

## University of Groningen

### Human in the AI loop via xAI and Active Learning for Visual Inspection

Rožanec, Jože; Montini, Elias; Cutrona, Vincenzo; Papamartzivanos, Dimitrios; Klemencic, Timotej; Fortuna, Blaz; Mladenec, Dunja; Veliou, Entso; Giannetsos, Thanassis; Emmanouilidis, Christos

*Published in:*  
Artificial Intelligence in Manufacturing

*DOI:*  
[10.1007/978-3-031-46452-2\\_22](https://doi.org/10.1007/978-3-031-46452-2_22)

**IMPORTANT NOTE:** You are advised to consult the publisher's version (publisher's PDF) if you wish to cite from it. Please check the document version below.

*Document Version*  
Publisher's PDF, also known as Version of record

*Publication date:*  
2023

[Link to publication in University of Groningen/UMCG research database](#)

*Citation for published version (APA):*

Rožanec, J., Montini, E., Cutrona, V., Papamartzivanos, D., Klemencic, T., Fortuna, B., Mladenec, D., Veliou, E., Giannetsos, T., & Emmanouilidis, C. (2023). Human in the AI loop via xAI and Active Learning for Visual Inspection. In J. Soldatos (Ed.), *Artificial Intelligence in Manufacturing: Enabling Intelligent, Flexible and Cost-Effective Production Through AI* (pp. 381-406). Springer. [https://doi.org/10.1007/978-3-031-46452-2\\_22](https://doi.org/10.1007/978-3-031-46452-2_22)

#### Copyright

Other than for strictly personal use, it is not permitted to download or to forward/distribute the text or part of it without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license (like Creative Commons).

The publication may also be distributed here under the terms of Article 25fa of the Dutch Copyright Act, indicated by the "Taverne" license. More information can be found on the University of Groningen website: <https://www.rug.nl/library/open-access/self-archiving-pure/taverne-amendment>.

#### Take-down policy

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

*Downloaded from the University of Groningen/UMCG research database (Pure): <http://www.rug.nl/research/portal>. For technical reasons the number of authors shown on this cover page is limited to 10 maximum.*

# Human in the AI Loop via xAI and Active Learning for Visual Inspection



**Jože M. Rožanec, Elias Montini, Vincenzo Cutrona,  
Dimitrios Papamartzivanos, Timotej Klemenčič, Blaž Fortuna,  
Dunja Mladenčić, Entso Veliou, Thanassis Giannetsos,  
and Christos Emmanouilidis**

## 1 Introduction

Industrial revolutions have historically disrupted manufacturing by introducing automation into the production process. Increasing automation changed worker responsibilities and roles. While past manufacturing revolutions were driven from the optimization point of view, the Industry 5.0 concepts capitalize on the technological foundations of Industry 4.0 to steer manufacturing toward human-centricity [32, 67], adding resilience and sustainability among its key targets [29]. This change is part of a holistic understanding of the industry's societal role. In particular, the

---

J. M. Rožanec (✉)

Jožef Stefan Institute, Ljubljana, Slovenia

Qlector d.o.o., Ljubljana, Slovenia

e-mail: [joze.rozanec@ijs.si](mailto:joze.rozanec@ijs.si)

E. Montini · V. Cutrona

University of Applied Science of Southern Switzerland, Manno, Switzerland

D. Papamartzivanos · T. Giannetsos

Ubitech Ltd., Athens, Greece

T. Klemenčič

University of Ljubljana, Ljubljana, Slovenia

B. Fortuna

Qlector d.o.o., Slovenia

D. Mladenčić

Jožef Stefan Institute, Ljubljana, Slovenia

E. Veliou

University of West Attica, Aigaleo, Greece

C. Emmanouilidis

University of Groningen, Groningen, The Netherlands

© The Author(s) 2024

J. Soldatos (ed.), *Artificial Intelligence in Manufacturing*,

[https://doi.org/10.1007/978-3-031-46452-2\\_22](https://doi.org/10.1007/978-3-031-46452-2_22)

European Commission expects the industry to collaborate on achieving societal goals that transcend jobs and company growth.

Human-centric manufacturing within the Industry 5.0 aims to ensure that human well-being, needs, and values are placed at the center of the manufacturing process. Furthermore, it seeks to enable collaborative intelligence between humans and machines to enable co-innovation, co-design, and co-creation of products and services [62], thus allowing leveraging on their strengths to maximize individual and joint outcomes and their joint added value [34]. It is expected that synergies enabled within Industry 5.0 will still allow for high-speed and mass-personalized manufacturing but will shift repetitive and monotonous tasks to be more assigned to machines to capitalize more on the human propensity for critical thinking and give them to more cognitively demanding tasks [71].

The emerging shift in human roles goes beyond allowing them to move away from repetitive tasks to undertake other physical activities. As non-human actors, including artificial intelligence (AI) - enabled ones, undertake tasks that can be automated, humans are not necessarily excluded but may well play a higher added value and steering role, bringing their cognitive capabilities into the AI loop [33]. This includes active synergies between AI-enabled non-human entities and humans, resulting in novel work configurations [41]. Such configurations empower human actors in new roles rather than diminishing them [8]. As a consequence, it is increasingly recognized that involving instead of replacing the human from the AI loop not only elevates the role of humans in such work environments but also significantly enhances the machine learning process, and, therefore, the emergent capabilities of the AI-enabled actors [77]. As a result, such synergies involve humans and non-human entities who jointly contribute to shaping an emergent meta-human learning system, which in turn is more capable and powerful than human and non-human entities acting alone [70].

A possible realization of such human-machine collaboration emerges from two sub-fields of artificial intelligence: active learning and explainable artificial intelligence (XAI). Active learning is concerned with finding pieces of data that allow machine learning algorithms to learn better toward a specific goal. Human intervention is frequently required, e.g., to label selected pieces of data and enable such learning. On the other hand, XAI aims to make the machine learning models intelligible to the human person so that humans can understand the rationale behind machine learning model predictions. While active learning requires human expertise to teach machines to learn better, XAI aims to help humans learn better about how machines learn and think. This way, both paradigms play on the strengths of humans and machines to realize synergistic relationships between them.

Among the contributions of the present work are (i) a brief introduction to the state-of-the-art research on human-machine collaboration, key aspects of trustworthiness and accountability in the context of Industry 5.0, and research related to automated visual inspection; (ii) the development of a vision on how an AI-first human-centric visual inspection solution could be realized; and (iii) a description of experiments and results obtained in the field of automated visual inspection at the EU H2020 STAR project.

The rest of the work is structured as follows: Sect. 2 describes related work, providing an overview of human-machine collaboration, the industry 5.0 paradigm and human-centric manufacturing, state-of-the-art on automated quality inspection, and a vision of how human-machine collaboration can be realized in the visual inspection domain. In Sect. 3, relevant research contributions from the EU H2020 STAR project are outlined, offering concrete examples of humans and AI working in synergy. Finally, Sect. 4 provides conclusions and insights into future work.

## 2 Background

### 2.1 Overview on Human-Machine Collaboration

The advent of increasingly intelligent machines has enabled a new kind of relationship: the relationship between humans and machines. Cooperative relationships between humans and machines were envisioned back in 1960 [39, 65]. This work defines machines in a broad sense, considering intelligent systems that can make decisions autonomously and independently (e.g., automated, autonomous, or AI agents, robots, vehicles, and instruments) [39, 83, 114]. Relationships between humans and machines have been characterized through different theories, such as the Socio-Technical Systems theory (considers humans and technology shape each other while pursuing a common goal within an organization), Actor-Network theory (considers machines should be equally pondered by humans when analyzing a social system, considering the later as an association of heterogeneous elements), Cyber-Physical Social Systems theory (extends the Socio-Technical Systems theory emphasizing social dimensions where computational algorithms are used to monitor devices), the theory on social machines (considers systems that combine social participation with machine-based computation), and the Human-Machine Networks theory (considers humans and machines form interdependent networks characterized by synergistic interactions). The first three theories conceptualize humans and machines as a single unit, while the last two consider social structures mediated in human-machine networks. In particular, the Socio-Technical Systems theory considers humans and technology shape each other while pursuing a common goal within an organization. The Cyber-Physical Social Systems theory extends this vision, emphasizing social dimensions where computational algorithms are used to monitor and control devices. Moreover, the Actor-Network theory conceptualizes the social system as an association of heterogeneous elements and advocates that machines should be equally pondered to humans. The theory of social machines is interested in systems that combine social participation with machine-based computation. In contrast, the Human-Machine Networks theory considers humans and machines to form interdependent networks characterized by synergistic interactions. A thorough analysis of the abovementioned concepts can be found in [105].

Regardless of the particular theory, the goal remains the same: foster and understand mutualistic and synergistic relationships between humans and machines, where the strengths of both are optimized toward a common goal to achieve what was previously unattainable to each of them. To that end, individual roles must either be clearly defined or allow for a clear sliding of roles when a role can be shared among different types of actors. This will ensure a dynamic division of tasks, optimal use of resources, and reduced processing time. Machines are aimed at supporting, improving, and extending human capabilities. The joint outcomes of human-machine collaboration can result in systems capable of creativity and intuitive action to transcend mere automation. Communication is a critical aspect of every social system. Therefore, emphasis must be placed on the interaction interfaces between such actors. To make such interfaces effective, the concept of shared context or situation awareness between collaborating agents becomes essential and can be seen as a form of mutual understanding [33]. This shared context is enabled through interaction communication of different modalities, including direct verbal (speech, text) and non-verbal (gestures, action and intention recognition, emotions recognition). On the other hand, means must be designed so that humans can understand the machine's goals and rationale for acting to reach such goals in a human-like form. In this regard, human-machine interfaces to support multi-modal interaction play a crucial role. These aspects were also identified by Jwo et al. [53], who described the 3I (Intellect, Interaction, and Interface) aspects that must be considered for achieving human-in-the-loop smart manufacturing.

Beyond shared context, human-machine cooperation requires adequate communication and shared or sliding control [101]. To realize an effective bidirectional information exchange, theory and methods must address how data and machine reasoning can be presented intuitively to humans. Frameworks and models abstracting human cognitive capabilities [60] are key to achieving this. Aligning the design of interactive interfaces and support tools for human-machine interactions with such concepts can be critically important for making effective human-machine interfaces. Enhancements in the interactivity, multisensitivity, and autonomy of feedback functions implemented on such interfaces allow for deeper integration between humans and machines. Shared control can be articulated at operational, tactical, and strategic levels, affecting information-gathering, information-analysis, decision-making, and action implementation.

Human-machine interactions can be viewed from multiple perspectives, necessitating a thorough consideration of several factors influencing such collaborations. These factors encompass emotional and social responses, task design and assignment, trust, acceptance, decision-making, and accountability [23]. Notably, research indicates that machines in collaborative settings impact human behavior, resulting in a diminished emotional response toward them. Consequently, this reduced emotional response can foster more rational interactions. Moreover, studies reveal that humans perceive a team more favorably when machines acknowledge and admit their errors. Additionally, the absence of social pressure from humans can detrimen-

tally affect overall human productivity. Furthermore, concerning accountability for decision-making, humans tend to shift responsibility onto machines.

Trust, a critical aspect to consider, has been explored extensively. Studies demonstrate that trust in machines is closely linked to perceived aptness [23]. Instances of machine errors often lead to a loss of trust, particularly when machines act autonomously. However, if machines operate in an advisory capacity, trust can be amended over time. Additionally, research reveals that while humans value machine advice, they hesitate to relinquish decision-making authority entirely. Nevertheless, relying excessively on machines can result in sub-optimal outcomes, as humans may fail to identify specific scenarios that necessitate their attention and judgment. For further details about the abovementioned experiments and additional insights the reader may be interested on the works by Chugunova et al. [23].

## ***2.2 Industry 5.0 and Human-Centric Manufacturing***

### **2.2.1 New Technological Opportunities to Reshape the Human Workforce**

Digital transformation in production environments demands new digital skills and radically reshapes the roles of plant and machine operators [22, 104]. While Industry 4.0 emphasizes the use of technologies to interconnect different stages of the value chain and the use of data analytics to increase productivity, Industry 5.0 emphasizes the role of humans in the manufacturing context [16, 54]. Furthermore, it aims to develop means that enable humans to work alongside advanced technologies to enhance industry-related processes [68]. An extensive review of this paradigm and its components was written by Leng et al. [62]. Nevertheless, two components are relevant to this work: collaborative intelligence and multi-objective interweaving. Collaborative Intelligence is the fusion of human and AI [110]. In the context of Industry 5.0, the fusion of both types of intelligence entails the cognitive coordination between humans and AI in machines, enabling them to collaborate in the innovation, design, and creation of tailored products and services. Complementarity between humans and AI (see Table 1) leads to the more effective execution of such tasks than would be possible if relegated to humans or machines only [18, 51, 75, 86].

When analyzing complementarities, humans have the knowledge and skills to develop and train machines by framing the problems to be solved and providing feedback regarding their actions or outputs [50, 62, 82, 112]. Furthermore, humans can enrich machine outcomes by interpreting results and insights and deciding how to act upon them [113]. Machines amplify workers' cognitive abilities: they can track many data sources and decide what information is potentially relevant to humans. Furthermore, machines can excel at repetitive tasks and free humans from such a burden. Such complementarity is considered within the multi-objective interweaving nature of Industry 5.0, which enables optimizing multiple goals beyond process performance and social and environmental sustainability [13].

**Table 1** Overview on humans and AI complementarity (adapted from [50, 62, 82], and complemented with our observations). x: capability completely fulfilled; o: capability partially fulfilled

Capability		Humans	Machines
Strengths and capabilities	Leadership	x	
	Teamwork	x	
	Creativity	x	o
	Problem-solving	x	x
	Risk assessment	x	o
	Intuition	x	
	Interpretation	x	
	Empathy	x	
	Adapt behavior	x	o
	Learn from experience	x	o
	Speed		x
	Scalability		x
	Endurance		x
	Quantitative accuracy		x
	Process large amounts of data		x
	Process different kinds of data in parallel	x	
	Perform continuous operations		x
	Consistent decision-making		x
	Physical and cognitive abilities	x	o
Weaknesses	Prone to biases and errors	x	
	Affected by emotions	x	
	Affected by distractions	x	
	Prone to frauds and adversarial attacks	o	x
	Affected by fatigue	x	
	Limited to certain scope and goals		x
	Lack of emotional intelligence		x
	Lack of social skills		x

Moreover, research suggests that leading companies are beginning to recognize the benefit of using machines and automation systems to supplement human labor rather than replacing the human workforce entirely [5, 28]. While AI was already able to tackle certain tasks with super-human capability [24], it has recently shown progress in areas such as creativity (e.g., through generative models such as DALL·E 2 [84]) or problem-solving [20], opening new frontiers of human-machine collaboration, such as co-creativity [9, 64].

In addition to the direct human involvement described above, digital twins [74] are another way to incorporate human insights into the AI processes. By creating virtual models of human behavior and mental processes, more profound insights into how humans interact with the world and use this information to improve AI systems. Digital twins can also support explainability and transparency in AI

systems, making explaining how they arrive at their decisions easier [12]. Moreover, digital representations can be used to consider users' preferences in the AI system behaviors, e.g., type of support [48, 106].

### 2.2.2 Trustworthiness and Implications for AI-Driven Industrial Systems

Trustworthiness for systems and their associated services and characteristics is defined according to the International Organization for Standardization (ISO) as “*the ability to meet stakeholders' expectations in a verifiable way*” [3]. It follows that trustworthiness can refer to products, services, technology, and data and, ultimately, to organizations. Therefore, the concept of trustworthiness is directly applicable to AI-driven systems, particularly to human-centric AI-enabled solutions. However, it should be understood that trustworthiness is a multifaceted concept, incorporating distinct characteristics such as accountability, accuracy, authenticity, availability, controllability, integrity, privacy, quality, reliability, resilience, robustness, safety, security, transparency, and usability [3].

Some of these characteristics should be seen as emerging characteristics of AI-enabled systems, which are not solely determined by the AI's contribution to an overall solution. Focusing specifically on the AI components of such solutions, ethics guidelines published by the European Commission (EC) identifies seven key requirements for trustworthiness characteristics that must be addressed [35]. These include (i) human agency and oversight; (ii) technical robustness and safety; (iii) privacy and data governance; (iv) transparency; (v) diversity, non-discrimination, and fairness; (vi) societal and environmental well-being; and (vii) accountability. Regarding some of these characteristics, there is a direct correspondence between broader trustworthiness as documented according to ISO and the EC guidelines. Technical robustness, safety, privacy, transparency, and accountability are identified in both sources. Human agency and oversight are directly linked to controllability, and so is governance, which is also the prime focus of ISO recommendations [2]. Given the societal impacts that AI-induced outcomes can have, the EC has also highlighted diversity, non-discrimination, fairness, and societal and environmental well-being as key characteristics of trusted AI solutions. However, these aspects are also partly addressed as part of the broader concept of “freedom from risk”, which can be defined as the extent to which a system avoids or mitigates risks to economic status, human life, health, and well-being and or the environment [1].

The trustworthiness of an AI system can be affected by multiple factors. Some of them relate to cybersecurity. In particular, machine learning algorithms are vulnerable to poison and evasion attacks. During poisoning attacks, the adversary aims to tamper with the training data used to create the machine learning models and distort the AI model on its foundation [42, 96]. Evasion attacks are performed during inference, where the attacker crafts adversarial inputs that may seem normal to humans but drive the models to classify the inputs wrongly [49, 72]. Such an adversarial landscape poses significant challenges and requires a collaborative approach between humans and machines to build defenses that can lead to more



robust and trustworthy AI solutions. While human intelligence can be used for the human-in-the-loop adversarial generation, where humans are guided to break models [108], AI solutions can be trained to detect adversarial inputs and uncover potentially malicious instances that try to evade the AI models [10]. Furthermore, human-machine collaboration can be fostered to detect such attacks promptly.

Accountability refers to the state of being accountable and relates to allocated responsibility [3]. At the system level, accountability is a property that ensures that the actions of an entity can be traced uniquely to the entity [4]. However, when considering governance, accountability is the obligation of an individual or organization to account for its activities, accept responsibility for them, and disclose the results in a transparent manner [2]. Therefore, accountability is closely linked to transparency for AI-enabled systems, which is served via XAI and interpretable AI. XAI and interpretable AI ensure that AI systems can be trusted when analyzing model outcomes that impact costs and investments or whenever their outputs provide information to guide human decision-making. Accuracy generally refers to the closeness of results and estimates to true values but, in the context of AI, further attains the meaning appropriate for specific machine learning tasks. Any entity that is what it claims to be is said to be characterized by authenticity, with relevant connotations for what AI-enabled systems claim to deliver. Such systems may furthermore be characterized by enhanced availability to the extent that they are usable on demand. Other characteristics such as integrity, privacy, and security attain additional meaning and importance in AI-driven systems and are further discussed in the next section. They can contribute to and affect the overall quality, reliability, resilience, robustness, and safety, whether the unit of interest is a component, a product, a production asset, or a service, with implications for individual workers all the way to the organization as a whole. When considering accountability for AI systems from the legal perspective, the EU AI Act [36] in its current form considers developers and manufacturers responsible for AI failures or unexpected outcomes. Nevertheless, the concept of accountability will evolve based on the issues found in practice and the corresponding jurisprudence that will shape the learning on how different risks, contexts, and outcomes must be considered in the industry context [46].

## ***2.3 Automated Quality Inspection***

### **2.3.1 The Role of Robotics**

The increasing prevalence of human-robot collaboration in diverse industries showcases the efforts to enhance workplace productivity, efficiency, and safety through the symbiotic interaction of robots and humans [44]. In manufacturing, robots are employed for repetitive and physically demanding tasks, enabling human workers to allocate their skills toward more intricate and creative endeavors. This collaborative

partnership allows for the fusion of human and robot capabilities, maximizing the overall outcomes.

The successful implementation of human-robot interaction owes credit to collaborative robots, commonly called cobots [56]. These advanced robots have sophisticated sensors and programming that facilitate safe and intuitive human interaction. This collaboration improves productivity and fosters a work environment where humans and robots can coexist harmoniously. This approach harmoniously merges robots' precision and accuracy with human workers' adaptability and dexterity.

Robotic integration in product quality control has become widespread across diverse industries and production sectors. Robots offer exceptional advantages within quality inspection processes, including precise repeatability and accurate movements [17]. They possess the capability to analyze various product aspects such as dimensions, surface defects, color, texture, and alignment, ensuring adherence to predefined standards. Robots' superior accuracy and efficiency make them an ideal choice for quality control applications.

To facilitate quality testing, robots are equipped with a range of sensors. These sensors enable precise measurement, detection, and sorting operations. Robots with cameras utilize advanced machine vision techniques to analyze image and video streams and identify anomalies like cracks, scratches, and other imperfections [107]. Subsequently, defective items are segregated from conforming ones, elevating overall production quality. The industry is witnessing an increasing adoption of 3D vision systems, particularly in applications requiring object grasping and precise information about object position and orientation.

Specially designed robots, such as coordinate measuring machines, are employed for dimensional and precision measurements. These robots feature high-precision axis encoders and accurate touch probes, enabling them to detect part measurements and consistently evaluate adherence to quality standards [61].

The active learning paradigm can be applied to enable efficient and flexible learning in robots. This can be particularly useful in resource-constrained industrial environments, where data scarcity and limited human knowledge prevail, acquiring essential data through unsupervised discovery becomes imperative [26]. Active learning demonstrates extensive applicability in robotics, encompassing prioritized decision-making, inspection, object recognition, and classification. Within quality control, active learning algorithms optimize machine learning models' defect detection and quality assessment training process. By actively selecting informative samples for labeling, active learning minimizes labeling efforts, augments model training efficiency, and ultimately enhances the accuracy and performance of quality control systems.

An intriguing domain of investigation pertains to the advancement of intuitive and natural interfaces that foster seamless communication and interaction between humans and robots. This entails the exploration of innovative interaction modalities, encompassing speech, gestures, and facial expressions, or even using augmented reality to customize the robots' appearance and foster better interaction with humans [59]. Other key research areas involve developing adaptive and flexible robotic

systems that dynamically adapt their behavior and actions to the prevailing context and the human collaborator's preferences, achieving low processing times [78]. These could be critical to enable real-time human intent recognition, situational awareness, and decision-making, all aimed at augmenting the adaptability and responsiveness of robots during collaborative tasks.

### 2.3.2 Artificial Intelligence—Enabled Visual Inspection

Visual inspection is frequently used to assess whether the manufactured product complies with quality standards and allows for the detection of functional and cosmetic defects [21]. It has historically involved human inspectors in determining whether the manufactured pieces are defective. Nevertheless, the human visual system excels in a world of variety and change, while the visual inspection process requires repeatedly observing the same product type. Furthermore, human visual inspection suffers from poor scalability and the fact that it is subjective, creating an inherent inspector-to-inspector inconsistency. The quality of visual inspection can be affected by many factors. See [94] classified them into five categories, whether they are related to the (i) task, (ii) individual, (iii) environment, (iv) organization, or (v) social aspects.

To solve the issues described above, much effort has been invested in automated visual inspection by creating software capable of inspecting manufactured products and determining whether they are defective. Cameras are used to provide visual input. Different approaches have been developed to determine whether a defect exists or not.

Automated optical quality control may target visual features as simple as colors, but more complex ones are involved in crack detection, the orientation of threads, defects in bolts [52] and metallic nuts [14]. Through automated optical inspection, it is also possible to detect defects on product surfaces of wide-ranging sizes [19, 103, 117]. Furthermore, it is also possible to target the actual manufacturing process, for example, welding [102], injection molding [66], or assembly of manufactured components [37]. Additionally, automated visual inspection applies to remanufacturing products at the end of their useful life [91].

State-of-the-art (SOTA) automated visual inspection techniques are dominated by deep learning approaches, achieving high-performance levels [6]. Among the many types of learning from data for visual inspection, unsupervised, weakly supervised, and supervised methods can be named. Unsupervised methods aim to discriminate defective manufactured pieces without labeled data. The weakly supervised approach assumes that data has an inherent cluster structure (instances of the same class are close to each other) and that the data lies in a manifold (nearby data instances have similar predictions). Therefore, it leverages a small amount of annotated data and unlabeled data to learn and issue predictions. Finally, supervised methods require annotated data and usually perform best among the three approaches. Often, labeled data are unavailable in sufficient range and numbers to enable fully supervised learning and additional exemplar images can be

produced through data augmentation [55]. In addition, multiple strategies have been developed to reduce the labeled data required to train and enhance a given classifier. Among them are active learning, generative AI, and few-shot learning. In the context of visual inspection, active learning studies how to select data instances that can be presented to a human annotator to maximize the models' learning. Generative AI aims to learn how to create data instances that resemble a particular class. Finally, few-shot learning aims to develop means by which the learner can acquire experience to solve a specific task with only a few labeled examples. To compensate for the lack of labeled data, it can either augment the dataset with samples from other datasets or use unlabeled data, acquire knowledge on another dataset, or algorithm (e.g., by adapting hyperparameters based on prior meta-learned knowledge) [109].

Regardless of the progress made in automated visual inspection, many challenges remain. First, there is no universal solution for automated visual inspection: solutions and approaches have been developed to target a specific product. Flexibility to address the inspection of multiple manufactured products with a single visual inspection system is a complex challenge and remains an open issue [21, 25, 80]. Second, unsupervised machine learning models do not require annotating data and may provide a certain level of defect detection when associating data clusters to categories (e.g., types of defects or no defects). Furthermore, given that no prior annotation of expected defects is required, they are suitable when various defects exist. Nevertheless, their detection rates are lower than those obtained by supervised machine learning models. Therefore, it should be examined use case by use case whether the unsupervised machine learning models are a suitable solution. Third, data collection and annotation are expensive. While data collection affects unsupervised machine learning models, data collection and annotations directly impact supervised machine learning approaches. While multiple strategies have been envisioned to overcome this issue (e.g., generative models, active learning, and few-shot learning), data collection and annotation remain an open challenge. Finally, better explainability techniques and intuitive ways to convey information to humans must be developed to understand whether the models learn and predict properly.

## ***2.4 Realizing Human-Machine Collaboration in Visual Inspection***

While much progress has been made in automated visual inspection, the authors recognize that most solutions are custom and developed for a particular product type. Developing systems that could adapt to a broad set of products and requirements remain a challenge. In human-centered manufacturing, it is critical to rethink and redesign the role of humans in the visual inspection process. The role of humans in automated visual inspection is shifting away from repetitive and manual tasks to roles with more cognitive involvement, which can still not be replicated by machines

and AI. In the simplest case, this involves humans labeling acquired image samples to guide the machine learning process [98]. However, the role of humans extends beyond data labeling and may involve interaction loops between humans and AI as part of the machine learning process [79].

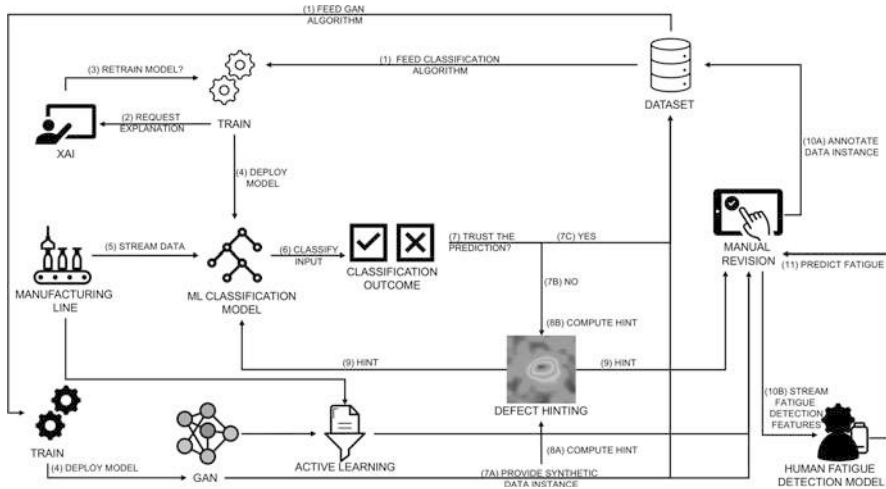
In this regard, two machine learning paradigms are particularly important: active learning and XAI. On the one side, active learning is an AI paradigm that seeks the intervention of an oracle (usually a human person) to help the machine learning model learn better toward an objective. XAI, on the other side, aims to explain the rationale behind a machine learning model action or prediction. Doing so enables a fruitful dialogue between humans and machines by providing insights into the machines' rationale and decision-making process.

Active learning for classification is based on the premises that unlabeled data (either collected or generated) is abundant, the data labeling is expensive, and the models' generalization error can be minimized by carefully selecting new input instances with which the model is trained [95, 99]. Active learning for classification has traditionally focused on the data (selecting or generating the data without further consideration for the model at hand) and the model's learning (e.g., considering the uncertainty at the predicted scores). Nevertheless, approaches have been developed to consider both dimensions and provide a holistic solution. One of them is the Robust Zero-Sum Game (RZSG) framework [119], which attempts to optimize both objectives at once, framing the data selection as a robust optimization problem to find the best weights for unlabeled data to minimize the actual risk, reduce the average loss (to achieve greater robustness to outliers) and minimize the maximal loss (increasing the robustness to imbalanced data distributions). Another perspective has been considered by Zajec et al. [118] and Križnar et al. [57], who aim to select data based on insights provided by XAI methods and, therefore, benefit from direct insights into the model's learning dynamics. Regardless of the approach, Wu et al. [111] propose that three aspects must be considered when searching for the most valuable samples: informativeness (contains rich information that would benefit the objective function), representativeness (how many other samples are similar to it), and diversity (the samples do not concentrate in a particular region but rather are scattered across the whole space). Strategies will be conditioned by particular requirements (e.g., whether the data instances are drawn from a pool of samples or a data stream). For a detailed review of active learning, the reader may be interested in some high-quality surveys of this domain. In particular, the works by Settles [95] and Rožanec et al. [87] can serve as an introduction to this topic. Furthermore, the surveys by Fu et al. [38] and Kumar et al. [58] provide an overview of querying strategies in a batch setting; the survey by Lughofer [69] give an overview of active learning in online settings, and the study by Ren et al. [85] describes active learning approaches related to deep learning models.

While AI models have the potential to automate many tasks and achieve super-human performance levels, in most cases, such models are opaque to humans: their predictions are mostly accurate, but no intuition regarding their reasoning process is conveyed to humans. Understanding the rationale behind a model's prediction is of utmost importance, given it provides a means to assess whether

the predictions are based on accurate facts and intuitions. Furthermore, it is crucial to develop means to understand the model's reasoning process given the impact such techniques have on the real world, either in fully automated settings or when decision-making is delegated to humans. Such insights enable responsible decision-making and accountability. The subfield of AI research developing techniques and mechanisms to elucidate the models' rationale and how to present them to humans is known as XAI. While the field can be traced back to the 1970s [93], it has recently flourished with the advent of modern deep learning [115]. When dealing with XAI, it is important to understand what makes a good explanation. A good explanation must consider at least three elements [11]: (a) reasons for a given model output (e.g., features and their values, how strongly do features influence a forecast, whether the features at which the model looks at make sense w.r.t. the forecast, how did training data influence the model's learning), (b) context (e.g., the data on which the machine learning model was trained, the context on which inference is performed), and (c) how is the abovementioned information conveyed to the users (e.g., target audience, the terminology used by such an audience, what information can be disclosed to it). XAI can be valuable in enhancing human understanding with new (machine-based) perspectives. It can also help to understand whether the model is optimizing for one or few of all required goals and, therefore, identify an appropriate compromise between the different goals that must be satisfied for the problem at hand [30]. To assess the goodness of an explanation, aspects such as user satisfaction, the explanation persuasiveness, the improvement of human judgment, the improvement of human-AI system performance, the automation capability, and the novelty of explanation must be considered [92]. For a detailed review of XAI, the reader may consider the works of Arrieta et al. [11], Doshi-Velez et al. [30], and Schwalbe et al. [92]. The work of Bodria et al. [15] provides a comprehensive introduction to XAI black box methods, and the works of Doshi-Velez et al. [30], Hoffman et al. [45], and Das et al. [27] focus on insights about how to measure the quality of explanations.

Active learning and XAI can complement each other. Understanding the rationale behind a model prediction provides valuable insight to humans and can also be leveraged in an active learning setting. In the particular case of defect inspection, insights obtained by XAI techniques are usually presented in anomaly maps. Such anomaly maps highlight regions of the image the machine learning models consider to issue a prediction. The more perfect the learning of a machine learning model, the better those anomaly maps should annotate a given image indicating defective regions. Therefore, the insights obtained from those anomaly maps can be used in at least two ways. First, the anomaly maps can be handed to the oracle (human inspector), who, aided by the anomaly map and the image of the product, may realize better where the manufacturing errors are, if any. Second, anomaly maps can be used to develop novel models and active learning policies that allow for data selection, considering what was learned by the model and how the model perceives unlabeled data. This approach is detailed in Fig. 1, which depicts how an initial dataset is used to train machine learning models for defect classification or data generation. In the model training process, XAI can be used to debug and iterate the



**Fig. 1** Envisioned setup for an AI-first human-centric visual inspection solution

model until getting satisfactory results. The classification model is then deployed to perform inference on incoming product images from the manufacturing line. If the classification scores for certain classes are high enough, the product can be classified as good or defective. When the uncertainty around the predicted scores is not low enough, the case can be sent for manual revision. Insights obtained through XAI and unsupervised classification models can be used to hint to the human inspector where the defects may be located. Alternative data sources for the manual revision or data labeling process can be generative models (e.g., generative adversarial networks), which can be used to generate labeled synthetic data and validate the level of attention of a human inspector. When collecting data, active learning techniques can be used to select the most promising data instances from either generative models or incoming images from the manufacturing line, reducing the labeling effort. Finally, a separate model can monitor human inspectors to predict fatigue and performance. Such models can be a valuable tool to ensure workplace well-being and enhance work quality. Some of the results obtained within the STAR project are presented in Sect. 3.1.

In recent years, researchers have made significant progress in understanding and quantifying fatigue and recognizing its impact on human performance and overall well-being. Through AI techniques, new approaches have emerged to accurately estimate the fatigue levels of individuals during different tasks and in different contexts [7, 47]. One notable area of inquiry concerns the assessment of fatigue in the workplace. Understanding and managing worker fatigue has become essential given the increasing demands and pressures of modern work environments. AI models can consider various factors and features to assess employee fatigue levels accurately. These models can provide valuable insights for organizations looking to implement strategies and interventions to optimize productivity and ensure



employee well-being or to support workflows, including quality controls, such as identifying when operators need a break. Although laboratory experiments have been conducted in this area [63], industrial applications remain relatively restricted compared to other fields, such as driving [97].

### 3 Industrial Applications

This section briefly describes how some ideas presented in the previous sections have been realized within the EU H2020 STAR project. Three domains are considered: artificial intelligence for visual inspection, digital twins, and cybersecurity.

#### 3.1 Machine Learning and Visual Inspection

In the domain of visual inspection, multiple use cases were considered. The datasets were provided by two industrial partners: *Philips Consumer Lifestyle BV* (Drachten, The Netherlands) and *Iber-Oleff - Componentes Tecnicos Em Plástico, S.A.* (Portugal). The *Philips Consumer Lifestyle BV* manufacturing plant is considered one of Europe's most important Philips development centers and is devoted to producing household appliances. They provided us with three datasets corresponding to different products. The first one corresponded to logo prints on shavers. The visual inspection task required understanding whether the logo was correctly printed or had some printing defect (e.g., double printing or interrupted printing). The second one corresponded to decorative caps covering the shaving head's center, and it required identifying whether the caps were correctly manufactured or if some flow lines or marks existed. Finally, the third dataset was about toothbrush shafts transferring motion from the handle to the brush. It required identifying whether the handles were manufactured without defects or if big dents, small dents, or some stripes could be appreciated. *Iber-Oleff - Componentes Tecnicos Em Plástico, S.A.* provided us with another dataset about automobile air vents they manufacture. The air vents have three components of interest: housing, lamellas (used to direct the air), and plastic links (which keep the lamellas tied together). The visual inspection task required us to determine whether (a) the fork is leaning against the support and correctly positioned, (b) the plastic link is present, (c) the lamella 1 is present, and the link is correctly assembled, and (d) the lamella 3 is present, and the link is correctly assembled.

Through the research, the researchers aimed to develop a comprehensive AI-first and human-centric approach to automated visual inspection. In particular, they (i) developed machine learning models to detect defects, (ii) used active learning to enhance the models' learning process while alleviating the need to label data, (iii) used XAI to enhance the labeling process, (iv) analyzed how data augmentation techniques at embeddings and image level, along with anomaly maps can enhance



the machine learning discriminative capabilities, (v) how human fatigue can be detected and predicted in humans, and (vi) how to calibrate and measure models' calibration quality to provide probabilistic predictive scores.

Research at the EU H2020 STAR project confirmed that active learning could alleviate the need for data labeling and help machine learning models learn better based on fewer data instances [90]. Nevertheless, the effort saved depends on the pool of unlabeled images, the use case, and the active learning strategy. Data augmentation techniques at an image or embedding level have increased the models' discriminative performance [89]. Furthermore, complementing images with anomaly maps as input to supervised classification models has substantially improved discriminative capabilities [88]. The data labeling experiments showed decreased labeling accuracy by humans over time [86], which was attributed to human fatigue. While the future labeling quality can be predicted, it requires ground truth data. This can be acquired by showing synthetically generated images. Nevertheless, more research is required to devise new models that would consider other cues and predict human fatigue in data labeling without the requirement of annotated data. Finally, predictive scores alone provide little information to the decision-maker: predictive score distributions differ across different models. Therefore, performing probability calibration is paramount to ensure probability scores have the same semantics across the models. The research compared some of the existing probability calibration techniques and developed metrics to measure and assess calibration quality regardless of ground truth availability [90].

### ***3.2 Human-Digital Twins in Quality Control***

In the context of STAR, significant advancements have been made in developing human-digital twins (HDTs). In particular, the project has developed an infrastructure (Clawdite Platform [74]) that allows the effortless creation of replicas of human workers through instantiating their digital counterparts. These HDTs have diverse features, encompassing static characteristics, dynamic data, and behavioral and functional models [73].

To ensure a comprehensive representation of the human worker, STAR's HDT incorporates two crucial data types. Firstly, it assimilates physiological data collected from wearable devices. Secondly, it utilizes quasi-static data, which encapsulates characteristic attributes of the human, offering a holistic perspective on their traits. Central to STAR's HDT is an AI model designed to detect mental stress and physical fatigue. By leveraging physiological and quasi-static data, this AI model effectively gauges the stress and fatigue levels experienced by the human worker. This breakthrough in automated quality control holds remarkable significance, manifesting in two distinct ways:

- During user manual inspection, the HDT continuously monitors the quality control process, actively identifying instances where the worker may be under

significant mental or physical stress. In such cases, the system promptly suggests the worker take a break, ensuring their well-being and preventing any potential decline in performance.


- During the training of automatic quality assessment models, as the worker evaluates and labels pictures during the dataset creation, the system periodically assigns a confidence score to each label provided by the user. This confidence score is computed based on evaluating the worker's mental and physical stress levels estimated through the HDT's AI model. By considering these stress levels as an integral part of the quality evaluation process, the HDT provides valuable insights into the worker's state of mind and physical condition, allowing one to consider these features during the training of AI models for quality assessment and control.

The integration of the HDTs, supported by the Clawdite Platform, in STAR's operations signifies a significant step forward in human-AI collaboration. This innovative approach prioritizes human workers' well-being and empowers automated quality control systems, ensuring optimal productivity and efficiency in various industrial settings.

### ***3.3 Making AI Visual Inspection Robust Against Adversarial Attacks***

In the context of the STAR project, an AI architecture was created for evaluating adversarial tactics and defense algorithms intended to safeguard, secure, and make the environments of manufacturing AI systems more reliable. More specifically, it was focused on AI-based visual inspection and tackled multiple use cases provided by two industrial partners: *Philips Consumer Lifestyle BV* (Drachten, The Netherlands) and *Iber-Oleff - Componentes Tecnicos Em Plástico, S.A.* (Portugal). Current production lines are often tailored for the mass production of one product or product series in the most efficient way. Given its many advantages, AI is being increasingly adopted for quality inspection. Such models are usually trained considering some convolutional neural network (CNN), which then classifies whether a product is defective through inference upon receiving images captured by the inspection cameras. Nevertheless, such models can be attacked through adversarial data, leading AI models to wrongly classify the products (e.g., not detecting defects). For instance, the adversary may exploit a vulnerability in the visual inspection camera and compromise the integrity of the captured data by manipulating the operational behavior of this business resource.

Among the various experimental testbeds built in the context of the STAR project, the ones created with soother cherries provided by *Philips Consumer Lifestyle BV* were the most challenging. The cherry is the upper part of the soother. The high quality of the cherry must be guaranteed to avoid any harm to the babies. Therefore, detecting any adversarial attack is of primary importance, given the



**Fig. 2** Evaluation results of pairwise comparison of adversarial attacks and defenses

consequences of the attack can directly impact children’s health. The goal of the testbed was to quantify the impact of adversarial attacks on classification models performing a visual inspection and evaluate how effective the defenses against such attacks were. To build the testbed, the Adversarial Robustness Toolbox [81] was used. In the experiments, the following adversarial methods were used: Fast Gradient Sign Attack (FGSM) [40], DeepFool [76], NewtonFool [49], and Projected Gradient Descent (PGD) [72]. The aim was to utilize these well-documented adversarial methods to derive crafted instances that can be used to attack the baseline classification model. Experiments were performed with defense strategies, namely feature squeezing [116], JPEG compression [31], spatial smoothing [116], TotalVarMin [43], and adversarial training [100].

To gather insights regarding adversarial tactics and defenses, they were evaluated pairwise. This enabled us to identify adversarial training as the best defense strategy to enhance the robustness of the CNN models. The basic idea behind adversarial training is to create examples that will be used later in the training process, creating a model aware of adversarial vectors launched against the quality control system. The results of the pairwise evaluation of the attacks and defenses are summarized in Fig. 2. The results are grouped into four sets based on the attack strategy. A baseline classifier was initially trained for each of the four experiments (see tag “Training”) to get the perception of the accuracy level that the quality inspection algorithm can achieve. The baseline model achieved an accuracy between 93% and 98%. The “Attack” bar indicates the accuracy of the classifier when posed against the adversarial attack. The DeepFool, FGSM, and PGD attacks strongly affected the classifier, causing the model’s accuracy to drop below 30%. This was not the case for the NewtonFool attack, where the classifier’s accuracy dropped to 84%. When considering defense strategies, feature squeezing, JPEG compression, and spatial smoothing can defend against the DeepFool attack: for the given dataset, they

led to an accuracy of 98%. However, TotalVarMin failed to defend the model. All the defenses failed against the FGSM and the PGD attacks. Based on the acquired results of the pairwise evaluations, it became clear that no clear mapping exists between types of attacks and defenses. Therefore, it can be challenging for defenders to plan a strategy to cope against any attack successfully. This outcome advocates the criticality and the challenge of defending against adversarial AI attacks. While the off-the-shelf and state-of-the-art defenses cannot perform in a stable manner under different adversarial methods, the adversarial training approach seems robust. The results agree with the literature, advocating that Adversarial Training can be a robust solution that can cope with adversaries despite its simplicity. A more detailed description of the abovementioned work can be found in [10].

## 4 Conclusion

This work has briefly introduced state-of-the-art research on human-machine collaboration, perspectives on human-centric manufacturing, and the key aspects of trustworthiness and accountability in the context of Industry 5.0. It described research on automated quality inspection, considering the role of robotics, AI approaches, and solutions to visual inspection and how a fruitful human-machine collaboration can be developed in the visual inspection domain. Finally, it described the experience and results obtained through research performed in the EU H2020 STAR project.

The converging view from the literature analysis is that human-machine cooperation requires adequate communication and control realized through effective bidirectional information exchange. Studies have been performed to understand peoples' emotional and social responses in human-machine interactions, understand task design, and how humans' trust, acceptance, decision-making, and accountability are developed or impacted in the presence of machines. In the field of visual inspection, much research was invested in automating the task of visual inspection by developing machine learning models to detect product defects. Furthermore, many research efforts targeted the development of techniques for XAI related to machine vision. Visual aids and hints derived by XAI are conveyed to humans through heat maps. Similarly, insights obtained from unsupervised machine learning models are conveyed to humans as anomaly maps. While such approaches solve particular problems, little research describes how a human-in-the-loop approach could be developed for visual inspection in manufacturing settings. This research aims to bridge the gap by implementing existing and researching novel active learning techniques for data selection to enhance the learning of machine learning algorithms. It also explores how labeling requirements could be reduced by employing few-shot learning and active learning techniques. Furthermore, research was conducted to understand how XAI and unsupervised classification methods can be used to generate heat maps and anomaly maps to facilitate data labeling in the context of manual revision or data annotation tasks. Moreover, predictive

models were developed to predict how heat maps and anomaly maps should be adapted over time to bridge the gap between the information conveyed by machine learning algorithms and explainability techniques and human perception. In addition, experiments were performed to gain insights related to human fatigue monitoring in the context of visual inspection. The present work described a complete and modular infrastructure developed to instantiate HDT, and different AI models for perceived fatigue exertion and mental stress have been trained to derive relevant features for human-centered production systems. Finally, it describes some research on adversarial attacks and defenses to enhance the understanding of protecting visual inspection setups in manufacturing environments.

While the research presented above advances the understanding of developing a human-in-the-loop approach for visual inspection in manufacturing, many open issues remain to be solved. Further research is required to understand how adaptive humans perceive hinting and how the many solutions described above contribute to building trust between humans and machines. Furthermore, effort must be invested to quantify the benefits such solutions bring to a manufacturing plant when implemented. Future research will encompass the integration of these solutions, aiming to achieve a comprehensive and synergistic implementation. The research will aim to develop new approaches that interleave active learning and XAI. Furthermore, novel few-shot learning solutions will be considered to allow for greater flexibility of the visual inspection while reducing data labeling requirements to a minimum. Finally, integrating AI visual inspection models and the HDT is expected to significantly augment the efficacy of quality inspection processes during user manual assessment and AI model training.

**Acknowledgments** This work was supported by the Slovenian Research Agency and the European Union's Horizon 2020 program project STAR under grant agreement number H2020-956573.

## References

1. 25022:2015 I: System and software quality requirements and evaluation (square) - measurement of quality in use (2015)
2. 38500:2015 I: Information technology — governance of it for the organization (2015)
3. 5723 IT: Trustworthiness - vocabulary (2022)
4. 7498-2:1989 I: Information processing systems - security architecture (1989)
5. Accenture: Reworking the revolution (2018). [https://www.accenture.com/\\_acnmedia/pdf-69/accenture-reworking-the-revolution-jan-2018-pov.pdf](https://www.accenture.com/_acnmedia/pdf-69/accenture-reworking-the-revolution-jan-2018-pov.pdf)
6. Aggour, K.S., Gupta, V.K., Ruscitto, D., et al.: Artificial intelligence/machine learning in manufacturing and inspection: A GE perspective. *MRS Bull.* **44**(7), 545–558 (2019)
7. Aguirre, A., Pinto, M.J., Cifuentes, C.A., et al.: Machine learning approach for fatigue estimation in sit-to-stand exercise. *Sensors* **21**(15), 5006 (2021)
8. Amershi, S., Cakmak, M., Bradley Knox, W., Kulesza, T.: Power to the people: the role of humans in interactive machine learning. *AI Mag.* **35**, 105–120 (2014). <https://doi.org/10.1609/aimag.v35i4.2513>

9. Anantrasirichai, N., Bull, D.: Artificial intelligence in the creative industries: a review. *Artif. Intell. Rev.*, 1–68 (2022)
10. Anastasiou, T., Karagiorgou, S., Petrou, P., et al.: Towards robustifying image classifiers against the perils of adversarial attacks on artificial intelligence systems. *Sensors* **22**(18), (2022). <https://doi.org/10.3390/s22186905>, <https://www.mdpi.com/1424-8220/22/18/6905>
11. Arrieta, A.B., Díaz-Rodríguez, N., Del Ser, J., et al.: Explainable artificial intelligence (XAI): Concepts, taxonomies, opportunities and challenges toward responsible AI. *Inf. Fusion* **58**, 82–115 (2020)
12. Bansal, G., Nushi, B., Kamar, E., et al.: Beyond accuracy: The role of mental models in human-AI team performance. In: *Proceedings of the AAAI Conference on Human Computation and Crowdsourcing*, pp. 2–11 (2019)
13. Bettoni, A., Montini, E., Righi, M., et al.: Mutualistic and adaptive human-machine collaboration based on machine learning in an injection moulding manufacturing line. *Procedia CIRP* **93**, 395–400 (2020)
14. Bharti, S., McGibney, A., O’Gorman, T.: Edge-enabled federated learning for vision based product quality inspection. In: *2022 33rd Irish Signals and Systems Conference (ISSC)*, pp. 1–6 (2022). <https://doi.org/10.1109/ISSC55427.2022.9826185>
15. Bodria, F., Giannotti, F., Guidotti, R., et al.: Benchmarking and survey of explanation methods for black box models. *Data Mining Knowl. Discovery*, 1–60 (2023)
16. Breque, M., De Nul, L., Petridis, A.: *Industry 5.0: Towards a sustainable, human-centric and resilient European industry* (2021)
17. Brito, T., Queiroz, J., Piardi, L., et al.: A machine learning approach for collaborative robot smart manufacturing inspection for quality control systems. *Procedia Manuf.* **51**, 11–18 (2020)
18. Cai, C.J., Reif, E., Hegde, N., et al.: Human-centered tools for coping with imperfect algorithms during medical decision-making. In: *Proceedings of the 2019 Chi Conference on Human Factors in Computing Systems*, pp. 1–14 (2019)
19. Cao, G., Ruan, S., Peng, Y., et al.: Large-complex-surface defect detection by hybrid gradient threshold segmentation and image registration. *IEEE Access* **6**, 36235–36246 (2018). <https://doi.org/10.1109/ACCESS.2018.2842028>
20. Cao, L.: A new age of AI: Features and futures. *IEEE Intell. Syst.* **37**(1), 25–37 (2022)
21. Chin, R.T., Harlow, C.A.: Automated visual inspection: A survey. *IEEE Trans. Pattern Anal. Mach. Intell.* **4**(6), 557–573 (1982)
22. Chuang, S.: Indispensable skills for human employees in the age of robots and AI. *Eur. J. Train. Dev.* (ahead-of-print) (2022)
23. Chugunova, M., Sele, D.: We and it: An interdisciplinary review of the experimental evidence on human-machine interaction. *Center for law & economics working paper series* 12 (2020)
24. Ciregan, D., Meier, U., Schmidhuber, J.: Multi-column deep neural networks for image classification. In: *2012 IEEE Conference on Computer Vision and Pattern Recognition*, pp 3642–3649. *IEEE* (2012)
25. Czimmermann, T., Ciuti, G., Milazzo, M., et al.: Visual-based defect detection and classification approaches for industrial applications—a survey. *Sensors* **20**(5), 1459 (2020)
26. Daniel, C., Viering, M., Metz, J., et al.: Active reward learning. In: *Robotics: Science and Systems* (2014)
27. Das, A., Rad, P.: Opportunities and challenges in explainable artificial intelligence (XAI): A survey. Preprint (2020). [arXiv:2006.11371](https://arxiv.org/abs/2006.11371)
28. Deloitte: The rise of the social enterprise. 2018 Deloitte global human capital trends (2018). [https://www2.deloitte.com/content/dam/insights/us/articles/HCTrends2018/2018-HCTrends\\_Rise-of-the-social-enterprise.pdf](https://www2.deloitte.com/content/dam/insights/us/articles/HCTrends2018/2018-HCTrends_Rise-of-the-social-enterprise.pdf)
29. Directorate-General for Research and Innovation, European Commission, Breque, M., De Nul, L., Petridis, A.: *Industry 5.0: towards a sustainable, human centric and resilient European industry*. Publications Office of the European Union (2021). <https://op.europa.eu/en/publication-detail/-/publication/468a892a-5097-11eb-b59f-01aa75ed71a1/language-en>

30. Doshi-Velez, F., Kim, B.: Towards a rigorous science of interpretable machine learning. Preprint (2017). arXiv:170208608
31. Dziugaite, G.K., Ghahramani, Z., Roy, D.M.: A study of the effect of jpg compression on adversarial images. Preprint (2016). arXiv:160800853
32. EESC: Industry 5.0 (2018). <https://www.eesc.europa.eu/en/agenda/our-events/events/industry-50>, Accessed: 24 May 2023
33. Emmanouilidis, C., Pistofidis, P., Bertonecelj, L., et al.: Enabling the human in the loop: Linked data and knowledge in industrial cyber-physical systems. *Annu. Rev. Control* **47**, 249–265 (2019). <https://doi.org/10.1016/j.arcontrol.2019.03.004>
34. Emmanouilidis, C., Waschull, S., Bokhorst, J.A., et al.: Human in the AI Loop in Production Environments, vol. 633 IFIP. Springer Science and Business Media Deutschland GmbH, pp. 331–342 (2021). [https://doi.org/10.1007/978-3-030-85910-7\\_35](https://doi.org/10.1007/978-3-030-85910-7_35)
35. European Commission: Ethics guidelines for trustworthy AI (2019)
36. European Commission: Laying down harmonised rules on artificial intelligence (artificial intelligence act) and amending certain union legislative acts (2021)
37. Frustaci, F., Spagnolo, F., Perri, S., et al.: Robust and high-performance machine vision system for automatic quality inspection in assembly processes. *Sensors* **22**, 2839 (2022). <https://doi.org/10.3390/s22082839>, <https://www.mdpi.com/1424-8220/22/8/2839>
38. Fu, Y., Zhu, X., Li, B.: A survey on instance selection for active learning. *Knowl. Inf. Syst.* **35**(2), 249–283 (2013)
39. Gerber, A., Derckx, P., Döppner, D.A., et al.: Conceptualization of the human-machine symbiosis—a literature review. In: Proceedings of the 53rd Hawaii International Conference on System Sciences (2020)
40. Goodfellow, I.J., Shlens, J., Szegedy, C.: Explaining and harnessing adversarial examples. In: Bengio Y, LeCun Y (eds) 3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7–9, 2015, Conference Track Proceedings (2015)
41. Grønsund, T., Aanestad, M.: Augmenting the algorithm: Emerging human-in-the-loop work configurations. *J. Strat. Inf. Syst.* **29**, 101614 (2020). <https://doi.org/10.1016/j.jsis.2020.101614>
42. Gu, T., Dolan-Gavitt, B., Garg, S.: Badnets: Identifying vulnerabilities in the machine learning model supply chain. Preprint (2017). arXiv:170806733
43. Guo, C., Rana, M., Cisse, M., et al.: Countering adversarial images using input transformations. In: International Conference on Learning Representations (2018). <https://openreview.net/forum?id=SyJ7CIWcb>
44. Heyer, C.: Human-robot interaction and future industrial robotics applications. In: 2010 IEEE/RSJ International Conference on Intelligent Robots and Systems, pp 4749–4754. IEEE (2010)
45. Hoffman, R.R., Mueller, S.T., Klein, G., et al.: Metrics for explainable ai: Challenges and prospects. Tech. rep., DARPA Explainable AI Program (2018)
46. Hohma, E., Boch, A., Trauth, R., et al.: Investigating accountability for artificial intelligence through risk governance: A workshop-based exploratory study. *Front. Psychol.* **14**, 86 (2023)
47. Hooda, R., Joshi, V., Shah, M.: A comprehensive review of approaches to detect fatigue using machine learning techniques. *Chronic Dis. Transl. Med.* **8**(1), 26–35 (2022)
48. Hu, Z., Lou, S., Xing, Y., et al.: Review and perspectives on driver digital twin and its enabling technologies for intelligent vehicles. *IEEE Trans. Intell. Veh.* (2022)
49. Jang, U., Wu, X., Jha, S.: Objective metrics and gradient descent algorithms for adversarial examples in machine learning. In: Proceedings of the 33rd Annual Computer Security Applications Conference. Association for Computing Machinery, New York, NY, USA, ACSAC '17, pp. 262–277 (2017). <https://doi.org/10.1145/3134600.3134635>
50. Jarrahi, M.H.: Artificial intelligence and the future of work: Human-AI symbiosis in organizational decision making. *Bus. Horizons* **61**(4), 577–586 (2018)
51. Jarrahi, M.H., Davoudi, V., Haeri, M.: The key to an effective AI-powered digital pathology: Establishing a symbiotic workflow between pathologists and machine. *J. Pathol. Inf.* **13**, 100156 (2022)



52. John Rajan, A., Jayakrishna, K., Vignesh, T., et al.: Development of computer vision for inspection of bolt using convolutional neural network. *Mater. Today Proc.* **45**, 6931–6935 (2021). <https://doi.org/10.1016/j.matpr.2021.01.372>, <https://www.sciencedirect.com/science/article/pii/S2214785321004636>. International Conference on Mechanical, Electronics and Computer Engineering 2020: Materials Science
53. Jwo, J.S., Lin, C.S., Lee, C.H.: Smart technology–driven aspects for human-in-the-loop smart manufacturing. *Int. J. Adv. Manuf. Technol.* **114**, 1741–1752 (2021)
54. Kaasinen, E., Anttila, A.H., Heikkilä, P., et al.: Smooth and resilient human–machine teamwork as an industry 5.0 design challenge. *Sustainability* **14**(5), 2773 (2022)
55. Kim, T.H., Kim, H.R., Cho, Y.J.: Product inspection methodology via deep learning: An overview. *Sensors* **21**(15), 5039 (2021). <https://doi.org/10.3390/s21155039>
56. Kosuge, K., Hirata, Y.: Human-robot interaction. In: 2004 IEEE International Conference on Robotics and Biomimetics, pp. 8–11. IEEE (2004)
57. Križnar, K., Rožanec, J.M., Fortuna, B., et al.: Explainable artificial intelligence meets active learning: A novel gradcam-based active learning strategy, submitted (2023)
58. Kumar, P., Gupta, A.: Active learning query strategies for classification, regression, and clustering: A survey. *J. Comput. Sci. Technol.* **35**(4), 913–945 (2020)
59. Lambert, A., Norouzi, N., Bruder, G., et al.: A systematic review of ten years of research on human interaction with social robots. *Int. J. Human Comput. Interact.* **36**(19), 1804–1817 (2020)
60. Langley, P.: Interactive cognitive systems and social intelligence. *IEEE Intell. Syst.* **32**, 22–30 (2017). <https://doi.org/10.1109/MIS.2017.3121556>
61. Leach, R., Bourell, D., Carmignato, S., et al.: Geometrical metrology for metal additive manufacturing. *CIRP Ann.* **68**(2), 677–700 (2019)
62. Leng, J., Sha, W., Wang, B., et al.: Industry 5.0: Prospect and retrospect. *J. Manuf. Syst.* **65**, 279–295 (2022)
63. Leone, A., Rescio, G., Siciliano, P., et al.: Multi sensors platform for stress monitoring of workers in smart manufacturing context. In: 2020 IEEE International Instrumentation and Measurement Technology Conference (I2MTC), pp 1–5. IEEE (2020)
64. Liapis, A., Yannakakis, G.N., Alexopoulos, C., et al.: Can computers foster human users' creativity? theory and praxis of mixed-initiative co-creativity. *Digit. Cult. Educ.* **8**, (2016)
65. Licklider, J.C.R.: Man-computer symbiosis. *IRE Trans. Human Fact. Electron.* **HFE-1**(1), 4–11 (1960). <https://doi.org/10.1109/THFE2.1960.4503259>
66. Liu, J., Guo, F., Gao, H., et al.: Defect detection of injection molding products on small datasets using transfer learning. *J. Manuf. Process.* **70**, 400–413 (2021). <https://doi.org/10.1016/j.jmapro.2021.08.034>
67. Longo, F., Padovano, A., Umbrello, S.: Value-oriented and ethical technology engineering in industry 5.0: A human-centric perspective for the design of the factory of the future. *Appl. Sci.* **10**(12), 4182 (2020)
68. Lu, Y.: The current status and developing trends of industry 4.0: A review. *Inf. Syst. Front.*, 1–20 (2021)
69. Lughofer, E.: On-line active learning: A new paradigm to improve practical useability of data stream modeling methods. *Inf. Sci.* **415**, 356–376 (2017)
70. Lyytinen, K., Nickerson, J.V., King, J.L.: Metahuman systems = humans + machines that learn. *J. Inf. Technol.* (2020). <https://doi.org/10.1177/0268396220915917>
71. Maddikunta, P.K.R., Pham, Q.V., Prabadevi, B., et al.: Industry 5.0: A survey on enabling technologies and potential applications. *J. Ind. Inf. Integr.* **26**, 100257 (2022)
72. Madry, A., Makelov, A., Schmidt, L., et al.: Towards deep learning models resistant to adversarial attacks. In: 6th International Conference on Learning Representations, ICLR 2018, Vancouver, BC, Canada, April 30–May 3, 2018, Conference Track Proceedings (2018). OpenReview.net. <https://openreview.net/forum?id=rJzIBfZAb>
73. Montini, E., Bettoni, A., Ciavotta, M., et al.: A meta-model for modular composition of tailored human digital twins in production. *Procedia CIRP* **104**, 689–695 (2021)



74. Montini, E., Cutrona, V., Bonomi, N., et al.: An iiot platform for human-aware factory digital twins. *Procedia CIRP* **107**, 661–667 (2022)
75. Montini, E., Cutrona, V., Dell’Oca, S., et al.: A framework for human-aware collaborative robotics systems development. *Procedia CIRP* (2023)
76. Moosavi-Dezfooli, S.M., Fawzi, A., Frossard, P.: DeepFool: a simple and accurate method to fool deep neural networks. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 2574–2582 (2016)
77. Mosqueira-Rey, E., Hernández-Pereira, E., Alonso-Ríos, D., et al.: Human-in-the-loop machine learning: a state of the art. *Artif. Intell. Rev.* (2022). <https://doi.org/10.1007/s10462-022-10246-w>
78. Mukherjee, D., Gupta, K., Chang, L.H., et al.: A survey of robot learning strategies for human-robot collaboration in industrial settings. *Robot. Comput. Integr. Manuf.* **73**, 102231 (2022)
79. Müller, D., März, M., Scheele, S., et al.: An interactive explanatory AI system for industrial quality control. In: *Thirty-Sixth AAAI Conference on Artificial Intelligence, AAAI 2022, Thirty-Fourth Conference on Innovative Applications of Artificial Intelligence, IAAI 2022, The Twelfth Symposium on Educational Advances in Artificial Intelligence, EAAI 2022 Virtual Event, February 22–March 1, 2022*, pp. 12580–12586. AAAI Press (2022)
80. Newman, T.S., Jain, A.K.: A survey of automated visual inspection. *Comput. Vis. Image Understand.* **61**(2), 231–262 (1995)
81. Nicolae, M.I., Sinn, M., Tran, M.N., et al.: Adversarial robustness toolbox v1.0.0. Preprint (2018). [arXiv:180701069](https://arxiv.org/abs/180701069)
82. Paul, S., Yuan, L., Jain, H.K., et al.: Intelligence augmentation: Human factors in ai and future of work. *AIS Trans. Human Comput. Interact.* **14**(3), 426–445 (2022)
83. Rahwan, I., Cebrian, M., Obradovich, N., et al.: Machine behaviour. *Machine Learning and the City: Applications in Architecture and Urban Design*, pp. 143–166 (2022)
84. Ramesh, A., Dhariwal, P., Nichol, A., et al.: Hierarchical text-conditional image generation with CLIP latents. Preprint (2022). [arXiv:220406125](https://arxiv.org/abs/220406125)
85. Ren, P., Xiao, Y., Chang, X., et al.: A survey of deep active learning. *ACM Comput. Surv. (CSUR)* **54**(9), 1–40 (2021)
86. Rožanec, J.M., Karel, K., Montini, E., et al.: Predicting operators’ fatigue in a human in the artificial intelligence loop for defect detection in manufacturing. In: *Proceedings of the 2023 IFAC World Congress* (2023)
87. Rožanec, J.M., Fortuna, B., Mladenović, D.: The future of data mining. chapter 6: Active learning (2022). <https://doi.org/10.52305/KCIN5931>
88. Rožanec, J.M., Zajec, P., Theodoropoulos, S., et al.: Robust anomaly map assisted multiple defect detection with supervised classification techniques. Preprint (2022). [arXiv:221209352](https://arxiv.org/abs/221209352)
89. Rožanec, J.M., Zajec, P., Theodoropoulos, S., et al.: Synthetic data augmentation using GAN for improved automated visual inspection. Preprint (2022). [arXiv:221209317](https://arxiv.org/abs/221209317)
90. Rožanec, J.M., Bizjak, L., Trajkova, E., et al.: Active learning and novel model calibration measurements for automated visual inspection in manufacturing. *J. Intell. Manuf.*, 1–22 (2023)
91. Saiz, F.A., Alfaro, G., Barandiaran, I.: An inspection and classification system for automotive component remanufacturing industry based on ensemble learning. *Information* **12**(12), (2021)
92. Schwalbe, G., Finzel, B.: A comprehensive taxonomy for explainable artificial intelligence: a systematic survey of surveys on methods and concepts. *Data Mining Knowl. Disc.*, 1–59 (2023)
93. Scott, A.C., Clancey, W.J., Davis, R., et al.: Explanation capabilities of production-based consultation systems. Tech. rep., Stanford Univ CA Dept Of Computer Science (1977)
94. See, J.E.: Visual inspection: a review of the literature. Sandia Report SAND2012-8590, Sandia National Laboratories, Albuquerque, New Mexico (2012)
95. Settles, B.: Active learning literature survey. Tech. rep., University of Wisconsin-Madison Department of Computer Sciences (2009)
96. Shokri, R., et al.: Bypassing backdoor detection algorithms in deep learning. In: *2020 IEEE European Symposium on Security and Privacy (EuroS&P)*, pp 175–183. IEEE (2020)

97. Sikander, G., Anwar, S.: Driver fatigue detection systems: A review. *IEEE Trans. Intell. Transp. Syst.* **20**(6), 2339–2352 (2018)
98. Silva, B., Marques, R., Faustino, D., et al.: Enhance the injection molding quality prediction with artificial intelligence to reach zero-defect manufacturing. *Processes* **11**, (2023). <https://doi.org/10.3390/pr11010062>
99. Sugiyama, M., Kawanabe, M.: Active Learning, pp 183–214. MIT Press (2012)
100. Szegedy, C., Zaremba, W., Sutskever, I., et al.: Intriguing properties of neural networks. In: Bengio, Y., LeCun, Y. (eds.) 2nd International Conference on Learning Representations, ICLR 2014, Banff, AB, Canada, April 14–16, 2014, Conference Track Proceedings (2014)
101. Tang, F., Mohammed, M., Longazo, J.: Experiments of human-robot teaming under sliding autonomy. In: 2016 IEEE International Conference on Advanced Intelligent Mechatronics (AIM), pp 113–118 (2016). <https://doi.org/10.1109/AIM.2016.7576752>
102. Tripicchio, P., Camacho-Gonzalez, G., D’Avella, S.: Welding defect detection: coping with artifacts in the production line. *Int. J. Adv. Manuf. Technol.* **111**, 1659–1669 (2020). <https://doi.org/10.1007/s00170-020-06146-4>
103. Tsai, D.M., Jen, P.H.: Autoencoder-based anomaly detection for surface defect inspection. *Adv. Eng. Inf.* **48**, (2021). <https://doi.org/10.1016/j.aei.2021.101272>
104. Tschang, F.T., Almirall, E.: Artificial intelligence as augmenting automation: Implications for employment. *Acad. Manag. Perspect.* **35**(4), 642–659 (2021)
105. Tsvetkova, M., Yasseri, T., Meyer, E.T., et al.: Understanding human-machine networks: a cross-disciplinary survey. *ACM Comput. Surv. (CSUR)* **50**(1), 1–35 (2017)
106. van Berkel, N., Skov, M.B., Kjeldskov, J.: Human-AI interaction: intermittent, continuous, and proactive. *Interactions* **28**(6), 67–71 (2021)
107. Villalba-Diez, J., Schmidt, D., Gevers, R., et al.: Deep learning for industrial computer vision quality control in the printing industry 4.0. *Sensors* **19**(18), 3987 (2019)
108. Wallace, E., Rodriguez, P., Feng, S., et al.: Trick me if you can: Human-in-the-loop generation of adversarial examples for question answering. *Trans. Assoc. Comput. Linguist.* **7**, 387–401 (2019)
109. Wang, Y., Yao, Q., Kwok, J.T., et al.: Generalizing from a few examples: A survey on few-shot learning. *ACM Comput. Surv. (CSUR)* **53**(3), 1–34 (2020)
110. Wilson, H.J., Daugherty, P.R.: Collaborative intelligence: Humans and AI are joining forces. *Harv. Bus. Rev.* **96**(4), 114–123 (2018)
111. Wu, D.: Pool-based sequential active learning for regression. *IEEE Trans. Neural Networks Learn. Syst.* **30**(5), 1348–1359 (2018)
112. Wu, J., Huang, Z., Hu, Z., et al.: Toward human-in-the-loop AI: Enhancing deep reinforcement learning via real-time human guidance for autonomous driving. *Engineering* **21**, 75–91 (2023). <https://doi.org/10.1016/j.eng.2022.05.017>, <https://www.sciencedirect.com/science/article/pii/S2095809922004878>
113. Wu, X., Xiao, L., Sun, Y., et al.: A survey of human-in-the-loop for machine learning. *Fut. Gener. Comput. Syst.* **135**, 364–381 (2022). <https://doi.org/10.1016/j.future.2022.05.014>, <https://www.sciencedirect.com/science/article/pii/S0167739X22001790>
114. Xiong, W., Fan, H., Ma, L., et al.: Challenges of human—machine collaboration in risky decision-making. *Front. Eng. Manag.* **9**(1), 89–103 (2022)
115. Xu, F., Uszkoreit, H., Du, Y., et al.: Explainable AI: A brief survey on history, research areas, approaches and challenges. In: CCF International Conference on Natural Language Processing and Chinese Computing, pp 563–574. Springer (2019)
116. Xu, W., Evans, D., Qi, Y.: Feature squeezing: Detecting adversarial examples in deep neural networks. Preprint (2017). arXiv:170401155
117. Yun, J.P., Shin, W.C., Koo, G., et al.: Automated defect inspection system for metal surfaces based on deep learning and data augmentation. *J. Manuf. Syst.* **55**, 317–324 (2020)
118. Zajec, P., Rožanec, J.M., Theodoropoulos, S., et al.: Few-shot learning for defect detection in manufacturing, submitted (2023)
119. Zhu, D., Li, Z., Wang, X., et al.: A robust zero-sum game framework for pool-based active learning. In: The 22nd International Conference on Artificial Intelligence and Statistics, pp 517–526. PMLR (2019)

**Open Access** This chapter is licensed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license and indicate if changes were made.

The images or other third party material in this chapter are included in the chapter's Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the chapter's Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder.

