International Journal of Computer Integrated Manufacturing

ISSN: 0951-192X (Print) 1362-3052 (Online) Journal homepage: www.tandfonline.com/journals/tcim20

AI-assisted auditing review of mechanical engineering drawings: a Perception–Cognition–Collaboration reference architecture and its implementation Perspectives

Xin Chen & Kai Cheng

To cite this article: Xin Chen & Kai Cheng (07 Aug 2026): AI-assisted auditing review of mechanical engineering drawings: a Perception–Cognition–Collaboration reference architecture and its implementation Perspectives, International Journal of Computer Integrated Manufacturing, DOI: 10.1080/0951192X.2026.2714117

To link to this article: https://doi.org/10.1080/0951192X.2026.2714117

© 2026 The Author(s). Published by Informa UK Limited, trading as Taylor & Francis Group.

Published online: 07 Aug 2026.

Submit your article to this journal

Article views: 474

View related articles

View Crossmark data

Full Terms & Conditions of access and use can be found at
https://www.tandfonline.com/action/journalInformation?journalCode=tcim20

INTERNATIONAL JOURNAL OF COMPUTER INTEGRATED MANUFACTURING
https://doi.org/10.1080/0951192X.2026.2714117

Taylor & Francis by informa

OPEN ACCESS Check for updates

AI-assisted auditing review of mechanical engineering drawings: a Perception–Cognition–Collaboration reference architecture and its implementation Perspectives

Xin Chen[a*] and Kai Cheng[b]

[a]School of Intelligent Science and Information Engineering, Shenyang University, Shenyang, China; [b]Department of Mechanical and Aerospace Engineering, Brunel University London, Uxbridge, UK

ABSTRACT

Mechanical drawing audits are essential to product design and manufacturing but remain labor-intensive and error-prone, particularly when checking tolerances, surface roughness, materials, treatments and assembly requirements. Although computer vision, document understanding, large language models (LLMs), knowledge representation and multi-agent systems (MAS) can support individual tasks, existing efforts remain fragmented. This paper proposes an engineering-oriented approach to AI-assisted mechanical drawing auditing and a three-layer Perception – Cognition – Collaboration (P – C – C) pre-validation reference architecture. Perception transforms carrier-specific evidence into provenance-bearing entities and relations. Cognition aligns them with a versioned Standards-and-Requirements Knowledge Graph and applies approved deterministic numerical and logical rules outside the LLM. Collaboration coordinates audit tasks, preserves conflicts, and routes critical or uncertain findings to authorized engineers. An audit of the Merge_places Computer Vision Dataset reveals limitations in schema consistency, lineage, class semantics and clause grounding. De-identified drawing regions illustrate clearance calculation, detection of a controlled tolerance mismatch, and the routing of GD&T and bill-of-materials (BOM) requests when evidence is incomplete. These bounded cases clarify interface behavior and define validation requirements for provenance, deterministic checking, interoperability, uncertainty, cybersecurity and human governance before engineering deployment. The framework integrates heterogeneous evidence, executable rules and accountable human decisions within a traceable and selectively automated audit workflow.

ARTICLE HISTORY
Received 19 November 2025
Accepted 31 July 2026

KEYWORDS
Engineering drawing auditing review; AI-assisted compliance checking; P–C–C reference architecture; standards-grounded reasoning; multi-agent systems; design automation

1. Introduction

Engineering design and manufacturing are central to modern product development, while persistent pressures include efficiency, cost, shorter delivery times, and increasingly stringent precision and quality requirements in the design and manufacturing process. Among the bottlenecks encountered throughout the process, the auditing review of mechanical engineering drawings remains critical and labor-intensive, particularly for precision products and machine design. Engineering drawings communicate design intent through dense geometric data, textual annotations, symbols, dimensions and tolerances that support design for manufacturing and design for assembly. Manual auditing requires experienced engineers to verify drawings against evolving geometric product specification (GPS) standards, engineering standards and project-specific requirements; it is therefore time-consuming and susceptible to errors and inconsistent interpretations. Such errors can propagate into manufacturing and assembly, leading to rework, safety risks and legal disputes. Intelligent augmentation of drawings auditing review is consequently relevant to productivity, standards compliance and product quality, but any automation must preserve traceability and accountable engineering judgement.

The pursuit of automated compliance checking is not new; it has been an active research area for several decades and a long-standing aspiration of design and manufacturing communities. Early work translated prescriptive requirements into computer-executable logic, but such systems were difficult to maintain when source information or requirements were ambiguous (Amor and Dimyadi 2021; Aslam and Umar 2024). Model-based representations, including computer-aided design (CAD) and model-based definition (MBD),

CONTACT Kai Cheng xin.chen@brunel.ac.uk; kai.cheng@brunel.ac.uk
*currently working as a visiting academic at Brunel University London.

© 2026 The Author(s). Published by Informa UK Limited, trading as Taylor & Francis Group.
This is an Open Access article distributed under the terms of the Creative Commons Attribution License (http://creativecommons.org/licenses/by/4.0/), which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited. The terms on which this article has been published allow the posting of the Accepted Manuscript in a repository by the author(s) or with their consent.

subsequently enabled richer object-level checking, while related building information modelling research demonstrated both the value and the cost of semantic enrichment (Yahya, Breslin, and Intizar Ali 2021). These adjacent-domain results inform workflow and rule-representation mechanisms, but they do not constitute performance evidence for mechanical GPS, GD&T, PMI, tolerance or assembly auditing.

The rise of deep learning opened new possibilities for automating parts of the auditing process, especially for legacy projects that rely on two-dimensional drawings rather than fully annotated MBD or ISO 10303 STEP AP242 models. More recently, generative AI and LLMs have been investigated for interpreting technical standards and design specifications, while MAS offers an organizational paradigm for coordinating multi-specialty auditing tasks. These technologies may contribute to the next generation of AI-assisted drawing auditing, but component-level success does not by itself establish end-to-end compliance accuracy, safe decision authority or industrial readiness.

This paper presents an engineering-oriented approach to AI-assisted auditing review of mechanical engineering drawings. The approach is particularly focused on using integrated AI and precision engineering principles for assisting the engineering design and manufacturing of precision machines. Therefore, three-tier reference architecture is developed and formulated from the complementary roles of computer vision, standards-grounded reasoning and controlled multi-agent collaboration. The AI-assisted auditing review focuses on auditing capability, evidence type, limitation and industrial readiness, while precision engineering focuses on specifying typed interfaces, deterministic rule gates, uncertainty propagation, abstention, conflict escalation and human release for mechanical drawing auditing. Together, auditable integration requirements and mechanical-domain decision boundaries provide testable specifications for implementation and validation. The remainder of the paper presents the implementation perspectives of the approach particularly focusing on key enabling technologies, the evolution and limits of existing evidence, the P–C–C reference architecture, and furthermore, the data, benchmark, interoperability, governance and validation conditions required for implementations.

The evidence set was assembled from DOI-indexed bibliographic and publisher records and updated through July 2026. Searches paired mechanical drawing, technical drawing, CAD, MBD, PMI, QIF, GD&T, tolerance, audit, verification, document understanding, knowledge graph, large language model, multi-agent, interoperability and human-AI terms; a targeted International Journal of Computer Integrated Manufacturing search added recent manufacturing work on semantic capability reasoning, task allocation, collaboration quality, Industry 5.0, governance and digital integration. Records were retained when their identity was reconcilable and their reported input, output, evidence type and limitation could be coded. The synthesis asks which capabilities support a traceable audit finding, where cross-layer uncertainty or domain-transfer assumptions break that chain, and which contracts and validation conditions are required for implementation. It is a structured integrative review rather than an exhaustive database census or meta-analysis, and no performance ranking is inferred across incompatible studies.

2. Foundations for enabling AI-Assisted engineering drawings auditing review

The foundation for supporting AI-assisted auditing review of engineering drawings includes five technology streams and their integration, i.e. the mechanical engineering drawing perception; CAD/MBD and quality-information exchange; knowledge representation and semantic rules; LLM and vision-language assistance; and multi-agent and human-AI collaboration. Table 1 compares their reported respective evidence, required mechanical-auditing artifacts, and maturity boundaries and positions the proposed P–C–C synthesis against them.

2.1. *Knowledge representation and semantic technologies*

Knowledge representation forms the semantic backbone of intelligent auditing systems. Knowledge graphs (KGs) support the representation of heterogeneous mechanical codes, standards and design entities by making relationships explicit (Paulheim 2017). Ontologies define a controlled conceptualization of a domain and can support reasoning about design decisions and their implications (Noy and McGuinness 2001). The Web Ontology Language (OWL) and Resource Description Framework (RDF) provide formal languages for knowledge modelling and data exchange (Hitzler et al. 2012; Tomaszuk and Hyland-Wood 2020). These technologies can represent designs and requirements, but the formalism alone does not guarantee ontology coverage, correct entity alignment, current standard editions or valid executable rules.

Table 1. Evidence streams, auditing artifacts and maturity boundaries.

Technology stream	Reported evidence and evaluation setting	Required mechanical-auditing artifact or interface	Principal limitation and maturity boundary (representative sources)
Mechanical drawing perception	Task-specific studies report symbol, text, layout, region and graph extraction from selected drawing types or entity classes.	Candidate entities and relations with source locators, transcription, units, alternatives, provenance and calibrated uncertainty	Evidence remains predominantly component-level and task-specific; extraction performance does not by itself establish complete feature association, clause applicability or compliance (Bickel, Goetz, and Wartzack 2024; Jamieson, Moreno-García, and Elyan 2024; Lin et al. 2023; Reddy et al. 2026; Villena Toro, Wiberg, and Tarkian 2023).
CAD, MBD and QIF exchange	Structured CAD models preserve geometry and feature relations, while QIF supports exchange across product definition, inspection planning, execution and results.	Carrier-specific adapters feeding a canonical evidence record and DKG while retaining feature identity and revision provenance	Structured exchange reduces dependence on raster recognition but does not establish PMI completeness, semantic correctness or executable compliance; schema-conformance and round-trip tests remain necessary (Cao et al. 2020; Michaloski et al. 2016; Z. Zhang, Jaiswal, and Rai 2018).
Knowledge representation and semantic rules	Studies report ontology construction, rule formalization and model querying in bounded manufacturing or adjacent-domain applications.	A provenance-bearing Drawing Knowledge Graph (DKG) linked to a separately governed Standards-and-Requirements Knowledge Graph (SRKG), with clause identity, edition, scope, applicability, approved constraints and tests	Mechanical GPS, datum, modifier, unit, exception and supersession semantics require domain-specific formalization and expert validation. The integrated DKG/SRKG contract remains unvalidated in mechanical drawing auditing (Amor and Dimyadi 2021; Järvenpää et al. 2023; Noy and McGuinness 2001; Paulheim 2017; Yahya, Breslin, and Intizar Ali 2021; J. Zhang and El-Gohary 2017).
LLM and vision-language assistance	Studies report prototypes for candidate retrieval, terminology normalization, requirement parsing, CAD assistance and explanation drafting.	Source-linked candidate clauses and operands supplied to version-controlled deterministic checks; generated content has no independent decision authority	Probabilistic outputs may omit exceptions, alter values or cite unsupported sources; bounded tool access, structured outputs, source verification and risk-appropriate human authorization are required (Deng, Khan, and Ahmet Erkoyuncu 2024; Fuchs et al. 2024; Mustapha 2025; Picard et al. 2023; Yang and Zhang 2024).
Multi-agent and human-AI collaboration	Studies report task decomposition, monitoring, reassignment and collaboration-quality concepts, but do not provide end-to-end mechanical-audit evidence.	Typed candidate findings, dependency and conflict records, independent re-checks, safe routes, overrides and accountable human release	End-to-end evidence for integrated mechanical auditing remains unavailable; correlated errors, orchestration overhead, workload, authority allocation, cost and safety remain unquantified in this application (Briseno et al. 2024; Guo et al. 2024; Kokotinis et al. 2023; Peruzzini, Prati, and Pellicciari 2024; Petzoldt, Harms, and Freitag 2023).
Proposed P–C–C reference architecture	An engineering-oriented, traceability-centred reference architecture is specified and examined through five bounded conditions: drawing-internal numerical agreement, controlled numerical conflict, missing standards prerequisites, feature-control-frame evidence and incomplete BOM context. Evaluation is limited to manual transcriptions and controlled test values rather than end-to-end system execution.	Carrier-aware evidence, DKG/SRKG separation, deterministic authority for hard constraints, typed cross-role findings, conflict handling, explicit safe states and accountable human release	The contribution is an architectural and evidence-governance synthesis supported by bounded checks. Component accuracy, integration benefit, scalability, industrial performance and GD&T compliance require independent end-to-end validation.

2.2. AI foundations: deep Learning and natural Language processing

Deep learning and natural language processing (NLP) underpin candidate perceptual and language-assistance capabilities in AI-assisted auditing review. Transformer architectures introduced attention mechanisms for long-range dependencies (Vaswani et al. 2017), and document models can integrate visual, textual and spatial information (Huang et al. 2022). Self-supervised learning can reduce reliance on labelled data, although it does not remove the need for task-specific validation on engineering drawings (Y. Xu et al. 2021). Foundation models such as BERT (Devlin et al. 2019) and GPT-family models (Brown et al. 2020; Radford et al. 2018) support contextual language processing, but their performance on precise standards interpretation depends on retrieval, controlled terminology, rule validation and domain evidence. Convolutional neural networks (Krizhevsky, Sutskever, and Hinton 2017; Simonyan and Zisserman 2015), residual networks (K. He et al. 2016), Faster R-CNN (Ren et al. 2017) and YOLO (Redmon et al. 2016) support image and symbol detection. These methods can contribute observations to the Perception layer; detection or language fluency alone does not establish compliance.

2.3. CAD and technical Document processing

In computer-aided design (CAD), graph neural networks (GNNs) can represent geometric and topological structures (Z. Wu et al. 2021). Graph convolutional networks (Kipf and Welling 2017) and graph attention networks (Veličković et al. 2018) model nodes as geometric entities and edges as spatial or topological relationships such as connectivity, containment and adjacency. This relational representation can complement pixel-based perception, although errors in vectorization, association or graph construction propagate to downstream reasoning. Optical character recognition (OCR) extracts textual and numerical observations from drawings, where orientation, fonts, symbols and dense graphical context remain difficult. CRAFT (Baek et al. 2019), end-to-end sequence recognition (Shi, Bai, and Yao 2017) and differentiable binarization (M. Liao et al. 2020) have improved general text detection and recognition, but engineering deployment also requires exact numerical, unit, sign, decimal and entity-association checks.

2.4. Product definition technologies and Applications

In mechanical domains, MBD and semantically annotated exchange packages can preserve geometry together with product and manufacturing information, while the Quality Information Framework (QIF) connects product definition, inspection planning, execution and results within a digital quality thread (Michaloski et al. 2016). Such structured sources can reduce dependence on raster recognition, but exchange success does not prove semantic completeness, correct PMI association, rule computability or revision consistency. Experience from heterogeneous BIM and IFC environments illustrates why carrier-specific conformance and round-trip loss tests remain necessary when information crosses tools (Azhar, Khalfan, and Maqsood 2015; Lai and Deng 2018; Sacks et al. 2018).

Research on requirement transformation and automated checking has developed useful mechanisms for clause segmentation, semantic classification, logic representation and model-to-rule alignment (Nawari 2012; J. Zhang and El-Gohary 2017; P. Zhou and El-Gohary 2016). These studies are concentrated largely in Architecture, Engineering and Construction (AEC) and cannot be reused as mechanical rules without reconstructing the ontology, applicability conditions, tabulated GPS semantics, datum relations, units, exceptions and standard editions. In the proposed architecture, language models may assist retrieval and candidate formalization, but only approved rules are executed by a version-controlled deterministic engine. QIF or Product Lifecycle Management (PLM) adapters preserve source and lifecycle provenance rather than serving as evidence that a drawing is compliant.

2.5. Synthesis of foundational enabling technologies

In summary, mechanical engineering drawing auditing may be supported by the integration of semantic representation, perceptual evidence acquisition, structured rule execution and controlled collaboration. Knowledge graphs can provide an explicit representational backbone; deep learning and NLP can contribute candidate observations and language assistance; GNNs and OCR can bridge geometric, relational and textual evidence; and CAD, MBD, STEP AP242 and QIF can improve semantic interoperability. These capabilities establish useful components, but the literature must still be examined for interface compatibility, uncertainty treatment, clause traceability, human responsibility and end-to-end evidence.

3. Literature review: the evolution of automated auditing

The literature can be organized as a progression from rule-based and product-definition approaches to data-driven perception, language-model assistance, semantic enrichment and multi-agent coordination. Table 1 compares the reported evidence, required mechanical-audit artifacts, maturity boundaries, and representative sources for these streams and the proposed P–C–C synthesis.

3.1. Early approaches and the model-based product definition Paradigm

Early automated compliance checking relied on hard-coded rules that translated prescriptive drawing or tolerance requirements into executable logic (Malsane et al. 2015). Object-oriented design models, MBD and, by analogy, BIM supplied more structured representations, but semantic enrichment remained a persistent challenge. NLP and knowledge graphs were subsequently used to extract requirements and link them to model entities (Bloch and Sacks 2020; Nakhaee, Elshani, and Wortmann 2024; Wang et al. 2023). Regional fire-code implementations also illustrate interface and adoption questions (Ismail et al. 2023), while domain applications have examined knowledge-graph checking and computable compliance workflows (Dimyadi and Amor 2013; Peng and Liu 2023). These studies support explicit representation and rule execution, but BIM object schemas, building-code rules and AEC benchmarks cannot be transferred directly to mechanical GD&T, PMI or assembly auditing without reconstructing the domain entities, constraints and reference truth.

3.2. Deep Learning for visual analysis and CAD

Although MBD supports model-centric workflows, a substantial portion of existing engineering assets remains in two-dimensional drawing form. Drawing-specific research has examined symbol recognition, engineering-diagram digitization, structured information extraction, GD&T detection, sketch analysis, title-block processing and technical-drawing text detection (Bickel, Goetz, and Wartzack 2024; Jamieson, Moreno-García, and Elyan 2024; Joffe et al. 2024; Kashevnik et al. 2023; Khan et al. 2025; Lin et al. 2023; Reddy et al. 2026; Schlagenhauf, Netzer, and Hillinger 2023; Stürmer, Graumann, and Koch 2024). Related graph and document methods represent layout, geometry and topological relations (Beltagy, Peters, and Cohan 2020; Cao et al. 2020; Carrara, Nousias, and Borrmann 2024; Gori, Monfardini, and Scarselli 2005; Luo et al. 2022; Seff et al. 2020; X. Xu et al. 2025; Y. Xu et al. 2020; Z. Zhang, Jaiswal, and Rai 2018; W. Zhang et al. 2023), while pixel-based segmentation provides another route for complex drawing layers (Moreno-Garcia, Johnston, and Garkuwa 2020). These advances provide task-specific perception-layer evidence; detection or transcription accuracy does not by itself establish feature association, rule applicability, clause correctness or end-to-end compliance.

3.3. Large Language models and multi-agent systems

LLMs have been investigated for CAD automation, design activity and standards-grounded assistance (Chen et al. 2024; Deng, Khan, and Ahmet Erkoyuncu 2024; Kumar et al. 2025; Li et al. 2025; Madireddy et al. 2025; Mustapha 2025; Prohofsky and Lande 2024; Schüpbach et al. 2025; Yang and Zhang 2024; L. Zhang et al. 2025). Vision-language and multimodal studies extend this assistance to sketches, drawings and engineering design inputs (Khan et al. 2024; Z. Liao et al. 2024; H. Liu et al. 2023; Picard et al. 2023; Zhu et al. 2023). MAS research supplies task-decomposition and coordination concepts, including software-agent co-learning and adjacent compliance-review cases (Guo et al. 2024; J. He, Treude, and Lo 2024; Jin and Levitt 1996; Macal and North 2010; Qian et al. 2024; Wan et al. 2025). These methods can reduce interface effort by proposing retrieval queries, terminology mappings, candidate logic or explanations, but fluent output is not a formal derivation and agreement among agents is not independent corroboration when models, prompts, retrieval or upstream evidence are shared.

3.4. Semantic enrichment and knowledge representation

Bridging the semantic gap between engineering design data and requirements remains a central challenge. Early web-based design and manufacturing systems exposed implementation and responsiveness issues that remain relevant to distributed engineering services (Cheng, Pan, and Harrison 2001; Toussaint and Cheng 2002). Knowledge graphs and semantic-web methods can represent clauses, design entities and inferred relations, while ontology engineering can formalize functional and manufacturing knowledge (Chidara, Cheng, and Gallear 2025; Kitamura and Mizoguchi 2004; Pauwels, Zhang, and Lee 2017; Y. Zhou, Lin, and She 2021). Adjacent BIM studies demonstrate semantic-service, geospatial and embodied-impact reasoning but also the cost of domain modelling and entity alignment (Hou, Li, and Rezgui 2015; Karan and Irizarry 2015; J. Wu, Nousias, and Borrmann 2025). Knowledge-graph construction research further emphasizes extraction, learning and evaluation requirements (Bian 2025; Choi and Jung 2025). Mechanical transfer therefore requires separate Drawing and Standards-and-Requirements Knowledge Graphs with explicit provenance, applicability, standard editions, exceptions and inference tests.

3.5. Document understanding and engineering standards compliance

Document-understanding architectures such as DocFormer, Donut, DiT, LayoutLM and related text detectors combine visual, textual and spatial evidence (Appalaraju et al. 2021; Kim et al. 2022; Li et al. 2022; Y. Xu et al. 2020; X. Zhou et al. 2017). General language and vision foundations, including RoBERTa, EfficientNet, Xception, SSD, early document-recognition networks and deep-learning reference architectures, supply reusable components but not mechanical compliance evidence (Chollet 2017; Goodfellow, Bengio, and Courville 2016; LeCun et al. 1998; W. Liu et al. 2016; Y. Liu et al. 2019; Tan and Le 2019). AEC ontology, digital-model, interoperability, evidence-mapping and adoption research illustrate both the value and the governance cost of structured compliance information (Acierno et al. 2017; Beach, Hippolyte, and Rezgui 2020; Borrmann et al. 2018; Hosseini et al. 2018; Jiang et al. 2022; Pazlar and Turk 2008). Work on regulation interpretation and information extraction provides adjacent-domain mechanisms for requirement segmentation, computability analysis and candidate rule generation (Bommarito, Martin Katz, and Detterman 2018; Fuchs et al. 2024; R. Zhang and El-Gohary 2018; Zhong et al. 2020; P. Zhou and El-Gohary 2017). All of these methods still require controlled standards sources, edition tracking, expert adjudication, deterministic tests and explicit indeterminate states before a mechanical compliance decision is released.

3.6. Synthesis and research trends

Automated compliance checking has evolved from rigid rule-based systems toward data-driven, semantic and collaborative approaches. Rule systems offer inspectable execution but are costly to author and maintain; MBD and CAD representations improve structure but depend on semantic completeness; deep learning supports component extraction but inherits data and generalization limitations; LLMs support flexible language assistance but can produce unsupported or numerically incorrect outputs; and multi-agent coordination can distribute work but introduces communication, conflict, latency and correlated-error risks. The critical gap is therefore not the absence of individual technologies, but the absence of a validated mechanical evidence chain connecting source observations, relations, versioned requirements, executed checks, uncertainty, conflict disposition and human release.

Several challenges constrain large-scale industrial adoption. Public, well-annotated mechanical drawing datasets remain limited, and proprietary restrictions impede external validation and benchmark comparability. Diverse carriers, including DWG, DXF, PDF, native CAD, STEP AP242 and QIF, create different semantic and information-loss risks. Explainability requires source and clause provenance rather than generated narrative alone, and trustworthy operation requires calibrated uncertainty, failure containment, cybersecurity, version control and human accountability. These issues are treated as design and evaluation requirements in Sections 4 and 5 rather than as problems assumed to be solved by the proposed architecture.

Recent research indicates movement toward hybrid approaches that combine data-driven perception with explicit knowledge representations and rule execution. Manufacturing studies on semantic capability matching, collaborative task allocation, collaboration quality and human-automation symbiosis further show that reasoning, coordination, evaluation and human responsibility must be designed together (Järvenpää et al. 2023; Kokotinis et al. 2023; Peruzzini, Prati, and Pellicciari 2024; Petzoldt, Harms, and Freitag 2023). The P–C–C contribution is accordingly framed as a reference architecture that makes cross-layer evidence contracts, deterministic gates, abstention, conflict escalation and mechanical-domain boundaries explicit. Its feasibility and performance remain hypotheses for implementation and external validation.

4. A proposed P–C–C (Perception-Cognition-Collaboration) reference architecture

Building on this synthesis, the P–C–C reference architecture organizes mechanical drawing auditing into Perception, Cognition and Collaboration layers. Each layer defines a distinct responsibility for evidence acquisition, standards-grounded checking or controlled coordination; Table 2 specifies their inputs, outputs, authority limits and escalation routes (Michaloski et al. 2016).

Figure 1 summarizes the information flow. Raw drawings or models enter controlled ingestion and Perception; candidate entities and relations pass through schema and uncertainty gates; Cognition aligns approved evidence with versioned standards and deterministic rules; and Collaboration routes typed findings, conflicts and incomplete evidence to authorized human review.

4.1. Perception layer

The Perception layer forms the evidential foundation of the framework by transforming raw mechanical drawings or model content into structured candidate observations for subsequent checks. Mechanical drawings combine geometric entities, dimensions, tolerances, surface-texture symbols, welding marks, material notes and view relations. The layer must preserve each observation's source location, extraction route, candidate class or value, unit, quality flag, and uncertainty in the DKG. Passing a Perception gate means that an observation is structurally usable; it does not mean that the drawing is

Table 2. Proposed P–C–C layer contracts, authority limits and canonical safe routes.

Layer	Preconditions / required input	Proposed structured output	Permitted authority and prohibited inference	Canonical safe route / terminal disposition
Perception	Controlled drawing or model source, revision, carrier, access context and quality state	DKG candidates with entity IDs, regions or geometry, values, units, relations, locators, provenance and uncertainty	May apply carrier, schema, coordinate, unit, range and selective-prediction gates; may not infer standards compliance	NEEDS_RECOGNITION_REVIEW or NEEDS_EVIDENCE
Cognition	Structurally screened DKG observations, an authorized and versioned SRKG requirement and verified operands	Candidate clause-grounded claim containing applicability, executed rule result, severity, evidence links, uncertainty and execution trace	Approved deterministic rules may evaluate hard constraints; language models may retrieve or rank candidate clauses, parse candidate structures and draft explanations, but cannot supply missing semantics, clauses or decision authority	NEEDS_RULE_REVIEW or NEEDS_EVIDENCE; ABSTAINED if a required rule or evidence prerequisite remains unavailable
Collaboration	Typed candidate claims, dependencies, specialist or tool identities and versions, conflict classes, and audit policy	Screened candidate-claim set pending authorized release, with dependencies, conflicts, overrides and evidence history retained	May check completeness, source authority, duplicates and conflicts and rerun approved rules; only the authorized engineering signatory may release a consequential claim	NEEDS_CONTEXT, CONTESTED, or NEEDS_HUMAN_RELEASE; RELEASED and REWORK require the recorded human gate. ABSTAINED may follow an unresolved prerequisite or a recorded human decision.

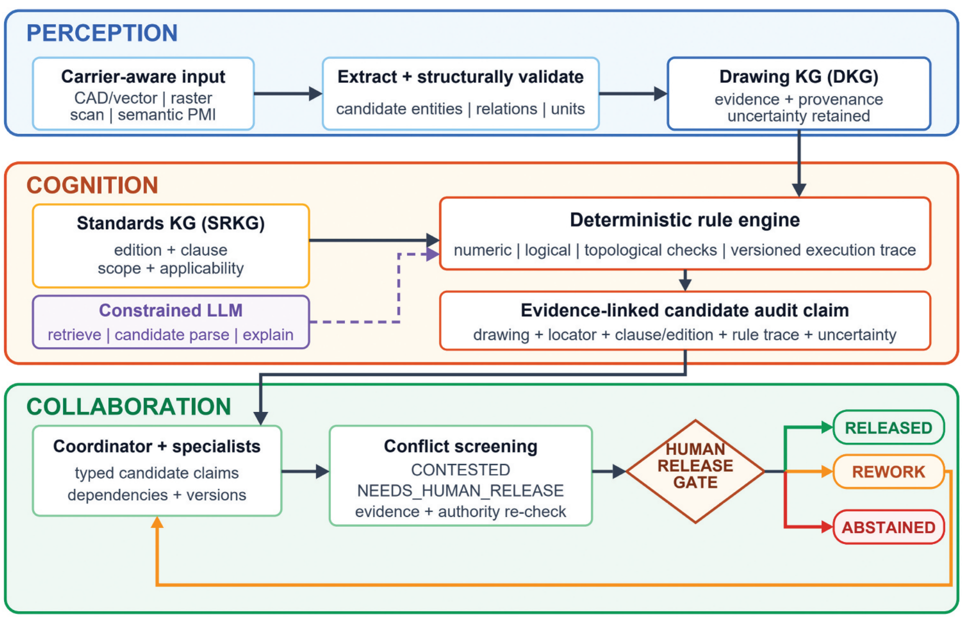

Figure 1. Evidence-preserving P–C–C reference architecture for mechanical drawing audit.

compliant or that the observation is semantically complete.

Semantic segmentation and layout models may identify drawing regions such as parts, sectional views and annotation blocks. Object detectors may locate GD&T, machining, weld and surface-texture symbols, while OCR may extract dimension values, material labels and manufacturing notes (W. Liu et al. 2016; Redmon and Farhadi 2018). These outputs are integrated into the DKG, where nodes represent candidate entities and edges represent candidate spatial, topological and semantic relations. Schema, coordinate, unit, cross-view and selective-prediction checks must precede transfer to Cognition, and ambiguous classes or associations must remain explicit or be routed to recognition review.

4.2. Cognition layer

The Cognition Layer aligns DKG observations with a separately maintained Standards-and-Requirements Knowledge Graph (SRKG). The SRKG records the issuing authority, document and clause identifier, edition, effective date, supersession relation, contractual scope, controlled terminology, applicability conditions, formal constraint, test cases and expert-approval status. Requirement-extraction and code-graph research provides mechanisms for candidate formalization (J. Zhang and El-Gohary 2017; P. Zhou and El-Gohary 2017; Y. Zhou, Lin, and She 2021), but mechanical rules must be reconstructed and validated against the applicable ISO, ASME, GB/T, customer and enterprise requirements. Representative controlled sources include ASME Y14.5 and the applicable editions of ISO 286, ISO 1101 and ISO 5459 (ASME 2018; ISO 2010a; ISO 2010b; ISO 2017; ISO 2024). A clause-grounded claim links the drawing entity and evidence locator to the applicable requirement, executed rule, severity, uncertainty state and decision status.

An LLM has a constrained supporting role in Cognition. It may formulate retrieval queries, rank candidate clauses, normalize terminology, parse a requirement into a candidate schema or draft an explanation from already verified evidence (Fuchs et al. 2024; Yang and Zhang 2024). It is not the authority for numerical or logical conformance. Candidate rules are admitted only after source verification, schema checking, unit and symbol checks, boundary test cases, and domain approval. Numerical comparisons, Boolean conditions, topology checks and unit conversions are executed by a version-controlled deterministic engine. Missing context, competing clauses, unsupported constructs or an explanation without a drawing locator, controlled clause and executed test produce an indeterminate or rule-review state.

4.3. Collaboration layer

The Collaboration Layer coordinates typed audit tasks and evidence without granting an orchestrator or report generator authority to invent missing facts. Task-allocation and multi-agent research provides mechanisms for decomposition, monitoring, reassignment and communication (Briseno et al. 2024; Graham, Decker, and Mersic 2003; Guo et al. 2024; Park and Sugumaran 2005). Manufacturing studies further show that task allocation, collaboration quality and human-automation symbiosis must be evaluated through workload, quality, safety, flexibility, cost and process state rather than agent count alone (Kokotinis et al. 2023; Peruzzini, Prati, and Pellicciari 2024; Petzoldt, Harms, and Freitag 2023). The Coordinator manages dependencies, timeouts, version identity, duplicate claims and completion state, while specialist roles are activated only when their required evidence is available.

Every specialist returns a typed finding containing a claim identifier, drawing-entity identifier, evidence locator, clause and edition, applicability assumptions, executed result, severity, uncertainty state, software or agent version, and recommended disposition. Conflicts are classified as fact, rule-applicability, scope, result or authority conflicts. Resolution follows evidence completeness, authoritative source and edition, deterministic re-check, independent domain review, and authorized human release; majority voting cannot overrule an unresolved safety-critical conflict. Material, process, structural-strength, tolerance-stack or assembly-feasibility tasks are not inferred from a drawing alone when their required CAD/CAE, bill-of-material, load, material, process or operating data are absent. Table 3 summarizes proposed role authority, required evidence, prohibited overreach and escalation paths.

4.4. Framework characteristics and implementation Perspectives

The P–C–C framework retains a modular architecture in which each layer has a distinct responsibility and communicates through explicit interfaces. This modularity may support incremental replacement or addition of models, agents, and standards, provided that schema versions, provenance, conformance tests and backward compatibility are maintained. Graph-based representations can improve traceability because relationships among drawing observations and requirements are explicit, but interpretability still depends on correct entity alignment, rule coverage and evidence completeness.

The distributed agent structure may support parallel auditing across drawing components or versions, but scalability and efficiency must be measured rather than assumed. Human-AI interaction preserves accountable expert authority for severe, ambiguous, conflicting or out-of-scope findings. Agentic-AI concepts may inform collaboration design, but they do not remove the need for bounded authority and empirical human-factors evidence (Yan 2025). The appropriate balance is risk-adaptive: validated low-risk tasks may be sampled, while high-risk or low-confidence tasks require human release. Review burden, override rates, alert fatigue, automation bias and skill degradation must be evaluated alongside technical accuracy.

By integrating perception, rule-grounded cognition and controlled collaboration, the proposed P–C–C reference architecture provides a candidate technical structure for future AI-assisted mechanical drawing auditing systems supporting precision-machine design and manufacturing. Its contribution is to make evidence flow, interface obligations, safe states and human

Table 3. Proposed audit roles, authority, evidence and escalation.

Role	Authorized task and authority limit	Evidence required	Disposition or escalation
Coordinator	Maintain scope, dependencies, versions, duplicate handling and completion state; do not adjudicate engineering truth by vote	Audit manifest, task graph, claim identifiers, evidence links and severity policy	Block incomplete output; reassign failed work; escalate unresolved conflicts
Source and perception assurance reviewer	Check carrier route, file integrity, schema, image quality, locators and uncertainty attached to candidate observations	Source files and revisions; parser or model identity; quality flags; and calibration evidence when probabilistic confidence is used	Request reprocessing or source confirmation; use NEEDS_RECOGNITION_REVIEW
Dimension, tolerance and GD&T reviewer	Check transcription, units, feature association, FCF syntax, datum relations, tolerances and fits only within supplied evidence	Verified drawing entities and view/feature relations; applicable clauses and deterministic operands when a rule-based decision is requested	Return NEEDS_EVIDENCE or NEEDS_RULE_REVIEW when signs, units, relations, applicability or editions are uncertain
Assembly and bill of materials (BOM) reviewer	Check visible identifiers, quantities, callouts and cross-document links without inferring missing parts from incomplete context	Visible source regions and revision identifiers for transcription; complete BOM scope, item associations and authoritative component-presence records for an error or omission decision	Return NEEDS_CONTEXT; preserve contested cross-document findings
Standards and rule curator	Maintain authorized sources, editions, applicability metadata, formal constraints, approvals and regression tests	Authorized standards source, supersession record, rule implementation and test suite	Return NEEDS_RULE_REVIEW for stale, ambiguous or unapproved rules
Context-dependent domain engineering reviewer	Assess material, process, strength or CAE questions only when the required model and operating context are supplied	Material/process records or CAD/CAE model, loads, boundary conditions and approved method	Return NEEDS_EVIDENCE or record the task as out of scope rather than infer missing engineering context
Conflict and assurance reviewer	Classify evidence, rule, scope, result and authority conflicts and perform an independent deterministic re-check	Competing claims, provenance, rule identity, severity and override history	Return NEEDS_RULE_REVIEW when no approved rule is available; escalate safety-critical or unresolved disagreement to the Authorized engineering signatory
Data governance and security steward	Control access, lineage, licences, retention, de-identification, model/rule package signatures and incident response	Authority records, access logs, data rights, model/rule signatures and security policy	Quarantine unauthorized or malicious input; invalidate or roll back affected workflow state
Authorized engineering signatory	Review consequential or sampled findings and release only claims contained in the approved machine-readable record	For a proposed release: complete evidence object; applicable clause where required by the claim; residual risk, alternatives, conflicts and override history	Release, return for rework or abstain with a recorded reason

responsibility explicit. Automation, interoperability, reliability and standardization remain evaluation objectives rather than established properties.

5. Discussion

5.1. Evolution and Framework overview

The evolution of AI-assisted auditing for mechanical engineering drawings reflects convergence among document vision, natural language processing, knowledge representation and distributed artificial intelligence. No single paradigm fully addresses compliance checking because each audit claim requires a controlled source, an observed entity or value, its relation and context, an applicable versioned requirement, an executed check, uncertainty and conflict disposition, and human-verifiable evidence. The P–C–C reference architecture organizes these dependencies without assuming that integration has already succeeded.

Historically, rule-based systems offered reproducible execution but were expensive to author and maintain, while object-oriented product models improved access to structured design evidence. AEC model-exchange, BIM and digital-twin reviews illustrate both the value of structured records and persistent implementation barriers (Eastman et al. 2010; Kaur et al. 2025; Opoku et al. 2021; Volk, Stengel, and Schultmann 2014). Process-image and knowledge-graph research shows how observations can be linked to structured state, while construction surveys emphasize extraction and evaluation limits (Bian 2025; Choi and Jung 2025; Pfitzner, Braun, and Borrmann 2024). Recent manufacturing syntheses position artificial intelligence and digital-twin-enabled automation within integrated-manufacturing systems (Webb, Tokhi, and Alkan 2025; S. Zhou and Bacanin 2024). These adjacent results support interface design and evaluation questions, but their schemas and benchmark outcomes are not mechanical compliance evidence.

The proposed P–C–C framework synthesizes these paradigms into a cohesive reference architecture. The Perception layer converts source material into candidate structured observations; the Cognition layer combines versioned knowledge with deterministic rules and bounded language assistance; and the Collaboration layer distributes domain tasks and reconciles only evidence-compatible findings under a defined human-release policy. This design separates responsibilities, enables uncertainty and conflict to remain visible and permits later testing of parallel execution. Compared with a linear rule pipeline, it adds feedback, abstention, provenance and controlled human escalation, but its end-to-end performance, scalability and reliability remain to be established.

To operationalize the framework, the data flow and interface definitions are specified across layers. The Perception layer outputs a DKG containing candidate features, dimensions, datums, symbols, annotations, views and relations, each with an identifier, source location, value or geometry, unit where applicable, extraction route, quality flag and uncertainty. The Cognition layer aligns these objects with versioned SRKG clauses and approved formal constraints. Each inference retains the clause identifier, edition, applicability record, executed operands, intermediate result, decision state and evidence locator. The Collaboration layer receives typed claims, detects duplicates and conflicts, and routes unresolved or high-risk findings to human release. These contracts are design requirements for controlling error propagation; their thresholds and effectiveness require empirical calibration.

5.2. Challenges and implementation considerations

Several failure modes can propagate across the architecture. Degraded scans, non-standard symbols, OCR sign or decimal errors, lost feature associations, stale standard editions, missing engineering context, conflicting requirements, unsupported LLM inferences, duplicate agent claims, timeouts, and report-generation drift can each convert a plausible intermediate output into an incorrect finding. The workflow therefore retains candidate alternatives and source regions, validates schema and units, checks requirement edition and applicability, executes approved constraints deterministically, blocks incomplete reports and uses explicit NEEDS_EVIDENCE, NEEDS_RECOGNITION_REVIEW, NEEDS_RULE_REVIEW, NEEDS_CONTEXT, CONTESTED, NEEDS_HUMAN_RELEASE and ABSTAINED states to route incomplete, conflicting or release-pending findings (Briseno et al. 2024; Graham, Decker, and Mersic 2003; Park and Sugumaran 2005; X. Xu and Cai 2020).

Table 4 summarizes the principal failure modes, their effects, detection controls, safe responses, responsible roles and evidence required to estimate residual risk. Risk scores require implementation and operational evidence.

Industrial deployment also depends on data governance, interpretability, workflow fit, cybersecurity, responsibility and economics. Recent manufacturing research links digital governance and security controls to Industry 4.0 maturity and privacy risk, while explainable-AI reviews show that interpretable outputs still require task-specific evaluation (Puthanveettil Madathil et al. 2025; Saunila et al. 2024; Tang, Zeng, and Zeng 2025). Mechanical drawings can contain intellectual property, export-controlled details, signatures, personal identifiers and customer information. A deployment therefore requires lawful authority, data minimization, access control, encryption, retention and deletion rules, source and licence records, de-identification, signed and versioned model and rule packages, least-privilege services, authenticated agent messages, incident response, and rollback. Shadow-mode evaluation should precede decision support, and selective automation should be limited to narrowly specified tasks with stable input distributions, validated thresholds, drift monitoring and periodic independent audit. Human-only and partially automated baselines must include false-positive burden, missed critical violations, review time, retraining, governance and downtime costs.

Standardization and interoperability remain decisive for adoption because mechanical workflows use heterogeneous carriers such as DWG, DXF, PDF, native CAD, STEP AP242 and QIF. A practical pathway requires carrier-specific adapters, a versioned canonical evidence record, DKG and SRKG interface contracts, conformance and round-trip tests, and controlled PLM write-back rather than an assumed universal schema. Future benchmarks should evaluate symbol and region detection with class-aware localization measures; text, dimensions and relations with character, word, numerical, unit and relation accuracy; rule retrieval and grounding with clause recall, citation precision, and applicability accuracy; deterministic execution with constraint accuracy and unsupported-inference rates; collaboration with coverage, conflict, escalation and omission measures; and end-to-end auditing with severity-stratified recall, false-negative rate, precision, false-positive burden and evidence traceability. Calibration, risk-coverage, human-review burden, latency, throughput, compute, energy or cost, and timeout behavior should be reported separately rather than collapsed into a single score.

Methods developed for AEC compliance offer useful mechanisms but not direct mechanical evidence. BIM and mechanical auditing differ in source carriers,

Table 4. Qualitative failure-control matrix for the proposed P–C–C workflow.

Failure mode	Potential effect	Detection / control	Safe response / owner	Evidence required to characterize residual risk
Degraded scan or non-standard symbol	Missed or wrong entity	Image-quality gate; candidate retention; out-of-distribution (OOD) check	Reprocess or route to NEEDS_RECOGNITION_REVIEW; Source and perception assurance reviewer	Stratified detection and selective-risk study
OCR numeral, sign, decimal or unit error	Wrong tolerance / dimension operand	Independent parser or transcription cross-check where available; range, unit and source-locator checks	Block rule execution; Dimension, tolerance and GD&T reviewer confirms source	Numeric exact-match and critical-error study
Lost CAD / PMI feature association	Rule applied to wrong feature	Native-semantic, cross-view and feature-relation validation	Reject link or request source model; Source and perception assurance reviewer + Dimension, tolerance and GD&T reviewer	Relation ground truth and round-trip tests
Missing, stale or wrong standard edition	Invalid or superseded decision	SRKG edition, effective-date, supersession and regression checks	NEEDS_RULE_REVIEW; Standards and rule curator	Version-conflict and update-regression tests
Ambiguous applicability or conflicting requirements	Inconsistent findings	Applicability record; authority hierarchy; independent re-check	CONTESTED; Conflict and assurance reviewer + Authorized engineering signatory	Expert-adjudicated conflict corpus
Unsupported LLM inference or fabricated citation	Plausible but ungrounded claim	Controlled-source retrieval; structured output; source-grounding check; deterministic re-check of operands and rules	Reject and log the candidate; NEEDS_RULE_REVIEW if no grounded alternative is available; Standards and rule curator + Conflict and assurance reviewer	Unsupported-source and citation-correctness rates
Duplicate or contradictory specialist/tool claims	Double count or inconsistent report	Claim IDs; provenance graph; duplicate and conflict checks	Merge only evidence-identical duplicates; otherwise mark CONTESTED and escalate; Coordinator + Conflict and assurance reviewer	Injected-conflict and arbitration study
Missing clause, feature-datum binding or measurement evidence	Unsupported standards or physical-conformance claim	Task-precondition schema; clause, edition, applicability, feature-datum-binding and measurement-evidence gates	Route missing or unvalidated rules to NEEDS_RULE_REVIEW and missing feature-datum or measurement evidence to NEEDS_EVIDENCE; Dimension, tolerance, and GD&T reviewer + Standards and rule curator; record ABSTAINED while a required prerequisite remains unresolved	Clause, binding, and measurement precondition-completeness audit
Missing bill of materials, revision or engineering context	Unsupported BOM-error, assembly, strength or process claim	Complete bill of materials, revision, component-presence, model, load and process-context checks	NEEDS_CONTEXT or out of scope; Assembly and bill of materials (BOM) reviewer or Context-dependent domain engineering reviewer	Context-precondition -completeness audit
Tool timeout or partial workflow	Incomplete audit shown as complete	Heartbeats; dependency state; completion manifest	Block release; resume or reassign; Coordinator	Load, timeout, and recovery tests
Prompt injection or malicious drawing text	Leakage or altered behavior	Content isolation; least privilege; instruction-data separation; red-team tests	Quarantine input and invalidate or roll back affected workflow state; Data governance and security steward	Threat-model and adversarial-test results
Report-generation drift	Unsupported conclusion in report	Machine-readable report-to-released-claim comparison	Block and regenerate; Coordinator + Authorized engineering signatory	Claim-consistency test suite

object schemas, rule forms, compliance subjects, evidence granularity and acceptable decision authority. Clause provenance, rule formalization, ontology alignment, task decomposition, audit logging and spatial-data interoperability may transfer as principles (Hjelmager et al. 2008; X. Xu and Cai 2020), whereas BIM/IFC schemas, building-code rule libraries and AEC benchmark results do not transfer as mechanical implementations or performance evidence. Mechanical adaptation requires reconstruction around CAD/MBD, STEP AP242, graphical and semantic PMI, datum reference systems, tolerance zones, fits, units, cross-view relations, assemblies and versioned ISO, ASME, GB/T, customer and enterprise requirements.

Table 5 distinguishes transferable principles from mechanisms requiring mechanical reconstruction and maps each transfer decision to a corresponding interoperability or validation test.

5.3. *Future research directions and socio-technical outlook*

Future research should investigate how LLM assistance and structured knowledge can reduce rule-authoring effort without transferring decision authority away from approved sources and deterministic constraints. Multimodal representations that connect symbols, geometry, text, DfM and DfA may improve evidence acquisition, but they require source-grouped datasets and

Table 5. Adjacent-domain transfer and interoperability gates for mechanical drawing audit.

Adjacent-domain mechanism or design principle	Mechanical transfer status	Mechanical adaptation / interoperability path	Evidence required before operational use
Clause provenance, authority, version and applicability	Principle transferable	Map applicable ISO, ASME, GB/T, customer and enterprise requirements; record issuing authority, edition, effective date, supersession, contractual scope and applicability in the SRKG	Clause-retrieval, applicability, authority-and-edition, supersession and update-regression tests
Requirement parsing and candidate formalization	Mechanical adaptation required	Extend beyond prose to tables, figures, GPS/PMI symbols, modifiers, datum systems, units, defaults and normative exceptions; retain each LLM output as a source-linked candidate without decision authority	Expert rule annotations, executable boundary cases and source-citation checks
BIM/IFC object schemas and exchange mappings	Mechanical reconstruction required	Rebuild around adapters for PDF, DWG/DXF, native CAD, STEP AP242, QIF and graphical or semantic PMI, feeding a canonical evidence record and DKG	Schema conformance, feature identity, semantic completeness and round-trip loss tests
Spatial and topological evidence linking	Mechanical adaptation required	Add mechanical feature association, datum sequence, tolerance zone, view projection, dimension chain, fit and assembly relations	Expert relation ground truth and cross-view or model-consistency tests
Domain ontology and rule library	Mechanical reconstruction required	Align DKG observations with versioned SRKG requirements; distinguish asserted observations, inferred relations, and approved rules; maintain rule test suites	Coverage, consistency, expert agreement and rule-regression evidence
Task decomposition and orchestration	Mechanical adaptation required	Define mechanical role inputs, permitted tools, typed outputs, authority boundaries, dependencies, timeouts, conflict classes and safe hand-offs	Ablation, conflict, latency, recovery, and human-escalation experiments
Human oversight and audit logging	Principle transferable	Map severity tiers, sampling, overrides, competence, release authority, accountability and record retention to the engineering quality system; treat liability allocation and regulatory acceptance as organization- and jurisdiction-specific deployment gates	Workflow studies with practicing engineers; workload, disagreement, override, alert-fatigue and skill-retention measures; organization- and jurisdiction-specific governance review
Cross-tool APIs and PLM write-back	Implementation-specific adaptation	Connect source adapters to the canonical evidence record and DKG/SRKG contracts through conformance and round-trip gates before controlled PLM write-back; preserve original locators and revisions	Vendor and version matrix, API contract tests, rollback, audit-trail verification and data-loss analysis
AEC benchmark results	Not transferable as performance evidence	Create independent mechanical benchmarks with expert-adjudicated GD&T, PMI, tolerance and assembly reference truth and relevant human and technical baselines	External mechanical validation using source-family-grouped partitions, uncertainty quantification and failure analysis

relation-level ground truth. Multi-agent systems should be tested for task allocation, conflict detection, latency, correlated errors and human escalation rather than assumed to function as proactive teammates. These directions should be evaluated through explicit baselines, failure injection, calibration and independent mechanical data.

The socio-technical dimension remains central to any transition from research prototypes to industry-grade systems. Co-development among academia, standards bodies, software vendors, data owners and practicing engineers is needed for shared data definitions, transparent validation and governance. Continuous oversight must include standards and model version control, access and retention controls, cybersecurity, incident response, rollback, audit logs, and named responsibility for data, rules, systems, security and final engineering release. Deployment should progress from shadow-mode observation to bounded decision support; selective automation should be considered only for narrowly defined tasks with validated thresholds, monitored drift and periodic independent audit.

5.4. Conceptual grounding and empirical validation requirements

This study evaluates the conceptual grounding, interface requirements and evidence boundaries of the reference architecture; it does not report an implemented end-to-end system. Literature mapping and worked conditions define testable requirements for subsequent implementation and validation.

The Merge_places Computer Vision Dataset (Symbol Detection 2024) contains 742 paired JPEG drawings and TXT annotation files, with all images stored as 2000×2000 RGB JPEGs. The TXT files contain 37,364 valid five-token normalized boxes across 12 anonymous numerical class identifiers and one 11-token polygon-like record, indicating a mixed annotation contract. The absence of a class map, annotation guide, source/licence record, predefined split, engineering relations, compliance labels and model outputs limits the present archive analysis to schema, lineage, duplication and evidence-routing behaviour.

Byte-level duplicate checking identified 111 byte-identical image pairs (222 files); 75 pairs were associated with different label-file bytes. Because source names also indicate multi-page and transformed variants, future perception experiments should consolidate duplicate contents, group source families, adjudicate label conflicts and freeze an external test set before partitioning. Cases without verified class semantics, applicable requirements, engineering relations or compliance labels would require rule review and abstention before any compliance claim.

Empirical validation should compare carrier-aware perception with raster-only processing, constrained LLM assistance with retrieval-only and unconstrained generation, deterministic rule execution with language-model adjudication, and the complete P-C-C workflow with sequential, single-agent, rule-based, partially automated and experienced-human baselines. The tests should include degraded scans, altered numerals or units, ambiguous symbols, missing views, stale or conflicting standards, unavailable agents, malicious drawing text, and missing engineering context. Results should report confidence intervals and severity-stratified error budgets and should identify outcomes that would disconfirm an architectural assumption. Until such evidence is available, the framework remains a pre-validation reference architecture and research agenda.

Table 6 specifies a prospective evaluation design for Perception, Cognition, Collaboration, human-governed operation and end-to-end auditing, including independent units, baselines, expert-adjudicated reference truth, measures, stratification and uncertainty analysis. Archive-scale studies should use source-family-aware sampling and report failure and abstention behavior separately.

Figure 2 maps a proposed route in which the available source and annotation records would pass through schema and lineage screens. Unresolved semantic or clause prerequisites would be assigned to the

Table 6. Prospective validation matrix for the P-C-C architecture.

Layer or task	Independent unit / reference truth	Baselines	Primary measures	Required assurance and stratification
Perception: symbols / regions	Source-grouped drawing; expert-adjudicated boxes or regions	Native parser where available; OCR / detector baselines	mAP at stated IoU; macro-F1; per-class precision / recall; localization error	Carrier, scan quality, object scale, class frequency and source family; source-family bootstrap CI
Perception: text / dimensions / relations	Drawing entity; verified transcription, units and relation graph	OCR and layout models; rule-based association	CER; WER; numeric exact match; unit accuracy; entity-relation F1	Numeral versus text subset, orientation, density, image quality and source family; explicit denominator and source-family bootstrap CI
Cognition: retrieval / grounding	Entity-clause pair; expert applicability judgement	Keyword / BM25; retrieval-only; less-constrained LLM	Clause recall@k; citation precision; applicability precision / recall; unsupported-source rate	Standard family and edition, clause type, exceptions, and conflicts; source-family CI and expert disagreement
Cognition: rule execution	Executable case; verified operands and expected result	Manual expert calculation; available deterministic checker; unconstrained generative baseline	Constraint-decision accuracy; numerical error; invalid-execution rate; unsupported-inference rate for generative baselines	Severity, missing inputs, units, boundary values and exceptions; case-family CI
Collaboration	Multi-claim case; adjudicated tasks and conflicts	Sequential Perception-Cognition (P-C); single agent; no-conflict-policy ablation	Task-completion rate; duplicate- and missed-claim rates; conflict-detection precision/recall; disposition agreement with the adjudicated reference; appropriate-escalation rate	Role/agent count, conflict type, task coupling, failure injection and source family; CI and adjudicator disagreement
End-to-end audit	Drawing or assembly; multi-expert adjudicated violations	Human-only; deterministic rule-based; partially automated; P-C-C ablations	Critical-violation recall; severity-specific false-negative rate; precision; false-positive findings per drawing; complete-evidence-chain rate; review time	Severity, carrier, source family, standard and complexity; source-family CI, expert disagreement and an external test set
Selective / human-governed operation	Claim and release decision, including abstained cases	No-abstention policy; mandatory manual review; alternative preregistered thresholds	Risk-coverage; selective risk/accuracy; ECE/ Brier when probabilistic confidence scores are available; human-review rate and time; override and erroneous-release rates	Preregistered risk tiers, thresholds, and release policy; calibration plots; post-release sampling
Efficiency / scale	Drawing or assembly job under stated concurrency	Single pipeline; alternative agent schedules	Latency; throughput; peak memory and GPU utilization; energy and cost; inter-role messages and, where applicable, model tokens; timeout and recovery rates	Offline versus interactive target, drawing density, role/agent count, hardware/software configuration and recovery policy; repeated-run uncertainty

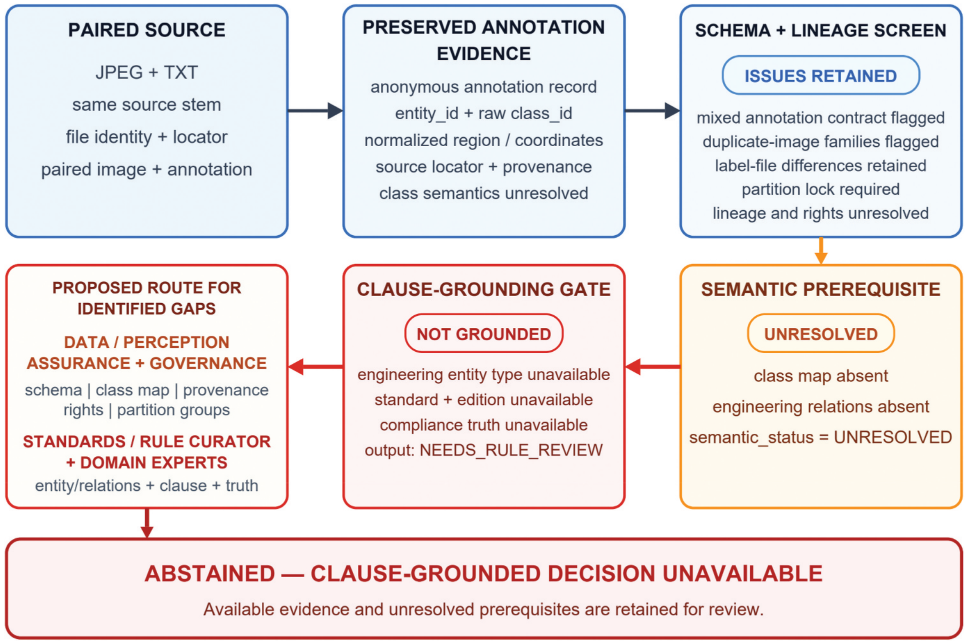

Figure 2. Evidence-boundary and safe-abstention workflow for the merge_places Computer Vision Dataset.

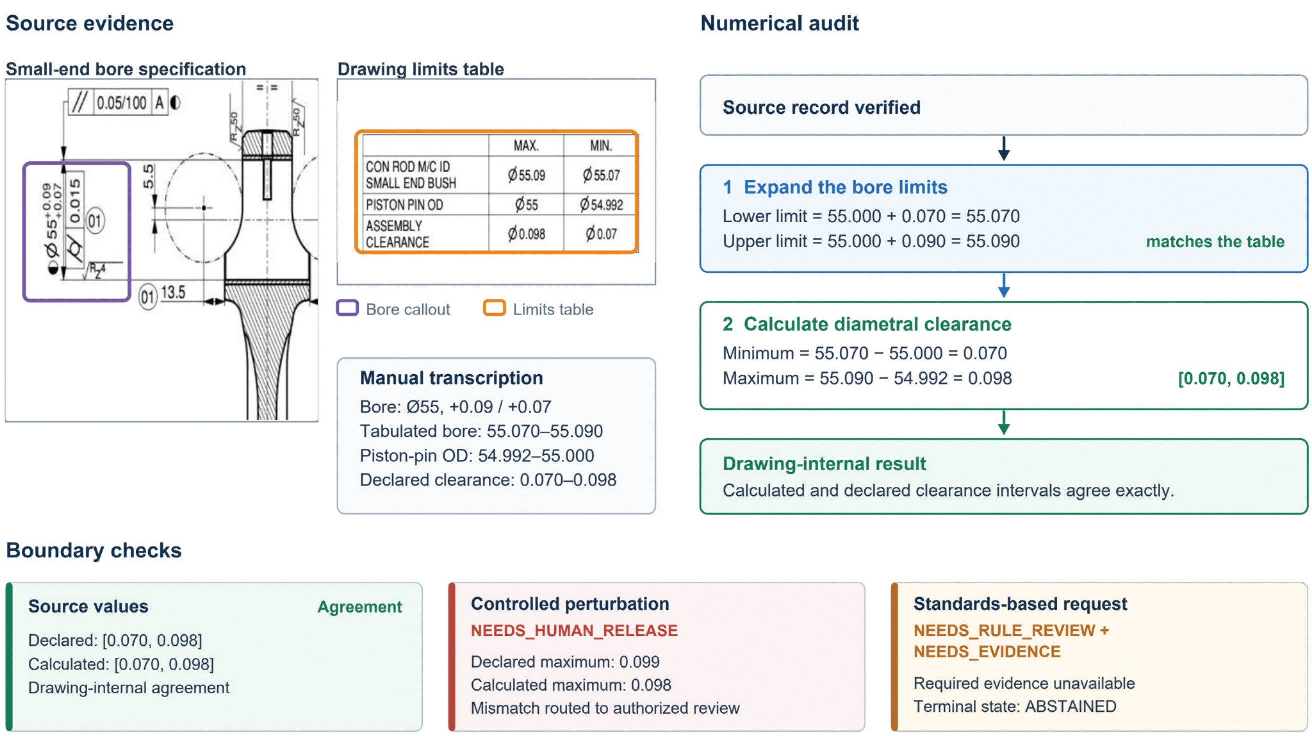

Figure 3. Drawing-internal diametral-clearance audit.

corresponding assurance, governance, standards and domain-review roles before an abstained decision is recorded.

A case-insensitive search for CONNECTING-ROD or CONROD identified six files representing four unique source families. The eligibility criteria required co-located small-end bore limits, mating piston-pin outside-diameter limits and a drawing-declared diametral-clearance interval; one family met all three. Figure 3 presents two de-identified regions from that drawing. This targeted selection supports the worked arithmetic case but is not an archive-wide sample.

The bore callout was transcribed as 55 with deviations +0.09/+0.07; the drawing table gives a bore interval of [55.070, 55.090], a piston-pin outside-diameter interval of [54.992, 55.000] and a declared diametral-clearance interval of [0.070, 0.098]. In the source drawing units, worst-case subtraction gives Cmin = 55.070–55.000 = 0.070 and Cmax = 55.090–54.992 = 0.098, reproducing the declared interval after trailing-zero normalization.

The same arithmetic was evaluated under two additional conditions. A controlled change of the declared maximum to 0.099 creates a mismatch; the corresponding proposed safe route is NEEDS_HUMAN_RELEASE pending authorized review. A standards-based request without class mapping or clause grounding is assigned the proposed routes NEEDS_RULE_REVIEW and NEEDS_EVIDENCE, with ABSTAINED as the terminal state if those prerequisites remain unresolved. These checks complement the source-value condition; archive-wide accuracy and prevalence require the study design in Table 6.

Figure 4 presents two evidence-boundary cases selected from six candidate source families (16 raster records; 11 unique byte contents) that were screened for examples distinct from the clearance check. One case contains a visible feature-control frame without clause and measurement context; the other contains a bill of materials (BOM) numbering transition without authoritative numbering and revision context. Selection was purposive and does not estimate archive-wide prevalence.

Figure 4a supports transcription of the visible feature-control-frame tokens but not a clause or measurement conclusion. Its proposed safe routes are NEEDS_RULE_REVIEW and NEEDS_EVIDENCE, with ABSTAINED as the terminal state if the required prerequisites remain unresolved. Figure 4b records the visible BOM rows 63, 62, 60 and 59 and the local 62-to-60 transition; its proposed safe route is NEEDS_CONTEXT. The authoritative numbering policy, revision history and component-presence record are required before classifying that transition as an error. Table 7 summarizes the corresponding evidence requirements, permissible findings and proposed safe routes.

Table 7 compares the evidence, check, permissible finding and proposed safe route across the three Figure 3 conditions and the two Figure 4 observations. Dimensional values are retained in source drawing units, and the Figure 4b typeset BOM is a source-linked transcription. The five purposively selected conditions define evidence and decision boundaries rather than archive-wide frequencies.

6. Conclusions

In this paper, the progression of AI-assisted auditing review of mechanical engineering drawings is critically highlighted from deterministic rules to data-driven perception, semantic reasoning and coordinated intelligence, particularly in relation to each respective technology with the specific segments of the auditing review process chain.

The proposed P-C-C reference architecture links source-controlled evidence, versioned requirements,

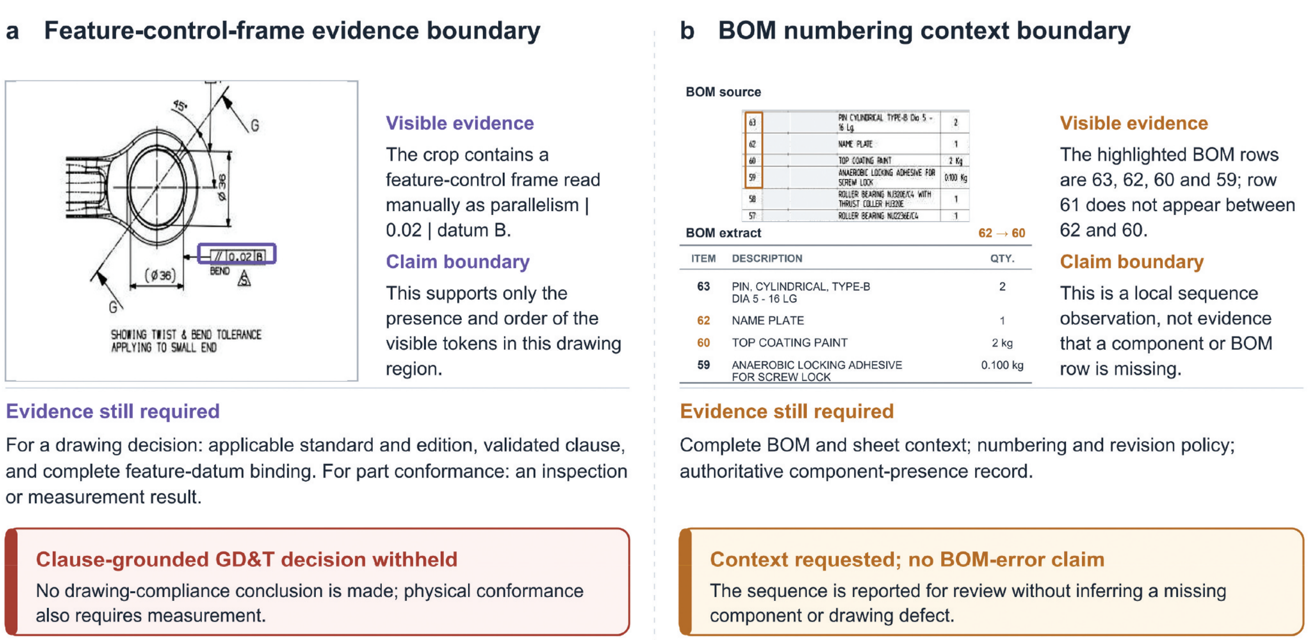

Figure 4. Source-grounded evidence boundaries in the merge_places Computer Vision Dataset.

Table 7. Evidence boundaries and proposed safe routes across Figures 3 and 4.

Condition	Source or test condition	Check or evidence gap	Permissible finding or proposed safe route	Not supported by this condition
Figure 3: source values	De-identified bore and limits-table crops; bore [55.070, 55.090], pin [54.992, 55.000], and declared clearance [0.070, 0.098] in source drawing units	Exact subtraction gives [0.070, 0.098], equal to the declared interval after trailing-zero normalization	Permissible finding: drawing-internal numerical agreement for the transcribed values.	Part conformance, standards compliance, measurement acceptance, system-level accuracy, archive prevalence or generalization
Figure 3: controlled perturbation	Test input changes only the declared maximum from 0.098 to 0.099; the source drawing is unchanged	Calculated maximum remains 0.098, so the test input and calculation disagree	Proposed safe route: NEEDS_HUMAN_RELEASE; an agreement finding is not eligible for release pending authorized review.	A naturally observed defect, defect rate, model performance or autonomous collaboration
Figure 3: standards-based request	The arithmetic evidence is present, but class mapping, the applicable standard and edition, clause grounding and measurement evidence if physical conformance is requested are unavailable	The numerical check cannot establish a standards or physical-conformance decision	Proposed safe route: NEEDS_RULE_REVIEW + NEEDS_EVIDENCE; terminal state if prerequisites remain unresolved: ABSTAINED.	GD&T or clause compliance, physical conformance, manufacturing acceptance or industrial suitability
Figure 4a: feature-control -frame (FCF) observation	A de-identified crop shows a feature-control frame manually transcribed as parallelism | 0.02 | datum B	Applicable standard and edition, validated clause, complete feature/datum binding and any physical measurement record are unavailable	Permissible finding: the visible token sequence only. Proposed safe route: NEEDS_RULE_REVIEW + NEEDS_EVIDENCE; terminal state if prerequisites remain unresolved: ABSTAINED.	Drawing-clause compliance from the missing clause/ binding evidence or physical conformance from the absent measurement evidence
Figure 4b: BOM observation	The source crop and source-linked transcription show rows 63, 62, 60 and 59; within this visible sequence, row 61 does not appear between 62 and 60	The authoritative numbering policy, revision/deletion history and component-presence record are unavailable	Permissible finding: the local visible sequence 63, 62, 60 and 59. Proposed safe route: NEEDS_CONTEXT; a BOM-error claim is not eligible for release without authoritative context.	A missing BOM row or component, a BOM error, nonconformance, archive prevalence or system performance

deterministic rule gates, uncertainty, conflict escalation, the authorized engineering approval and release, aiming at fulfilling the goals of design for manufacturing and design for assembly in a perfect applicable precision engineering manner. Its interfaces define independently testable responsibilities for Perception, Cognition and Collaboration.

The archived auditing review on precision engineering design exemplars shows how provenance, deterministic arithmetic, evidence prerequisites and abstention can be represented within selected P-C-C interfaces. Implementation studies with independent mechanical design benchmarks, validated standards mappings, interoperability tests, uncertainty analysis, cybersecurity evaluation, engineer baselines and workflow evaluation are required to establish system-level performance and precision engineering usage.

The research approach and results presented further demonstrate that the AI-assisted auditing review of mechanical engineering drawings can combine source-linked perception, standards-grounded rule execution and accountable collaboration within a traceable quality workflow applicable to the precision engineering design process. It has great application potentials for future engineering manufacturing industries.

Author contributions

CRediT: **Xin Chen:** Conceptualization, Data curation, Formal analysis, Funding acquisition, Investigation, Methodology, Resources, Software, Validation, Visualization, Writing – original draft, Writing – review & editing; **Kai Cheng:** Conceptualization, Data curation, Formal analysis, Investigation, Methodology, Resources, Software, Supervision, Validation, Visualization, Writing – review & editing.

Disclosure statement

No potential conflict of interest was reported by the author(s)

Funding

This study was supported by the grant from the basic research projects of educational department of Liaoning province [Grant No. LJ212411035018] and China Scholarship Council (CSC) under Grant [202408210337]. The authors gratefully acknowledge the support provided by the funding agencies and host institutions.

Data availability statement

The data presented in this study are available in [Merge_places Computer Vision Dataset] at [https://universe.roboflow.com/symboldetection-4xwmm/merge_places], accessed on 18 July 2025.

References

Acierno, Marta, Stefano Cursi, Davide Simeone, and Donatella Fiorani. 2017. "Architectural Heritage Knowledge Modelling: An Ontology-Based Framework for Conservation Process." *Journal of Cultural Heritage* 24:124–133. https://doi.org/10.1016/j.culher.2016.09.010.

Amor, Robert, and Johannes Dimyadi. 2021. "The Promise of Automated Compliance Checking." *Developments in the Built Environment* 5:100039. https://doi.org/10.1016/j.dibe.2020.100039.

Appalaraju, Srikar, Bhavan Jasani, Bhargava Urala Kota, Yusheng Xie, and R. Manmatha. 2021. "DocFormer: End-To-End Transformer for Document Understanding." 2021 IEEE/CVF International Conference on Computer Vision (ICCV), Montreal, QC, Canada, 973–983. IEEE. https://doi.org/10.1109/iccv48922.2021.00103.

Aslam, Bilal, and Tariq Umar. 2024. "Automated Code Compliance Checking Through Building Information Modelling." *Proceedings of the Institution of Civil Engineers —Structures and Buildings* 177 (9): 822–837. https://doi.org/10.1680/jstbu.22.00214.

ASME (American Society of Mechanical Engineers). 2018. *Dimensioning and Tolerancing*. ASME Y14.5-2018 (R2024). New York: ASME. https://www.asme.org/codes-standards/find-codes-standards/y14-5-dimensioning-tolerancing/2018/pdf.

Azhar, Salman, Malik Khalfan, and Tayyab Maqsood. 2015. "Building Information Modelling (BIM): Now and Beyond." *Construction Economics and Building* 12 (4): 15–28. https://doi.org/10.5130/ajceb.v12i4.3032.

Baek, Youngmin, Bado Lee, Dongyoon Han, Sangdoo Yun, and Hwalsuk Lee. 2019. "Character Region Awareness for Text Detection." 2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), Long Beach, CA, USA, 9357–9366. IEEE. https://doi.org/10.1109/cvpr.2019.00959.

Beach, Thomas H., Jean-Laurent Hippolyte, and Yacine Rezgui. 2020. "Towards the Adoption of Automated Regulatory Compliance Checking in the Built Environment." *Automation in Construction* 118:103285. https://doi.org/10.1016/j.autcon.2020.103285.

Beltagy, Iz, Matthew E. Peters, and Arman Cohan. 2020. "Longformer: The Long-Document Transformer." arXiv preprint. https://doi.org/10.48550/ARXIV.2004.05150.

Bian, Haonan. 2025. "LLM-Empowered Knowledge Graph Construction: A Survey." *arXiv Preprint*. https://doi.org/10.48550/ARXIV.2510.20345.

Bickel, Sebastian, Stefan Goetz, and Sandro Wartzack. 2024. "Symbol Detection in Mechanical Engineering Sketches: Experimental Study on Principle Sketches with Synthetic Data Generation and Deep Learning." *Applied Sciences* 14 (14): 6106. https://doi.org/10.3390/app14146106.

Bloch, Tanya, and Rafael Sacks. 2020. "Clustering Information Types for Semantic Enrichment of Building Information Models to Support Automated Code Compliance Checking." *Journal of Computing in Civil Engineering* 34 (6): 04020040. https://doi.org/10.1061/(ASCE)CP.1943-5487.0000922.

Bommarito, Michael J., Daniel Martin Katz, and Eric M. Detterman. 2018. "LexNLP: Natural Language Processing and Information Extraction for Legal and Regulatory Texts." *SSRN Electronic Journal*. https://doi.org/10.2139/ssrn.3192101.

Borrmann, André, Markus König, Christian Koch, and Jakob Beetz. 2018. *Building Information Modeling: Technology Foundations and Industry Practice*. Cham: Springer International Publishing. https://doi.org/10.1007/978-3-319-92862-3.

Briseno, Julian de Gortari, Roko Parać, Leo Ardon, Marc Roig Vilamala, Daniel Furelos-Blanco, Lance Kaplan, Vinod K. Mishra, et al. 2024. "TeamCollab: A Framework for Collaborative Perception-Cognition-Communication-Action." 2024 27th International Conference on Information Fusion (FUSION), Venice, Italy, 1–8. IEEE. https://doi.org/10.23919/fusion59988.2024.10706435.

Brown, Tom B., Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, et al. 2020. "Language Models are Few-Shot Learners." *arXiv Preprint*. https://doi.org/10.48550/ARXIV.2005.14165.

Cao, Weijuan, Trevor Robinson, Yang Hua, Flavien Boussuge, Andrew R. Colligan, and Wanbin Pan. 2020. "Graph Representation of 3D CAD Models for Machining Feature Recognition with Deep Learning." Volume 11A: 46th Design Automation Conference (DAC), Virtual, Online, V11AT11A003. American Society of Mechanical Engineers. https://doi.org/10.1115/detc2020-22355.

Carrara, Andrea, Stavros Nousias, and André Borrmann. 2024. "VectorGraphNET: Graph Attention Networks for Accurate Segmentation of Complex Technical Drawings." *Journal of Computing in Civil Engineering* 39 (6). https://doi.org/10.1061/JCCEE5.CPENG-6508.

Chen, Nanjiang, Xuhui Lin, Hai Jiang, and Yi An. 2024. "Automated Building Information Modeling Compliance Check Through a Large Language Model Combined with Deep Learning and Ontology." *Buildings* 14 (7): 1983. https://doi.org/10.3390/buildings14071983.

Cheng, K., P. Y. Pan, and D. K. Harrison. 2001. "Web-Based Design and Manufacturing Support Systems: Implementation Perspectives." *International Journal of Computer Integrated Manufacturing* 14 (1): 14–27. https://doi.org/10.1080/09511920150214875.

Chidara, Alexander, Kai Cheng, and David Gallear. 2025. "Ontology-Based Modelling and Analysis of Sustainable Polymer Systems: PVC Comparative Polymer and Implementation Perspectives." *Polymers* 17 (19): 2612. https://doi.org/10.3390/polym17192612.

Choi, Seungmin, and Yuchul Jung. 2025. "Knowledge Graph Construction: Extraction, Learning, and Evaluation." *Applied Sciences* 15 (7): 3727. https://doi.org/10.3390/app15073727.

Chollet, François. 2017. "Xception: Deep Learning with Depthwise Separable Convolutions." 2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR), Honolulu, HI, USA, 1800–1807. IEEE. https://doi.org/10.1109/CVPR.2017.195.

Deng, Haoxuan, Samir Khan, and John Ahmet Erkoyuncu. 2024. "An Investigation on Utilizing Large Language Model for Industrial Computer-Aided Design Automation." *Procedia CIRP* 128:221–226. https://doi.org/10.1016/j.procir.2024.07.049.

Devlin, Jacob, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. "BERT: Pre-Training of Deep Bidirectional Transformers for Language Understanding."

Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers), Minneapolis, MN: 4171–4186. Association for Computational Linguistics. https://doi.org/10.18653/v1/N19-1423.

Dimyadi, Johannes, and Robert Amor. 2013. "Automated Building Code Compliance Checking—Where is it At?" Proceedings of the 19th CIB World Building Congress, 172–185. https://doi.org/10.13140/2.1.4920.4161.

Eastman, Charles M., Young-Su Jeong, Rafael Sacks, and Ilan Kaner. 2010. "Exchange Model and Exchange Object Concepts for Implementation of National BIM Standards." *Journal of Computing in Civil Engineering* 24 (1): 25–34. https://doi.org/10.1061/(ASCE)0887-3801(2010)24:1(25).

Fuchs, Stefan, Michael Witbrock, Johannes Dimyadi, and Robert Amor. 2024. "Using Large Language Models for the Interpretation of Building Regulations." *arXiv Preprint*. https://doi.org/10.48550/ARXIV.2407.21060.

Goodfellow, Ian, Yoshua Bengio, and Aaron Courville. 2016. *Deep Learning*. Cambridge, MA: MIT Press. https://www.deeplearningbook.org/ .

Gori, Marco, Gabriele Monfardini, and Franco Scarselli. 2005. "A New Model for Learning in Graph Domains." Proceedings of the 2005 IEEE International Joint Conference on Neural Networks, Montreal, QC, Canada, vol. 2, 729–734. IEEE. https://doi.org/10.1109/IJCNN.2005.1555942.

Graham, John R., Keith S. Decker, and Michael Mersic. 2003. "DECAF - A Flexible Multi Agent System Architecture." *Autonomous Agents and Multi-Agent Systems* 7 (1–2): 7–27. https://doi.org/10.1023/A:1024120703127.

Guo, Taicheng, Xiuying Chen, Yaqi Wang, Ruidi Chang, Shichao Pei, Nitesh V. Chawla, Olaf Wiest, and Xiangliang Zhang. 2024. "Large Language Model Based Multi-Agents: A Survey of Progress and Challenges." *arXiv Preprint*. https://doi.org/10.48550/ARXIV.2402.01680.

He, Junda, Christoph Treude, and David Lo. 2024. "LLM-Based Multi-Agent Systems for Software Engineering: Literature Review, Vision and the Road Ahead." *ACM Transactions on Software Engineering and Methodology* 34 (5): 1–30. https://doi.org/10.1145/3712003.

He, Kaiming, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2016. "Deep Residual Learning for Image Recognition." 2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR), Las Vegas, NV, USA, 770–778. IEEE. https://doi.org/10.1109/cvpr.2016.90.

Hitzler, Pascal, Markus Krötzsch, Bijan Parsia, Peter F. Patel-Schneider, and Sebastian Rudolph. 2012. *OWL 2 Web Ontology Language Primer*. Second Edition) ed. Cambridge, MA: World Wide Web Consortium. https://www.w3.org/TR/owl2-primer/.

Hjelmager, Jan, Harold Moellering, Antony Cooper, Tatiana Delgado, Abbas Rajabifard, Petr Rapant, David Danko, et al. 2008. "An Initial Formal Model for Spatial Data Infrastructures." *International Journal of Geographical Information Science* 22 (11–12): 1295–1309. https://doi.org/10.1080/13658810801909623.

Hosseini, M. Reza, Mojtaba Maghrebi, Ali Akbarnezhad, Igor Martek, and Mehrdad Arashpour. 2018. "Analysis of Citation Networks in Building Information Modeling Research." *Journal of Construction Engineering and Management* 144 (8): 04018064. https://doi.org/10.1061/(asce)co.1943-7862.0001492.

Hou, Shangjie, Haijiang Li, and Yacine Rezgui. 2015. "Ontology-Based Approach for Structural Design Considering Low Embodied Energy and Carbon." *Energy & Buildings* 102:75–90. https://doi.org/10.1016/j.enbuild.2015.04.051.

Huang, Yupan, Tengchao Lv, Lei Cui, Yutong Lu, and Furu Wei. 2022. "LayoutLMv3: Pre-Training for Document AI with Unified Text and Image Masking." Proceedings of the 30th ACM International Conference on Multimedia, Lisboa, Portugal, 4083–4091. ACM. https://doi.org/10.1145/3503161.3548112.

Ismail, Aimi Sara, Kherun Nita Ali, Noorminshah A. Iahad, Mukhtar A. Kassem, and Najib Taher Al-Ashwal. 2023. "BIM-Based Automated Code Compliance Checking System in Malaysian Fire Safety Regulations: A User-Friendly Approach." *Buildings* 13 (6): 1404. https://doi.org/10.3390/buildings13061404.

ISO (International Organization for Standardization). 2010a. *Geometrical Product Specifications (GPS)—ISO Code System for Tolerances on Linear Sizes—Part 1: Basis of Tolerances, Deviations and Fits*. ISO 286-1:2010. 2nd ed. International Organization for Standardization. https://www.iso.org/standard/45975.html.

ISO (International Organization for Standardization). 2010b. *Geometrical Product Specifications (GPS)—ISO Code System for Tolerances on Linear Sizes—Part 2: Tables of Standard Tolerance Classes and Limit Deviations for Holes and Shafts*. ISO 286-2:2010. 2nd ed. International Organization for Standardization. https://www.iso.org/standard/54915.html.

ISO (International Organization for Standardization). 2017. *Geometrical Product Specifications (GPS)—Geometrical Tolerancing—Tolerances of Form, Orientation, Location and Run-Out*. ISO 1101:2017. 4th ed. International Organization for Standardization. https://www.iso.org/standard/66777.html.

ISO (International Organization for Standardization). 2024. *Geometrical Product Specifications (GPS)—Geometrical Tolerancing—Datums and Datum Systems*. ISO 5459:2024. 3rd ed. International Organization for Standardization. https://www.iso.org/standard/87855.html.

Jamieson, Laura, Carlos Francisco Moreno-García, and Eyad Elyan. 2024. "A Review of Deep Learning Methods for Digitisation of Complex Documents and Engineering Diagrams." *Artificial Intelligence Review* 57 (6): 136. https://doi.org/10.1007/s10462-024-10779-2.

Järvenpää, Eeva, Niko Siltala, Otto Hylli, Hasse Nylund, and Minna Lanz. 2023. "Semantic Rules for Capability Matchmaking in the Context of Manufacturing System Design and Reconfiguration." *International Journal of Computer Integrated Manufacturing* 36 (1): 128–154. https://doi.org/10.1080/0951192X.2022.2081361.

Jiang, Feng, Ling Ma, Tim Broyd, Weiya Chen, and Hanbin Luo. 2022. "Digital Twin Enabled Sustainable Urban Road Planning." *Sustainable Cities and Society* 78:103645. https://doi.org/10.1016/j.scs.2021.103645.

Jin, Yan, and Raymond E. Levitt. 1996. "The Virtual Design Team: A Computational Model of Project Organizations." *Computational & Mathematical Organization Theory* 2 (3): 171–195. https://doi.org/10.1007/bf00127273.

Joffe, Isaac, Yuchen Qian, Mohammad Talebi-Kalaleh, and Qipei Mei. 2024. "A Computer Vision Framework for

Structural Analysis of Hand-Drawn Engineering Sketches." *Sensors* 24 (9): 2923. https://doi.org/10.3390/s24092923.

Karan, Ebrahim P., and Javier Irizarry. 2015. "Extending BIM Interoperability to Preconstruction Operations Using Geospatial Analyses and Semantic Web Services." *Automation in Construction* 53:1–12. https://doi.org/10.1016/j.autcon.2015.02.012.

Kashevnik, Alexey, Nikolay Shilov, Nikolay Teslya, Fudail Hasan, Andrey Kitenko, Veronika Dukareva, Marat Abdurakhimov, Alexander Zingarevich, and Dmitry Blokhin. 2023. "An Approach to Engineering Drawing Organization: Title Block Detection and Processing." *IEEE Access*: 1. https://doi.org/10.1109/ACCESS.2023.3244603.

Kaur, Ranjit, Benesi John Mwambegele, Ashish George Abraham, Shammas Abdul Basheer, and Siddharth Garia. 2025. "A Comprehensive Review on Building Information Modelling (BIM), Its Implementations and Applications." *Discover Civil Engineering* 2 (1): 177. https://doi.org/10.1007/s44290-025-00342-5.

Khan, Muhammad Tayyab, Lequn Chen, Ye Han Ng, Wenhe Feng, Nicholas Yew Jin Tan, and Seung Ki Moon. 2024. "Fine-Tuning Vision-Language Model for Automated Engineering Drawing Information Extraction." *arXiv Preprint*. https://doi.org/10.48550/ARXIV.2411.03707.

Khan, Muhammad Tayyab, Zane Yong, Lequn Chen, Jun Ming Tan, Wenhe Feng, and Seung Ki Moon. 2025. "Automated Parsing of Engineering Drawings for Structured Information Extraction Using a Fine-Tuned Document Understanding Transformer." *arXiv Preprint*. https://doi.org/10.48550/ARXIV.2505.01530.

Kim, Geewook, Teakgyu Hong, Moonbin Yim, Jeongyeon Nam, Jinyoung Park, Jinyeong Yim, Wonseok Hwang, Sangdoo Yun, Dongyoon Han, and Seunghyun Park. 2022. "OCR-Free Document Understanding Transformer." In *Computer Vision—ECCV 2022*, 498–517. Cham: Springer. https://doi.org/10.1007/978-3-031-19815-1_29.

Kipf, Thomas N., and Max Welling. 2017. "Semi-Supervised Classification with Graph Convolutional Networks." 5th International Conference on Learning Representations (ICLR 2017), Conference Track Proceedings, Toulon, France. OpenReview.net. https://openreview.net/forum?id=SJU4ayYgl.

Kitamura, Yoshinobu, and Riichiro Mizoguchi. 2004. "Ontology-Based Systematization of Functional Knowledge." *Journal of Engineering Design* 15 (4): 327–351. https://doi.org/10.1080/09544820410001697163.

Kokotinis, George, George Michalos, Zoi Arkouli, and Sotiris Makris. 2023. "On the Quantification of Human-Robot Collaboration Quality." *International Journal of Computer Integrated Manufacturing* 36 (10): 1431–1448. https://doi.org/10.1080/0951192X.2023.2189304.

Krizhevsky, Alex, Ilya Sutskever, and Geoffrey E. Hinton. 2017. "ImageNet Classification with Deep Convolutional Neural Networks." *Communications of the ACM* 60 (6): 84–90. https://doi.org/10.1145/3065386.

Kumar, Sumit, Sarthak Kapoor, Harsh Vardhan, and Yao Zhao. 2025. "Generative AI for CAD Automation: Leveraging Large Language Models for 3D Modelling." *arXiv Preprint*. https://doi.org/10.48550/ARXIV.2508.00843.

Lai, Huahui, and Xueyuan Deng. 2018. "Interoperability Analysis of IFC-Based Data Exchange Between Heterogeneous BIM Software." *Journal of Civil Engineering & Management* 24 (7): 537–555. https://doi.org/10.3846/jcem.2018.6132.

LeCun, Y., L. Bottou, Y. Bengio, and P. Haffner. 1998. "Gradient-Based Learning Applied to Document Recognition." *Proceedings of the IEEE* 86 (11): 2278–2324. https://doi.org/10.1109/5.726791.

Li, Jiahao, Weijian Ma, Xueyang Li, Yunzhong Lou, Guichun Zhou, and Xiangdong Zhou. 2025. "CAD-Llama: Leveraging Large Language Models for Computer-Aided Design Parametric 3D Model Generation." *arXiv Preprint*. https://doi.org/10.48550/ARXIV.2505.04481.

Li, Junlong, Yiheng Xu, Tengchao Lv, Lei Cui, Cha Zhang, and Furu Wei. 2022. "DiT: Self-Supervised Pre-Training for Document Image Transformer." Proceedings of the 30th ACM International Conference on Multimedia, Lisboa, Portugal, 3530–3539. ACM. https://doi.org/10.1145/3503161.3547911.

Liao, Minghui, Zhaoyi Wan, Cong Yao, Kai Chen, and Xiang Bai. 2020. "Real-Time Scene Text Detection with Differentiable Binarization." *Proceedings of the AAAI Conference on Artificial Intelligence* 34 (7): 11474–11481. https://doi.org/10.1609/aaai.v34i07.6812.

Liao, Zhichao, Fengyuan Piao, Di Huang, Xinghui Li, Yue Ma, Pingfa Feng, Heming Fang, and Long Zeng. 2024. "Freehand Sketch Generation from Mechanical Components." Proceedings of the 32nd ACM International Conference on Multimedia, New York: 6755–6764. Association for Computing Machinery. https://doi.org/10.1145/3664647.3681046.

Lin, Yi-Hsin, Yu-Hung Ting, Yi-Cyun Huang, Kai-Lun Cheng, and Wen-Ren Jong. 2023. "Integration of Deep Learning for Automatic Recognition of 2D Engineering Drawings." *Machines* 11 (8): 802. https://doi.org/10.3390/machines11080802.

Liu, Haotian, Chunyuan Li, Qingyang Wu, and Yong Jae Lee. 2023. "Visual Instruction Tuning." *Advances in Neural Information Processing Systems* 36: 34892–34916. https://doi.org/10.52202/075280-1516.

Liu, Wei, Anguelov Dragomir, Erhan Dumitru, Szegedy Christian, Reed Scott, Fu Cheng-Yang, and C. Berg. Alexander. 2016. "SSD: Single Shot MultiBox Detector." In *Computer Vision—ECCV 2016*, edited by Bastian Leibe, Jiří Matas, Nicu Sebe, and Max Welling, 21–37. Lecture Notes in Computer Science 9905. Cham: Springer. https://doi.org/10.1007/978-3-319-46448-0_2.

Liu, Yinhan, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. "RoBERTa: A Robustly Optimized BERT Pretraining Approach." *arXiv Preprint*. https://doi.org/10.48550/ARXIV.1907.11692.

Luo, Siwen, Yihao Ding, Siqu Long, Josiah Poon, and Soyeon Caren Han. 2022. "Doc-GCN: Heterogeneous Graph Convolutional Networks for Document Layout Analysis." *arXiv Preprint*. https://doi.org/10.48550/ARXIV.2208.10970.

Macal, Charles M., and Michael J. North. 2010. "Tutorial on Agent-Based Modelling and Simulation." *Journal of Simulation* 4 (3): 151–162. https://doi.org/10.1057/jos.2010.3.

Madireddy, Soumya, Lu Gao, Zia Ud Din, Kinam Kim, Ahmed Senouci, Zhe Han, and Yunpeng Zhang. 2025. "Large Language Model-Driven Code Compliance Checking

in Building Information Modeling." *Electronics* 14 (11): 2146. https://doi.org/10.3390/electronics14112146.
Malsane, Sagar, Jane Matthews, Steve Lockley, Peter E. D. Love, and David Greenwood. 2015. "Development of an Object Model for Automated Compliance Checking." *Automation in Construction* 49:51–58. https://doi.org/10.1016/j.autcon.2014.10.004.
Michaloski, John, Tom Hedberg, Hui Huang, and Thomas Kramer. 2016. *End-To-End Quality Information Framework (QIF) Technology Survey*. NISTIR 8127. Gaithersburg, MD: National Institute of Standards and Technology. https://doi.org/10.6028/NIST.IR.8127.
Moreno-Garcia, Carlos Francisco, Pam Johnston, and Bello Garkuwa. 2020. "Pixel-Based Layer Segmentation of Complex Engineering Drawings Using Convolutional Neural Networks." 2020 International Joint Conference on Neural Networks (IJCNN), Glasgow, UK, 1–7. IEEE. https://doi.org/10.1109/ijcnn48605.2020.9207479.
Mustapha, K. B. 2025. "A Survey of Emerging Applications of Large Language Models for Problems in Mechanics, Product Design, and Manufacturing." *Advanced Engineering Informatics* 64:103066. https://doi.org/10.1016/j.aei.2024.103066.
Nakhaee, Ali, Diellza Elshani, and Thomas Wortmann. 2024. "A Vision for Automated Building Code Compliance Checking by Unifying Hybrid Knowledge Graphs and Large Language Models." In *Scalable Disruptors*, edited by Philipp Eversmann, Christoph Gengnagel, Julian Lienhard, Mette Ramsgaard Thomsen, and Jan Wurm, 445–457. Cham: Springer. https://doi.org/10.1007/978-3-031-68275-9_36.
Nawari, Nawari O. 2012. "BIM Standard in Off-Site Construction." *Journal of Architectural Engineering* 18 (2): 107–113. https://doi.org/10.1061/(asce)ae.1943-5568.0000056.
Noy, Natalya F., and Deborah L. McGuinness. 2001. *Ontology Development 101: A Guide to Creating Your First Ontology*. Technical Report KSL-01-05. Stanford, CA: Stanford Knowledge Systems Laboratory. https://protege.stanford.edu/publications/ontology_development/ontology101.pdf.
Opoku, De-Graft Joe, Srinath Perera, Robert Osei-Kyei, and Maria Rashidi. 2021. "Digital Twin Application in the Construction Industry: A Literature Review." *Journal of Building Engineering* 40:102726. https://doi.org/10.1016/j.jobe.2021.102726.
Park, Sooyong, and Vijayan Sugumaran. 2005. "Designing Multi-Agent Systems: A Framework and Application." *Expert Systems with Applications* 28 (2): 259–271. https://doi.org/10.1016/j.eswa.2004.10.006.
Paulheim, Heiko. 2017. "Knowledge Graph Refinement: A Survey of Approaches and Evaluation Methods." *Semantic Web* 8 (3): 489–508. https://doi.org/10.3233/SW-160218.
Pauwels, Pieter, Sijie Zhang, and Yong-Cheol Lee. 2017. "Semantic Web Technologies in AEC Industry: A Literature Overview." *Automation in Construction* 73:145–165. https://doi.org/10.1016/j.autcon.2016.10.003.
Pazlar, Tomaž, and Žiga Turk. 2008. "Interoperability in Practice: Geometric Data Exchange Using the IFC Standard." *ITcon* 13:362–380. https://www.itcon.org/papers/2008_24.content.00881.pdf.
Peng, Junlong, and Xiangjun Liu. 2023. "Automated Code Compliance Checking Research Based on BIM and Knowledge Graph." *Scientific Reports* 13 (1): 7065. https://doi.org/10.1038/s41598-023-34342-1.
Peruzzini, Margherita, Elisa Prati, and Marcello Pellicciari. 2024. "A Framework to Design Smart Manufacturing Systems for Industry 5.0 Based on the Human-Automation Symbiosis." *International Journal of Computer Integrated Manufacturing* 37 (10–11): 1426–1443. https://doi.org/10.1080/0951192X.2023.2257634.
Petzoldt, Christoph, Matthias Harms, and Michael Freitag. 2023. "Review of Task Allocation for Human-Robot Collaboration in Assembly." *International Journal of Computer Integrated Manufacturing* 36 (11): 1675–1715. https://doi.org/10.1080/0951192X.2023.2204467.
Pfitzner, Fabian, Alexander Braun, and André Borrmann. 2024. "From Data to Knowledge: Construction Process Analysis Through Continuous Image Capturing, Object Detection, and Knowledge Graph Creation." *Automation in Construction* 164:105451. https://doi.org/10.1016/j.autcon.2024.105451.
Picard, Cyril, Kristen M. Edwards, Anna C. Doris, Brandon Man, Giorgio Giannone, Md Ferdous Alam, and Faez Ahmed. 2023. "From Concept to Manufacturing: Evaluating Vision-Language Models for Engineering Design." *arXiv Preprint*. https://doi.org/10.48550/ARXIV.2311.12668.
Prohofsky, David, and Micah Lande. 2024. "Affordances of Large Language Models in Design Activity." 2024 ASEE Annual Conference & Exposition, Portland, Oregon, USA, ASEE Conferences. https://doi.org/10.18260/1-2--46532.
Puthanveettil Madathil, Abhilash, Xichun Luo, Qi Liu, Charles Walker, Rajeshkumar Madarkar, and Yi Qin. 2025. "A Review of Explainable Artificial Intelligence in Smart Manufacturing." *International Journal of Production Research* 63 (23): 8654–8697. https://doi.org/10.1080/00207543.2025.2513574.
Qian, Chen, Yufan Dang, Jiahao Li, Wei Liu, Zihao Xie, Yifei Wang, Weize Chen, et al. 2024. "Experiential Co-Learning of Software-Developing Agents." Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), Bangkok: 5628–5640. Association for Computational Linguistics. https://doi.org/10.18653/v1/2024.acl-long.305.
Radford, Alec, Karthik Narasimhan, Tim Salimans, and Ilya Sutskever. 2018. *Improving Language Understanding by Generative Pre-Training*. San Francisco, CA: OpenAI. https://cdn.openai.com/research-covers/language-unsupervised/language_understanding_paper.pdf.
Reddy, Tadigotla Narendra, Nitesh Kumar, Nachappa Pemmanda Ponnappa, Nagasiri Mohana, Prakash Vinod, Mervin A. Herbert, and Shrikantha S. Rao. 2026. "Intelligent GD&T Symbol Detection in Mechanical Drawings: A Comparative Study of YOLOv11, Faster R-CNN, and RetinaNet for Quality Assurance." *Journal of Intelligent Manufacturing* 37 (7): 2941–2959. https://doi.org/10.1007/s10845-025-02669-3.
Redmon, Joseph, Santosh Divvala, Ross Girshick, and Ali Farhadi. 2016. "You Only Look Once: Unified, Real-Time Object Detection." 2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR), Las Vegas, NV, USA, 779–788. IEEE. https://doi.org/10.1109/cvpr.2016.91.
Redmon, Joseph, and Ali Farhadi. 2018. "YOLOv3: An Incremental Improvement." *arXiv Preprint*. https://doi.org/10.48550/ARXIV.1804.02767.

Ren, Shaoqing, Kaiming He, Ross Girshick, and Jian Sun. 2017. "Faster R-CNN: Towards Real-Time Object Detection with Region Proposal Networks." *IEEE Transactions on Pattern Analysis & Machine Intelligence* 39 (6): 1137–1149. https://doi.org/10.1109/tpami.2016.2577031.

Sacks, Rafael, Charles Eastman, Ghang Lee, and Paul Teicholz. 2018. *BIM Handbook: A Guide to Building Information Modeling for Owners, Designers, Engineers, Contractors, and Facility Managers*. 3rd ed. Hoboken, NJ: Wiley. https://doi.org/10.1002/9781119287568.

Saunila, Minna, Juhani Ukko, Mina Nasiri, and Patrizia Garengo. 2024. "The Role of Digital Governance in the Integration of Performance Measurement Systems Uses and Industry 4.0 Maturity." *International Journal of Computer Integrated Manufacturing* 37 (7): 809–822. https://doi.org/10.1080/0951192X.2023.2257918.

Schlagenhauf, Tobias, Markus Netzer, and Jan Hillinger. 2023. "Text Detection on Technical Drawings for the Digitization of Brown-Field Processes." *Procedia CIRP* 118:372–377. https://doi.org/10.1016/j.procir.2023.06.064.

Schüpbach, Aurel, Raul San Miguel, Julian Ferchow, and Mirko Meboldt. 2025. "From Text to Design: A Framework to Leverage LLM Agents for Automated CAD Generation." *Proceedings of the Design Society* 5:1893–1902. https://doi.org/10.1017/pds.2025.10203.

Seff, Ari, Yaniv Ovadia, Wenda Zhou, and Ryan P. Adams. 2020. "SketchGraphs: A Large-Scale Dataset for Modeling Relational Geometry in Computer-Aided Design." *arXiv Preprint*. https://doi.org/10.48550/ARXIV.2007.08506.

Shi, Baoguang, Xiang Bai, and Cong Yao. 2017. "An End-To-End Trainable Neural Network for Image-Based Sequence Recognition and Its Application to Scene Text Recognition." *IEEE Transactions on Pattern Analysis & Machine Intelligence* 39 (11): 2298–2304. https://doi.org/10.1109/tpami.2016.2646371.

Simonyan, Karen, and Andrew Zisserman. 2015. "Very Deep Convolutional Networks for Large-Scale Image Recognition." 3rd International Conference on Learning Representations (ICLR 2015), San Diego, CA: 1–14. Computational and Biological Learning Society. https://arxiv.org/abs/1409.1556 .

Stürmer, Jan Marius, Marius Graumann, and Tobias Koch. 2024. "From Engineering Diagrams to Graphs: Digitizing P&Ids with Transformers." *arXiv Preprint*. https://doi.org/10.48550/ARXIV.2411.13929.

SymbolDetection. 2024. "Merge_places Computer Vision Dataset" (dataset). Roboflow Universe. Accessed July 18, 2025. https://universe.roboflow.com/symboldetection-4xwmm/merge_places.

Tan, Mingxing, and Quoc V. Le. 2019. "EfficientNet: Rethinking Model Scaling for Convolutional Neural Networks." *arXiv Preprint*. https://doi.org/10.48550/ARXIV.1905.11946.

Tang, Zeyu, Chang Zeng, and Yali Zeng. 2025. "Research on Data Security in Industry 4.0 Manufacturing Industry Against the Background of Privacy Protection Challenges." *International Journal of Computer Integrated Manufacturing* 38 (5): 636–648. https://doi.org/10.1080/0951192X.2024.2319656.

Tomaszuk, Dominik, and David Hyland-Wood. 2020. "RDF 1.1: Knowledge Representation and Data Integration Language for the Web." *Symmetry* 12 (1): 84. https://doi.org/10.3390/sym12010084.

Toussaint, J., and K. Cheng. 2002. "Design Agility and Manufacturing Responsiveness on the Web." *Integrated Manufacturing Systems* 13 (5): 328–339. https://doi.org/10.1108/09576060210429784.

Vaswani, Ashish, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. 2017. "Attention is All You Need." *arXiv Preprint*. https://doi.org/10.48550/ARXIV.1706.03762.

Veličković, Petar, Guillem Cucurull, Arantxa Casanova, Adriana Romero, Pietro Liò, and Yoshua Bengio. 2018. "Graph Attention Networks." 6th International Conference on Learning Representations (ICLR 2018), Conference Track Proceedings, Vancouver, Canada. OpenReview.net. https://openreview.net/forum?id=rJXMpikCZ.

Villena Toro, Javier, Anton Wiberg, and Mehdi Tarkian. 2023. "Optical Character Recognition on Engineering Drawings to Achieve Automation in Production Quality Control." *Frontiers in Manufacturing Technology* 3:1154132. https://doi.org/10.3389/fmtec.2023.1154132.

Volk, Rebekka, Julian Stengel, and Frank Schultmann. 2014. "Building Information Modeling (BIM) for Existing Buildings — Literature Review and Future Needs." *Automation in Construction* 38:109–127. https://doi.org/10.1016/j.autcon.2013.10.023.

Wan, Hanlong, Weili Xu, Michael Rosenberg, Jian Zhang, and Aysha Siddika. 2025. "Automatic Building Code Review: A Case Study." *arXiv Preprint*. https://doi.org/10.48550/ARXIV.2510.02634.

Wang, Yukang, Yue Liu, Haozhe Cai, Jia Wang, and Xiaoping Zhou. 2023. "An Automated Fire Code Compliance Checking Jointly Using Building Information Models and Natural Language Processing." *Fire* 6 (9): 358. https://doi.org/10.3390/fire6090358.

Webb, Louie, Osman M. Tokhi, and Bugra Alkan. 2025. "State of the Art and Future Directions of Digital Twin-Enabled Smart Assembly Automation in Discrete Manufacturing Industries." *International Journal of Computer Integrated Manufacturing* 38 (8): 1126–1160. https://doi.org/10.1080/0951192X.2024.2387775.

Wu, Jiabin, Stavros Nousias, and André Borrmann. 2025. "Design Healing Framework for Automated Code Compliance." *Automation in Construction* 171:106004. https://doi.org/10.1016/j.autcon.2025.106004.

Wu, Zonghan, Shirui Pan, Fengwen Chen, Guodong Long, Chengqi Zhang, and Philip S. Yu. 2021. "A Comprehensive Survey on Graph Neural Networks." *IEEE Transactions on Neural Networks and Learning Systems* 32 (1): 4–24. https://doi.org/10.1109/tnnls.2020.2978386.

Xu, Xin, and Hubo Cai. 2020. "Semantic Approach to Compliance Checking of Underground Utilities." *Automation in Construction* 109:103006. https://doi.org/10.1016/j.autcon.2019.103006.

Xu, Xuesong, Gang Xiao, Li Sun, Xia Zhang, Peixi Wu, Yuanming Zhang, and Zhenbo Cheng. 2025. "Associative Reasoning for Engineering Drawings Using an Interactive Attention Mechanism." *Automation in Construction* 170:105942. https://doi.org/10.1016/j.autcon.2024.105942.

Xu, Yiheng, Minghao Li, Lei Cui, Shaohan Huang, Furu Wei, and Ming Zhou. 2020. "LayoutLM: Pre-Training of Text and Layout for Document Image Understanding." Proceedings of the 26th ACM SIGKDD International Conference on

Knowledge Discovery & Data Mining, Virtual Event, CA, USA, 1192–1200. ACM. https://doi.org/10.1145/3394486.3403172.
Xu, Yiheng, Tengchao Lv, Lei Cui, Guoxin Wang, Yijuan Lu, Dinei Florencio, Cha Zhang, and Furu Wei. 2021. "LayoutXLM: Multimodal Pre-Training for Multilingual Visually-Rich Document Understanding." *arXiv Preprint*. https://doi.org/10.48550/ARXIV.2104.08836.
Yahya, Muhammad, John G. Breslin, and Muhammad Intizar Ali. 2021. "Semantic Web and Knowledge Graphs for Industry 4.0." *Applied Sciences* 11 (11): 5110. https://doi.org/10.3390/app11115110.
Yan, Lixiang. 2025. "From Passive Tool to Socio-Cognitive Teammate: A Conceptual Framework for Agentic AI in Human-AI Collaborative Learning." arXiv preprint. https://doi.org/10.48550/ARXIV.2508.14825.
Yang, Fan, and Jiansong Zhang. 2024. "Prompt-Based Automation of Building Code Information Transformation for Compliance Checking." *Automation in Construction* 168:105817. https://doi.org/10.1016/j.autcon.2024.105817.
Zhang, Jiansong, and Nora M. El-Gohary. 2017. "Integrating Semantic NLP and Logic Reasoning into a Unified System for Fully-Automated Code Checking." *Automation in Construction* 73:45–57. https://doi.org/10.1016/j.autcon.2016.08.027.
Zhang, Licheng, Bach Le, Naveed Akhtar, Siew-Kei Lam, and Tuan Ngo. 2025. "Large Language Models for Computer-Aided Design: A Survey." *arXiv Preprint*. https://doi.org/10.48550/ARXIV.2505.08137.
Zhang, Ruichuan, and Nora M. El-Gohary. 2018. "A Clustering Approach for Analyzing the Computability of Building Code Requirements." In *Construction Research Congress 2018*, edited by Chao Wang, Christofer Harper, Yongcheol Lee, Rebecca Harris, and Charles Berryman, 86–95. Reston, VA: American Society of Civil Engineers. https://doi.org/10.1061/9780784481264.009.
Zhang, Wentai, Joe Joseph, Yue Yin, Liuyue Xie, Tomotake Furuhata, Soji Yamakawa, Kenji Shimada, and Levent Burak Kara. 2023. "Component Segmentation of Engineering Drawings Using Graph Convolutional Networks." *Computers in Industry* 147:103885. https://doi.org/10.1016/j.compind.2023.103885.
Zhang, Zhibo, Prakhar Jaiswal, and Rahul Rai. 2018. "FeatureNet: Machining Feature Recognition Based on 3D Convolution Neural Network." *Computer-Aided Design* 101:12–22. https://doi.org/10.1016/j.cad.2018.03.006.
Zhong, Botao, Xuejiao Xing, Hanbin Luo, Qirui Zhou, Heng Li, Timothy Rose, and Weili Fang. 2020. "Deep Learning-Based Extraction of Construction Procedural Constraints from Construction Regulations." *Advanced Engineering Informatics* 43:101003. https://doi.org/10.1016/j.aei.2019.101003.
Zhou, Peng, and Nora El-Gohary. 2016. "Domain-Specific Hierarchical Text Classification for Supporting Automated Environmental Compliance Checking." *Journal of Computing in Civil Engineering* 30 (4): 04015057. https://doi.org/10.1061/(ASCE)CP.1943-5487.0000513.
Zhou, Peng, and Nora El-Gohary. 2017. "Ontology-Based Automated Information Extraction from Building Energy Conservation Codes." *Automation in Construction* 74:103–117. https://doi.org/10.1016/j.autcon.2016.09.004.
Zhou, Shengzong, and Nebojsa Bacanin. 2024. "Artificial Intelligence in Advanced Manufacturing." *International Journal of Computer Integrated Manufacturing* 37 (4): 383–384. https://doi.org/10.1080/0951192X.2024.2327226.
Zhou, Xinyu, Cong Yao, He Wen, Yuzhi Wang, Shuchang Zhou, Weiran He, and Jiajun Liang. 2017. "EAST: An Efficient and Accurate Scene Text Detector." *arXiv Preprint*. https://doi.org/10.48550/ARXIV.1704.03155.
Zhou, Yucheng, Jiarui Lin, and Zhongtian She. 2021. "Automatic Construction of Building Code Graph for Regulation Intelligence." In *ICCREM 2021*, edited by Yaowu Wang, Dayu Zhang, and Geoffrey Q. P. Shen, 248–254. Reston, VA: American Society of Civil Engineers. https://doi.org/10.1061/9780784483848.028.
Zhu, Deyao, Jun Chen, Xiaoqian Shen, Xiang Li, and Mohamed Elhoseiny. 2023. "MiniGPT-4: Enhancing Vision-Language Understanding with Advanced Large Language Models." *arXiv Preprint*. https://doi.org/10.48550/ARXIV.2304.10592.