by an increasingly dense regulatory landscape governing data processing, cybersecurity, product security, accountability, traceability, interoperability, and safety-relevant deployment. A PRISMA ScR-informed scoping review is used to examine how European Union regulation influences artificial intelligence adoption across four representative domains: energy and smart grids, smart buildings, mobility and transport systems, and industrial and manufacturing environments. The analysis draws on primary legal sources, the peer-reviewed literature, and policy and standards-related materials, and is structured around three dimensions: regulatory barriers, technological and architectural challenges, and cross-sector implications for governance, innovation, and competitiveness. The results show that regulation functions simultaneously as a constraint and an enabling condition. It increases compliance burden, raises integration complexity, and slows deployment in higher risk settings, while promoting trustworthy artificial intelligence, stronger cybersecurity, lifecycle governance, clearer accountability, and more interoperable digital infrastructures. The central finding is that regulation is not external to artificial intelligence adoption in cyber–physical systems, but actively shapes the design space within which such systems can be developed, integrated, validated, and scaled. Future progress therefore depends on regulation-aware systems engineering, stronger implementation guidance, and cross-sector reference architectures capable of aligning legal compliance with technical architecture and operational value creation.

Keywords: artificial intelligence; cyber–physical systems; European Union; regulation; AI Act; GDPR; NIS2; Cyber Resilience Act; Data Act; Data Governance Act; cybersecurity; interoperability; compliance by design

Check for updates

Academic Editor:
Aryya Gangopadhyay

Received: 26 April 2026
Revised: 15 May 2026
Accepted: 16 May 2026
Published: 19 May 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 Attribution (CC BY) license.

1. Introduction

Cyber–physical systems are becoming a defining technological foundation of the contemporary digital economy. Across energy systems, buildings, mobility infrastructures, industrial production environments, and urban services, these systems integrate sensing, communication, software, and physical processes into increasingly interconnected

operational environments. Artificial intelligence is accelerating this transformation by enabling cyber–physical systems to move beyond fixed automation towards prediction, adaptive control, anomaly detection, optimisation, and more autonomous forms of decision support [1–4]. Artificial intelligence is therefore no longer an isolated digital capability. It is increasingly embedded in real-time socio-technical environments whose failures may have physical, economic, safety, and societal consequences [4–6].

Within the European Union, the adoption of artificial intelligence in cyber–physical systems is shaped by a dense and overlapping regulatory environment. Regulation EU 2024/1689, commonly referred to as the Artificial Intelligence Act, establishes a harmonised legal framework for the placing on the market, putting into service, and use of artificial intelligence systems within the Union [7]. At the same time, Regulation EU 2016/679 on data protection continues to govern the processing of personal data [8]. Directive EU 2022/2555 establishes cybersecurity obligations for essential and important entities [9]. Regulation EU 2024/2847 introduces horizontal cybersecurity requirements for products with digital elements [10]. Regulation EU 2023/2854 addresses fair access to and use of data [11]. Regulation EU 2022/868 seeks to facilitate trusted data sharing [12]. Regulation EU 2019/881 establishes a Union framework for cybersecurity certification [13]. In machinery-rich environments, Regulation EU 2023/1230 further connects product safety and digital intelligence [14]. Artificial intelligence adoption in cyber–physical systems therefore takes place under multiple interacting legal conditions rather than under a single technology specific regime [2,3,5,15].

This regulatory density is especially significant in cyber–physical systems because such systems are characterised by close coupling between digital functions and physical operations. Artificial intelligence in these environments may influence the timing, quality, and safety of decisions affecting energy dispatch, indoor climate control, industrial machinery, mobility coordination, logistics, and critical infrastructures. Compliance cannot therefore be treated as a downstream legal exercise. Data governance, traceability, logging, human oversight, cybersecurity, software update mechanisms, identity and access control, and resilience against cyber threats must increasingly be incorporated into system design from the outset [16–19]. Regulation consequently shapes not only whether artificial intelligence can be adopted, but also how cyber–physical systems must be architected, documented, procured, validated, and governed [5,17,20].

At the same time, the regulatory role of the European Union should not be understood solely as restrictive. Compliance obligations may slow experimentation, increase documentation burdens, and raise development and deployment costs [20–22]. Yet they also establish conditions for trustworthy artificial intelligence, security by design, stronger accountability, and greater interoperability across digital infrastructures [2,3,23]. This dual character is particularly important in cyber–physical systems, where public trust, safety, reliability, and resilience are not secondary considerations but central prerequisites for adoption. The European regulatory model therefore creates a persistent tension between short-term innovation friction and long-term institutional and technical robustness.

Despite a growing body of literature on artificial intelligence governance in individual sectors, the present knowledge base remains fragmented. Existing studies have examined artificial intelligence governance in domains such as smart grids, building energy management, smart cities, digital twins, autonomous systems, and industrial automation, often focusing on technological capabilities, optimisation potential, or domain-specific regulatory constraints [2,3,5,23,24]. While these studies provide valuable insights, they remain largely confined to individual sectors and do not sufficiently explain how a shared European Union regulatory environment shapes artificial intelligence adoption across cyber–physical systems with different operational logics, risk profiles, institutional arrangements, and

market conditions [4,5]. As a result, the interaction between regulatory structures and system-level implementation remains insufficiently synthesised. A cross-sector perspective is therefore needed to identify both recurring regulatory pressures across cyber–physical systems and the distinct ways in which these pressures are translated into architectural, organisational, and operational consequences in different domains [2–5].

The analysis therefore examines how European Union regulation influences the adoption of artificial intelligence in cyber–physical systems across four representative domains: energy and smart grids, smart buildings and building automation, mobility and transport systems, and industrial and manufacturing environments. These domains were selected as analytically representative rather than exhaustive. Together, they capture variation in infrastructure criticality, safety exposure, data intensity, institutional complexity, and installed base conditions, while maintaining comparability as mature classes of cyber–physical systems in which artificial intelligence is increasingly embedded [2–4,24] and where obligations related to data protection, cybersecurity, accountability, product safety, and sectoral governance are especially consequential.

Three analytical dimensions structure the synthesis. The first concerns regulatory barriers, including compliance burden, legal fragmentation, uncertainty around accountability, and the burden of high-risk governance. The second concerns technological and architectural challenges, including privacy-preserving data design, secure system stacks, logging and auditability, human oversight, interoperability, and lifecycle governance. The third concerns cross-sector implications for governance, innovation, and competitiveness, with particular attention to how legal requirements shape organisational capability, market structure, and deployment pathways.

The contribution lies in establishing a cross-sector and regulation-aware architectural perspective on artificial intelligence adoption in cyber–physical systems. More specifically, the analysis shows how European Union legal requirements are translated into system design constraints, integration requirements, operational governance conditions, and differentiated adoption pathways across representative cyber–physical domains.

The analysis is guided by three research questions. First, to what extent does the current European Union digital regulatory framework provide a coherent basis for governing AI-enabled cyber–physical systems in which digital failures may create physical, operational, safety, or societal harm? Second, how are obligations from the Artificial Intelligence Act, the General Data Protection Regulation, NIS2, the Cyber Resilience Act, the Data Act, the Data Governance Act, and product- or sector-specific safety rules translated into system-level requirements for safety, cybersecurity, accountability, transparency, adaptability, data governance, and traceability? Third, how do these regulatory translation effects differ across energy and smart grids, smart buildings and building automation, mobility and transport systems, and industrial and manufacturing environments when artificial intelligence is used at observational, advisory, or operational layers?

The remainder of the manuscript is organised as follows. Section 2 presents the review design, search strategy, eligibility criteria, source selection logic, and synthesis approach. Section 3 defines the conceptual foundations of artificial intelligence in cyber–physical systems. Section 4 maps the European Union regulatory landscape relevant to these systems. Section 5 examines the principal regulatory barriers to adoption. Section 6 analyses the technological and architectural implications of compliance. Section 7 compares adoption conditions across the four selected domains. Section 8 synthesises cross-sector implications for governance, innovation, and competitiveness. Section 9 develops policy and research recommendations. Section 10 discusses the broader significance of the findings. Section 11 concludes.

2. Methodology

A structured scoping review approach informed by PRISMA ScR logic was adopted in order to map an interdisciplinary evidence base on the relationship between European Union regulation and the adoption of artificial intelligence in cyber–physical systems. This approach was selected because the topic spans legal instruments, policy documents, technical studies, the architecture-oriented literature, standards-related materials, and sector-specific analyses distributed across multiple operational domains. The aim was not to test a narrowly defined causal hypothesis or aggregate results from a homogeneous empirical literature, but to identify, organise, and synthesise the principal regulatory, technological, and cross-sector patterns shaping the adoption of artificial intelligence in cyber–physical systems within the European Union. The review process therefore followed a structured scoping logic adapted to the heterogeneous character of the evidence base.

The review was delimited in two principal ways. First, the analysis focused on artificial intelligence as embedded in cyber–physical systems rather than artificial intelligence as a purely digital or standalone software phenomenon. The emphasis therefore fell on contexts in which artificial intelligence interacts with sensing infrastructures, communication layers, digital representations, control functions, operational technologies, and human supervisory arrangements in environments where physical processes are materially affected. Second, the review was bounded to the European Union regulatory context. Non-European sources were considered only where they clarified technical, conceptual, or comparative issues directly relevant to the main analytical objective.

The scope of the review was organised around four representative cyber–physical domains: energy and smart grids, smart buildings and building automation, mobility and transport systems, and industrial and manufacturing systems. These domains were selected because they capture different combinations of infrastructure criticality, safety exposure, data intensity, organisational complexity, and installed base conditions while still sharing core cyber–physical characteristics. Each exhibits a growing role for artificial intelligence in functions such as forecasting, optimisation, anomaly detection, coordination, and operational decision support. Each is also materially shaped by European Union rules concerning artificial intelligence, data protection, cybersecurity, data access, product security, and operational accountability.

The evidence base included peer-reviewed journal articles, conference papers, and official European Union regulations and directives. Primary legal sources were treated as indispensable because several of the most relevant instruments are recent and are represented most authoritatively in official legal texts rather than in secondary commentary.

The review period for the technical and policy literature was centred on the period from 2015 onwards. This boundary was selected to capture the rise of contemporary artificial intelligence deployment in operational environments together with the major regulatory developments associated with the General Data Protection Regulation and the subsequent evolution of European digital legislation. More recent legal and implementation-related materials up to 2026 were included where relevant because several major instruments affecting cyber–physical systems were enacted, entered into force, or moved into implementation during the period from 2022 to 2025.

Searches were conducted across Scopus, Web of Science, and IEEE Xplore using a structured query logic that originally combined five concept blocks: cyber–physical systems, representative application domains, artificial intelligence, regulatory and governance terms, and the European Union context. However, the concept block on cyber–physical systems limited the search results to fewer than five sources and was therefore excluded from the final search query. This low retrieval rate was interpreted as an indication that the relevant literature is often indexed through domain-specific terminology rather than

through the generic term cyber–physical systems. Studies relevant to the review were more commonly described using terms such as smart grids, smart buildings, transport systems, mobility systems, robotics, manufacturing, digital twins, operational technology, or industrial automation. Excluding the generic cyber–physical systems block therefore improved recall while retaining conceptual relevance through the domain blocks and inclusion criteria. The cyber–physical framing was subsequently applied during screening and synthesis rather than enforced as a mandatory keyword condition.

The final search query took the following form: ("smart grid*" OR "smart building*" OR mobility OR transport OR industr* OR manufacturing OR Robotics) AND ("artificial intelligence" OR AI OR "machine learning" OR "digital twin*") AND (regulation* OR governance OR compliance OR law*) AND ("European Union" OR EU). The database searches were complemented by targeted retrieval of primary legal materials from EUR Lex and official European Union institutional repositories. This combined search strategy was designed to retrieve the academic and technical literature corresponding to the representative cyber–physical domains analysed in the review, while ensuring that the relevant European legal and institutional framework was captured through authoritative primary sources.

Sources were included when they met at least one of the following criteria: they addressed European Union regulation relevant to artificial intelligence-enabled cyber–physical systems; they examined the design, deployment, governance, or adoption of artificial intelligence in cyber–physical or closely related operational systems; they analysed regulatory, cybersecurity, data governance, safety, accountability, interoperability, or human oversight implications for artificial intelligence-enabled systems; or they provided sector-specific evidence from energy and smart grids, smart buildings and building automation, mobility and transport systems, or industrial and manufacturing environments. Primary legal sources, official European Union institutional materials, and peer-reviewed studies were included when they were necessary to interpret regulatory obligations, implementation conditions, or architectural consequences.

Sources were excluded when they focused on artificial intelligence as a standalone digital technology without cyber–physical or operational relevance; addressed cyber–physical systems without meaningful treatment of artificial intelligence; discussed abstract artificial intelligence ethics without regulatory, architectural, or deployment linkage; examined non-European legal environments without clear comparative relevance to the European Union context; presented purely algorithmic or model performance studies without reference to deployment context, governance, system integration, compliance, or operational use; or addressed sectors outside the four selected domains unless they provided transferable insight into cross-sector regulation, cybersecurity, data governance, safety, or interoperability.

The search and selection process is summarised in Figure 1. Records identified through database searching and targeted source retrieval were screened in stages through duplicate removal, title and abstract review, and full-text assessment. In addition to conventional academic sources, primary legal texts and selected implementation-oriented materials were incorporated where they were necessary to interpret regulatory obligations, implementation conditions, or institutional guidance relevant to artificial intelligence adoption in cyber–physical systems. The final synthesis corpus therefore combined the academic literature with authoritative legal and policy materials appropriate to the interdisciplinary character of the review question.

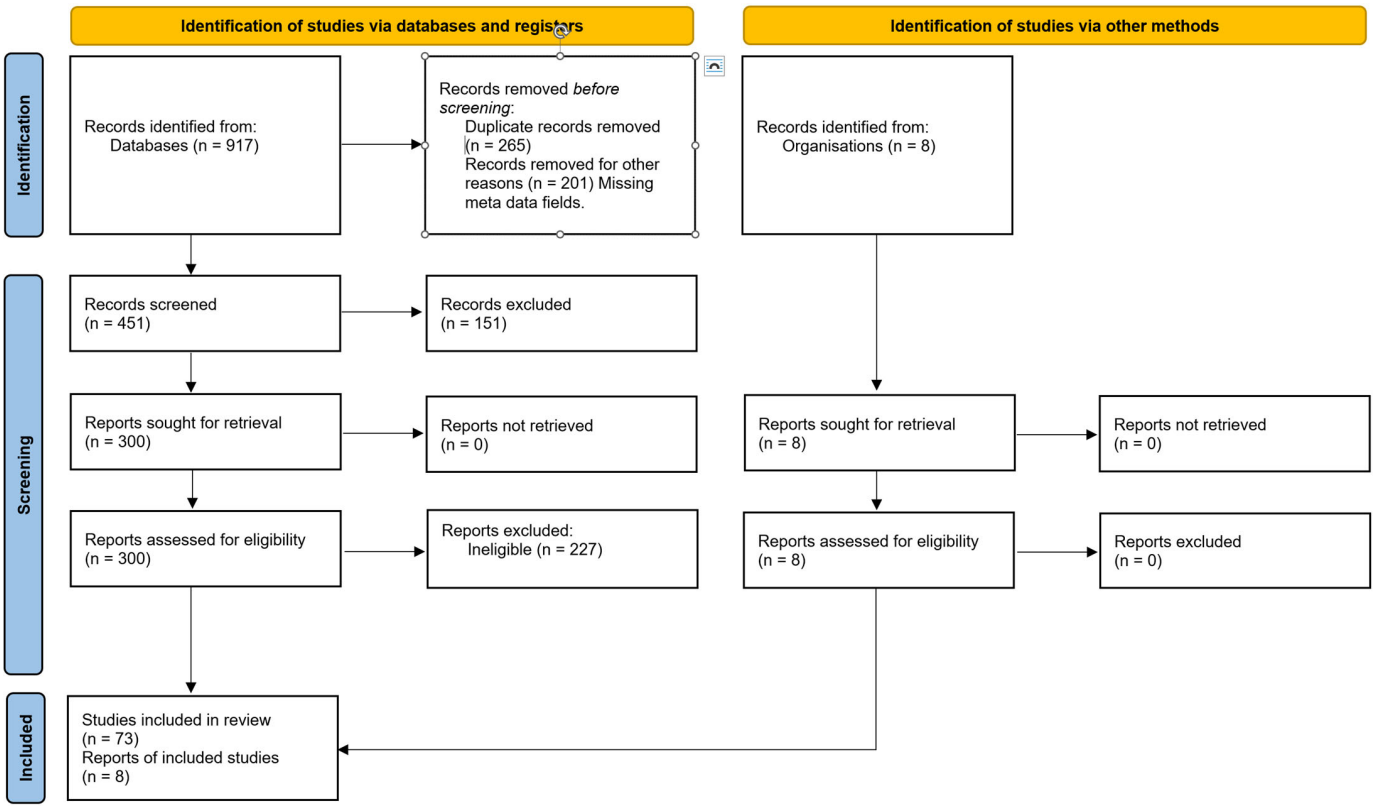


Figure 1. PRISMA ScR search and selection process.

Data extraction was organised through a structured charting approach rather than formal quality scoring. For each included source, the review recorded bibliographic information, source type, domain relevance, regulatory instruments addressed, artificial intelligence application context, cyber–physical characteristics, identified barriers, architectural implications, and cross-sector insights. Particular attention was given to how sources described the interaction between legal requirements and system design. This included the treatment of data minimisation, logging and traceability, explainability, human oversight, resilience requirements, software update obligations, conformity assessment, interoperability, and organisational accountability.

The synthesis combined deductive and inductive thematic analysis. Deductively, the analysis was structured around three main dimensions: regulatory barriers, technological and architectural challenges, and cross-sector implications for governance, innovation, and competitiveness. Inductively, subthemes emerged within these dimensions, including lawful data processing, high-risk artificial intelligence classification, cybersecurity obligations, documentation burdens, privacy-preserving architecture, legacy integration, certification pathways, human oversight design, and governance asymmetries across sectors and actor types. This combined strategy made it possible to retain a coherent article architecture while remaining responsive to variation across domains, source types, and implementation contexts.

Given the scoping character of the review and the heterogeneity of the source base, no formal quality score was applied across all included materials. Instead, evidentiary weight was differentiated by source type. Primary legal texts and official implementation materials were treated as authoritative for regulatory interpretation. Peer-reviewed studies were prioritised for technical, sector-specific, and implementation-related claims. This differentiated treatment of evidence was necessary in a field where legal authority, technical feasibility, and institutional interpretation must be considered together. Based on this differentiated treatment of evidence, the review required an analytical structure

capable of linking legal obligations, technical design requirements, and domain-specific implementation conditions.

To strengthen the analytical consistency of the review, the synthesis was structured around seven evaluation criteria derived from the recurring obligations and implementation concerns identified across the legal and technical evidence base. These criteria are safety, cybersecurity, accountability, transparency, adaptability, data governance, and traceability. They were selected because they capture the principal normative and architectural requirements that arise when artificial intelligence is embedded in cyber–physical systems whose outputs may influence physical processes, infrastructure operation, machinery behaviour, user environments, or public services. The criteria do not function as independent legal categories. Rather, they provide an analytical bridge between legal obligations and systems engineering consequences.

Safety refers to the capacity of the regulatory framework to address physical harm, unsafe operation, service disruption, and safety-relevant interaction between artificial intelligence functions and physical processes. Cybersecurity refers to the capacity to protect connected products, organisational infrastructures, data pipelines, software components, update mechanisms, and operational technology environments against cyber threats. Accountability refers to the allocation of responsibility among developers, manufacturers, providers, deployers, operators, integrators, and human supervisors. Transparency refers to the availability of meaningful information about system purpose, functioning, limitations, and operational implications. Adaptability refers to the ability of the regulatory and architectural framework to support software updates, model retraining, vulnerability remediation, changing operational contexts, and evolving standards without undermining compliance. Data governance refers to lawful, controlled, and interoperable access to data across personal, operational, device-generated, and cross-organisational contexts. Traceability refers to the ability to reconstruct data flows, model versions, system decisions, human interventions, incidents, and update histories across the lifecycle of artificial intelligence-enabled cyber–physical systems.

These criteria were applied in two steps. First, the major European Union legal instruments were examined according to the regulatory functions they perform across the seven criteria. Second, the four cyber–physical domains were compared according to how these criteria are translated into dominant architectural consequences, adoption bottlenecks, and governance requirements. This approach makes it possible to move beyond a descriptive mapping of legal instruments and towards a systematic assessment of whether the current regulatory framework provides a coherent basis for governing artificial intelligence-enabled cyber–physical systems. Table 1 summarises the analytical criteria used to structure the synthesis.

Several limitations should be recognised. The regulatory environment remains dynamic, and implementation guidance, standards, delegated acts, and sector-specific compliance practices may continue to evolve after the review period. The literature availability is also uneven across domains, with some sectors currently better documented than others. Document-based review methods are far better suited to mapping formal regulatory and technical discourse than to capturing tacit organisational experience from deployment practice. Despite these limitations, the review provides a structured analytical basis for understanding how European Union regulatory requirements are translated into system-level design constraints, architectural choices, and differentiated adoption conditions in cyber–physical systems.

Table 1. Analytical criteria structuring the synthesis.

Evaluation Criterion	Normative Concern	Architectural Manifestation in Artificial Intelligence-Enabled Cyber–Physical Systems	Main European Union Regulatory Link
Safety	Prevention of physical harm, unsafe operation, and service disruption	Safety cases, risk classification, validation, safe fallback, human intervention points	Artificial Intelligence Act, Machinery Regulation, sector-specific safety rules
Cybersecurity	Protection against attacks on products, organisations, data pipelines, models, and operational infrastructures	Secure-by-design products, vulnerability handling, update security, identity and access management, incident response	NIS2, Cyber Resilience Act, Cybersecurity Act
Accountability	Allocation of responsibility across providers, manufacturers, deployers, operators, integrators, and supervisors	Role allocation, documentation, governance workflows, approval chains, responsibility matrices	Artificial Intelligence Act, product safety law, liability-related frameworks
Transparency	Meaningful information about purpose, functioning, limits, and operational implications	User information, operator dashboards, explainability layers, documentation, model cards where relevant	Artificial Intelligence Act, General Data Protection Regulation, sector guidance
Adaptability	Capacity to remain compliant during software updates, retraining, changing use contexts, and evolving standards	Lifecycle governance, version control, change management, rollback, post-market monitoring	Artificial Intelligence Act, Cyber Resilience Act, Machinery Regulation
Data governance	Lawful and controlled use of personal, operational, device-generated, and shared data	Data minimisation, access rights, retention rules, edge–cloud partitioning, data space interfaces	General Data Protection Regulation, Data Act, Data Governance Act
Traceability	Ability to reconstruct decisions, data flows, model versions, incidents, and interventions	Logging, provenance, audit trails, time synchronisation, model and data lineage	Artificial Intelligence Act, NIS2, Cyber Resilience Act, General Data Protection Regulation

3. Conceptual Foundations: Artificial Intelligence in Cyber–Physical Systems

Cyber–physical systems are integrated operational environments in which computational processes interact continuously with physical processes through sensing, communication, control, and actuation. Their defining characteristic is not merely the presence of digital technology, but the tight coupling between software-based functions and material infrastructures whose behaviour unfolds in real time. Sensors generate data about physical conditions, networks coordinate that information, software interprets system states, and control actions influence machines, devices, buildings, infrastructures, or mobility environments. Cyber–physical systems therefore occupy the space in which digital decision processes and physical consequences become inseparable [4,17,25].

This distinguishes cyber–physical systems from stand-alone information systems. A conventional digital platform may support record keeping, communication, or recommendation without directly affecting the physical environment. By contrast, cyber–physical systems are implicated in physical state transitions. They regulate temperatures and ventilation conditions in buildings, coordinate distributed energy flows, govern the movement of vehicles, manage industrial machinery, monitor infrastructures, and support human de-

cision making in settings where timing, reliability, and safety are essential. The implications of artificial intelligence are therefore substantially different in cyber–physical systems than in purely digital environments. In cyber–physical systems, erroneous prediction, opaque optimisation, or compromised data pipelines may produce not only computational error but also service disruption, equipment damage, unsafe operating conditions, or wider societal harm [4,6,26].

Artificial intelligence intensifies this coupling by expanding the range of functions that digital systems can perform. Traditional automation in cyber–physical systems typically relies on predefined control logic, fixed thresholds, deterministic rules, and relatively stable assumptions about operating conditions. Artificial intelligence introduces the capacity to infer patterns from data, adapt to changing conditions, detect subtle anomalies, generate predictions, optimise across multiple objectives, and support or automate complex decision processes. This does not displace conventional control logic. Rather, it adds probabilistic, model-based, and data-driven capabilities to existing cyber–physical architectures. As a result, artificial intelligence changes both the functional potential and the governance requirements of cyber–physical systems [3,16,27].

From a functional perspective, the role of artificial intelligence in cyber–physical systems can be understood through five recurring capabilities. Artificial intelligence strengthens perception by enabling richer interpretation of sensor data, including event classification, state estimation, and anomaly detection. It supports prediction by forecasting future states such as energy demand, equipment degradation, occupancy dynamics, or process deviation. It enables optimisation by coordinating resources, schedules, and control actions across multiple interacting objectives. It strengthens adaptive control by allowing system behaviour to respond more flexibly to dynamic conditions. It also contributes to decision support by helping human operators interpret complex system states, prioritise interventions, and evaluate operational alternatives [2,3,28].

These capabilities should not be understood as the insertion of a model into an otherwise unchanged technical environment. Artificial intelligence in cyber–physical systems is embedded within broader architectural arrangements that include sensor infrastructures, edge and cloud resources, communication protocols, digital representations, control loops, device identities, software update pathways, and human oversight mechanisms. The relevant unit of analysis is therefore not the isolated model, but the socio-technical system within which the model is trained, integrated, monitored, documented, secured, and used [5,16,27]. This distinction is essential because the practical implications of artificial intelligence depend on the wider system conditions under which it operates.

Several design properties are especially important in this context. One is real-time or near-real-time responsiveness. In many cyber–physical settings, system value depends on acting within narrow temporal windows. Another is safety significance. Even where a system is not formally classified as safety-critical, operational failure may still affect physical assets, essential services, or human wellbeing. A third is interdependence. Cyber–physical systems rarely operate as isolated units. They are connected to vendors, infrastructures, operators, and adjacent systems through data flows and control relationships. A fourth is legacy integration. Many operational environments combine new artificial intelligence-enabled components with long-lived assets, proprietary interfaces, and fragmented governance arrangements. A fifth is accountability complexity. Decisions and outcomes are often distributed across software providers, system integrators, infrastructure operators, maintenance teams, and human supervisors. These properties explain why the governance of artificial intelligence in cyber–physical systems cannot be reduced to general concerns about transparency or fairness alone. The central issue is how legal, technical,

and organisational requirements interact in environments where digital outputs can have material consequences [4,29,30].

For analytical clarity, three layers of artificial intelligence integration in cyber–physical systems can be distinguished. The observational layer uses artificial intelligence primarily for monitoring, diagnostics, forecasting, and pattern recognition without direct control action. The advisory layer provides recommendations or decision support to human operators or supervisory systems. The operational layer directly influences orchestration routines, control decisions, or automated intervention. These layers are not always sharply separated in practice, but the distinction is useful because governance requirements become more demanding as systems move from observation to advice and from advice to operational influence [28,31,32].

A further conceptual issue concerns the relationship between cyber–physical systems and socio-technical governance. Cyber–physical systems are often described in engineering terms, yet in practice, they are embedded within institutional routines, procurement structures, professional roles, sectoral regulation, and public expectations. Artificial intelligence adoption therefore depends not only on technical feasibility, but also on governance capacity. Organisations must determine who validates models, who supervises outputs, how incidents are escalated, how updates are managed, how data rights are interpreted, and how competing objectives such as efficiency, safety, privacy, and resilience are balanced. This is particularly relevant in Europe, where the regulatory model places strong emphasis on accountability, human oversight, transparency, and the protection of fundamental rights [16,33,34]. Artificial intelligence in cyber–physical systems must therefore be engineered as part of a governed operational arrangement rather than as a detached optimisation module [3,16,17].

Within this framing, artificial intelligence adoption in cyber–physical systems can be defined as the process through which artificial intelligence-based functions are developed, integrated, validated, deployed, governed, and scaled within environments where digital outputs influence physical processes or materially shape operational decisions. Adoption is therefore broader than procurement or installation. It includes technical integration, organisational alignment, legal compliance, operator trust, and lifecycle management. Many initiatives in cyber–physical systems do not fail because the model performs poorly in principle. They fail because they cannot be integrated safely into workflows, cannot satisfy documentation and accountability demands, cannot secure adequate data rights, or cannot be maintained under cybersecurity and product governance constraints [2,20,22].

This conceptual framing establishes the basis for the analysis that follows. In cyber–physical systems, artificial intelligence is not simply a software capability added to an existing technical stack. It is part of a larger operational arrangement in which data governance, cybersecurity, product integrity, human oversight, and institutional accountability become inseparable from functional performance. Cyber–physical systems therefore provide a particularly revealing lens for analysing how regulatory conditions are translated into concrete design choices, integration patterns, and operational consequences.

4. European Union Regulatory Landscape Relevant to Artificial Intelligence in Cyber–Physical Systems

The European Union regulatory landscape relevant to artificial intelligence in cyber–physical systems is not organised around a single legal instrument. It is composed of overlapping layers of regulation that address distinct but interacting dimensions of digital and physical system governance. At the centre of this landscape is Regulation EU 2024/1689, the Artificial Intelligence Act, which lays down harmonised rules on artificial intelligence across the Union. Artificial intelligence-enabled cyber–physical systems are also

shaped by Regulation EU 2016/679 on data protection, Directive EU 2022/2555 on cybersecurity, Regulation EU 2024/2847 on horizontal cybersecurity requirements for products with digital elements, Regulation EU 2023/2854 on access to and use of data, Regulation EU 2022/868 on data governance, Regulation EU 2019/881 on cybersecurity certification, and in machinery-rich contexts, Regulation EU 2023/1230 on machinery. Artificial intelligence adoption in cyber–physical systems must therefore be understood as taking place within a layered compliance environment in which artificial intelligence law, data law, cybersecurity law, product security, and sector-specific operational obligations interact [2–5].

This layered structure is especially significant in cyber–physical systems because such systems combine software intensive intelligence with material processes, connected devices, operational technologies, and institutional responsibilities. A single deployment may therefore fall simultaneously within multiple legal categories. An artificial intelligence-enabled industrial platform may trigger obligations under the Artificial Intelligence Act because of its artificial intelligence functionality, under the Cyber Resilience Act because it includes products with digital elements, under the Data Act because device-generated data must be made accessible under-specified conditions, under NIS2 because the operating organisation is an essential or important entity, and under the Machinery Regulation because the system influences the safe functioning of equipment [4,5,35]. The regulatory significance therefore lies not only in the existence of multiple instruments, but in the cumulative way their requirements shape the design, deployment, and governance of cyber–physical systems [2,3,5].

4.1. The Artificial Intelligence Act as the Central Horizontal Framework

The Artificial Intelligence Act provides the first comprehensive horizontal legal framework for artificial intelligence in the European Union [7,33,34,36]. Its importance for cyber–physical systems lies in the fact that it regulates artificial intelligence not as abstract software alone, but in relation to intended use, risk, market placement, safety significance, and operational context. This is especially important where artificial intelligence is embedded in infrastructure, machinery, transport, buildings, and industrial operations. The act introduces a governance model in which documentation, logging, human oversight, risk management, and conformity-related obligations become integral to engineering and deployment planning [17,20,22].

For cyber–physical systems, the high-risk categories are particularly consequential. Artificial intelligence used as a safety component in machinery, artificial intelligence used in critical infrastructure contexts, and artificial intelligence deployed in operational environments with material safety or societal implications may fall within high-risk obligations [2,21,37]. This matters because cyber–physical systems often use artificial intelligence in monitoring, optimisation, diagnostics, control support, and resource coordination in environments where physical safety, service continuity, and resilience are essential. High-risk governance therefore affects not only software developers but also system integrators, infrastructure operators, equipment manufacturers, and deployers [20,22,38].

The practical consequence is substantial. High-risk artificial intelligence governance requires structured risk management, technical documentation, logging capabilities, data governance measures, transparency and information duties, human oversight arrangements, robustness expectations, and post-market monitoring [17,22,39]. In cyber–physical systems, these obligations must be incorporated into environments that often already include legacy equipment, proprietary control systems, safety procedures, and multiple organisational actors. The relevant unit of compliance is therefore rarely the artificial intelligence model in isolation. It is the wider socio-technical arrangement within which the model is embedded and through which its outputs acquire operational effect [5,16,17].

Implementation capacity is also relevant. The transition to conformity assessment is likely to create bottlenecks, particularly because notified body capacity may remain limited relative to expected demand [20,22,38]. This is especially significant for smaller firms and specialised integrators seeking to deploy high-risk artificial intelligence-enabled systems in operational settings. The Artificial Intelligence Act is intended to support trustworthy adoption and legal certainty, but in the short to medium term, it also raises the threshold for market entry and deployment in precisely those contexts where cyber–physical consequences are most significant [20–22].

4.2. The General Data Protection Regulation and the Governance of Personal Data

The General Data Protection Regulation remains one of the foundational instruments shaping artificial intelligence adoption in cyber–physical systems. Its relevance arises wherever artificial intelligence-enabled environments process personal data directly or indirectly [3,8,18]. In buildings, this may include occupancy-related sensing, access patterns, or behavioural traces [40]. In mobility systems, it may include location data, identifiable travel records, or vehicle-related user data [5]. In industrial contexts, worker-adjacent monitoring data may create similar issues. Even where data are not obviously personal at first sight, combinations of sensor streams and contextual information in cyber–physical environments may create identifiability or profiling risks [3,8,41].

The significance of the regulation is not limited to lawful processing. It also shapes architecture. Principles such as data minimisation, purpose limitation, storage limitation, security of processing, and privacy by design influence where data are processed, how long they are retained, whether they are aggregated or pseudonymised, and how training and inference pipelines are structured [3,16,18]. In this sense, the regulation is not simply a downstream compliance burden. It acts directly on sensing strategies, data retention logic, edge versus cloud decisions, and governance arrangements in data-rich cyber–physical systems [3,16,18].

The interaction with artificial intelligence is particularly significant because artificial intelligence often depends on rich, continuous, and context-sensitive data. The stronger the demand for behavioural granularity and historical depth, the greater the tension between model ambition and lawful governance [3,27,42]. In cyber–physical systems, this tension becomes especially visible because data are generated continuously across environments that often contain human presence, movement, and routine. The General Data Protection Regulation therefore affects not only whether data may be used, but how artificial intelligence ready data environments can be designed in the first place [3,16,18].

4.3. NIS2 and the Cybersecurity Governance of Operating Organisations

Directive EU 2022/2555, commonly known as NIS2, establishes measures for a high common level of cybersecurity across the Union. Its importance for artificial intelligence-enabled cyber–physical systems lies in the fact that it governs organisations rather than products alone. Essential and important entities in covered sectors must implement cybersecurity risk management measures and comply with governance and incident reporting obligations [6,9]. Because many cyber–physical systems operate in sectors such as energy, transport, manufacturing of critical products, water, health, and public administration, NIS2 becomes a major legal condition for the organisational environments in which artificial intelligence is deployed.

This is conceptually important because it shifts attention from the model to the operational setting within which that model is embedded. A technically sophisticated artificial intelligence component cannot be meaningfully compliant in practice if it is deployed within an organisation lacking appropriate cybersecurity governance, incident handling

capacity, supply chain control, or resilience measures [2,6,43]. In cyber–physical systems, where artificial intelligence often depends on live data streams, connected assets, remote interfaces, and distributed infrastructures, the organisational layer is especially consequential. Artificial intelligence adoption must therefore be analysed through governance capacity as well as through technical functionality [6,20,43].

NIS2 is also important because it reinforces the idea that cybersecurity in cyber–physical systems is not an optional enhancement added after deployment. It is a precondition for operational legitimacy in sectors where continuity, resilience, and incident preparedness are already matters of public and institutional concern [6,26,43]. The organisational environment into which artificial intelligence is introduced is therefore part of the regulatory condition of adoption.

4.4. The Cyber Resilience Act and Security Requirements for Products with Digital Elements

Regulation EU 2024/2847, the Cyber Resilience Act, establishes horizontal cybersecurity requirements for products with digital elements. This regulation is highly relevant because cyber–physical systems are typically composed of connected devices, embedded software, gateways, controllers, interfaces, and updateable digital components distributed across physical environments [4,6,10]. Unlike NIS2, which addresses organisational obligations, the Cyber Resilience Act focuses on the security characteristics of the products themselves across their lifecycle.

For artificial intelligence-enabled cyber–physical systems, the implication is clear. Artificial intelligence rarely exists without supporting digital products such as sensors, controllers, edge devices, communication modules, supervisory software, or embedded platforms. The security posture of these components directly affects the trustworthiness, integrity, and resilience of artificial intelligence-enabled operation [2,4,6]. By introducing horizontal requirements for products with digital elements, the act brings product cybersecurity into the core of system design and market placement.

This is particularly important in cyber–physical environments where compromised devices can produce cascading operational consequences. It also broadens the security challenge beyond conventional software vulnerability management. Artificial intelligence-enabled systems may be affected by issues related to model integrity, training data manipulation, insecure update pipelines, and the interaction between software vulnerability and safety-relevant control environments [4,6,44]. The Cyber Resilience Act therefore strengthens the requirement that artificial intelligence in cyber–physical systems must be grounded in lifecycle security rather than in model performance alone [4,6,10,29].

4.5. The Data Act and the Question of Access to Device-Generated Data

Regulation EU 2023/2854, the Data Act, lays down harmonised rules on fair access to and use of data. For cyber–physical systems, this is especially significant because these systems generate large volumes of operational and device-related data through sensors, equipment, connected products, and digital services. Artificial intelligence adoption depends heavily on access to such data for training, optimisation, diagnostics, forecasting, and service improvement. The Data Act therefore affects the practical conditions under which developers, operators, users, and third parties may gain lawful access to relevant data resources [5,11,18,45].

The strategic significance of the act lies in its potential to reduce data lock-in and rebalance control over machine-generated data. In cyber–physical environments, data are often concentrated in vendor-controlled platforms or fragmented across organisations, devices, and service layers. This can constrain artificial intelligence experimentation, reinforce proprietary dependence, and limit interoperability [2,5,46]. By creating harmonised rules

on access to and use of data, the Data Act contributes to a broader regulatory effort to make operational data more usable under fair conditions.

At the same time, legal access does not automatically create technical usability. Even where data can be accessed lawfully, artificial intelligence adoption still depends on semantics, metadata quality, time alignment, interface stability, and secure exchange. The Data Act therefore creates enabling conditions for more open data environments, but its practical effect in cyber–physical systems depends on whether those legal rights can be translated into technically workable data infrastructures [5,18,46].

4.6. The Data Governance Act and Trusted Data Sharing

Regulation EU 2022/868, the Data Governance Act, establishes a framework for European data governance intended to facilitate trusted data sharing across sectors. Its relevance to cyber–physical systems lies in the fact that artificial intelligence adoption increasingly depends not only on data held by a single operator but also on the ability to combine or reuse data across institutional and sectoral boundaries. This is especially visible in smart grids, transport ecosystems, industrial value chains, and urban infrastructures, where data relevant to artificial intelligence may be dispersed across public actors, infrastructure operators, device providers, and private firms [2,5,12,18].

The Data Governance Act should therefore be understood as an enabling layer within the wider regulatory landscape. Whereas the General Data Protection Regulation protects personal data and the Data Act structures rights around access and use, the Data Governance Act contributes to the institutional conditions for trusted sharing and reuse. This matters for federated architectures, common data spaces, and more integrated forms of artificial intelligence-enabled coordination [5,18,42].

At the same time, the existence of a legal framework for sharing does not by itself resolve semantic, technical, and organisational interoperability problems. The act creates legal possibility, but technical implementation remains a substantial challenge. Its significance therefore lies in enabling governance conditions that must still be matched by architecture, standardisation, and operational coordination [5,18,46].

4.7. The Cybersecurity Act and the Role of Certification

Regulation EU 2019/881, the Cybersecurity Act, strengthens the role of ENISA [2,6,13] and establishes a European cybersecurity certification framework for information and communication technology products, services, and processes. For artificial intelligence-enabled cyber–physical systems, certification is significant because assurance increasingly functions as a mechanism for demonstrating trustworthiness in procurement-sensitive, critical, or safety-relevant environments. Certification does not resolve all risks associated with artificial intelligence-enabled operation, but it contributes to a broader regulatory logic in which security, evidence, and assurance become strategic assets rather than optional additions [4,6,47].

In cyber–physical systems, this relevance is magnified by system interdependence. Artificial intelligence-enabled infrastructures are often assembled from multiple components, vendors, and service layers, which makes assurance difficult without recognised schemes, technical standards, or certification pathways. The Cybersecurity Act therefore supports the emergence of more structured trust mechanisms in complex operational environments where confidence in digital components increasingly affects the adoptability of artificial intelligence [5,6,47].

4.8. Product Safety, Machinery, and the Safety-Relevant Dimension of Artificial Intelligence in Cyber–Physical Systems

Regulation EU 2023/1230 on machinery highlights that the governance of artificial intelligence in cyber–physical systems cannot be reduced to digital regulation alone. Where artificial intelligence is embedded in or materially influences machinery and related products, product safety law remains a major determinant of lawful deployment. This is especially relevant in industrial and manufacturing systems, but the logic extends more broadly to environments in which actuation, automation, and physical equipment interact with artificial intelligence-enabled decision making [14,35,48–50].

The broader point is that artificial intelligence in cyber–physical systems becomes legally significant not only because it processes data or supports decisions, but because it may affect the safe functioning of physical systems. The interaction between digital governance and product safety frameworks is therefore a central feature of the European regulatory landscape. This is one reason why a cyber–physical perspective is analytically stronger than a purely software-centred view of artificial intelligence regulation. It makes visible the point at which digital functionality, operational consequence, and legal responsibility intersect [29,30,35].

4.9. A Multi-Layered Compliance Environment

Taken together, these instruments show that artificial intelligence adoption in cyber–physical systems within the European Union occurs within a multi-layered compliance environment composed of at least six interacting dimensions, shown in Figure 2. The first is artificial intelligence governance, centred on the Artificial Intelligence Act. The second is personal data governance, centred on the General Data Protection Regulation. The third is organisational cybersecurity governance, centred on NIS2. The fourth is product and software security governance, shaped strongly by the Cyber Resilience Act and supported by the Cybersecurity Act. The fifth is data access and sharing governance, structured through the Data Act and the Data Governance Act. In safety-relevant environments, product and sector-specific law adds a further layer [4,5,35].

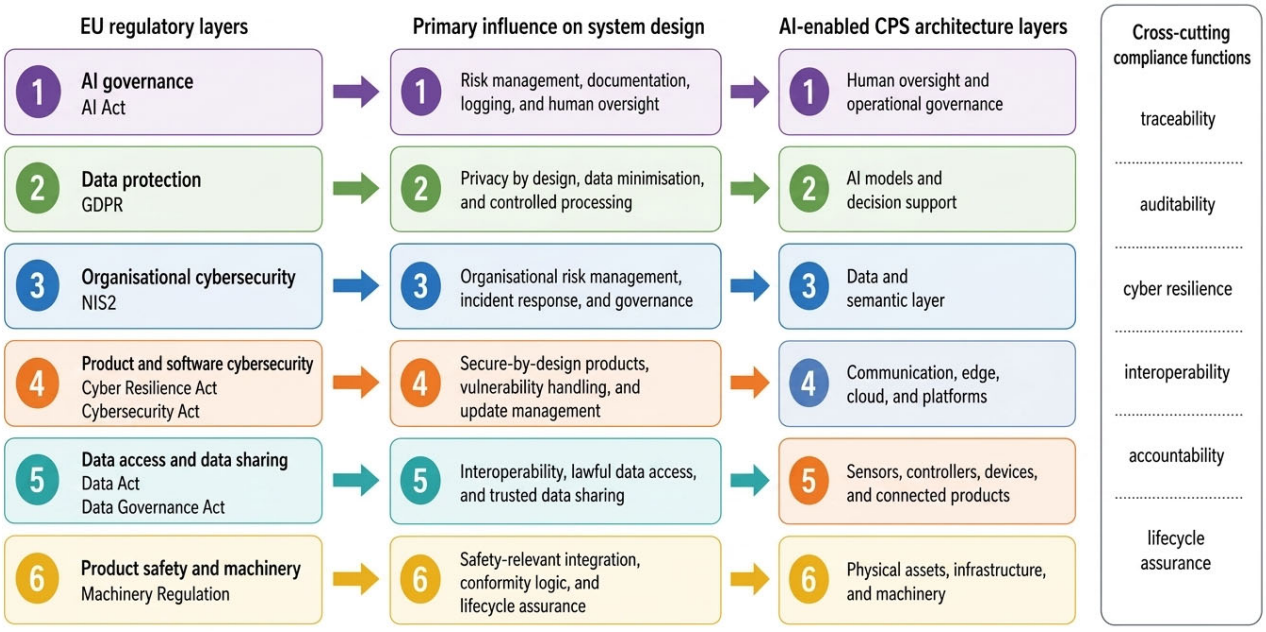


Figure 2. Multi-layered EU regulatory framework shaping AI-enabled cyber–physical systems.

The analytical significance of this structure lies in cumulative translation rather than simple accumulation. The relevant question is not merely how many instruments apply, but how their combined requirements are translated into the design, integration, governance, and assurance of artificial intelligence-enabled cyber–physical systems. Legal obligations influence where data may flow, how devices must be secured, how artificial intelligence functions must be documented, how responsibilities are allocated, what interfaces are required for lawful data access, and what forms of assurance may be needed by operators or procurers [5,17,20]. The European regulatory landscape therefore shapes the operational design space within which artificial intelligence-enabled cyber–physical systems can be developed, integrated, and scaled.

5. Regulatory Barriers to Artificial Intelligence Adoption in Cyber–Physical Systems

The adoption of artificial intelligence in cyber–physical systems within the European Union is shaped by a regulatory environment that creates both legal discipline and practical friction. The main challenge is not that European law simply prohibits artificial intelligence in cyber–physical contexts. The deeper issue is that adoption must occur under overlapping obligations concerning risk classification, documentation, data protection, cybersecurity, product security, access to data, and operational accountability. In cyber–physical systems, where digital outputs can materially influence physical processes, these obligations become especially consequential because compliance affects procurement, integration, validation, lifecycle management, and operational governance. The barriers examined below are analytically separated for clarity, even though in practice they reinforce one another [2,3,22].

5.1. High-Risk Governance and the Burden of the Artificial Intelligence Act

A first barrier arises from the governance model established by the Artificial Intelligence Act. Because many cyber–physical applications are embedded in environments linked to infrastructure, machinery, transport, buildings, and other operational settings where reliability and safety matter, a substantial share of artificial intelligence-enabled innovation may fall into contexts where legal classification and compliance obligations are materially demanding [21,22,37].

The burden does not arise only from the possibility of classification as high-risk. It arises from the operational consequence of that classification. High-risk governance requires risk management, technical documentation, logging capabilities, data governance measures, transparency, information duties, human oversight arrangements, and post-market monitoring. In a purely digital service environment, these requirements may be burdensome but still manageable within software development workflows. In cyber–physical systems, they must be incorporated into environments that include legacy equipment, physical assets, safety procedures, proprietary interfaces, and multiple organisational actors. This makes compliance more difficult and more expensive [20–22].

The barrier is especially acute for small and medium-sized enterprises, specialist system integrators, municipal actors, and infrastructure operators with limited internal regulatory capacity [20,22,51,52]. For such actors, the Artificial Intelligence Act may raise the threshold for experimentation and deployment by making it necessary to establish compliance competence before operational value can be realised. The expected scarcity of conformity assessment capacity further increases uncertainty and may delay market entry or deployment in precisely those cases where cyber–physical consequences are most relevant [21,22,38].

5.2. Data Protection Constraints and the Complexity of Lawful Processing

A second barrier concerns the General Data Protection Regulation and the challenge of lawful data processing in cyber–physical environments. Artificial intelligence adoption often depends on rich, continuous, and context-sensitive data flows. Yet data protection law requires lawful basis, purpose limitation, data minimisation, storage limitation, and security of processing. These principles are normatively clear but operationally difficult in environments where data are generated continuously across devices, spaces, and infrastructures [3,8,16].

A building management system may not be designed originally around fine-grained consent logic. A transport platform may require detailed event data for optimisation while simultaneously facing strong privacy constraints. An industrial analytics system may collect machine and worker-adjacent data in ways that create uncertainty about controller responsibilities, purpose boundaries, and proportionality. In such cases, data protection does not merely regulate data after collection. It limits the design space for sensing strategies and data architectures from the outset [3,8,16,41].

The effect is often architectural. Organisations may move towards edge processing, stronger aggregation, pseudonymisation, stricter retention logic, or reduced model ambition where personal data risk becomes too difficult to manage. Such design shifts may improve privacy and trust, but they may also make artificial intelligence development more costly, less flexible, or less accurate in environments where granular data would otherwise support stronger performance [3,27,42]. The barrier is therefore one of structured tension between data richness and lawful governance.

5.3. Cybersecurity Obligations as a Barrier to Deployment Readiness

A third barrier arises from the cybersecurity obligations that now shape both organisations and digital products in Europe. NIS2 imposes cybersecurity risk management and incident reporting duties on many entities operating cyber–physical systems. The Cyber Resilience Act imposes lifecycle security requirements on products with digital elements. The Cybersecurity Act strengthens the role of certification. Together, these instruments make cybersecurity a legal and operational precondition for deployment rather than a secondary technical quality concern [4,6,9,10,13,43].

For artificial intelligence adoption, this creates friction because artificial intelligence-enabled cyber–physical systems depend on connected sensors, controllers, edge devices, gateways, supervisory software, and communication infrastructures. If these components are insecure, the trustworthiness of the artificial intelligence function becomes operationally irrelevant. A strong forecasting or control support model embedded in a weak digital stack may remain unsuitable for deployment in regulated environments because the surrounding system fails to meet required resilience expectations [2,4,32].

This challenge is amplified in legacy-rich sectors. Many cyber–physical systems operate with ageing field devices, fragmented operational technology networks, and vendor-specific interfaces not originally designed for contemporary cybersecurity requirements. Artificial intelligence adoption may therefore be postponed not because the artificial intelligence function lacks value, but because the surrounding infrastructure is not yet secure or governable enough to support lawful and resilient deployment [2,43,44]. In advanced settings, further challenges also arise from artificial intelligence-specific cyber threats such as adversarial inputs, model extraction, and training data manipulation, which extend beyond traditional software vulnerability models [32,44,53].

5.4. Responsibility, Liability, and Human Oversight Uncertainty

A fourth barrier concerns uncertainty surrounding responsibility, liability, and human oversight in artificial intelligence-enabled cyber–physical systems. European regulation increasingly insists on accountability, traceability, and meaningful human oversight, particularly in higher risk contexts. In practical deployment settings, however, it is often difficult to determine how responsibility should be distributed between model developers, component manufacturers, system integrators, infrastructure operators, maintenance providers, and human supervisors [29,30,54].

Cyber–physical systems intensify this challenge because decisions and outcomes are distributed across socio-technical chains rather than produced by a single actor. An artificial intelligence-enabled anomaly detection system may provide alerts that influence operator intervention. A predictive optimisation system may alter schedules that affect physical assets. A control support module may generate recommendations accepted by supervisory software rather than by a human directly. In such cases, formal demand for human oversight does not automatically resolve the practical question of who is responsible when harm occurs, performance degrades, or an incident unfolds [30,54,55].

Manufacturer and provider responsibility becomes especially difficult where cybersecurity defects, insecure update mechanisms, or poorly governed model changes contribute to physical harm or operational disruption. Traditional product liability logic was developed around relatively stable products whose defectiveness could often be assessed at the point of market placement. Artificial intelligence-enabled cyber–physical systems challenge this logic because their risk profile may change after deployment through software updates, model retraining, vulnerability discovery, data drift, interface changes, or new forms of connectivity. Responsibility therefore needs to be understood as a lifecycle question rather than only as a point of market placement question. This requires clearer allocation of duties for secure updates, post-deployment monitoring, incident response, change documentation, and the preservation of traceability across manufacturers, providers, deployers, integrators, and operators.

This uncertainty slows adoption because risk-averse operators may prefer conventional automation or manual procedures over artificial intelligence-enabled arrangements whose governance implications remain unclear. The barrier is therefore not only liability in a narrow legal sense. It is also institutional uncertainty around role allocation, escalation, validation, update approval, model drift management, and intervention authority [28,29,31].

5.5. Fragmentation Across Horizontal, Sectoral, and Product-Specific Rules

A fifth barrier arises from the fragmented and layered character of the regulatory environment itself. Artificial intelligence-enabled cyber–physical systems are rarely governed by one legal regime alone. They may be subject simultaneously to the Artificial Intelligence Act, the General Data Protection Regulation, NIS2, the Cyber Resilience Act, the Data Act, the Data Governance Act, the Cybersecurity Act, and sector- or product-specific frameworks such as the Machinery Regulation [3,5,6]. This creates a cumulative compliance burden, especially for cross-domain or system of systems environments.

The challenge is not simply the number of rules. It is the difficulty of translating multiple legal logics into a coherent implementation pathway. One regulation may focus on data rights, another on organisational cybersecurity, another on product lifecycle obligations, and another on artificial intelligence documentation and oversight. Each has its own scope, terminology, institutional context, and compliance logic [3,5,6,56]. In cyber–physical systems, where deployment depends on the integration of devices, software, services, and operational practices, this fragmentation slows decision making and increases transaction costs [5,20,22].

The problem is especially acute in cross-sector cyber–physical applications, such as a smart building linked to grid services or an industrial site using artificial intelligence-enabled energy optimisation connected to external infrastructure. Such systems may not fit neatly into one sectoral compliance pathway. Fragmentation therefore favours larger actors with stronger internal legal and regulatory capacity and discourages modular experimentation by smaller organisations [22,51,52].

5.6. Interoperability Gaps and the Absence of Harmonised Implementation Pathways

A sixth barrier concerns interoperability and the continued absence of fully harmonised implementation pathways. European regulation increasingly recognises the importance of interoperability, data sharing, and common digital infrastructures. Yet many cyber–physical systems remain fragmented across proprietary interfaces, vendor-specific formats, legacy protocols, and institutionally siloed data environments. Compliance often presupposes a degree of system visibility and control that such architectures do not provide [2,5,23].

Traceability requires access to logs and interfaces. Data governance requires clarity about flows and rights. Cybersecurity requires consistent control over assets and update mechanisms. Human oversight requires understandable system behaviour and identifiable intervention points. Where cyber–physical environments are technically fragmented, meeting these expectations becomes more difficult and more expensive. Organisations may understand the general legal direction, yet lack concrete reference architectures, accepted standards profiles, or sector-specific guidance showing how compliance should be operationalised [17,25,27].

The absence of harmonised implementation pathways also affects market dynamics. It can lock adoption into ad hoc or vendor-driven approaches, reduce procurement flexibility, and increase long-term dependency on closed ecosystems. In cyber–physical systems, interoperability is not a matter of convenience alone. It is tied to resilience, auditability, maintainability, and governability [2,5,23].

5.7. Cumulative Adoption Friction

Taken together, the barriers examined above show that the principal regulatory challenge for artificial intelligence adoption in cyber–physical systems in Europe is cumulative adoption friction, as illustrated in Figure 3. High-risk governance, lawful data processing constraints, cybersecurity obligations, accountability demands, regulatory fragmentation, and interoperability gaps do not operate independently. They interact across legal, technical, organisational, and market dimensions. A system that is difficult to secure is also difficult to certify, difficult to trust, and difficult to govern. A system with unclear data rights is also difficult to optimise, difficult to document, and difficult to scale. A system with fragmented interfaces is also difficult to audit, supervise, and maintain.

A practical example can be seen in an AI-enabled smart building that is connected to local energy flexibility services. The building operator may wish to use occupancy data, indoor climate data, energy prices, weather forecasts, and grid signals to optimise HVAC operation and shift flexible demand. The use case appears technically straightforward, but cumulative adoption friction emerges quickly. GDPR constrains the collection and reuse of occupancy-related data. The Cyber Resilience Act affects the connected sensors, gateways, and controllers used to collect and act on data. NIS2 may become relevant if the building belongs to an essential or important entity. The Data Act may influence access to device-generated data held by vendors. The AI Act may become more relevant if the optimisation function affects safety-relevant building services or becomes part of higher risk operational decision making. At the same time, legacy building management systems

may lack secure update mechanisms, standardised interfaces, or sufficient logging. The adoption barrier is therefore not one single rule, but the combined requirement to make the use case lawful, secure, interoperable, explainable, traceable, and operationally governable.

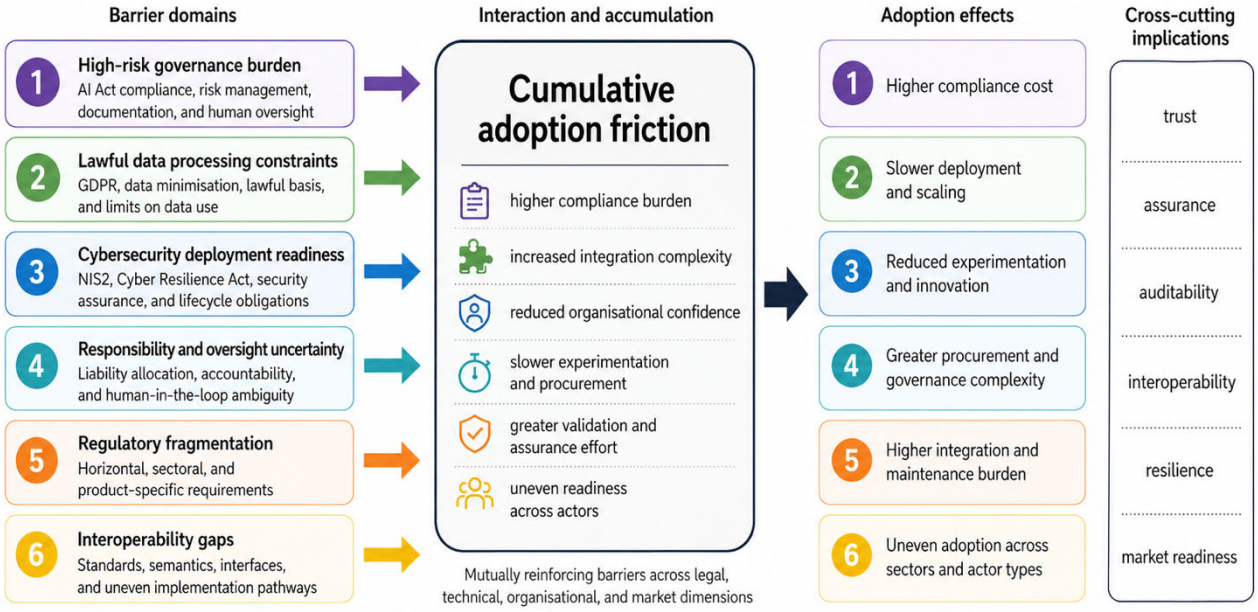


Figure 3. Cumulative adoption friction created by interacting regulatory barriers in AI-enabled cyber–physical systems.

This cumulative character is especially important in cyber–physical systems because deployment depends not only on model performance, but on the coordinated reliability of data pipelines, connected devices, software components, human workflows, and organisational responsibilities. The regulatory challenge therefore lies less in isolated legal prohibition than in the combined burden of making artificial intelligence-enabled cyber–physical systems sufficiently lawful, secure, explainable, maintainable, and governable for operational deployment. This is also the point at which regulatory barriers must be examined as design constraints, since the practical effect of law is ultimately realised through architecture.

The dual role of regulation depends on the maturity of the regulatory instrument, the availability of implementation guidance, the capacity of the adopting organisation, and the layer at which artificial intelligence is integrated. Regulation primarily functions as a constraint when obligations are new, implementation guidance is incomplete, conformity pathways are uncertain, or the adopting actor lacks legal, cybersecurity, and systems engineering capacity. This is especially visible where AI functions move from observational monitoring to advisory decision support or operational influence, because the need for documentation, oversight, traceability, and assurance increases with the potential for physical or societal consequence. Regulation functions more clearly as an enabling condition when it reduces uncertainty, clarifies responsibility, improves cybersecurity baselines, supports data access, strengthens interoperability, and creates trust among operators, users, procurers, and public authorities. The same legal instrument may therefore constrain short-term experimentation while enabling longer term adoption by making AI-enabled cyber–physical systems more secure, accountable, and institutionally acceptable.

6. Technological and Architectural Challenges Under European Union Regulatory Conditions

The barriers identified in the previous section do not remain external to system design, but are translated into architectural constraints that shape how artificial intelligence can be implemented within cyber–physical systems. The technological and architectural challenges facing artificial intelligence adoption in cyber–physical systems within the European Union must therefore be understood as the design-level consequences of the regulatory environment described above. In many technology debates, law is treated as an external constraint applied after engineering decisions have been made. In cyber–physical systems, this sequencing is increasingly untenable. European regulation influences where data may be processed, how systems must be secured, how updates must be managed, how decisions must be documented, and how human oversight must be enabled. Architecture becomes the medium through which compliance is operationalised [5,16,17].

The central challenge is that artificial intelligence-enabled cyber–physical systems must now satisfy several expectations simultaneously. They must be functionally effective, operationally reliable, secure against cyber threats, capable of supporting traceability and auditability, compatible with data governance requirements, and sufficiently understandable to support meaningful human supervision where required. These expectations are individually demanding and collectively transformative. They move artificial intelligence-enabled cyber–physical systems away from opportunistic integration and towards more disciplined, layered, and governable design [3,16,17]. The subsections that follow therefore examine how regulatory obligations are translated into concrete architectural demands.

6.1. Privacy-Preserving Data Architectures

A first major challenge concerns the design of privacy-preserving data architectures. Data protection law does not merely regulate data after collection. It influences what data should be collected in the first place, for what purposes, under what lawful basis, and with what safeguards. In cyber–physical systems, this is difficult because artificial intelligence performance often improves with richer, more granular, and more continuous data streams, while privacy law pushes towards data minimisation, purpose limitation, and controlled retention. The result is an architectural tension between analytical ambition and lawful data governance [3,16,18].

This tension often appears in decisions about local versus central processing. Edge based processing can reduce the need to transfer raw data to central platforms and may therefore support stronger privacy protection and lower exposure of sensitive data. Cloud-based architectures can support richer model training, broader system integration, and cross-site optimisation, but they intensify legal and governance complexity [3,18,27]. Many artificial intelligence-enabled cyber–physical systems therefore move towards hybrid architectures in which some data are processed locally for responsiveness or privacy reasons while selected features, aggregates, or model artefacts are transmitted centrally. Hybridisation, however, increases design complexity and requires careful specification of data boundaries, retention logic, and responsibility allocation.

Model development adds further difficulty. Artificial intelligence in cyber–physical systems often depends on historical operational data for training, validation, and ongoing improvement. Yet data protection rules may constrain how those data are assembled, combined, and reused. This encourages the use of pseudonymisation, aggregation, federated learning, privacy-enhancing computation, or synthetic data. Such approaches may reduce legal exposure, but they may also introduce trade-offs in model accuracy, engineering effort, and maintainability [3,27,42].

6.2. Secure-by-Design and Resilient System Stacks

A second major challenge concerns the requirement for secure-by-design system stacks. Cybersecurity obligations increasingly mean that artificial intelligence functions cannot be treated as model centric subsystems. They must be integrated into secure operational stacks that include device security, communication security, identity and access management, vulnerability handling, software update mechanisms, and incident response integration [6,10,17].

A forecasting model, optimisation engine, or anomaly detection component may be technically strong and still be unsuitable for deployment if it depends on insecure gateways, weak authentication, uncontrolled interfaces, or unsupported edge devices. In cyber–physical systems, where sensors, controllers, and supervisory software interact continuously with physical processes, the weakest connected component can undermine the entire operational arrangement [4,6,44]. The engineering challenge is therefore to design resilient end-to-end stacks rather than isolated artificial intelligence modules [4,6,44].

This becomes especially difficult in legacy environments. Many energy systems, buildings, mobility infrastructures, and industrial facilities still rely on field devices, controllers, and protocols that were not designed for contemporary cybersecurity expectations. Artificial intelligence adoption in such contexts often requires middleware, segmentation strategies, protocol wrappers, or staged retrofit architectures. These compensatory measures can enable transition, but they also increase technical debt and make assurance more difficult [2–4].

6.3. Logging, Traceability, and Auditability as Architectural Requirements

A third challenge concerns logging, traceability, and auditability. Artificial intelligence law, cybersecurity obligations, and data governance expectations all reinforce the need for architectures capable of evidencing system behaviour, access events, model versions, data flows, and incident history [5,17,57]. In cyber–physical systems, logging is no longer a secondary engineering convenience. It is a core architectural function.

Systems must increasingly be able to reconstruct how an artificial intelligence-related output was generated, what data were used, which model version was active, what subsequent human or automated actions followed, and how security-relevant events were recorded. This is difficult because cyber–physical systems are distributed. Data may be generated at the edge, transformed in middleware, enriched in platforms, and consumed by dashboards, supervisory tools, or control support services. If these layers are procured from different vendors or developed independently, traceability fragments quickly [5,17,57].

The technical challenge is therefore to build observability across the full artificial intelligence and cyber–physical chain rather than within isolated components. This requires version control, provenance tracking, time synchronisation, coherent interface design, access records, and event logging that is meaningful for both engineering diagnosis and regulatory accountability. Auditability is also essential for lifecycle management. Without structured logs and deployment histories, it becomes difficult to distinguish between model drift, data quality problems, sensor faults, integration errors, or malicious interference [57–59].

6.4. Human Oversight and the Design of Intervention Points

A fourth challenge concerns the translation of legal expectations about human oversight into operationally credible system design. Human oversight is frequently invoked as a safeguard, yet regulation does not automatically specify how oversight should be implemented in real cyber–physical workflows [28,31]. In practice, oversight requires interfaces, escalation paths, interpretive aids, role allocation, decision thresholds, and clear

intervention rights. It also requires that humans be placed in positions where intervention is meaningful rather than symbolic.

This is especially difficult in cyber–physical systems because timing and complexity matter. In some settings, humans cannot realistically inspect every model output. In others, the artificial intelligence output may only become visible after it has already influenced a physical process or scheduling logic. The design challenge is therefore to define appropriate intervention points rather than merely to require human presence. This may involve confidence thresholds, exception handling rules, staged autonomy, role-specific dashboards, and distinctions between observational, advisory, and operational use [28,31,58].

Cognitive load must also be considered. Cyber–physical environments often generate high volumes of alerts, diagnostics, and system states. Poorly designed oversight interfaces may cause operators either to defer excessively to artificial intelligence or to ignore it because of information overload. Meaningful oversight therefore depends on operational interpretability, timing, and task relevance rather than on generic transparency alone [28,31,57,60].

6.5. Explainability, Transparency, and Operational Interpretability

A fifth challenge concerns explainability and operational interpretability. European regulation creates strong pressure towards meaningful information, documentation, and accountability, especially where systems affect safety, rights, or significant operational outcomes. In cyber–physical systems, artificial intelligence components must often be intelligible enough to support trust, troubleshooting, oversight, and regulatory review [57,59,61,62].

The difficulty is that many high-performing artificial intelligence methods do not naturally align with operational interpretability. In predictive maintenance, demand forecasting, adaptive control support, or mobility optimisation, organisations may face a trade-off between model complexity and governability. Simpler models may be easier to explain and validate but may capture less operational nuance. More complex models may perform better in some conditions but be harder to justify, monitor, or debug [57,59,61]. In cyber–physical systems, where outputs may influence physical processes and interdependent infrastructures, this becomes a major architectural and governance challenge.

Operational interpretability is also role-dependent. An explanation sufficient for data scientists may be useless for building operators, grid controllers, facilities teams, or maintenance engineers. Explainability therefore often needs to be layered. One layer may support technical validation, another operator interaction, and another audit or regulatory review. Designing such layered interpretability requires coordination between model developers, domain experts, interface designers, and compliance stakeholders [3,57,63].

6.6. Data Access, Interoperability, and Federated System Integration

A sixth challenge concerns data access and federated system integration. Data sharing law may support greater access to and reuse of data, but it does not remove the underlying technical difficulty of integrating data across proprietary products, organisational boundaries, and heterogeneous infrastructures. Cyber–physical systems often span devices, vendors, protocols, and institutional actors that were never designed to participate in a shared artificial intelligence ready data environment [5,23,46].

Artificial intelligence-enabled cyber–physical systems therefore require more than raw data access. They require coherent semantics, metadata quality, time alignment, identity management, interface stability, and secure exchange mechanisms. Without these qualities, data may be legally accessible yet operationally unusable [5,23,46]. In multi-actor environments such as smart grids, transport systems, building portfolios, and indus-

trial value chains, the challenge is to move from fragmented data possession towards governed interoperability.

Federated architectures may reduce some privacy and control concerns by avoiding full centralisation, but they increase demands on coordination, standardisation, and assurance. Artificial intelligence in such environments must often operate under partial observability, uneven data quality, asynchronous updates, and differing institutional rules. Interoperability is therefore not a peripheral technical matter. It is a condition for lawful, scalable, and governable deployment [18,23,42].

6.7. Lifecycle Governance, Software Updates, and Continuous Compliance

A seventh challenge concerns lifecycle governance. Artificial intelligence-enabled cyber–physical systems are not static products. They evolve through software updates, model retraining, vulnerability remediation, changing intended uses, and operational adaptation. Regulation increasingly reflects this reality. Cybersecurity law addresses lifecycle security for products with digital elements. Artificial intelligence governance reinforces expectations around post-market monitoring. Together, these developments mean that compliance is no longer a one-time approval event [4,6,64].

This creates a demanding engineering requirement. Organisations must be able to update software and models without undermining safety, interoperability, documentation coherence, or operator trust. In cyber–physical systems, updates may affect devices in the field, control support logic, data schemas, and interfaces simultaneously. Poorly managed updates can break validated assumptions and create operational instability. Continuous compliance therefore becomes a lifecycle discipline requiring change management, observability, testing, rollback capability, and version governance [4,58,64].

Continuous software updates also create a normative tension between adaptability and legal certainty. Cyber–physical systems must be patchable and updatable because insecure or obsolete software can create unacceptable cybersecurity and safety risks. At the same time, updates may alter validated behaviour, modify interfaces, change model performance, or affect assumptions used in conformity assessment and operational approval. The architectural requirement is therefore not simply to allow updates, but to govern them through version control, impact assessment, testing, approval workflows, rollback mechanisms, and updated documentation. In artificial intelligence-enabled cyber–physical systems, continuous compliance depends on treating every material update as part of the regulated lifecycle of the system.

The challenge is especially acute where models are retrained or adapted using live operational data. Such adaptation may improve performance, but it can also complicate reproducibility, accountability, and safety assurance. This creates tension between technical adaptability and regulatory stability. Cyber–physical systems require both, yet achieving both simultaneously demands architectural maturity that many organisations do not yet possess [20,58,64].

6.8. Legacy Integration and the Cost of Regulation-Aware Retrofitting

An eighth challenge concerns legacy integration. Many European cyber–physical systems were not originally designed for artificial intelligence, layered cybersecurity governance, or contemporary data regulation. They have evolved through incremental automation, sector-specific procurement, and vendor-dependent maintenance practices. Artificial intelligence adoption therefore often takes the form of retrofitting rather than greenfield design [3,65,66].

Retrofitting requires organisations to insert artificial intelligence functions into environments that may lack modern telemetry, secure software update channels, standardised

interfaces, or coherent asset inventories. It may also require them to document and validate system behaviour more rigorously than before. The cost of such retrofitting is not purely financial. It also includes organisational learning, procurement redesign, contractual renegotiation, and reallocation of responsibility across legal, information technology, operational technology, security, and operations teams. This is one reason why regulation-aware artificial intelligence adoption progresses unevenly across sectors and organisations [3,51,52].

6.9. Regulation as an Architectural Design Force

Taken together, these challenges show that European regulation operates in cyber–physical systems not only as a compliance framework but also as architectural constraints that shape their design, as shown in Figure 4. Artificial intelligence governance influences logging, documentation, transparency, oversight structures, and post-deployment monitoring. Data protection influences data placement, minimisation, and retention. Cybersecurity law influences device selection, update mechanisms, identity management, and network design. Data access law influences interface design, interoperability strategy, and conditions for lawful reuse. Product and safety law influences how artificial intelligence may be embedded into machinery and other physically consequential systems [35,48,49].

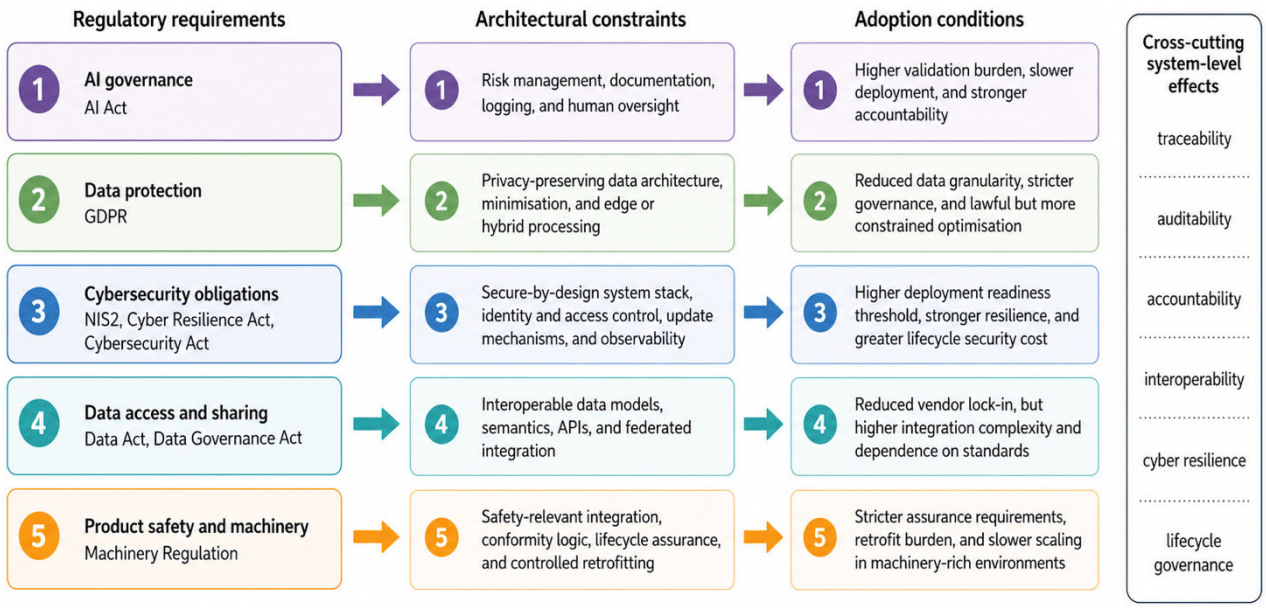


Figure 4. Translation of regulatory requirements into architectural constraints and adoption conditions in cyber–physical systems.

The practical implication is not that regulation simply suppresses innovation. In several respects, it drives more mature engineering by forcing attention to trust, resilience, lifecycle management, and accountability. Yet it also raises the threshold for adoption and places disproportionate pressure on actors that lack the resources to implement regulation-aware architectures [20,51,52]. These design pressures do not manifest uniformly. Their practical significance depends on the type of cyber–physical system into which artificial intelligence is being introduced, the nature of the surrounding infrastructure, and the institutional conditions under which deployment occurs. These differences become clearer through cross-sector comparison.

7. Cross-Sector Domain Analysis

The regulatory and architectural pressures described above do not manifest uniformly across cyber–physical systems, but are translated differently depending on operational

context, infrastructure characteristics, safety exposure, data conditions, and institutional arrangements. A cross-sector analysis is therefore necessary to identify both recurring regulatory logics and differentiated implementation patterns. The four selected domains provide this comparative basis by representing distinct but comparable classes of artificial intelligence-enabled cyber–physical systems. The purpose of the comparison is not merely descriptive. It is to show how the same regulatory environment produces different dominant pressures, architectural priorities, and adoption bottlenecks across domains. The comparison applies the three AI integration layers introduced in Section 3. Observational AI primarily supports monitoring, forecasting, diagnostics, and anomaly detection. Advisory AI supports human or supervisory decision making by recommending actions or prioritising interventions. Operational AI directly influences orchestration, control, automation, or intervention routines. This distinction is important because the governance implications of AI-enabled CPS intensify as AI moves closer to physical effect. Observational AI mainly raises concerns around data governance, privacy, cybersecurity, and transparency. Advisory AI adds stronger concerns around accountability, explanation, human oversight, and decision responsibility. Operational AI raises the highest demands for safety, traceability, resilience, lifecycle assurance, and clearly defined intervention authority.

7.1. Energy and Smart Grids

Energy and smart grids represent one of the clearest cases of artificial intelligence-enabled cyber–physical systems under strong regulatory influence. The sector combines large-scale physical infrastructure, distributed sensing and control, high reliability requirements, and growing dependence on digital coordination as renewable integration, electrification, flexibility services, and distributed energy resources expand. Artificial intelligence is increasingly used in forecasting, congestion management, fault detection, predictive maintenance, distributed resource coordination, market-facing optimisation, and network planning support [1,2,21,23]. This makes the domain especially useful for examining how regulation is translated into infrastructure-level requirements for resilience, traceability, interoperability, and institutional trust.

The observational layer is prominent in load forecasting, fault detection, asset monitoring, and renewable generation prediction. The advisory layer appears where artificial intelligence supports congestion management, flexibility activation, maintenance prioritisation, or operator decision support. The operational layer becomes relevant where artificial intelligence directly influences dispatch, switching support, distributed energy resource coordination, or automated control routines. Regulatory pressure increases substantially across these layers. Observational applications mainly require secure and lawful data pipelines, while advisory and operational applications require stronger traceability, validation, human oversight, and resilience because they may influence infrastructure stability and continuity of supply.

In this domain, the Artificial Intelligence Act matters because artificial intelligence functions can become closely linked to the management and operation of critical infrastructure. The legal importance of risk classification is therefore particularly strong. Grid-related artificial intelligence is often not merely advisory in a generic business sense. It can influence operational decision making in environments where continuity of service, infrastructure resilience, and system stability are essential. Documentation, traceability, human oversight, and risk management obligations therefore have unusual weight [2,17,21].

NIS2 is also highly salient because energy falls directly within the sectors covered by the Directive. Cybersecurity governance is a formal organisational requirement rather than a matter of good practice alone [2,6,43]. Artificial intelligence deployment in this context is therefore inseparable from institutional cybersecurity maturity. Data governance and

interoperability are likewise central because energy-related artificial intelligence often depends on combining data across transmission, distribution, distributed resources, weather systems, prosumers, and market platforms. The sector has historically been fragmented across operators, voltage levels, technical standards, and national practices. This makes energy an especially clear example of how data access law and standardisation ambitions may function as long-term enablers while leaving substantial short-term implementation challenges unresolved [1,2,23].

A distinctive feature of the energy domain is that the main tension is often not privacy alone, although personal data may matter in demand-side applications. The more dominant tension is between innovation and infrastructural trustworthiness. Operators must justify artificial intelligence adoption in environments where failure can have cascading effects. Explainability, resilience, validation discipline, and operational accountability therefore carry unusual weight [2,17,21].

7.2. Smart Buildings and Building Automation

Smart buildings and building automation represent a different but equally important class of artificial intelligence-enabled cyber–physical systems. Artificial intelligence is used in heating, ventilation, air conditioning, indoor climate optimisation, occupancy-aware control, predictive maintenance, load shifting, anomaly detection, and integration with wider energy management functions [3,47,65]. The promise of improved energy efficiency, comfort, and flexibility is closely tied to sensing density and digital control capability. This domain is especially useful for examining how privacy-sensitive sensing, fragmented installed bases, and retrofit constraints reshape the practical conditions of adoption.

The observational layer is visible in occupancy estimation, anomaly detection, equipment monitoring, indoor environmental analytics, and energy consumption forecasting. The advisory layer appears when artificial intelligence recommends set point changes, maintenance actions, demand response participation, or comfort–energy trade-offs to facilities teams. The operational layer emerges when artificial intelligence directly adjusts HVAC, lighting, storage, or load shifting strategies. The regulatory profile changes across these layers. Observational applications are dominated by privacy, transparency, and lawful data processing, while advisory and operational applications raise stronger accountability, safety, cybersecurity, and traceability demands because artificial intelligence outputs may affect occupants, indoor climate, energy flexibility, and building service continuity.

The regulatory profile of smart buildings differs from that of smart grids because privacy and personal data governance are often more immediate. Buildings generate operational data about presence, routines, indoor environmental responses, access patterns, and behavioural regularities. Even where the intended purpose is technical optimisation, these data can become personal or at least privacy-sensitive when linked to identifiable occupants or user groups. Data protection therefore becomes especially salient. The main architectural consequence is that building-level artificial intelligence often requires stronger attention to lawful basis, minimisation, retention boundaries, and privacy-preserving data handling than system-level infrastructure applications might require [3,16,41].

At the same time, smart buildings are affected by cybersecurity obligations in two ways. First, building systems increasingly depend on products with digital elements such as gateways, sensors, controllers, and connected management software, which makes the Cyber Resilience Act highly relevant for the product layer. Second, building portfolios associated with essential or important entities may fall under NIS2 obligations. Artificial intelligence adoption in building automation is therefore shaped by a combination of privacy law, product cybersecurity, and in some cases organisational cybersecurity governance [3,6,47].

Another distinctive feature of the building domain is strong dependence on retrofit realities. Many building environments were not designed as digital-first cyber–physical systems. They often combine modern artificial intelligence-capable platforms with older building management systems, vendor-specific subsystems, limited semantic consistency, and constrained maintenance resources. Even where the artificial intelligence use case is compelling, legacy integration, fragmented data models, and unclear governance over vendors and facility management functions can slow deployment considerably [3,65,67]. This makes smart buildings a particularly clear example of regulation interacting with installed base complexity.

7.3. Mobility and Transport Systems

Mobility and transport systems represent a third major cyber–physical domain in which artificial intelligence adoption is strongly shaped by regulation. These systems include traffic management platforms, connected and automated mobility services, fleet coordination, infrastructure monitoring, multimodal optimisation, and safety related assistance or control functions [45,68–71]. They combine real-time coordination, distributed data sources, public space operation, and direct interaction with human movement and safety. This makes them especially useful for examining the combined importance of safety, public legitimacy, time-critical coordination, and cross-organisational responsibility.

The observational layer includes traffic monitoring, infrastructure diagnostics, risk detection, passenger flow analysis, and fleet state estimation. The advisory layer includes routing recommendations, incident response support, driver assistance, maintenance prioritisation, and operational planning. The operational layer appears where artificial intelligence affects traffic signalling, automated vehicle behaviour, fleet orchestration, or safety-relevant transport control. The movement from observation to operation is particularly consequential in mobility because timing, public space governance, safety, and accountability interact directly. Operational artificial intelligence in transport therefore requires stronger evidence of safety, cybersecurity, auditability, and human intervention capability than monitoring or forecasting applications.

The defining regulatory characteristic of mobility-related artificial intelligence is the combination of safety relevance and public space governance. Transport-related artificial intelligence may intersect with the Artificial Intelligence Act where systems are used in contexts linked to critical infrastructure, safety components, or other high-impact operational settings. It may also intersect strongly with data protection law where location data, video feeds, behavioural traces, or identifiable travel patterns are involved. In public space contexts, mobility systems can further raise concerns about surveillance, proportionality, transparency, and acceptable forms of automated analysis [18,41,63].

From an architectural perspective, transport systems are unusually demanding because timing and coordination requirements are often strict. Artificial intelligence outputs may only have value if they can be incorporated rapidly into signalling, routing, scheduling, fleet orchestration, or hazard response. Yet legal expectations around auditability, oversight, and secure operation still apply. Mobility therefore illustrates a classic cyber–physical tension between responsiveness and governability. A transport system that is too slow, opaque, or difficult to validate may be unacceptable operationally. A system that is too weakly documented, privacy intrusive, or insecure may be unacceptable legally or politically [41,54,57].

Mobility systems are also characterised by high interdependence. Vehicles, roadside infrastructure, municipal systems, platform providers, communications networks, and public authorities may all be involved in a single artificial intelligence-enabled service. Responsibility allocation is therefore more complex than in single-owner industrial

settings. Interoperability and trusted data exchange become central, which makes mobility a key domain for understanding the practical significance of data sharing law and cross-organisational digital coordination [18,45,68].

7.4. Industrial and Manufacturing Systems

Industrial and manufacturing systems form a fourth representative domain and highlight another distinctive variant of artificial intelligence-enabled cyber–physical governance. Artificial intelligence is used here for predictive maintenance, quality control, process optimisation, energy management, anomaly detection, scheduling, machine coordination, and increasingly the support of autonomous or semi-autonomous industrial operations [44,72–75]. The core cyber–physical feature is the tight coupling between software-based intelligence and machinery, production lines, robotics, industrial networks, and safety-relevant operational processes. This makes the domain especially useful for examining how product safety, lifecycle assurance, operational technology constraints, and retrofitting pressures shape adoption.

The observational layer includes condition monitoring, anomaly detection, quality inspection, energy performance monitoring, and predictive maintenance. The advisory layer appears when artificial intelligence recommends maintenance interventions, production rescheduling, process adjustments, or safety actions. The operational layer emerges where artificial intelligence directly affects machine coordination, robotic behaviour, production control, or safety-relevant automation. This progression is especially important in manufacturing because operational artificial intelligence may interact directly with machinery, workers, and production continuity. As artificial intelligence moves closer to machine control, product safety, lifecycle assurance, cybersecurity, update governance, and responsibility allocation become more demanding.

The regulatory character of industrial systems differs from that of buildings and mobility because product safety and machinery governance become especially prominent. Where artificial intelligence is embedded in or materially influences machinery and related products, the Machinery Regulation becomes directly relevant alongside horizontal artificial intelligence and cybersecurity rules [35,48,49,74]. The key challenge is therefore not only whether the artificial intelligence works analytically, but whether its integration into machinery-rich environments remains safe, controllable, and lawfully documented.

Industrial environments are also strongly affected by the Cyber Resilience Act because they depend on products with digital elements throughout the production environment, including sensors, controllers, industrial gateways, embedded software, and connected supervisory systems. In many cases, NIS2 is also relevant where operators fall within covered sectors or categories. Industrial artificial intelligence adoption is therefore shaped by a dual expectation of secure products and secure organisations [4,6,44]. The technical threshold for deployment is raised because artificial intelligence functions must be integrated into operational technology environments with long lifecycles, strong uptime requirements, and limited tolerance for experimentation.

A distinctive feature of the industrial domain is the weight of legacy and installed base dependence. Manufacturing environments often contain long-lived machinery, proprietary ecosystems, and process-specific interfaces. Artificial intelligence adoption may therefore require extensive retrofit engineering, middleware layers, staged deployment logic, and careful segregation between new digital intelligence and older operational assets. This makes industrial systems a particularly strong example of regulation-aware retrofitting [4,52,73].

7.5. Comparative Synthesis Across the Four Domains

Comparison across the four domains reveals both common pressures and clear differences. In all four cases, artificial intelligence adoption is shaped by the Artificial Intelligence Act, cybersecurity obligations, data governance requirements, and the need for stronger documentation, logging, and human oversight [2–5]. Also in all four cases, architecture becomes the point at which regulation is translated into operational form. The major difference lies in which regulatory pressure becomes dominant and how that pressure is translated into the primary adoption bottleneck. Table 2 consolidates this comparison by applying the same analytical categories across the four domains.

Table 2. Cross-sector comparison of regulatory pressures, AI integration layers, architectural consequences, and adoption bottlenecks in AI-enabled cyber–physical systems.

Domain	Dominant AI Integration Layers	Main Regulatory Instruments	Dominant Evaluation Criteria	Architectural Consequences	Main Adoption Bottleneck
Energy and smart grids	Observational and advisory, with emerging operational use in flexibility, dispatch support, and distributed resource coordination	AI Act, NIS2, Cyber Resilience Act, Data Act, Cybersecurity Act, sector energy rules	Safety, cybersecurity, traceability, accountability, interoperability	Secure and traceable software stacks, resilient data pipelines, operator oversight, interoperability across grid actors	High reliability expectations, cybersecurity maturity, legacy infrastructures, and institutional risk aversion
Smart buildings and building automation	Observational and advisory are common, operational use appears in HVAC, lighting, demand response, and energy management control	GDPR, AI Act, Cyber Resilience Act, Data Act, NIS2 where applicable	Data governance, transparency, cybersecurity, adaptability, accountability	Privacy-preserving sensing, edge–cloud partitioning, retrofit-compatible integration, secure connected products	Privacy-sensitive data, fragmented building systems, vendor lock-in, limited facilities governance capacity
Mobility and transport systems	Observational and advisory are widespread, operational use appears in traffic management, fleet orchestration, and automated mobility functions	AI Act, GDPR, NIS2, Cyber Resilience Act, Data Act, sector transport safety rules	Safety, transparency, accountability, cybersecurity, traceability	Real-time auditability, explainable decision support, secure data exchange, intervention points	Balancing responsiveness, safety, public legitimacy, privacy, and cross-organisational responsibility
Industrial and manufacturing systems	Observational and advisory are mature, operational use appears in robotics, process control, scheduling, and safety related automation	AI Act, Cyber Resilience Act, Machinery Regulation, NIS2 where applicable, Cybersecurity Act	Safety, cybersecurity, adaptability, accountability, traceability	Secure OT integration, lifecycle update governance, machinery safety validation, staged retrofitting	Long-lived machinery, proprietary OT systems, safety assurance, update control, compliance cost

The comparison shows that the same European regulatory landscape does not produce uniform adoption effects. Instead, each domain translates the shared regulatory framework into a different combination of dominant pressure, architectural requirement, and adoption bottleneck.

In energy and smart grids, the dominant pressure arises from critical infrastructure resilience, cybersecurity governance, and interoperability across distributed operational environments. The main architectural consequence is the need for highly reliable, traceable, and securely integrated software stacks capable of supporting infrastructure-level coordination. The corresponding adoption bottleneck lies in the difficulty of introducing artificial intelligence into environments where reliability, resilience, and institutional trust are non-negotiable [2,21,23].

In smart buildings and building automation, the dominant pressure arises from privacy-sensitive sensing environments combined with fragmented installed base conditions. The main architectural consequence is the need for privacy-preserving data handling,

stronger data governance boundaries, and retrofit-compatible integration across heterogeneous building management systems. The corresponding adoption bottleneck lies in the tension between data-rich optimisation ambitions and the practical realities of lawful data use, legacy subsystems, and constrained facilities governance [3,41,65].

In mobility and transport systems, the dominant pressure arises from the combination of safety relevance, public space operation, and strong demands for timely coordination. The main architectural consequence is the need for systems that remain both operationally responsive and sufficiently auditable, secure, and governable under public and institutional scrutiny. The corresponding adoption bottleneck lies in balancing speed, coordination, oversight, and legitimacy in environments where artificial intelligence may directly affect movement, safety, and public trust [54,57,69,70].

In industrial and manufacturing systems, the dominant pressure arises from machinery safety, product cybersecurity, lifecycle assurance, and dependence on long-lived operational technology environments. The main architectural consequence is the need for carefully controlled integration of artificial intelligence into machinery-rich settings through secure products, staged retrofit logic, and robust lifecycle governance. The corresponding adoption bottleneck lies in reconciling innovation with safety assurance, legacy operational technology constraints, and the technical cost of regulation-aware retrofitting [4,49,73,74].

The cross-sector pattern is therefore not one of uniform legal effect but of differentiated regulatory translation. The same horizontal legal instruments shape artificial intelligence adoption differently depending on whether the core cyber–physical challenge is infrastructure coordination, privacy-sensitive optimisation, public movement under time-critical conditions, or machinery-linked process control [4,5]. This comparative pattern provides the basis for the broader governance, innovation, and competitiveness implications considered in the next section. Figure 5 provides a visual synthesis of these domain-specific regulatory translation patterns.

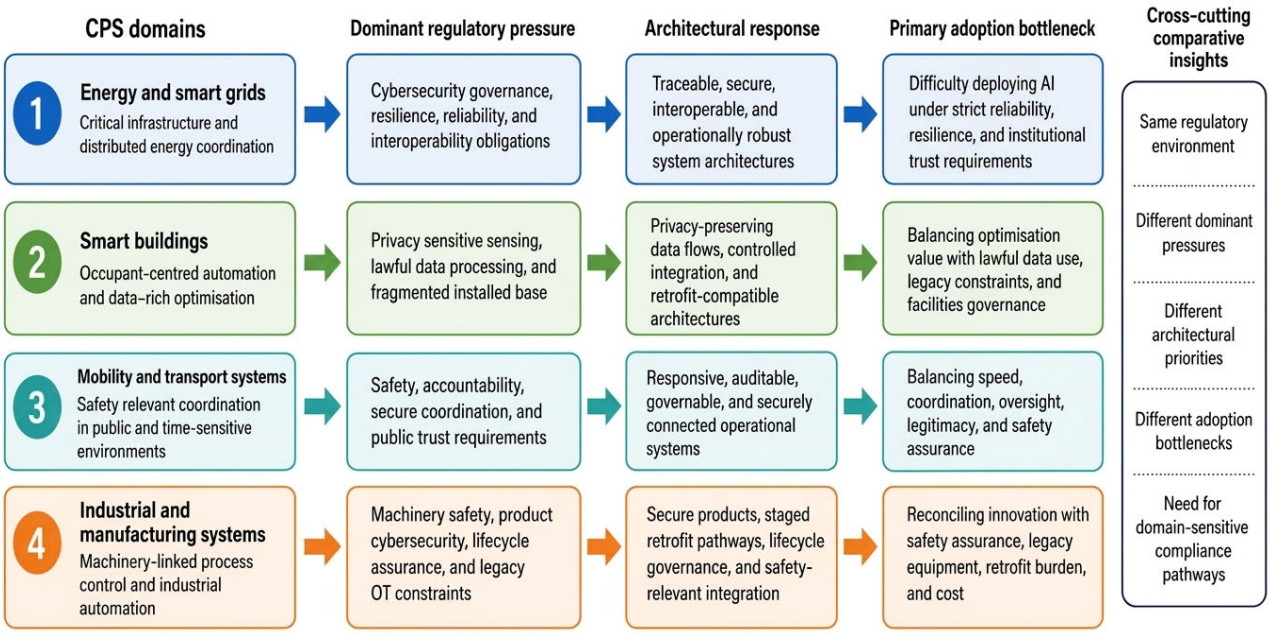


Figure 5. Cross-sector differentiation of regulatory pressures, architectural responses, and adoption bottlenecks.

8. Cross-Sector Implications for Governance, Innovation, and Competitiveness

The cross-sector comparison shows that the impact of European Union regulation on artificial intelligence adoption in cyber–physical systems cannot be reduced to a simple binary of constraint versus enablement. Across energy, buildings, mobility, and industrial systems, regulation functions as a structural force that reshapes governance arrangements, innovation pathways, procurement logic, and competitive positioning. The most important cross-sector implication is that artificial intelligence adoption increasingly depends on the capacity of organisations to align legal interpretation, technical architecture, cybersecurity maturity, and operational responsibility within a coherent implementation model [4,5,20,22]. The implications discussed below arise directly from the differentiated regulatory translation observed across the four domains.

A first implication concerns governance capacity. The domain comparison shows that compliance demands are not absorbed equally across actors. Large utilities, industrial firms, infrastructure operators, and well-resourced technology vendors are generally better positioned to absorb the burdens associated with artificial intelligence governance, cybersecurity, product security, and data management. They are more likely to possess internal legal teams, cybersecurity functions, systems engineering expertise, and structured procurement processes [20,22,52]. Smaller firms, municipalities, facilities teams, and specialist integrators often operate with far more limited capacity [38,51,52]. The result is not merely uneven compliance readiness, but uneven adoptability. Artificial intelligence becomes easier to deploy where governance capability already exists [20,22,52].

A second implication concerns innovation pathways. The domain analysis shows that European cyber–physical innovation is unlikely to follow the model of rapid, lightly governed experimentation that characterises some less regulated digital environments [76–78]. Documentation, traceability, secure integration, lawful data governance, and meaningful human oversight push innovation towards slower and more structured development trajectories. Innovation becomes more architecture intensive, more compliance aware, and more dependent on coordination between engineering, legal, cybersecurity, and operations functions [16,33,77]. In this context, successful innovation depends not only on model quality but on the ability to engineer governable deployment environments [16,17,22].

A third implication concerns procurement and market structure. As the four domains illustrate, regulation increasingly makes documentation quality, interoperability, secure lifecycle management, and assurance capability part of the practical value proposition of artificial intelligence-enabled systems. Procurement decisions are therefore likely to place greater weight on certification readiness, product security, observability, updateability, and evidence of governance maturity [6,17,47]. This creates opportunities for vendors capable of demonstrating regulation-aware design, but it also risks reinforcing market concentration if only a limited set of suppliers can bear the fixed cost of compliance [22,51,52].

A fourth implication concerns trust and legitimacy. In cyber–physical systems, adoption depends not only on technical efficiency but also on confidence that artificial intelligence-enabled operation is safe, resilient, accountable, and aligned with public and institutional expectations [16,33,34]. This is especially visible in mobility and building contexts, where artificial intelligence may be proximate to users, occupants, or citizens, but it is equally relevant in energy and industry, where failure can produce broader infrastructural and economic consequences [2,4,26]. The European approach seeks to institutionalise such trust through legal requirements rather than through voluntary self-governance alone. This may slow short-term deployment, but it may also create more durable conditions for acceptance in critical and semi-critical environments [16,33,34].

A fifth implication concerns competitiveness at the European level. There is a genuine risk that cumulative compliance obligations may slow the pace of artificial intelligence adoption relative to jurisdictions that permit faster deployment with fewer ex ante controls [20,76,77]. This concern is particularly relevant for smaller firms facing high fixed compliance costs [21,51,52]. Yet the cross-sector analysis also points to an opposing possibility. If Europe succeeds in aligning artificial intelligence governance with robust architectures, secure products, interoperable data environments, and credible assurance mechanisms, it may establish a distinctive competitive advantage in trustworthy cyber–physical artificial intelligence [5,23,33]. Such an advantage would rest not on speed alone, but on reliability, safety, compliance readiness, and suitability for regulated and physically consequential operational settings.

Taken together, these implications suggest that European Union regulation is gradually redefining the basis on which artificial intelligence becomes adoptable in cyber–physical systems. Technical performance remains important, but it is no longer sufficient. Organisational maturity, cybersecurity discipline, documentation capability, lifecycle governance, and interoperability strategy are becoming equally central. The adoption of artificial intelligence in cyber–physical systems is therefore shifting from a model-centred challenge to an institutional and architectural capability challenge [5,16,20]. This shift also frames the broader interpretive discussion that follows, particularly the question of how the European regulatory model reshapes the relationship between innovation, trust, and operational governability.

9. Policy and Research Recommendations

The analysis points to a need for more coordinated and implementation-oriented action if the European Union is to support both trustworthy governance and viable adoption of artificial intelligence in cyber–physical systems. The challenge is not to weaken regulatory ambition. The challenge is to make that ambition more operationally actionable across diverse sectors and organisational contexts.

A first recommendation concerns the development of sector-specific implementation guidance for high-impact artificial intelligence-enabled cyber–physical systems. The Artificial Intelligence Act establishes a common horizontal framework, but the practical meaning of documentation, oversight, risk management, traceability, and post-market monitoring differs substantially between smart grids, buildings, mobility systems, and machinery-rich industrial environments. More detailed domain-specific guidance would reduce interpretive ambiguity and make compliance more predictable for operators, vendors, and public authorities [2,22,53].

A second recommendation concerns stronger alignment between horizontal regulation and technical reference architectures. Many organisations understand the general legal direction but lack concrete engineering patterns for compliant implementation. European institutions, standards bodies, agencies, and sectoral consortia should therefore support the development of reference architectures for regulation-aware artificial intelligence-enabled cyber–physical systems. Such architectures should address data segmentation, identity and access management, secure update pathways, logging and provenance, human oversight interfaces, and interoperability design. They would not replace legal interpretation, but they would provide much-needed implementation scaffolding [17,25,27].

A third recommendation concerns targeted support for small and medium-sized enterprises, municipalities, and smaller infrastructure operators. These actors often lack the internal capacity to interpret and operationalise complex regulatory requirements even when they possess strong domain knowledge and valuable use cases. Support mechanisms could include documentation templates, compliance toolkits, access to expert

advisory structures, shared testing and validation environments, and where appropriate, financial support for conformity-related processes. Without such support, the regulatory environment risks reinforcing market concentration and slowing socially valuable adoption in less resourced settings [38,51,52].

A fourth recommendation concerns wider use of regulatory sandboxes and supervised experimentation environments for cyber–physical artificial intelligence. Because many domains combine safety, infrastructure, and public interest considerations, real-world learning is essential but difficult. Structured sandboxes can make it possible to test regulation-aware artificial intelligence under controlled conditions while preserving accountability and oversight. They are particularly valuable in cross-domain settings where systems span energy, buildings, transport, and industrial operations. Sandboxes should not function as zones of regulatory suspension, but as instruments for governance informed innovation [33,77,78].

A fifth recommendation concerns stronger integration of cybersecurity by design and lifecycle security into public funding, procurement, and digital transition programmes. The Cyber Resilience Act and NIS2 both reinforce the principle that artificial intelligence adoption in cyber–physical systems cannot rest on insecure digital foundations. Public support programmes should therefore prioritise architectures that are secure, updateable, observable, and maintainable over time rather than rewarding deployment speed alone [4,6,43].

A sixth recommendation concerns interoperability and data governance. The Data Act and the Data Governance Act create an important legal basis for fairer data access and more trusted sharing, but technical implementation remains uneven. Greater effort is needed to develop common semantic models, secure exchange mechanisms, application programming interfaces, digital identity schemes, and sector-compatible federated data infrastructures that can support artificial intelligence-enabled cyber–physical coordination [18,42,46].

A final recommendation concerns research. Future work should move beyond isolated legal commentary and isolated technical benchmarking and instead focus on regulation-aware systems engineering. There is a need for empirical studies examining how organisations actually integrate artificial intelligence governance, cybersecurity obligations, and data rules into cyber–physical design and operation. There is also a need for comparative work on domain-specific implementation pathways, on meaningful human oversight under real operational conditions, and on the long-term cost implications of compliant versus weakly governed deployment. Research linking law, architecture, and operational performance will be especially valuable [5,16,22].

10. Discussion

The preceding analysis shows that the adoption of artificial intelligence in cyber–physical systems within the European Union is undergoing a structural shift from a predominantly technology-driven process to a regulation-shaped systems engineering challenge. In earlier phases of digitalisation, the central questions often concerned technical feasibility, data availability, and the expected value of prediction, optimisation, or automation [1,79,80]. These questions remain important, but they are no longer sufficient. Artificial intelligence adoption in cyber–physical systems is now also determined by legal architecture, organisational accountability, product security, and institutional trust. The central analytical implication is that regulation increasingly shapes not only what organisations are permitted to do, but what kinds of systems they are capable of building, integrating, and operating responsibly [5,16,20].

This helps explain why sector-specific studies, although valuable, often remain incomplete when treated in isolation. Smart grids, buildings, mobility systems, and industrial

environments each present distinctive governance problems, yet the cross-sector analysis reveals recurring patterns. Across all four domains, legal requirements are translated into demands for stronger architecture, stronger documentation, more robust cybersecurity, clearer human role allocation, and more disciplined lifecycle management [2–5]. What differs is the dominant point of translation. In some domains, privacy is more central. In others, infrastructure resilience, public legitimacy, or product safety becomes more decisive. The cyber–physical perspective is analytically powerful because it reveals both the shared regulatory logic and the differentiated operational consequences [3–5].

The European approach also becomes clearer when contrasted with alternative governance trajectories. In the United States, artificial intelligence governance is often more market-oriented and sector-specific, with greater reliance on flexible ex post oversight than on horizontally integrated ex ante requirements [81]. This can favour rapid experimentation, but may also produce fragmented governance conditions across domains [76–78]. In China, artificial intelligence development is shaped more strongly by state-directed coordination, centralised data access, and industrial policy-driven deployment [72,76,81]. This can support rapid implementation at scale, but places less emphasis on decentralised accountability, transparency, and fundamental rights [34,72,76]. The European Union follows a distinct path in which horizontal legal frameworks prioritise safety, accountability, trust, and the protection of rights [16,33,34]. The significance of this path for cyber–physical systems lies not merely in the existence of regulation, but in the way regulation directly shapes system design, integration, and operational governance [5,16,17].

A further implication concerns the relationship between regulation and innovation. Public debate often frames these as opposing forces, yet the present analysis suggests a more complex relationship. European regulation unquestionably raises short-term barriers. It increases transaction costs, slows some forms of experimentation, and advantages actors with stronger compliance capacity [22,51,52]. At the same time, it also pushes artificial intelligence adoption towards more mature, robust, and governable forms [17,33,77]. In cyber–physical systems, where digital error can become physical harm, service disruption, or security incident, this is not only a legal preference but an operational necessity [4,29,30]. The decisive question is therefore not whether regulation should matter, but whether its implementation burden can be made sufficiently predictable, proportionate, and actionable to support beneficial deployment [22,33,77].

The analysis also highlights the growing importance of organisational capability. Compliance cannot be treated as a downstream task delegated to legal functions after technical design has been completed. It is an institutional capability spanning legal interpretation, systems engineering, cybersecurity management, procurement, documentation, and operations [20,22,38]. Many barriers to adoption are therefore not failures of model quality, but failures of institutional integration. Artificial intelligence adoption in cyber–physical systems increasingly depends on the ability to establish cross-functional governance structures rather than on isolated technical expertise alone [16,20,38].

The dynamic character of the European regulatory landscape also creates a practical requirement for adaptable architectures. AI-enabled cyber–physical systems should not be designed only for the minimum obligations known at the time of deployment, because legal guidance, harmonised standards, delegated acts, cybersecurity certification schemes, and sector-specific implementation practices will continue to evolve. Adaptability requires modular architectures in which data governance, model management, cybersecurity controls, logging, human oversight, and update mechanisms can be modified without redesigning the entire system. It also requires explicit version governance for models, datasets, software components, interfaces, and documentation. From a systems engineering perspective, future-ready compliance is therefore supported by separation of concerns, standardised

interfaces, machine-readable logs, configurable access control, secure update pathways, rollback capability, and traceable change management. Such architectural properties reduce the risk that regulatory or standardisation changes will require disruptive retrofitting after deployment.

At the same time, the current regulatory environment still exhibits important limitations. Regulatory layering creates ambiguity, implementation pathways remain underspecified, and standardisation has not yet fully aligned with legal ambition [5,6,25]. Capacity asymmetries risk producing uneven adoption patterns that favour large actors and well-resourced vendors [20,51,52]. If these issues remain unresolved, they may slow beneficial deployment and reinforce market concentration. The central challenge is therefore not the existence of regulation, but the effectiveness with which it is translated into operational practice across different cyber–physical domains [5,16,22].

11. Conclusions

Artificial intelligence is emerging as a central capability in cyber–physical systems across energy, buildings, mobility, and industrial environments, where its value lies in enabling prediction, optimisation, anomaly detection, adaptive coordination, and more informed operational decision making. Within the European Union, however, its adoption is shaped by a multi-layered regulatory environment that extends well beyond the Artificial Intelligence Act alone. Data protection law, cybersecurity governance, product security rules, data access frameworks, and sector or product-specific safety obligations collectively define the conditions under which artificial intelligence can be integrated into operational cyber–physical systems.

The analysis shows that the main barriers to adoption arise not from outright legal prohibition, but from cumulative compliance friction. High-risk governance obligations, lawful data processing constraints, cybersecurity and lifecycle security requirements, fragmented regulatory layering, uncertainty around responsibility and human oversight, and underdeveloped interoperability pathways all raise the threshold for deployment. These barriers are especially consequential in cyber–physical systems because they affect environments in which digital outputs influence physical processes, infrastructures, or machinery.

At the same time, the analysis shows that regulation also functions as a long-term enabling condition. It promotes secure-by-design products, stronger data governance, clearer accountability, more disciplined lifecycle management, and greater institutional trust in artificial intelligence-enabled operation. Regulation therefore does not simply slow adoption. It reshapes the conditions under which adoption becomes possible. Artificial intelligence becomes operationally viable not when it merely demonstrates technical performance, but when it can be embedded lawfully, securely, and governably within cyber–physical systems.

The central conclusion is that European Union regulation has become constitutive of artificial intelligence adoption in cyber–physical systems. It defines the practical design space within which such systems can be developed, integrated, validated, and scaled. The systems most likely to succeed in the European context will be those designed from the outset around privacy-aware data architectures, secure digital components, traceable and auditable artificial intelligence functions, credible human oversight mechanisms, interoperable interfaces, and clearly assigned organisational responsibilities.

The future competitiveness of European cyber–physical artificial intelligence will depend on whether regulatory discipline can be translated into engineering strength. If legal ambition is matched by stronger implementation guidance, accessible compliance support, interoperable data infrastructures, and robust cross-sector reference architectures, the European Union may establish a distinctive advantage in trustworthy artificial intelligence

for critical and physically consequential systems. If not, fragmented implementation and high transition costs may slow adoption and weaken innovation capacity. The decisive task is therefore not deregulation, but more effective alignment between law, architecture, and operational practice.

Author Contributions: Conceptualisation, B.N.J.; methodology, Z.G.M.; validation, Z.G.M. and B.N.J.; formal analysis, B.N.J. and Z.G.M.; investigation, B.N.J. and Z.G.M.; resources, B.N.J.; data curation, Z.G.M.; writing—original draft preparation, B.N.J.; writing—review and editing, Z.G.M. and B.N.J.; visualisation, B.N.J. All authors have read and agreed to the published version of the manuscript.

Funding: This research received no external funding.

Data Availability Statement: No new data were created or analysed in this study. Data sharing is not applicable to this article.

Conflicts of Interest: The authors declare no conflicts of interest.

Abbreviations

The following abbreviations are used in this manuscript:

AI	Artificial Intelligence
AI Act	Artificial Intelligence Act
API	Application Programming Interface
BEMS	Building Energy Management System
CPS	Cyber–Physical System
CRA	Cyber Resilience Act
DGA	Data Governance Act
DER	Distributed Energy Resource
ENISA	European Union Agency for Cybersecurity
EU	European Union
EUR Lex	European Union Legal Database
GDPR	General Data Protection Regulation
HVAC	Heating, Ventilation, and Air Conditioning
ICT	Information and Communication Technology
IoT	Internet of Things
IT	Information Technology
ML	Machine Learning
NIS2	Network and Information Security Directive 2
OT	Operational Technology
PRISMA ScR	Preferred Reporting Items for Systematic Reviews and Meta-Analyses Extension for Scoping Reviews
SME	Small and Medium-sized Enterprise
TSO	Transmission System Operator
DSO	Distribution System Operator

References

1. Zahid, H.; Zulfiqar, A.; Adnan, M.; Iqbal, S.; Mohamed, S.E.G. A review on socio-technical transition pathway to European super smart grid: Trends, challenges and way forward via enabling technologies. *Results Eng.* **2025**, *25*, 104155. [CrossRef]
2. Jorgensen, B.N.; Gunasekaran, S.S.; Ma, Z.G. Impact of EU Laws on AI Adoption in Smart Grids: A Review of Regulatory Barriers, Technological Challenges, and Stakeholder Benefits. *Energies* **2025**, *18*, 3002. [CrossRef]
3. Jorgensen, B.N.; Ma, Z.G. Impact of EU Laws on the Adoption of AI and IoT in Advanced Building Energy Management Systems: A Review of Regulatory Barriers, Technological Challenges, and Economic Opportunities. *Buildings* **2025**, *15*, 2160. [CrossRef]
4. Ndiaye, N.G.; Waedt, K.; Dejon, N.; Gaber, C.; Marinakis, A.; Gizelis, C.A.; Gur, G.; Rennhard, M.; Zeddini, O.; Wary, J.-P.; et al. *Safety and Cybersecurity Under Emerging EU Legislation for Industry: A Use-Case Driven Perspective*; Springer Nature: Cham, Switzerland, 2026; pp. 304–321.

5. Jorgensen, B.N.; Ma, Z.G. Digital Twins Under EU Law: A Unified Compliance Framework Across Smart Cities, Industry, Transportation, and Energy Systems. *Electronics* **2025**, *14*, 4881. [CrossRef]
6. Metallinou, A.; Papadimitriou, A.; Zatti, S.; Giannopapa, C. An Evolving Regulatory Framework in Europe in the Field of Cybersecurity. In *Proceedings of the 58th IAA Symposium on Safety, Quality and Knowledge Management in Space Activities*; International Astronautical Federation (IAF): Paris, France, 2025; pp. 415–437. [CrossRef]
7. European Parliament. Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 Laying down Harmonised Rules on Artificial Intelligence and Amending Regulations (EC) No 300/2008, (EU) No 167/2013, (EU) No 168/2013, (EU) 2018/858, (EU) 2018/1139 and (EU) 2019/2144 and Directives 2014/90/EU, (EU) 2016/797 and (EU) 2020/1828 (Artificial Intelligence Act). Available online: http://data.europa.eu/eli/reg/2024/1689/oj (accessed on 26 February 2026).
8. European Parliament. Regulation (EU) 2016/679 of the European Parliament and of the Council of 27 April 2016 on the Protection of Natural Persons with Regard to the Processing of Personal Data and on the Free Movement of Such Data, and Repealing Directive 95/46/EC (General Data Protection Regulation). Available online: http://data.europa.eu/eli/reg/2016/679/oj (accessed on 26 February 2026).
9. European Parliament. Directive (EU) 2022/2555 of the European Parliament and of the Council of 14 December 2022 on Measures for a High Common Level of Cybersecurity Across the Union, Amending Regulation (EU) No 910/2014 and Directive (EU) 2018/1972, and Repealing Directive (EU) 2016/1148 (NIS 2 Directive). Available online: http://data.europa.eu/eli/dir/2022/2555/oj (accessed on 26 February 2026).
10. European Parliament. Regulation (EU) 2024/2847 of the European Parliament and of the Council of 23 October 2024 on Horizontal Cybersecurity Requirements for Products with Digital Elements and Amending Regulations (EU) No 168/2013 and (EU) 2019/1020 and Directive (EU) 2020/1828 (Cyber Resilience Act). Available online: http://data.europa.eu/eli/reg/2024/2847/oj (accessed on 26 February 2026).
11. European Parliament. Regulation (EU) 2023/2854 of the European Parliament and of the Council of 13 December 2023 on Harmonised Rules on Fair Access to and Use of Data and Amending Regulation (EU) 2017/2394 and Directive (EU) 2020/1828 (Data Act). Available online: http://data.europa.eu/eli/reg/2023/2854/oj (accessed on 26 February 2026).
12. European Parliament. Regulation (EU) 2022/868 of the European Parliament and of the Council of 30 May 2022 on European Data Governance and Amending Regulation (EU) 2018/1724 (Data Governance Act). Available online: http://data.europa.eu/eli/reg/2022/868/oj (accessed on 26 February 2026).
13. European Parliament. Regulation (EU) 2019/881 of the European Parliament and of the Council of 17 April 2019 on ENISA (the European Union Agency for Cybersecurity) and on Information and Communications Technology Cybersecurity Certification and Repealing Regulation (EU) No 526/2013 (Cybersecurity Act). Available online: http://data.europa.eu/eli/reg/2019/881/oj (accessed on 26 February 2026).
14. European Parliament. Regulation (EU) 2023/1230 of the European Parliament and of the Council of 14 June 2023 on Machinery and Repealing Directive 2006/42/EC of the European Parliament and of the Council and Council Directive 73/361/EEC. Available online: http://data.europa.eu/eli/reg/2023/1230/oj (accessed on 26 February 2026).
15. Jørgensen, B.N.; Ma, Z.G. Impact of EU Regulations on AI Adoption in Smart City Solutions: A Review of Regulatory Barriers, Technological Challenges, and Societal Benefits. *Information* **2025**, *16*, 568. [CrossRef]
16. Horzyk, A.M. Implications of European By-Design Laws on Data-Driven System Design. In *Proceedings of the 2024 International Joint Conference on Neural Networks (IJCNN)*; IEEE: Piscataway, NJ, USA, 2024; pp. 1–8.
17. Vereno, D.; Polanec, K.; Neureiter, C. Compliance by Design for Cyber-Physical Energy Systems: The Role of Model-Based Systems Engineering in Complying with the EU AI Act. In *Proceedings of the 12th International Conference on Model-Based Software and Systems Engineering*; SciTePress: Setúbal, Portugal, 2024; Volume 1, pp. 365–370. [CrossRef]
18. Marcelo Corrales, C.; Iheanyi, N.; Nadia, M.; Nikolay, M.; Mark, F. Data Protection and Data Security by Design in Mobility Data Spaces: Insights from the MobiSpaces Project. *Eur. Data Prot. Law Rev.* **2024**, *10*, 257–276. [CrossRef]
19. Jørgensen, B.N.; Ma, Z.G. Cybersecurity and Resilience of Smart Grids: A Review of Threat Landscape, Incidents, and Emerging Solutions. *Appl. Sci.* **2026**, *16*, 981. [CrossRef]
20. Kauffman, M.E.; Soares, M.N.; Chao, K.-M.; Long, T.; Manzato, W.J.J. EU AI Act regulation: A study of non-European Union manufacturers' compliance preparedness. *J. Manuf. Technol. Manag.* **2026**, 1–20. [CrossRef]
21. Heymann, F.; Parginos, K.; Bessa, R.J.; Galus, M. Operating AI systems in the electricity sector under European's AI Act—Insights on compliance costs, profitability frontiers and extraterritorial effects. *Energy Rep.* **2023**, *10*, 4538–4555. [CrossRef]
22. Wagner, M.; Song, Q.; Borg, M.; Engstrom, E.; Lysek, M. AI Act high-risk AI compliance challenge and industry impact: A multiple case study. *Inf. Softw. Technol.* **2026**, *194*, 108067. [CrossRef]
23. Jorgensen, B.N.; Ma, Z.G. Digital Twin of the European Electricity Grid: A Review of Regulatory Barriers, Technological Challenges, and Economic Opportunities. *Appl. Sci.* **2025**, *15*, 6475. [CrossRef]
24. Fakhraian, E.; Semanjski, I.; Semanjski, S.; Aghezzaf, E.-H. Towards Safe and Efficient Unmanned Aircraft System Operations: Literature Review of Digital Twins' Applications and European Union Regulatory Compliance. *Drones* **2023**, *7*, 478. [CrossRef]

25. Christoph, S.; Abdelkader Magdy, S. Overview of AI standardization. In *IDIMT-2023: New Challenges for ICT and Management*; Trauner Verlag: Linz, Austria, 2023; pp. 143–149. [CrossRef]
26. Jovanovic, A.S. Stress-Testing the Resilience of Critical Infrastructures Exposed to Polycrises Triggered by Emerging Risks. *Int. J. Disaster Risk Sci.* **2025**, *16*, 594–617. [CrossRef]
27. Kotilainen, P.; Makitalo, N.; Systa, K.; Mehraj, A.; Waseem, M.; Mikkonen, T.; Murillo, J.M. Allocating distributed AI/ML applications to cloud-edge continuum based on privacy, regulatory, and ethical constraints. *J. Syst. Softw.* **2025**, *222*, 112333. [CrossRef]
28. Treiss, A.; Walk, J.; Kuehl, N. An Uncertainty-Based Human-in-the-Loop System for Industrial Tool Wear Analysis. In *Proceedings of the Joint European Conference on Machine Learning and Knowledge Discovery in Databases*; Springer International Publishing: Cham, Switzerland, 2021; pp. 85–100.
29. Borg, J.; McGeorge, D. Managing Changing Product Liability Obligations Emerging from New Proposed EU Directive. In *Proceedings of the International Conference on Computer Safety, Reliability, and Security*; Springer Nature: Cham, Switzerland, 2024; pp. 188–199.
30. Nikolinakos, N.T. *Adapting the EU Civil Liability Regime to the Digital Age: Artificial Intelligence, Robotics, and Other Emerging Technologies*; Springer: Berlin/Heidelberg, Germany, 2024; Volume 68, pp. 1–637.
31. Andrew, K.; Renee, P.-M.; Guillaume, S.; Mathilde, L.; Charles, D. Exploring Cooperation and Collaboration in Human AI Teaming (HAT)—EA SA AI Concept paper V2.0. *Transp. Res. Procedia* **2025**, *88*, 297–304. [CrossRef]
32. Hosseini, A.M.; Sauter, T.; Kastner, W. Trustworthy AI for Security Decision-Making in ICSs: Towards Compliance with the EU AI Act. In *Proceedings of the 2025 IEEE 30th International Conference on Emerging Technologies and Factory Automation (ETFA)*; IEEE: Piscataway, NJ, USA, 2025. [CrossRef]
33. Cancela-Outeda, C. The EU's AI act: A framework for collaborative governance. *Internet Things* **2024**, *27*, 101291. [CrossRef]
34. Casonato, C.; Marchetti, B. The proposal for a regulation of the European Union on artificial intelligence: Preliminary notes. *BioLaw J.* **2021**, *2021*, 415–437. [CrossRef]
35. Mahler, T. Smart Robotics in the EU Legal Framework. *Oslo Law Rev.* **2024**, *11*, 1–18. [CrossRef]
36. Burgstaller, P. The Use of AI and Its Legal Boundaries in the EU. In *Proceedings of the 2023 1st International Conference on Optimization Techniques for Learning (ICOTL)*; IEEE: Piscataway, NJ, USA, 2023; pp. 1–5.
37. Dzhusupova, R.; Bosch, J.; Olsson, H.H.; van den Brand, M. From Regulation to Implementation: EU AI Act Compliance in Engineering and Construction of the Oil & Gas and Petrochemical Sectors. In *Proceedings of the 2025 IEEE International Conference on Metrology for eXtended Reality, Artificial Intelligence and Neural Engineering (MetroXRAINE)*; IEEE: Piscataway, NJ, USA, 2025; pp. 1060–1065.
38. Ramona, K. EU AI Act Demands Assurance of AI Competence What Companies in Industry and Manufacturing Need to Know Now to Properly Train their Employees. *ZWF Z. Fuer Wirtsch. Fabr.* **2025**, *120*, 752–754. [CrossRef]
39. Ledda, A.; Fizzano, M.R.; Stefanelli, M.; D'Emilia, E.; Schena, E.; Massaroni, C.; Pinelli, M.; Giordani, D. Intelligence Applications in Personal Protective Equipment: Bibliometric Insights and Evolution Patterns (2020–2025). In *Proceedings of the 2025 International Workshop on Biomedical Applications, Technologies and Sensors (BATS)*; IEEE: Piscataway, NJ, USA, 2025; Volume 2025, pp. 1–6.
40. Bakar, A.A.; Yussof, S.; Ghapar, A.A.; Sameon, S.S.; Jørgensen, B.N. A Review of Privacy Concerns in Energy-Efficient Smart Buildings: Risks, Rights, and Regulations. *Energies* **2024**, *17*, 977. [CrossRef]
41. Andrew, M.; Lachlan, U. In cars (are we really safest of all?): Interior sensing and emotional opacity. *Int. Rev. Law Comput. Technol.* **2022**, *36*, 470–493. [CrossRef]
42. Dalgkitsis, A.; Koufakis, A.; Stutterheim, J.; Mifsud, A.; Atwani, P.; Gommans, L.; de Laat, C.; Papagianni, C.; Oprescu, A. Secure Collaborative Model Training with Dynamic Federated Learning in Multi-Domain Environments. In *Proceedings of the SC24-W: Workshops of the International Conference for High Performance Computing, Networking, Storage and Analysis*; IEEE: Piscataway, NJ, USA, 2024; pp. 755–759.
43. Tameem Hadi, F.; Mustafa, I.A.-K.; Luttfi, A.A.-H. Legal and Communication Challenges in Smart Grid Cybersecurity: Classification of Network Resilience Under Cyber Attacks Using Machine Learning. *J. Commun.* **2025**, *20*, 221–228. [CrossRef]
44. Evgeni, S.; Roumen, T.; Galya, P.; Kamelia, R. Analysis of Practical Machine Learning Scenarios for Cybersecurity in Industry 4.0. *Wseas Trans. Syst. Control* **2023**, *18*, 444–459. [CrossRef]
45. Katalin, T. Adapting the ITF's Transport Policy with a Focus on the EU's Strategy of Sustainable Mobility. *Period. Polytech. Transp. Eng.* **2025**, *53*, 222–225. [CrossRef]
46. Dürr, T.; Kaltenborn, U.; Staiger, D. Digital product passport for T&D enabling IoT integration and LCA. In *Proceedings of the 28th International Conference and Exhibition on Electricity Distribution (CIRED 2025)*; IET: London, UK, 2025; pp. 3163–3167.
47. Hägele, G.; Bouguerra, A.; Sarkheyli-Hägele, A. Towards the Certification of an Evacuation Assistance System Utilizing AI-based Approaches. In *Proceedings of the 2024 IEEE 35th International Symposium on Software Reliability Engineering Workshops (ISSREW)*; IEEE: Piscataway, NJ, USA, 2024; pp. 240–246.

48. Christophe, G.; Markus, M.; Christophe, G. Introduction to the European Machinery Regulation. In *European Digital Regulations*; Springer Nature: Cham, Switzerland, 2025; Volume 265, pp. 207–228.
49. Giulia De, C.; Damiano, T.; Rossana, P.; Chiara, V.; Lorenzo, B.; Daniela, P.B. Construction Site Equipment: Moving Towards the New Machinery Regulation. *Buildings* **2026**, *16*, 1391. [CrossRef]
50. Segate, R.V.; Daly, A. Encoding the Enforcement of Safety Standards into Smart Robots to Harness Their Computing Sophistication and Collaborative Potential: A Legal Risk Assessment for European Union Policymakers. *Eur. J. RISK Regul.* **2024**, *15*, 665–704. [CrossRef]
51. Pariso, P.; Picariello, M.; Marino, A. How AI Can Aim and Reduce the Annual Amount of Energy Consumption in European SMEs. *Macromol. Symp.* **2025**, *414*, 70045. [CrossRef]
52. Orosnjak, M.; Kedziora, S.; Desloges, M. Digital Maturity of SMEs in the EU: Leaders and Laggards of Luxembourg 's Manufacturing Ecosystem. *Technologies* **2025**, *13*, 541. [CrossRef]
53. Bunzel, N. Compliance Made Practical: Translating the EU AI Act into Implementable Security Actions. In *Proceedings of the 2025 IEEE/ACM International Workshop on Responsible AI Engineering (RAIE)*; IEEE: Piscataway, NJ, USA, 2025; pp. 69–73.
54. Gless, S.; Ligeti, K. Regulating driving automation in the European Union—Criminalliability on the road ahead? *New J. Eur. Crim. Law* **2024**, *15*, 33–57. [CrossRef]
55. Ma Carmen Núñez, Z. Towards a uniform European legal framework in the prevention of risks and civil liability in the field of intelligent automated driving. *Cuad. Derecho Transnacional* **2023**, *15*, 689–723. [CrossRef]
56. Przemysław, G.; Lucrezia, N.; Monika, Z. Harmonising risk assessments for high-risk AI systems under the GDPR and the AI Act. *J. Data Prot. Priv.* **2025**, *7*, 359–371. [CrossRef]
57. Canas, P.N.; Nieto, M.; Otaegui, O.; Rodriguez, I. A Methodology to Enhance Transparency for Trustworthy Artificial Intelligence for Cooperative, Connected, and Automated Mobility. *SAE Int. J. Connect. Autom. Veh.* **2025**, *8*, 125–139. [CrossRef]
58. Torens, C.; Junger, F.; Schirmer, S.; Nagarajan, P.; Schopferer, S.; Zhukov, D.; Dauer, J. Runtime monitoring of operational design domain to safeguard machine learning components. *CEAS Aeronaut. J.* **2025**, *16*, 973–991. [CrossRef]
59. Strem, N.; Dhami, D.S.; Kersting, K. Industrial-Scale Autonomous eXplainable Iterative Learning. In *Proceedings of the Pacific Rim International Conference on Artificial Intelligence*; Springer Nature: Singapore, 2026; Volume 16454 LNAI, pp. 183–193.
60. Jørgensen, B.N.; Ma, Z.G. Infostructure: A Scoping Review and Reference Architectural Framework for Situation Awareness in Future Power System Control Rooms. *Energies* **2026**, *19*, 1472. [CrossRef]
61. Zhu, Y.; Abali, B.E.; Vollmecke, C. Explainable AI analysis of printing parameter effects on life cycle as sessment for sustainable material-extrusion additive manufacturing. *Virtual Phys. Prototyp.* **2026**, *21*, e2632456. [CrossRef]
62. Kleoniki, P.; Ifigeneia, T. AI as an Environmental Challenge: Mapping Safeguards in EU Environment Al and Climate Law to Address the 'Silence' in the EU AI Act. *Eur. Energy Environ. Law Rev.* **2025**, *34*, 25–37. [CrossRef]
63. Katsoura, E.; Ferryman, J.; Stavropoulos, G.; Uhl, A.; Gierszal, H.; Kruszyński, A.; Tyczka, P.; Murray, S.; Martín Zamorano Barrios, M.; Murphy, E.; et al. Vision Paper: AutoBorder—Vehicle-Integrated Solutions for Drive-Through and Human-Centric Borders. In *Proceedings of the 2025 IEEE International Conference on Big Data (BigData)*; IEEE: Piscataway, NJ, USA, 2025; Volume 2025, pp. 4892–4900.
64. Kaakai, F.; Dmitriev, K.; Adibhatla, S.; Baskaya, E.; Bezzecchi, E.; Bharadwaj, R.; Brown, B.; Gentile, G.; Gingins, C.; Grihon, S.; et al. Toward a Machine Learning Development Lifecycle for Product Certification and Approval in Aviation. *SAE Int. J. Aerosp.* **2022**, *15*, 127–143. [CrossRef]
65. Doukari, O.; Bacnak, T.; Mahamedi, E. Stakeholder-informed renovation digital twin for multi-domain monitoring: Pilot implementation and acceptance evaluation. *Energy Build.* **2026**, *354*, 116979. [CrossRef]
66. Benkrid, A.; Zahra, O.; Szőke, R.; Shukla, A.; Szőke, I. Digital Transformation of the Nuclear Industry: Leveraging Robotics, AI, and Digital Twins for Standardised, Safe, and Efficient Operations. In *Proceedings of the 22nd International Conference on Informatics in Control, Automation and Robotics*; SciTePress: Setúbal, Portugal, 2025; Volume 2, pp. 526–535. [CrossRef]
67. Gondak, M.d.O.; Hluszko, C.; Moretti, V.; Zeni, G.L.F.; de Francisco, A.C. Guide to digital technologies for circularity in the project lifecycle: Perspective of the EN 15643-3:2012 standard. *Constr. Innov.* **2025**. [CrossRef]
68. Andrasko, J.; Hamul'ak, O.; Mesarcik, M.; Kerikmae, T.; Kajander, A. Sustainable Data Governance for Cooperative, Connected and Automated Mobility in the European Union. *Sustainability* **2021**, *13*, 10610. [CrossRef]
69. Neumann, T.; Lukasik, R. Integrating Autonomous Trucks into Human-Centric Operations: A Path to Safer and More Energy-Efficient Road Transport. *Energies* **2025**, *18*, 4219. [CrossRef]
70. Hadj-Mabrouk, H. A literature review on the applications of artificial intelligence to European rail transport safety. *IET Intell. Transp. Syst.* **2024**, *18*, 2291–2324. [CrossRef]
71. Lee, Y.-G.; Lee, C.-H.; Jeon, Y.-H.; Bae, J.-H. Transformative Impact of the EU AI Act on Maritime Autonomous Surface Ships. *Laws* **2024**, *13*, 61. [CrossRef]

72. Legat, C.; Li, R.; Hong, P.; de Meer, J.; Boell, M. Opportunities for Leveraging Artificial Intelligence in Manufacturing: A Comparative Perspective of China and Germany on Regulation and Standardization. *IFAC-PapersOnLine* **2025**, *59*, 422–427. [CrossRef]
73. Katja, G.; Benjamin, P.; Jonas, H.; Michael, E. Use of Machine Learning in Smart Factories A Risk Assessment. *ZWF Z. Fuer Wirtsch. Fabr.* **2024**, *119*, 356–361. [CrossRef]
74. Di Paco, F.; Burattini, L.; Gabbrielli, R.; Landi, L.; Marcelloni, F.; Marrazzini, L.; Palumbo, M.; Pirozzi, M. AISAFETY: An AI-based smart system for enhancing operator safety in production processes. *Saf. Sci.* **2026**, *199*, 107201. [CrossRef]
75. Florian, S.; Rouven, J.; Stefano, P.; Elisa, V.; Maurizio, G. Integrative inspection methodology for enhanced PCB remanufacturing using artificial intelligence. *Procedia CIRP* **2025**, *132*, 227–232. [CrossRef]
76. Barr, P.A.; Imroz, S.M. A Cross-Regional Review of AI Safety Regulations in the Commercial Aviation Industry. *Adm. Sci.* **2026**, *16*, 53. [CrossRef]
77. Justo-Hanani, R. Risk-based approach to EU AI act: Benefits and challenges of co-regulation. *Policy Des. Pract.* **2026**, *9*, 171–180. [CrossRef]
78. Mitja, K. Tragedy of the commons, civil drones and hybrid modes of technology regulation: A comparative law and economics perspective. *J. Law Mark. Innov.* **2024**, *3*, 179–211. [CrossRef]
79. Jankowska, A. Innovative Approaches to Transport Challenges in Port Cities Based on Literature Review. *TRANSNAV-Int. J. Mar. Navig. Saf. SEA Transp.* **2025**, *19*, 469–476. [CrossRef]
80. Matviienko, H.; Kucherkova, S.; Yanovska, V.; Hurochkina, V.; Ternovsky, V.; Kęsy, M. Governmental Management and Regulatory Measures for Advancing AI in the Ukrainian Energy Sector as a Basis for Rapid and Sustainable Development of the Ukrainian Economy. In *Proceedings of the 2023 13th International Conference on Advanced Computer Information Technologies (ACIT)*; IEEE: Piscataway, NJ, USA, 2023; pp. 303–307.
81. Jørgensen, B.N.; Ma, Z.G. Regulating AI in the Energy Sector: A Scoping Review of EU Laws, Challenges, and Global Perspectives. *Energies* **2025**, *18*, 2359. [CrossRef]

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.