



EXPLAINABLE ARTIFICIAL INTELLIGENCE IN INDUSTRIAL MACHINE VISION: A BIBLIOMETRIC ANALYSIS

Truong Thanh Cong^{1*}, Mai Thanh Thao¹, Vo xuan The¹

¹University of Finance - Marketing, Vietnam

*Corresponding author: Email: ttcong@ufm.edu.vn

ARTICLE INFO	ABSTRACT
DOI: 10.52932/jfmr.v4i4ene.1521 <i>Received:</i> June 09, 2026 <i>Accepted:</i> June 30, 2026 <i>Published:</i> July 25, 2026 Keywords: Bibliometric Analysis, Explainable Artificial Intelligence, Industrial Machine Vision, Industry 4.0 Article classification: Literature review JEL codes: O33, L60, C55	Purpose – This study maps the structural properties of explainable artificial intelligence (XAI) in the industrial machine vision research field at the corpus level. Explainability requirements for industrial machine vision systems have intensified since the European Union Artificial Intelligence Act classified high-risk image-based inspection systems as requiring traceable automated decisions, yet the resulting research field has not been systematically mapped. Design/ methodology/ approach – 526 peer-reviewed publications retrieved from Scopus and Web of Science for the period 2015 to March 2026 were analysed using six sequential bibliometric analyses: publication trend, source productivity, geographic distribution, keyword co-occurrence network, thematic mapping, and thematic evolution. All analyses were conducted using the bibliometrix package (version 4.3.0) in R (version 4.4.1). Findings – Annual output grew from 6 papers in 2017 to 158 in 2025 (compound annual growth rate: approximately 50%), with a 94% single-year acceleration in 2022-2023 coinciding with early EU AI Act commentary and the displacement of convolutional architectures by vision transformers. Among XAI method terms, Grad-CAM leads to 40 keyword occurrences, more than three times LIME (11) and nearly six times SHAP (7). China and India together account for 30.0% of the corpus (158 papers). Thematic mapping shows deep learning migrating from a motor to a basic theme between the sub-periods 2015-2021 and 2022-2026, while decision support systems and ethical artificial intelligence emerge as new motor themes.

Originality/ value – This study presents the first dual-database bibliometric analysis of XAI in industrial machine vision, providing a reference map for researchers entering the field, journal editors tracking disciplinary boundaries, and policymakers benchmarking research response to regulatory pressure. The findings indicate that the field has moved from performance demonstration toward regulatory-compliance integration, with post-hoc visual explanation as the dominant methodological choice.

1. Introduction

Machine vision systems now perform quality inspection at throughput rates no human inspector can match. Convolutional neural networks and vision transformers identify surface defects, dimensional deviations, and assembly errors across automotive, semiconductor, pharmaceutical, and food manufacturing with accuracy rates exceeding 90% in controlled benchmarks (Xia et al., 2020; Yi et al., 2017). This accuracy carries a structural cost: these models map input images to classification decisions through hundreds of millions of learned parameters, with no interpretable account of why a specific item was flagged as defective. In safety-critical production environments, that opacity creates two concrete problems. First, process engineers cannot diagnose systematic failure modes from model outputs alone, because the model provides no basis for tracing a false rejection or missed defect to a causal input feature. Second, accountability frameworks now require that decisions affecting product safety carry a traceable justification.

The European Union Artificial Intelligence Act (Regulation (EU) 2024/1689, 2024) classifies machine vision systems operating in critical industrial infrastructure as high-risk AI and mandates explainability for decisions affecting safety and fundamental rights. Its legislative trajectory was evident from 2022, when draft

provisions generated sustained commentary from automation and manufacturing sectors. This regulatory context provided institutional expression for a shift already underway in the research record: XAI had moved from a technical research interest to a deployment consideration in industrial machine vision.

Despite this growth, the structural properties of the field have not been mapped at the corpus level. Existing reviews address XAI broadly across all domains (Arrieta et al., 2020; Herrera, 2025; Nguyen et al., 2025), while bibliometric analyses of artificial intelligence in manufacturing examine industrial adoption without restriction to image-based inspection or to explainability as a primary focus. No comprehensive dual-database bibliometric study addresses the specific intersection of XAI methods and industrial machine vision. The dominant explanation methods deployed in production contexts, the geographic distribution of contributing institutions, the source journals concentrating this output, and the trajectory of thematic change over the field's growth period remain undocumented in systematic form.

This paper presents the first dual-database bibliometric analysis of XAI in industrial machine vision, drawing on 526 peer-reviewed publications retrieved from Scopus and Web of Science for the period 2015 to March 2026. Six sequential analyses characterise publication

growth, source and country productivity, keyword co-occurrence structure, and thematic evolution. The findings provide a reference map for researchers entering the field, for journal editors tracking disciplinary boundaries, and for policymakers benchmarking research response to regulatory pressure.

The remainder of this paper is organized as follows: Section 2 presents the theoretical foundation and review of related studies. Section 3 describes the search protocol and analysis framework. Section 4 reports the results and discussion across all six analyses. Section 5 states the conclusions and managerial implications.

2. Theoretical Foundation

Explainable Artificial Intelligence (XAI) refers to methods and techniques that make the outputs of machine learning and deep learning models interpretable to human users. The foundational framework proposed by Arrieta et al. (2020) defines XAI as a set of processes and methods that allow human users to understand and appropriately trust the results and outputs created by machine learning algorithms. XAI approaches are broadly classified into two categories: ante-hoc methods (inherently interpretable models such as decision trees and rule-based systems) and post-hoc methods (applied after training to explain otherwise opaque models). Post-hoc visual explanation methods dominate in image-centric domains; among these, Gradient-weighted Class Activation Mapping (Grad-CAM) produces class-activation maps overlaid on input images, LIME (Local Interpretable Model-agnostic Explanations) generates local approximations, and SHAP (SHapley Additive exPlanations) computes feature attribution values.

Industrial machine vision encompasses automated systems that acquire, process, and interpret image data for quality control,

defect detection, and process monitoring within manufacturing environments. The integration of deep learning into machine vision substantially increased classification accuracy while simultaneously reducing the interpretability of model decisions, creating a performance-interpretability tension that XAI methods are designed to address. This tension is sharpened in regulated industries where ISO standards, sector-specific certification requirements, and now the EU AI Act mandate that automated decisions be auditable.

Bibliometric analysis provides a systematic method for quantifying and mapping research output across a defined corpus. The methodology proposed by Donthu et al. (2021) distinguishes performance analysis (measuring the productivity and impact of research constituents) from science mapping (visualising the intellectual structure of a field). Callon's framework (Callon et al., 1991) extends science mapping through thematic mapping, positioning keyword clusters in centrality-density space to classify themes as motor, niche, basic, or emerging/declining. Bradford's Law of Scattering predicts that a small core of journals accounts for the majority of a field's output, with the remainder distributed across a much larger set of peripheral sources.

Related bibliometric studies on XAI have addressed the general XAI landscape (Arrieta et al., 2020; Herrera, 2025) and specific application domains such as additive manufacturing (Nguyen et al., 2025). Bibliometric analyses of artificial intelligence in manufacturing have examined broader industrial adoption patterns but have not restricted their scope to image-based inspection or explainability as a primary variable. The present study fills this gap by restricting the corpus to the intersection of XAI methods and industrial machine vision and applying dual-database retrieval to maximise coverage.

3. Research Methodology

3.1. Data Collection and Search Protocol

Data were retrieved from two complementary databases: Scopus and Web of Science (WoS). Both databases were searched because single-database coverage in applied AI research consistently misses 30-50% of eligible records, with each platform indexing distinct journal sets in engineering and computer science.

The search string combined three conceptual blocks: XAI methods, machine vision and visual inspection techniques, and industrial context. In Scopus, the query applied TITLE-ABS-KEY syntax:

TITLE-ABS-KEY(("explainable artificial intelligence" OR "explainable AI" OR XAI OR "interpretable machine learning" OR "model explainability" OR "model interpretability" OR SHAP OR LIME OR "saliency map" OR "Grad-CAM") AND ("machine vision" OR "industrial vision" OR "computer vision" OR "image processing" OR "image analysis" OR "visual inspection" OR "automated inspection" OR "defect detection" OR "quality inspection") AND (industrial OR industry OR manufacturing OR "Industry 4.0" OR factory OR production)).

The WoS equivalent applied TS = (Topic Search) syntax across the same three conceptual blocks. Restricting the second block to machine vision and inspection terms ensured the corpus centred on paper where explainability methods operate on image or sensor data within production systems, not on XAI applied to business analytics or supply chain modelling.

Both searches were filtered to English-language journal articles and reviews published from 2015 to March 2026. The start date aligns with the earliest traceable industrial machine vision papers applying explainability methods: foundational SHAP and Grad-CAM work entered industrial inspection literature

from 2015, preceding the formal DARPA XAI programme (2016) that subsequently accelerated the field. Searches were conducted on 10 March 2026.

Scopus returned to 575 records and WoS returned 294 records. After merging (869 total) and removing 245 duplicate entries identified by DOI matching and author-title comparison, 624 unique records remained.

3.2. Inclusion and Exclusion Criteria

Records were screened against pre-specified criteria applied sequentially to titles, abstracts, and author keywords. Inclusion required that a paper be a peer-reviewed journal article, conference paper, or review indexed in Scopus or WoS, written in English, addressing at least one XAI or model interpretability method, and reporting an application within an industrial, manufacturing, or factory context. Book chapters and editorials were excluded as they lack the peer-review and citation traceability required for bibliometric analysis. Non-English records and papers addressing XAI with no empirical or applied industrial component were excluded. Papers without accessible abstracts were excluded before screening.

Screening of 624 records excluded 98 papers, primarily because the industrial application context was absent (papers addressed XAI in healthcare, finance, or other non-manufacturing domains) or the explainability method was incidental rather than a primary focus. The final corpus comprised 526 publications.

3.3. Bibliometric Analysis Framework

The analysis framework was adapted from the guidelines proposed by Donthu et al. (2021). All analyses were conducted using the bibliometrix package (version 4.3.0) in R (version 4.4.1). Six sequential analyses were applied to the corpus.

Descriptive statistics characterised publication volume by year, journal source, and

country of corresponding authorship. Source productivity was assessed against Bradford’s Law of Scattering, which predicts that a small core of journals accounts for the majority of a field’s output.

Keyword co-occurrence analysis constructed a network in which nodes represent author keywords and edges represent co-occurrence within the same paper. The minimum co-occurrence threshold was set to 3. Network topology was visualised using the Fruchterman-Reingold layout algorithm.

Thematic mapping positioned keyword clusters in a two-dimensional space following Callon’s framework (Callon et al., 1991), defined by centrality (the degree of external connection between clusters) and density (the internal cohesion of keywords within a cluster). Maps were produced for the full corpus period and separately for 2015-2021 and 2022-2026.

The sub-period boundary at 2021/2022 reflects two concurrent structural shifts: COVID-19 supply chain disruptions from 2020 onward accelerated Industry 4.0 adoption, and vision transformer architectures displaced convolutional networks as the dominant paradigm in computer vision from 2021. Thematic evolution analysis tracked how

keyword clusters quantified the proportion of keywords that migrated between clusters, merged, or split across the boundary year.

4. Results and evaluation

4.1. Research results

Publication trends

The 526 papers retrieved span 2015 to March 2026. Publication output grew from 6 papers in 2017 to 158 in 2025, representing a compound annual growth rate (CAGR) of approximately 50% over the eight-year period. Growth was not uniform. From 2017 to 2022, output expanded gradually from 6 to 34 papers (average of fewer than five additional publications per year). Output growth accelerated from 2022: 34 papers in 2022 rose to 66 in 2023, a 94% year-on-year increase, and continued rising through 2024 (94 papers) and 2025 (158 papers). The 2022–2023 inflection corresponds with the mainstreaming of large vision models and the European Union’s preparations for the AI Act. By 2025, annual output exceeded 26 times the 2017 baseline. The 2026 count of 67 papers covers the period through 10 March 2026; the annualised projection exceeds 300 publications, suggesting the growth trajectory has not plateaued.




Figure 1. Annual Publication Trend – XAI in Industrial Machine Vision

Source analysis

The 15 most productive publication venues together account for 107 of the 526 papers (20.3%); the remaining 79.7% distributes across a large number of lower-productivity venues. This high scatter, in which no single source exceeds 2.7% of the corpus, is consistent with Bradford's Law of scattering in an interdisciplinary field. IEEE Access leads with 14 papers. The Journal of Intelligent Manufacturing (11 papers) is the highest-ranked domain-specific journal, followed by Applied Sciences-Basel (9 papers), Computers and Electronics in Agriculture (8 papers), Scientific Reports (7 papers), and Engineering Applications of Artificial Intelligence (6 papers). Conference proceedings series account for 38 of the 107 top 15 papers. Four open-access venues in the top 15 (IEEE Access, Applied Sciences-Basel, Scientific Reports, and Sensors) contribute 35 papers combined, reflecting the broader engineering shift toward open-access publishing (*see Appendix 1 online*).

Geographic distribution

The 15 most productive countries together account for 335 of the 526 papers (63.7%). China leads the corpus with 94 papers, followed by India with 64 papers; together these two countries contribute 158 papers (30.0% of the total). The USA ranks third with 34 papers, Germany fourth with 25, and South Korea fifth with 24. This five-country group produces 241 papers (45.8% of the corpus). Asian countries within the top 15 (China, India, South Korea, Bangladesh, Japan) account for 197 papers combined (37.5%). European countries (Germany, United Kingdom, Spain, Italy, Poland, Switzerland) contribute 79 papers (15.0%). Bangladesh, with 8 papers, is the highest-ranked lower-middle-income economy in the corpus (*see Appendix 2 online*).

Keyword co-occurrence analysis

The keyword co-occurrence network retains terms appearing in at least 3 publications and organises into three clusters. The most frequent keyword is "deep learning" (153 occurrences), followed by "explainable artificial intelligence" (111 occurrences). Among XAI-specific method terms, Grad-CAM ranks first at 40 occurrences; LIME and SHAP appear at 11 and 7 occurrences respectively. Among application-domain terms, "defect detection" is the most frequent at 36 occurrences, followed by "transfer learning" (20), "quality control" (8), and "predictive maintenance" (7). Cluster 1 groups the deep-learning methodology core: deep learning, convolutional neural network, computer vision, machine learning, feature extraction, defect detection, inspection, and quality control. Cluster 2 groups the XAI and interpretability strand: explainable artificial intelligence, decision making, image analysis, artificial intelligence, diagnosis, and image segmentation. Cluster 3 is a smaller, method-specific cluster centred on LIME, image processing, and forecasting (*see Appendix 3 online*).

Thematic mapping

Thematic maps were constructed using Callon's framework. In the full-period map (2015-2026), the basic quadrant contains the largest and most central clusters: deep learning, explainable artificial intelligence, and computer vision occupy a single high-frequency cluster with the greatest corpus-wide centrality, alongside convolutional neural network, saliency map, and inspection. These themes function as the structural backbone of the field. The motor quadrant holds four clusters: feature extraction and computational modelling, decision support systems and ethical artificial intelligence, predictive modelling, and fabric defect detection. The niche quadrant contains

specialised application strands including internal detection and solar cell applications, eco-friendly concrete, fault detection and intelligent manufacturing, and disease detection and IoT. The emerging or declining quadrant holds interpretable machine learning, additive manufacturing and process monitoring, SHAP and YOLO, LIME and black-box models, and UAV-based applications (*see Appendix 4 online*).

Thematic evolution

A Sankey diagram traces the evolution of thematic clusters across four time slices spanning the corpus period. A core deep-learning and computer-vision stream forms the widest single band across all four time slices. A convolutional neural network strand, visible as a distinct band in the earliest time slice, merges into the broader deep-learning stream by the third slice, consistent with the shift of CNN to a basic theme. XAI-specific themes (explainable artificial intelligence, Grad-CAM, and LIME) appear as separate, growing bands beginning in the second time slice and widen through the third and fourth slices. Decision support systems and ethical artificial intelligence appear as new bands only in the third and fourth slices (*see Appendix 5 online*).

4.2. Discussion of research results

Growth trajectory and regulatory context

The corpus grew at a CAGR of approximately 50% between 2017 and 2025, a rate that places XAI in industrial machine vision among the fastest-expanding applied AI subfields indexed in Scopus and WoS. The 2022-2023 inflection, in which annual output rose 94% in a single year, marks a structural acceleration rather than incremental growth. Two converging forces account for this inflection. The first is the enactment of the European Union Artificial Intelligence Act, which classifies machine vision systems used in critical industrial infrastructure as high-risk AI and mandates explainability

for decisions affecting safety and fundamental rights (Regulation (EU) 2024/1689, 2024). The second is the mainstreaming of large vision transformer architectures from 2021 onward (Dosovitskiy et al., 2021), which introduced architectures whose internal workings are less transparent than the convolutional architectures they began to displace.

Geographic concentration and source patterns

China (94 papers) and India (64 papers) together account for 30.0% of the corpus. Both countries operate large manufacturing sectors that generate direct demand for machine vision quality-control systems, and both have seen sustained growth in AI research investment across the study period. Germany and South Korea rank fourth and fifth and share a comparable industrial profile as major producers of automotive and semiconductor components. Bangladesh at eighth rank (8 papers) concentrates in agricultural machine vision applications (fabric defect detection and crop disease identification), suggesting that demand for XAI-enabled machine vision extends to applications where decision accountability carries regulatory or livelihood stakes. Source-level concentration mirrors the country-level pattern: publication activity is distributed rather than consolidated, consistent with Bradford's Law of scattering in interdisciplinary fields.

XAI method patterns: Post-hoc visualisation dominance

Grad-CAM leads XAI-specific terms at 40 occurrences, more than three times LIME (11) and nearly six times SHAP (7). This asymmetry reflects a domain-specific feature of industrial machine vision: the tasks are overwhelmingly image-based, and Grad-CAM generates class-activation maps overlaid on input images to produce explanations auditable by quality engineers without statistical expertise. SHAP

and LIME produce feature-importance scores suited to tabular or mixed-feature inputs and appear far less often because the corpus is image-centric. The two most-cited papers confirm this pattern: Yi et al. (2017, 139 citations) establish the deep learning baseline, and Xia et al. (2020, 131 citations) introduce Grad-CAM visual explanation methods to make model decisions auditable by process engineers.

Thematic evolution: Maturation and emerging frontiers

The shift of “deep learning” from the motor quadrant (2015-2021) to the basic quadrant (2022-2026) is the central bibliometric signal of this corpus. In Callon’s framework, motor themes are productive specialisations of a field: well-developed internally and well-connected externally. Basic themes are foundational, central to the field’s vocabulary but no longer sites of specialist differentiation. The repositioning of deep learning as basic signals that demonstrating deep learning works in an industrial context is no longer a sufficient research contribution. The emergence of “decision support systems” and “ethical artificial intelligence” as new motor themes in 2022-2026 reflects the regulatory shift. Rožanec et al. (2023), the most-cited paper in the corpus at 177 citations, exemplifies this shift with its focus on human-centric AI architecture for Industry 5.0 treating explainability as a design requirement for human-machine collaboration. The persistence of “interpretable machine learning” in the emerging or declining quadrant indicates that the tension between predictive performance and inherent interpretability remains unresolved (Kruschel et al., 2026).

5. Conclusions and managerial implications

5.1. Conclusions

XAI in industrial machine vision grew from fewer than 10 papers per year in 2017 to 158 in 2025, a compound annual growth rate

of approximately 50% sustained across eight years. The 2022-2023 inflection, in which annual output rose 94% in a single year, marks a structural transition driven by two concurrent forces: EU AI Act regulatory pressure created compliance-driven research incentives, and the shift from convolutional to vision transformer architectures introduced explanation demands that CNNs did not face.

Three structural properties define the field across 2022-2026. Post-hoc visual explanation dominates: Grad-CAM accounts for 40 keyword occurrences, more than three times LIME (11) and nearly six times SHAP (7), consistent with an image-centric task profile where saliency-map overlays serve production engineers without statistical expertise. Geographic output is concentrated but not uniform: China and India together contribute 30.0% of the corpus and the top five countries account for 45.8%, yet the remaining output spans more than 15 nations across four continents, including lower-income economies where agricultural machine vision applications drive demand. Thematically, deep learning has moved from motor to basic theme, and decision support systems and ethical AI have taken its place, signalling a research agenda now organised around deployment and compliance rather than performance demonstration.

5.2. Managerial implications

Practitioners deploying machine vision systems in regulated environments should prioritise Grad-CAM or equivalent saliency-map approaches for explainability, as these methods dominate the current evidence base and produce outputs directly interpretable by production engineers. Firms in the EU should align explainability documentation with the AI Act’s high-risk classification requirements, referencing the audit-trail standards emerging in the decision support systems research cluster.

Research funders and journal editors can use the method frequency data, thematic maps, and geographic distribution reported here to identify active frontiers and track disciplinary boundary dynamics. The primary open question is whether the post-hoc explanation norm persists as transformer-based architectures mature, or whether architectural interpretability becomes a tractable design target as the regulatory environment tightens.

5.3. Limitations of the study

First, the corpus is restricted to records indexed in Scopus and WoS; grey literature, technical reports, and conference proceedings not indexed in these databases are excluded. Second, keyword co-occurrence analysis depends on author-supplied keywords, introducing noise from variant terminology and papers without keywords. Third, citation counts reflect full lifetime accumulation as of March 2026, which advantages older papers by construction. Fourth, geographic distribution reflects institutional affiliation as recorded in

Scopus and WoS, not the physical location of the research or industry partner. Fifth, the corpus search was conducted on 10 March 2026; publications indexed after this date are not represented, and the 2026 annual count covers only the first 69 days of the year.

Funding Acknowledgment

This research received no specific grant from any funding agency.

Data Availability

Data is available on reasonable request. Interested researchers may contact the corresponding author at ttcong@ufm.edu.vn.

AI Usage Statement

During the preparation of this work, the author used Claude (Anthropic) in order to assist with language editing and improve the readability of the manuscript. After using this tool, the author reviewed and edited the content as needed and takes full responsibility for the content of the publication.

References

- Arrieta, A. B., Díaz-Rodríguez, N., Del Ser, J., Bennetot, A., Tabik, S., Barbado, A.,... & Herrera, F. (2020). Explainable Artificial Intelligence (XAI): Concepts, taxonomies, opportunities and challenges toward responsible AI. *Information fusion*, 58, 82-115. <https://doi.org/10.1016/j.inffus.2019.12.012>
- Callon, M., Courtial, J. P., & Laville, F. (1991). Co-word analysis as a tool for describing the network of interactions between basic and technological research. *Scientometrics*, 22(1), 155-205. <https://doi.org/10.1007/BF02019280>
- Donthu, N., Kumar, S., Mukherjee, D., Pandey, N., & Lim, W. M. (2021). How to conduct a bibliometric analysis: An overview and guidelines. *Journal of Business Research*, 133, 285-296. <https://doi.org/10.1016/j.jbusres.2021.04.070>
- Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T., Dehghani, M., Minderer, M., Heigold, G., Gelly, S., Uszkoreit, J., & Houlsby, N. (2021). An image is worth 16x16 words: Transformers for image recognition at scale (Version 2). *arXiv*. <https://doi.org/10.48550/ARXIV.2010.11929>
- European Parliament (2024). Regulation (EU) 2024/1689 of the European Parliament and of the Council. *Official Journal of the European Union*. <https://eur-lex.europa.eu/eli/reg/2024/1689/oj/eng>
- Herrera, F. (2025). Reflections and attentiveness on eXplainable Artificial Intelligence (XAI). The journey ahead from criticisms to human-AI collaboration. *Information Fusion*, 121, 103133. <https://doi.org/10.1016/j.inffus.2025.103133>
- Kruschel, S., Hambauer, N., Weinzierl, S., Zilker, S., Kraus, M., & Zschech, P. (2026). Challenging the performance-interpretability trade-off: An evaluation of interpretable machine learning models. *Business & Information Systems Engineering*, 68(1), 159-183. <https://doi.org/10.1007/s12599-024-00922-2>
- Nguyen, Q. T., Truong, T. C., Nguyen, T. P., Tran, T. N., Mesicek, J., & Pagac, M. (2025). Explainable artificial intelligence in additive manufacturing: A systematic review on method convergence and assessment of standardization gaps. *MM Science Journal*, 2025(4), 8638. https://doi.org/10.17973/MMSJ.2025_10_2025095

- Rožanec, J. M., Novalija, I., Zajec, P., Kenda, K., Tavakoli Ghinani, H., Suh, S., Bian, S., Veliou, E., Papamartzivanos, D., Giannetsos, T., Menesidou, S. A., Alonso, R., Cauli, N., Meloni, A., Recupero, D. R., Kyriazis, D., Sofianidis, G., Theodoropoulos, S., Fortuna, B.,... Soldatos, J. (2023). Human-centric artificial intelligence architecture for industry 5.0 applications. *International Journal of Production Research*, 61(20), 6847–6872. <https://doi.org/10.1080/00207543.2022.2138611>
- Xia, C., Pan, Z., Fei, Z., Zhang, S., & Li, H. (2020). Vision based defects detection for keyhole TIG welding using deep learning with visual explanation. *Journal of Manufacturing Processes*, 56, 845–855. <https://doi.org/10.1016/j.jmapro.2020.05.033>
- Yi, L., Li, G., & Jiang, M. (2017). An end-to-end steel strip surface defects recognition system based on convolutional neural networks. *Steel Research International*, 88(2), 1600068. <https://doi.org/10.1002/srin.201600068>