Deep Science Publishing, 2025
https://doi.org/10.70593/978-93-7185-753-6_3

Chapter 3: Artificial Intelligence Applications in Mission-Critical Domains

Swarup Panda
SRM Institute of Science and Technology, Kattankulathur, Tamil Nadu, India

1. Introduction to AI in Mission-Critical Domains

The historical development of Artificial Intelligence (AI) has been an incremental one, where cascades of many breakthroughs in technology have led to advances in intelligence for all-parallel, all-digital simulators for systems with pre-defined rules sets [1-2]. Notable exceptions exist which display truly general native intelligence, and which do not require nearly infinite compute resources, but none have yet been able to successfully navigate in physical environments to create lasting organizations, replicate themselves, and achieve self-sustainability quality, as these forms of life do [3-5]. In mimicking the intelligence of these complex naturally evolved Living Systems, AI systems have become general production-quality, affordable tools for high-volume data and signal processing tasks, and semi-autonomous decision-support systems for moderately larger artifice processes in many mission-critical domains undertaken by humans and human-centered organizations, supported by their technical systems.

AI has also been embedded within, and enabled dramatic advances in traditional systems of automation like speech processing, natural language processing, intelligent data and sensor fusion, decision making and intelligence analysis, and swarm robotics [6-8]. There are frameworks that have defined Digital Evolution principles for accelerating systems of man and machine within evermore autonomous cognitive, mission-critical systems, which provide society many of its current defense and governance missions, and created the crisis of trust in AI systems. Yet, these same frameworks warn that relied-upon, human-centered organizations are failing within dynamic global transformation, and therefore caution as to what extent; and so raises questions around how and when to place AI in control of tasks in the mission-space of essence without being backed-up by bounding strategic parameter bounding decision rules inherently incorporated with self-interest and mutual trust into these same organizations.

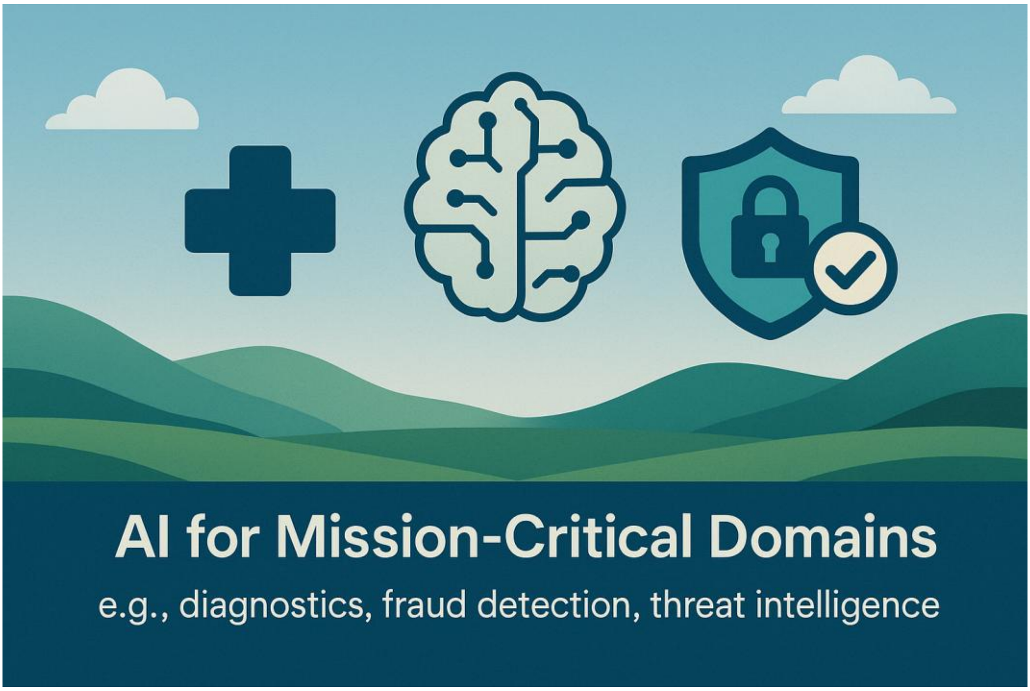
AI for Mission-Critical Domains
e.g., diagnostics, fraud detection, threat intelligence

2. Overview of Key Mission-Critical Domains

In this section, we provide details for four focused mission-critical domains where AI applications currently have a strong presence and importance. In particular, we explore healthcare diagnostics, fraud detection in finance, security in cybersecurity, and optimization in Supply Chain Management. These four domains were selected because of the strong potential for innovation in those application areas and also because of the specific technology directions that were mentioned above.

2.1. Healthcare Diagnostics

AI applications in the healthcare domain have become a major research topic during the past decade as the world has seen major advances in AI technologies that match or exceed human capabilities in areas like visual diagnosis, writing, and planning [7,9-10]. Healthcare data is typically full of sensitive information about the patients, and hence, AI technologies can only be applied after concept validation. However, once they are validated, the applications can significantly reduce the time of diagnosis and treatment for the patient as well as the cost for the healthcare providers. Recent years have seen the outbreak of a new disease caused by a deadly virus, which has made the importance of valid and efficient healthcare technologies urgently needed [1,11-14]. The urgent need for testing has also redirected the focus of many countries toward AI-based testing mechanisms. The need for healthcare diagnostics has been a subject of interest since time immemorial. In recent years, the testing technologies have also become largely technology-driven rather than operator-driven by using core sensor technologies. As a result, AI-based innovations specifically involving deep learning, computer vision, and natural language processing have seen increased popularity such as the use of computer vision algorithms in the analysis of pathology slides and sputum thermography.

In the area of healthcare diagnostics, we have also seen a massive investment by leading technology companies around the world as healthcare has been compared with the digital industry of the 1990s. Medical diagnostics is a core enabler of the larger healthcare ecosystem consisting of the hospital care, retail, home, and IT-enabled support [13,15-17]. These larger segments for which the impact of AI applications could be felt largely contain redundant operational overheads that could potentially benefit from the continuous monitoring and optimization capabilities offered by AI. In this chapter, we will elaborate on the many such AI applications in these segments. We will start with an elaborate discussion on healthcare diagnostics and follow with an extensive overview of other application domains in this larger ecosystem.

2.2. Financial Fraud Detection

Fraudulent behavior can take many different forms in financial transactions, typically driven by the motive of negative gain at the expense of a larger audience. It can be perpetrated by individuals or organizations, intentionally or unintentionally, remotely or physically [18-20]. The arena of financial transactions can include a myriad of services such as payments, loans, deposits, mortgages, credit and insurance, pensions, or security and commodity trading. The phenomenal growth of the Internet, especially with the ubiquitous connectivity offered through smart devices, has opened the floodgates of online financial transactions, resulting in the accompanying surge of financial crime. Financial crimes are continually evolving as criminals and con artists become tech-savvy, increasingly developing more sophisticated and complex schemes and techniques to exploit weaknesses in financial systems. This entails a formidable challenge of keeping one jump ahead of the fraudsters who are employing new and cleverer techniques. Although the techniques they use may be sophisticated, the actual motivation for their fraudulent behavior - deceit and theft - remains unchanged.

Financial crimes entail the loss of millions of dollars every year to financial institutions and businesses [19,21-22]. Actual estimates vary, but the costs of financial fraud are some of the highest in the economy, often impossible to quantify accurately. Given the vast size and scale of financial transactions being undertaken digitally today, it is not surprising that the phenomenon has attracted the attention of data scientists and machine learning researchers for developing intelligent solutions for fraud detection. AI-enabled solutions are capable of rapidly sifting through massive volumes of digital transactions, looking for indications and patterns of potential fraud on a real-time basis. They deploy intelligent algorithms capable of learning normal transactional behaviors for every account and are also able to flag potentially fraudulent transactions for further investigation by fraud analysts without creating a large volume of false positives or negatives.

2.3. Cybersecurity and Threat Intelligence

Artificial intelligence (AI) has found substantial success in mission-critical domains, particularly over the past decade [11,23-25]. In threat intelligence and cybersecurity, AI is applied to synthesize security intelligence from diverse datasets, and alerts are raised for threat actors with malicious intentions against the enterprise. These alerts often lead to further investigation by security operations center (SOC) personnel to decide if any actual action needs to be taken for the possible threat to become real. Generally, SOC experts derive the threat intelligence from different modalities of data such as records of network traffic, logs from various security applications, and telemetry from endpoint agents. This problem of assembling disparate datasets for intelligence, and specifically security alerts, is referred to as triage. The major success factors for AI application in threat intelligence and cybersecurity are a massive amount of unlabeled points and high label accuracy for supervised or semi-supervised usage. This high label accuracy often comes about through the involvement of myriad domain experts during the human operations for triage.

The generation of actionable insights essentially drives AI models in many other domains, including cybersecurity [26-28]. The different focus is, however, on the asset aspect of people and their behavior and intent. Therefore, the processes implementing the action or decisions are different from many other domains. In contrast to other domains with defined input–output mapping, cybersecurity AI essentially divides the inputs into those that require actions and executions to follow and those that do not require inputs for actionable outputs. The need for investments in this mission-critical domain and prudent executive decisions is very high, as the intention or the very word “intent” is front and center at different levels.

2.4. Supply Chain Management

Supply chains have become globally interlinked over recent decades due to the increased outsourcing of product components to regions of low labor cost, such as Southeast Asia [29-32]. This geographical separation between supply and demand has generated increasing risks of product supply disruption, which is typically referred to as “the bullwhip effect.” In particular, recent events have exacerbated supply chain risks and demonstrated the critical importance of supply chain resilience. The increasing demand volatility for key products, such as biopharmaceuticals, semiconductors, and food products, has further demonstrated the potential for balance sheet exposure to catastrophic loss from inefficient or fragile supply chains. In this context, the application of AI in various aspects of supply chain practice has gained increasing interest across both academia and industry.

Supply chain management (SCM) encompasses the operational and strategic planning of within- and cross-organizational management, decision-making, and coordination functions with the aim of designing, developing, integrating, and controlling supply and distribution networks to achieve specified goals, such as minimizing cycle time, total delivered cost, or inventory levels, and maximizing quality and customer service level. In addition to being a critical operating cost, efficient supply chains increasingly serve as a value-adding source of differentiation from competitors through addressing critical end-customer requirements, which often include product features, availability, flexibility, lead time, and risk mitigation. Recent advances in information and communication technologies have enabled companies to achieve worldwide integration and coordination of processes across the complete supply/demand chain associated with products or services.

3. AI Techniques Used in Diagnostics

While Artificial Intelligence (AI) serves various roles in a diagnostic solution, from information processing and analysis to task automation or computerized support, it is the individual AI techniques, such as Machine Learning (ML), Natural Language Processing (NLP), Knowledge Representation (KR), and Cognitive Architectures, which form the building blocks used by AI toolkits for providing unique diagnostic services [31,33-35]. The following sections outline the major AI technologies used in diagnostic tools, with a focus on their mission-critical domain applications.

ML is the process of using algorithms and statistical models to analyze and synthesize patterns from complex data inputs and generate generalizations, predictions, and recommendations. Patient evaluations vary significantly and generate large amounts of data; therefore, ML techniques are well suited to extract latent inaccurate correlations in the data and utilize them as proxies for the actual underlying cause. The key advantage of ML over traditional diagnostics is the avoidance of the need to create and maintain complex causal models to explain and evaluate various maladies. While unsupervised learning can be used to diagnose cases with previously unknown maladies, the more common application of ML in diagnostics is through supervised classification of patients into different diagnostics and supervised regression for obtaining severity measures of existing diagnostics. Non-linear functions generated through deep learning applied to raw data are showing remarkable improvement in classification accuracy and explainability. While initially applied mainly to winning cognitive tasks like facial emotion recognition, its ability to replace ML with simpler supervised classifiers has made deep learning the primary ML technique deployed in diagnostics.

Conversational Agents have grown significantly in their role of initial virtual clinical assistants and pre-evaluation filters through client-facing Natural Language Processing systems [36-38]. NLP is a subfield of AI that plays a major role in diagnostic solutions because of the utilization of unstructured free-form textual data. NLP enables the automation of all major steps involved in naturally communicating with a patient, from language understanding to language generation to user intent recognition to named entity recognition to relationship extraction. These steps are effectively the textual inputs and outputs of a typical diagnostic tool. As analytical and prediction functions using machine learning, deep learning, and Bayesian network techniques for diagnosing, recommending, and generating analysis reports are integrated into these conversational agents, AI-powered chatbots will transform into true virtual clinical assistants.

3.1. Machine Learning Algorithms

Machine learning emphasizes the learning component, which may be considered a statistical problem. Probabilistic models like Bayesian and maximum entropy models are ubiquitous in ML. The core of machine learning has always been supervised learning. This is due to the robustness of supervised learning methods, whose practical success has meant that it has largely marginalized or overshadowed the other ML topics like unsupervised or active learning. These methods have roots in Bayesian decision theory, statistical estimation theory, information theory, and optimum experimental design.

ML has a more practical focus than the traditional areas of statistics. A statistical analysis normally has three stages. The model is selected from a collection of candidates and fit to the data. Lastly, inferences like confidence intervals are made from the fitted model. In contrast, ML methodology is focused on the second aspect of statistical analysis, model fitting. ML uses model fitting to address challenges with data, such as large volume, high dimension, and missing values [1,39-41]. Again, compared with statistics, ML emphasizes a different criterion – predictive accuracy, as compared to likelihood or posterior sampling. The checkpoints of ML predictions are accuracy or error rate on test data, and therefore, evaluate the predictive performance via a test sample has also been integrated into the fitting methods via resampling, such as cross-validation or bootstrapping. In particular, choosing a better predictive model among several candidates has become such a common practice that it also has its own name, Model Selection.

Somewhat more formally, model selection is the key point of model selection, which is the fundamental models in supervised learning. All of these models comprise the fundamental toolbox of supervised learning, offering different theoretical properties, practical advantages, and pitfalls. Besides algorithm selection, careful tuning of the model parameters that go beyond the statistical sampling stage is essential to attain good prediction accuracy. These parameters typically control the size and complexity of the estimated model.

3.2. Natural Language Processing

Terms of definition characterizing a range of mental health fields suggest that apart from description, NLP can support therapy [42-44]. The open questions arise whether we could build intelligent agents that are able to conduct a meaningful communication with a user dealing with depression or substance abuse. Chatbots negotiating with a user with potential psychological problems are one of the approaches to answer validation whether we are able to code the human knowledge and interactions to the level that we are able to outsource and automate parts of these activities. The two colonial regions with the basic underlying goals and background motivation can be distinguished in the interdisciplinary research topics bridging informal areas of linguistics and psychiatry.

The first colonizers — the psychiatrists use chatbots to assess the user's features. The examples of psycholinguistic features commonly employed in user profiling involved in the assessments of psychopathology include toxic words counts, negative or positive emotion words counts, speech dismemberment. The second colonizers — the linguists exploit the psychological knowledge to enhance the functionality and improve the underlying algorithms used by the chatbots. The actions performed in all the cases use the curated sets of transcripts representing selected population assessments sharing the same psychological disorder, type-affected area, or user's profile and representing the underlying properties of interest and considered features.

3.3. Image Recognition Technologies

Image recognition technologies allow identifying objects in a user-provided photo or video stream. Algorithms trained on a large number of example images and video clips use a substantial amount of computational resources to classify video or photo images as one or more objects. Initially developed for circumventing the security risk involved in face recognition passwords for phones and user login, further enhancements were completed by tech giants to support object recognition.

The main weaknesses of these technologies, or deep learning algorithms, are their need for an extremely large number of example photos to train the algorithm and inability to handle unseen object types unless retrained [45-46]. Used in a prudent way, image recognizers—a sub-area of machine intelligence—can be extremely useful in particular cases. For instance, recently researchers matched the words in medical records with clinical-motor and audio-visual data. Then they successfully used these deep false recognition networks to discover symptoms of autism spectrum disorder from non-labeled videos that were previously deemed unfeasible. AI systems based on deep learning algorithms were successfully applied to multiple domains: analyzing the type of skin conditions in order to assist doctors in diagnosis; classifying medical images, such as the detection of diabetic retinopathy in retinal fundus images; detecting pneumonia from chest X-ray images; identifying breast cancer in mammography images.

4. AI in Fraud Detection

Fraud can be defined as the manipulation or distortion of data for the purpose of financial gain. In general, therefore, fraud is a gray-area concept that can be defined and detected with greater or lesser success only if a normative framework is defined. A legal system sets certain bounds within which transactions should take place. When the limits are transgressed, a criminal act is committed, and law enforcement agencies can take corrective action. Thus the application of AI to fraud detection, just like to other areas where human decision making plays an important role, such as legal practice or psychology, allows for higher accuracy in predictions and quicker and more effective recommendations for action to be taken. Fraud can be found in many spheres of human life. Fraud syndicates manipulate immigrant visas to the detriment of nation-states, terrorists destroy commercial property in order to extract ransom, and fingerprints and other data can be copied and falsified, causing great harm to the rightful owners of the credentials. Fraud at the corporate level can be in the form of underinsured assets or concealed liabilities, or the inactive corporate sponsor indemnifying the parent corporation for litigation costs relating to the operation of the indirectly held subsidiary’s business.

The nexus of AI and fraud detection is strongest in the domain of predictive analytics using supervised learning. A number of statistical techniques have, for years now, been trained on large lags of historical data from credit card transactions to identify whether any current transaction is likely to be fraudulent. AI-enabled predictive analytics thus inform those conducting the transaction, those issuing the credit card, and the transaction networks. In addition to anomaly detection techniques and predictive analytics, organizations have used AI for behavioral analysis of employees in the context of internal fraud detection. Organizations are constantly being subject to review and reassessment and are expected to explore advanced technologies to further increase compliance with mandates set forth by regulatory oversight. By leveraging AI, companies can close the gap in those risks while significantly reducing the time and resources it currently takes to accomplish or achieve compliance.

4.1. Anomaly Detection Techniques

The field of cybersecurity is faced with an increasing number of threats as we use the Internet for more and more different domains [18,47-50]. In particular, malware instances are increasing in numbers, and becoming more and more sophisticated thanks to AI applications that help their malware-related tasks and goals, such as automated malware creation and distribution, ransom-cracking, and the use of steganography. The use of the Internet for critical domains, in particular, begs for efficient cybersecurity solutions, pushing these areas to the forefront of traditional AI research.

4.1. Anomaly Detection Techniques. AI applications for anomaly detection in various data objects is a long-term goal of AI research. Traditionally, anomaly detection techniques are typically centered on monitoring logical system operations and detecting anomalies based on the possibility of or the impossibility of certain system events, either by using domain specific knowledge or utilizing logic-based heuristics. More formally, anomaly detection is often characterized in the following manner.

Let D denote a certain amount of traditional prior domain knowledge about the space of possible system states given by the various parameters of the system. Let a variable N be the total number of possible admissible scenarios of the monitored system based on the prior domain knowledge D. Let a variable T denote the total number of actually executed scenarios and let a E be the number of actually executed scenarios not contained in D, which are marked as alerting scenarios. Then anomaly detection attempts to find N, T, and a E such that N is maximized, T is minimized, and a E is minimized. Deviation of the monitored scenarios from the prior domain knowledge D is an indicator of potential malicious operation. The main principle is that machine learning attempts to define allowed and disallowed scenarios for the monitored system, both being based on prior domain knowledge.

4.2. Predictive Analytics

Predictive analytics reveals the future with certain probabilities for specific events. It allows businesses to factor in explanations for prediction errors. The essential building blocks for predictions depend on business knowledge and experience, existing workflows and operations, and mathematical models and advanced statistical techniques to combine information from forecasts, attributes, and other signals. Predictive analytics answers questions such as: Do existing sales pipelines flags, tools, and models for risk questions allow sales executives to accurately anticipate failures and provide savings for clients who are coming late? Do certain client behaviors that have been identified allow telecom companies to target and engage the right clients with campaigns that will retain them for the future? Are certain organizations within the insurance value chain leveraging loss models that guide pricing based on risk attributes? Do predictive models, tracking model performance for specific customer segments, deal sizes, product types, various customers, help companies overcome statistics, recognize warning signals, and continuously improve their modeling efforts?

The intent of predictive modeling is not to match all characteristics once and for all, but to obtain realistic grades for individuals for a very large number of purposes. The predictive model for a person answering such questions will normally be based not only on the current input information for that individual, but also on databases containing aggregated information derived from current and historical data from similar individuals. Only a small number of individuals can be observed through constant monitoring, and the most appropriate time to intervene is often key to success in an action, such as introducing an offset pack for a latent complaint or an attractive price environment for value reinforcement. In many organizations, the only monitoring process available is prospect list testing, combatting complaints for a special generation period and assessing the results.

4.3. Behavioral Analysis

Behavioral analysis, particularly of online movement and activity, provides businesses with detection processes and playbooks that limit losses and deter fraudsters [18-20]. For example, on receiving a log-in or account recovery request, a rules-based playbook checks the existence of the user within a specific geolocation, the timing and travel speed of the current log-in attempt, the current device being used to initiate the log-in attempt, if the request is made via an anonymous VPN, the denial of multiple login attempts, pattern instability of login attempts or access histories, and so on. If check passes, verification is requested and granted. If it fails, risk-based online verification options for user verification and multiple flexible alternative paths are examined.

In addition to its uses as a standalone capability, behavioral analysis is an important building block for other aspects of predictive analytics, employed as a signal for items such as transaction and account risk assessments. Instead of basing verification request and transaction permission decision-making on static account age restrictions and business classifications only, a risk score based on real-time measurement of such behavior analysis signals helps refine decision quality. The capability to enlarge the volume of "silent stop" hard log-out fraud and its related account risk assessments and acquire signals of money laundering or failed account takeover attempts is what potential users should envisage as future available behavioral risk scoring analytics upon sharply increased use of a much-enhanced account sign-up verification and log-in multi-layer automated verification solution.

5. Threat Intelligence and AI

Threat intelligence is a hands-on cybersecurity process that helps organizations at different levels of maturity implement and scale. This section reviews several key areas of the threat intelligence workflows that can benefit from the help of AI, and explores a range of existing and future AI roles in threat intelligence operations. Organizations today collect massive amounts of threat intelligence data, from open sources to paid-for commercial feeds, to internal telemetry feeds from their own systems and networks to track potential attacks. To make sense of this data, analysts must curate it, extract and catalog the data, identify which alerts deserve further action, get context on the organizations that are allegedly attacking them, and distill down to a summary of the information that is useful for teams within the organization. All these steps are well-suited for being made more efficient, if not fully automated, by AI.

AI is already being used with success for several of these operations, from curated collection to enrichment, and will likely enable full automation of many other steps in the coming years. Humans are still a critical part of the process due to the need to ensure that intelligence is derived and acted upon at sufficient quality levels, but AI can greatly augment their efforts. More frequent updates about the threat landscape, timely notifications, and richer, more contextualized interaction are other areas where AI can improve these processes. These solutions can help reduce time to prepare for, respond to, and recover from attacks.

5.1. Data Collection and Analysis

Data collection and analysis is a crucial part of cyber threat intelligence gathering. Their goal is to match possible digital patterns of criminal or hostile cyber activity to geospatial, physical, or online locations or entities involved in such activity. In other words, they are used to make intelligence predictions about risk scenarios. Many times organizations use open information from several different types of sources about the attackers as the basis for prediction. The most important open source information for threat intelligence is social media, domains, cyber criminal forums, hacktivist groups, malware leak blogs, and the dark web. In addition, local information from law enforcement agencies is also widely used.

The meta-trend of information gathering in physical threat intelligence is the advent of social media activity monitoring [19,21,22]. Social media are being increasingly used not only by political entities, but also by many other kinds of organizations, interest groups, or individuals engaged in nefarious activity. Some public or private firms specialize in collecting, processing, and analyzing social media data to match local or temporal spikes in social media activity with risk scenarios around the world. Such agreement helps threat prediction about physical places that may become targets of hostile actions, such as armed attacks, demonstrations, violent protests, kidnappings, or hijacking of individuals or dynamics of civil wars and insurrections that may evolve into armed conflict.

The recent introduction of generative large language models plus the datasets of knowledge digitization and organization about the world have made the information and digital pattern recognition process more efficient and reliable. Trained large language models, acting as mechanized threat analysts, can be used to make sense of social media, forums, chats, and blogs conversation and messaging by detecting discourse or discussions about risky behaviors and scenarios. In particular, large language models can analyze current and past information content for early signs and signals of increased interest or emerging trends in the physical or cyber-social space about events of significance and this effect can be linked to adverse physical or cyber social activity or events, either planned or ongoing.

5.2. Automated Threat Detection

Detection plays a critical role in threat intelligence. The speed with which a situation is detected can have a major impact on the necessary response time of response teams, while monitoring of an area can yield essential information and intelligence. Continuous monitoring of the cyberspace for potential malicious infiltration by external and internal actors, for the potential compromise of critical infrastructure and networks, and detection of anomalous behavior pointing to an existing compromise are all functions that lend themselves well to the application of AI.

AI algorithms excel at analyzing large data sets for indicating patterns and signatures of actors, either good or bad [6-8]. This capability can assist in the detection of intrusions into critical networks and services, as well as indications of compromise and anomalous behavior during or after a compromise. For industry sectors with a large attack surface, such as those associated networked control systems or the industrial control system implementation of Industry 4.0. Furthermore, large public datasets have been developed for these tasks, which are comprised of large training sets of flow files with factual labeling. More recently, the use of technology has been proposed for the protection of Industrial Internet of Things in the context of Industry 4.0 from cyber attacks. The method is described for using physical data from the IIoT to detect deviations from established patterns, which can be exploited by actors for intrusion and compromise.

5.3. Incident Response Automation

AI can also be applied to the incident handling phase of the incident response lifecycle. In this phase the response team works to contain and remediate the malicious attack. The incident response team is motivated to eliminate the immediate threat, but also to fully understand what happened, and how to prevent recurrence. The need to identify the specific attack used, the specific malware used, and the specific vulnerabilities exploited drives this deeper analysis. It is during this deeper phase of incident response that expert knowledge of specific malware families, and the expertise of those with experience in analyzing those malware families is required. Knowledge of the tactics, techniques, and procedures of specific APT actors is often critical in incident response. Many times, the ultimate goal of an adversary is not simply an incident, but instead is information needed to loot or attack further. So, collecting, analyzing, visualizing, and sharing knowledge is an important aspect of APT domain multitasking. And as in all domains of collaboration, achieving synergy in these actions is paramount, both within your own team, and between partner teams in the APT community.

There has been historical reluctance to wholly rely on AI in the domain of incident response activity, and repair decisions. This area also should have AI-enhanced tools, and approaches to optimize the value and speed of those incident response actions, while still including humans to ensure best fidelity in the final outcomes. Automating incident and communication of incident causes also can support dynamic decision-making processes to classify, escalate, and respond during a crisis while managing the risk. For example, machine-generated alerts and messages regarding network anomalous behavior are also of concern: if there is no explicit validation process, the environment can become saturated with alerts, making it difficult for human operators to take action, monitor critical systems, and respond to actual incidents.

6. Challenges in Implementing AI

Implementing AI in mission-critical applications has become increasingly feasible due to the growing availability of data resources, computing power, and algorithmic advancements. Despite these advantages, many organizations remain skeptical about adopting AI for these applications. Such concerns stem from the potential loss of control over important business processes and the reservations regarding the quantity and quality of data required in addition to validation and performance guarantees. Moreover, such efforts require a level of collaboration and integration across functional boundaries that extends beyond adopting other IT products and services.

As an example of these challenges, we refer to a few well-known incidents in the early development of AI, such as data poison incidents that resulted from the use of AI in the context of autonomous vehicles. Such events resulted from the training of AI systems on data collected from the real world and revealed the unsafe deployment of mission-critical systems trained on big data. Further emphasizing the importance of safety in AI, researchers and leaders made statements that deep learning has an intrinsic responsibility problem and that deep learning for vision is driving AI in the wrong direction. It is crucial that AI systems be rigorously validated in their deployment for mission-critical domains, and there exists an urgent need to impose stringent safety requirements in collaboration with the relevant stakeholders in such applications. System qualities, beyond measurement and validation of unidimensional criteria such as accuracy or bias, are of utmost importance across all stakeholders in these applications.

Most problems at this population level are multifactorial, and no single factor is responsible for failure. In addition to convergence and overfitting issues associated with large training data in deep learning, other serious issues include data biases, inappropriate selection of training data, the curse of dimensionality, adversarial training, and generation, attention mechanisms, and others. The selection of these parameters and their necessary tuning is not a simple engineering problem but a key aspect of the researcher and developer of neural network models. Efforts are ongoing to automate such choices, but a better understanding of these efforts is necessary for mission-critical applications.

6.1. Data Privacy Concerns

AI solutions are being rapidly put forward in a variety of applications, especially to augment or assist the human in the loop. However, there are some serious concerns associated with adopting AI for workplace applications, especially in mission-critical domains. Some of these concerns are unique for specific domains, whereas others are more generalized. In this chapter, we provide a brief overview of some of these more generalized concerns, in the following manner. We will first describe the issues, and then elaborate by looking at specific examples from AI in mission-critical domains, especially related to healthcare and best practices to address these concerns. Specifically, in the following sections, we delve into data privacy, data bias, integration of solutions into existing workflows.

First, data privacy for robust, trustworthy, and accurate AI models is required. Healthcare institutions get a lot of funds to hold on to the data privacy of their clients. Such data privacy laws demand that any non-compliant data disclosures that lead to compromise of patient privacy attract heavy penalties. As a result, not only is it difficult to obtain real verified labels for the images used to train AI models based on medical data, but the real data itself may not be available for model training. For example, a patient undergoing an MRI has his/her data stored at the healthcare institution. This healthcare institution in turn needs to remember the identity of the patient, and de-identify the data in order to share the data with researchers to develop AI models to assist in some MRI-related tasks.

6.2. Bias and Fairness in AI Models

The biased distribution of input datasets can lead AI to make broken decisions, and such bias can arise at any point in the pipeline of the AI project. For example, an aspect of the input data can be biased due to past social influence or historical mistrust, the biased sampling of the input data might occur while collecting the input data, the feature engineering approach may favor some group attributes, or the performance evaluation methodology may be flawed.

A major dilemma with AI model fairness is that a fair algorithm on one specified fairness criterion can be unfair on another specified fairness criterion, and that the fairness score changes with population attributes. Moreover, fairness is context-specific and attribute-specific. The current fairness notions attribute either too much power to the modeler or do not take population configurations into account. AI practitioners need to take every angle into account to create a fair AI model while doing bias mitigation based on a reasonable assumption.

Specific bias-mitigation recommendations utilize bias detectors, balanced datasets, adversarial unlearning during training to ensure nondiscrimination, and unbiasing techniques using model post-supervision to change predictions in a minimal manner after training or bias-removing techniques to process the input data or post-process the AI model in a fair way. AI experts propose changing the AI strategy more generally to take decision impact into account and incorporate fairness into the reward structure of algorithms. Besides the popular definition of privacy as confidentiality data protection, an AI bias-management strategy should include a specific privacy-preserving procedure.

6.3. Integration with Existing Systems

One of the most common challenges faced by organizations adopting AI is integrating it with existing systems. Companies have been investing vast amounts of money for years on legacy systems. These systems may not be updated frequently, but they hold substantial personal and institutional knowledge. They house mission-critical data, and any downtime they incur when AI is integrated with them must be counted in hours. This is especially true of systems such as payroll, which also process a substantial number of transactions daily. Integrating these systems with new AI systems is currently expensive. Almost every use case for AI requires a data pipeline, and building a new AI system from scratch for every pipeline is not only going to be expensive from a financial standpoint but will also take AI consultants many months to complete.

Once a company chooses a vendor's AI platform, they are generally locked into that vendor's technology ecosystem. While these vendors do offer access to the features of their products, they are limited in functionality compared to the core product. Due to the diverse and fragmented landscape of AI, a company may need to rely on different vendors for different features. Payment and advertising may be managed by one vendor while the others use different vendors for the rest, leading to inefficiencies that can be addressed to some extent with AI-specific design patterns. These design patterns also highlight a company's efforts to expose AI's weaknesses in the back-end systems that can be addressed in early phases of the AI project to deliver efficiency and reliability gain.

7. Case Studies in Healthcare Diagnostics

Healthcare is a high-stakes sector, which makes it a strong candidate for incorporating artificial intelligence (AI), especially in safety-critical applications, such as medical diagnostics, therapy assignment, and drug discovery. Patients' health and lives are of utmost importance, leading to rigorous standards for safety and effectiveness in the products and services used in the healthcare field. The likelihood of using medical device AI for improving the accuracy and time efficiency of healthcare providers for a well-defined and frequently used task when an answer will be seen by a second radiologist before it results in a protocol used on a patient has led to the acceleration of the clinical adoption of computer vision-inspired deep learning models for image analysis in radiology. Radiologists produce interpretable reports based on the images that they view and analyze. Predictive monitoring is also common in hospital settings, particularly for patients in critical care. Often with so many patients to observe, nurses rely on computerized alerts to signal them when a monitoring device detects potential patient deterioration. Alert fatigue, however, is a common problem induced by the overwhelming daily dissemination of alerts. Studies show that as few as 15-20% of these alerts are due to real events. Thus, patient decline detection is a fitting area for the adoption of machine learning models.

Mortality alerts that leverage a patient cohort's large amount of medical history data as they are receiving treatment at a healthcare facility would also contribute to the accuracy of the time to event estimation. Deferring just some of the patient mortality to allow clinical validation post-annotation may improve patient care and engagement.

7.1. AI in Radiology

Medical imaging stands at the forefront of Artificial Intelligence (AI) applications in medicine, and radiology is arguably the most mature translation of AI development in health care to date. Already, AI is deployed in clinical practice as decision-support tools for radiologists. With the ever-increasing incident volume of imaging studies, many academic radiology departments and commercial radiology practices are struggling to meet the needs of referring clinicians and patients. AI tools promise an assistive role on the horizon, with research efforts aimed at replacing or augmenting the diagnostic accuracy of radiologists. Some of the earliest commercial AI products in medicine were radiology-focused, including imaging-based tools for osteoporotic fracture prediction, colon polyp detection, skin cancer risk assessment, and others.

The applications for AI in radiology have been exhaustively reviewed and summarized. The specific indications and conditions include, but are not limited to: breast cancer detection and characterization in mammography, breast ultrasound, and breast magnetic resonance imaging; lung cancer detection and characterization in computed tomography of the chest; other thoracic / pulmonary applications in CT; cardiac assessment in non-invasive imaging arteries using electron-beam CT; enhancement in MRI of prostate and brain; colorectal disease detection and assessment in colonography and CT; musculoskeletal imaging in plain film and MRI; and neurological imaging in CT and MRI. The production and deployment of clinical AI tools require significant collaboration among industry, physicians from all medical specialties and disciplines, governing bodies, and regulatory institutions, who must balance the quest for expedient development and distribution of effective tools with the safeguarding of patient safety and wellbeing.

7.2. AI for Predictive Patient Monitoring

Predictive patient monitoring based on bio-signals like ECG, pulse-oximetry, respiration rate, inter-beat intervals, and pupil dilation can identify the physiological trajectory of patients during both home-based and hospital-based care. Such potentially actionable predictions include pre-symptomatic detection of adverse effects from surgery, trauma, and toxicology; evolution of long-term chronic diseases; the onset of episodic conditions; and life-threatening emergencies like cardiac arrest, stroke, and sepsis. These predictions are a critical step in the clinical workflow. They can aid clinical staff in their patient care and intervention planning. Timely alerts about patient degeneration can guide triaging during hazmat contaminations, mass-casualty events, and other emergencies involving hundreds of simultaneous patient presentations. A critical dimension of intelligent patient monitoring is long-term, ambulatory predictive monitors. Compared to static monocular closed-loop feedback, ambulatory health monitoring systems capture real-time temporal dynamics of involuntary biophysiological phenomena over extended periods of time. Intelligent assistive technologies are being developed for long-term and high-fidelity detection of monitor-measurable behaviors using various modalities like video, depth, radar, audio, and inertial sensors. Although some technologies are targeted towards the cognitively impaired, the primary applicability of such monitors lies in the elderly, and those with chronic co-morbidities like heart disease and diabetes. With an ageing population, hospitals have become critically understaffed and overburdened. Consequently, there are often long wait times for senior patients presenting with health issues. There is also a need for real-time differentiability between healthy non-patients and diseased patients, both before they enter the clinical system, and long after they've exited the system.

8. Case Studies in Fraud Detection

The financial, insurance, and retail industries, as well as public administration and local authorities, invest substantial resources in identifying fraudulent actions. Fraud involving credit cards, written checks, or tax returns affects all of us. More than twenty billion dollars a year is lost to fraud in the credit card industry alone, with banks, merchants, and credit issuers picking up a large share of these losses. The insurance industry also suffers billions of dollars in fraud losses annually. It is estimated that these losses equal between 10% and 25% of all claims. Tax fraud, especially refund fraud, is also a growing problem.

Still, how would we know which one of our neighbors is committing tax fraud? Almost by definition, fraud is a rare event in comparison to non-fraud commits. If the taxpayer base of a country were 250 million, a reasonable estimate of fraudulent behavior would be 1% (250,000). As even a fraction of that number is enough to generate academic interest, we must emphasize that our concern must be with screening but not executing the process. Making some mistakes is not a big issue as long as the model-specific adjustment of costs is not too high using a normalized confusion matrix. Payment systems, such as credit card issuers, and multi-institution tax systems can be used to design better support systems that adjust the model error costs.

In this work, we present two different implementations of models for fraud detection. The methodologies vary due to different internal and external constraints; however, they share the common goal of providing some degree of real-world operational capability. The first case concerns the detection of fraud in credit card transactions. The second, the detection of fraud in the automobile insurance industry. Though many aspects of these two applications differ, some important issues unify them.

8.1. AI in Credit Card Fraud Prevention

Plastic money or credit cards have become indispensable in our consumer and purchasing lives. Banks, financial and credit card companies, organizations, and merchants approve large amounts of electronic transactions every day for the purchase of products and services. While this cashless transaction mode has several advantages—such as ease of carrying, guaranteed payments, guaranteed purchase, and others—credit card transactions are prone to fraud and cyber-crimes. Theft of credit card information often leads to significant financial losses to banks and financial institutions. It is estimated that credit card fraud results in losses of around USD 27 billion each year. With the rapid growth of e-commerce, online purchases have also increased tremendously and fraudulent transactions have also increased significantly. This has led to an increased frequency of credit card fraud detection to protect against credit card fraud.

Credit card fraud is defined as unauthorized use of a credit card, where information, such as that regarding the card number, cardholder name, card validity period, card verification number, etc., itself, or while making a purchase of products and services online. Predicting fraudulent transactions is an extremely difficult exercise, as the majority of all the electronic card transactions are valid and accepted. This imbalance between valid transactions and invalid transactions leads to a high level of difficulty for credit card companies in spotting the fraudulent transactions and taking remedial actions in real-time for reducing losses. Though credit card frauds use various methods to dupe banks and customers, machine learning methods have proven vastly successful in predicting credit card fraud.

8.2. Insurance Fraud Detection Using AI

The insurance sector is at the forefront of adopting novel technologies for optimizing their business processes. A critical operation in the insurance industry that has recently seen a growing interest in using AI techniques is the detection of fraudulent claims. This activity has captured the attention of both academics and practitioners because of the importance of reducing the annual costs due to fraud. However, the volume of data involved and the intrinsic difficulties to detect low probability events such as fraud create opportunities for researchers and practitioners to develop advanced techniques to assist their analytical capabilities. For example, because of the low frequency of fraudulent claims, traditional supervised learning methods face huge difficulties in learning classification rules capable of generalizing beyond the historical data.

Two different ways have been proposed to handle the class imbalance problem: one involves the modification of the learning algorithms by experimenting with different types of cost-sensitive learning that assigns different costs to the two types of errors and the other one is the data balancing approach which modifies the training set given to the algorithm in different ways, by either oversampling the minority class or undersampling the majority class, in order to provide a more balanced data distribution. Despite many proposed techniques to address the class imbalance issue, the performance in practical cases did not always improve. One of the issues is that it is unclear how best to modify the original training data distribution to facilitate better performance of learning algorithms. In this context, a clear recommendation is to combine the data balancing approaches with the cost-sensitive algorithms. This requires a careful tuning process that, unfortunately, is often resource-consuming.

9. Case Studies in Cybersecurity

Cybersecurity is an area where AI in general, and machine learning (ML) in particular, have been deployed in multiple domains. Some of them include malware detection, phishing detection, intrusion detection/prevention, and vulnerability management. Cybersecurity stands out from other mission-critical domains because the application of AI was more based on economics than necessity, as most cybersecurity applications do not rely on AI for their effectiveness. Yet, AI is being sanctioned to play an important role in making cybersecurity applications better in regard to their accuracy, speed, or ability to deal with scale and complexity. These applications consist of security detection, response, and defense; and security operations, management, and compliance. While a few organizations have initiated the AI-in-cybersecurity effort, it is still early days in the usage and deployment of AI solutions that are effective in the real world. This chapter reviews the known history of AI in cybersecurity, and presents case studies based on past and present experiences in phishing detection, and malware detection and prevention. AI contains the potential for improving the work of cybersecurity professionals, as well as augment the practice of overall cybersecurity for everyday Internet citizens. The future is promising – organizations in all domains need to begin the journey of deploying AI capabilities into their cybersecurity arsenal. We believe that this need is paramount and compelling enough that compliance with this journey move from stretching the three pillars of data science – leadership, effectuation, and wait time – to enabling these three pillars to propel the security ecosystem into a new age. Without AI, however, this success may not be possible and is becoming increasingly hard to achieve.

9.1. AI-Driven Phishing Detection

In recent years, phishing attacks remained a major information security threat due to their effectiveness in harvesting sensitive information. As a result, phishing research has received considerable attention in the information security community. Historically an early warning system in a simplistic way, existing methods could lead to high false positive rates. For example, if a new URL was observed, then it would be suspected as being a phishing URL until being confirmed as legitimate in the near future. More common characteristics used in detection algorithms include keyword spotting, URL obfuscation, and IP blacklisting. Phishing detection from the perspective of a knowledge graph, however, does not seek to predict phishing URLs. Instead, it seeks to proactively warn users of phishing hazards associated with observed phishing URLs. The warnings are issued in real-time and predict user actions that may subsequently lead to a phishing incident, straining the relationship between phishing prevention and detection. In other words, an improved method for creating real-time, causal security warnings requires the application of artificial intelligence techniques.

Smartphone users, especially those of the younger generation, fill their devices with applications related to social engineering. Phishing attacks take advantage of the public's constant demand for easy access to friends and information. Attackers exploit user desires by creating counterfeit websites for popular social and news sites. Phishing apps target groups involved in hijacking sessions to intercept sensitive information between users and real service providers. Online social network applications encourage sharing information. Attackers rely on posting stolen information to lure organizations into unnecessary despair. Organizations must be constantly scanning for harm caused by their employees' naïve behavior. By injecting phishing emails into the business workflow, the company can detect how many emails were opened and how many attempts were made to click the links.

9.2. Malware Detection and Prevention

Malware is a class of software that is designed to disrupt computer functioning and to attack or exploit computers or networks. Despite years of trying to make systems less vulnerable to malware, including building higher security components into operating systems, implementing automated update reminders, blocking untrusted communications, and recognizing viruses through heuristic detection patterns, malware development has grown to unprecedented numbers. Malware detection and prevention absorb most of the resources allocated to cybersecurity research. AI and machine learning have been long-time contributors to malware mitigation, with research verified by deployment in production detection systems.

Machine learning has a strong foundation as a useful tool for some types of cybersecurity work, especially anomaly detection. Cybersecurity is observing a change in the traditional role of machine learning — starting in the early days as a supplementary approach to traditional detection and prevention techniques, with signatures and rules based on explicit conditions. Now, malware environment landscapes are too dynamic for traditional techniques. AI is applied in multiple stages of malicious programs' life cycle, including: Creation predicting, Propagation predicting and estimating damages, Detection monitoring activity on the local system or network model, Response blocking the activity, and Forensics detecting overlooked present and past artifacts or victim responses. Malware has the capability of penetrating intrusion detection systems, although these systems are essential for stopping those types of threats. In the future, we expect a substantial collaboration between AI and machine learning and cybersecurity.

10. Future Trends in AI for Mission-Critical Domains

With an increasing demand for improved performance, reliability, efficiency, and availability of mission-critical AI applications, users expect more from their AI systems. On the technical side, the potential list for nearly mature AI capabilities includes creditable explainable AI; more efficient, incremental continual learning; unsupervised and self-supervised learning as cost-effective alternatives to expensive human annotation; pervasive trustworthy AI powered by domain-aligned and reliable reasoning and simulation; autonomous netted and swarmed AI; and next-gen mixed-reality interfaces that break the human-user barrier. On the user side, growing ROI motivation from transforming organizational mission-critical functions into a closed-loop continuous autonomous processes driven by AI agents is a specific role model for the next maturity level of AI technology. Further, AI democratization at multi-levels, from national to industrial to institutional level, asks for and encourages parties interested in improving efficiency and effectiveness of public services to collaborate and cooperate with these governments for an AI service infrastructure.

However, the increasing AI autonomy, capabilities, and pervasiveness, augmented by the accompanying algorithmic biases, highlight AI risk management and insurance as the next hot global topics. AI ethics, equity, and fairness in system design, capabilities, application, and impact; trustworthy, controllable, transparent; and responsible, social-aware, and human supportive integrated concepts and integrated services will fill the growing demand for AI governance. AI regulation has been set on multidisciplinary multi-actor, multi-instrument national, international, and global agendas. New-generation permutable, dynamic, coalition-oriented, service-oriented group market mechanisms and infrastructures will be considered. Alternatively, with their federated learning, decentralized, cross-domain capabilities, AI may become a

Trojan horse for enabling such a journey toward an AI-enabled governance ecosystem.

10.1. Advancements in AI Technologies

Mission-critical domains are unfortunate early and frequent adopters and increasingly depend on AI in a full-stack, homogeneous stack. They have bettering investments in future-proofing their underlying foundations and systems to allow for continual innovation and future advancements in AI technologies. Moving forward, most of the largest boosts in AI capabilities will happen in these foundational technologies. These advancements will be increasingly generalizable and potent, allowing for more effective solutions and easier integration to diverse domain problems.

We anticipate rapid advancements particularly in three areas of underlying technological capabilities: scaling of foundation models using sparsity in standard AI, greater alignment and trustworthiness of AI systems by specializing foundation models with better dataset and methodology, and innovative new capabilities from efficient operators on AI systems using model-driven approaches through provable AI or breakthrough developments in software or hardware systems. These factors, coupled with improved interfaces and development economies such as annotation tooling, will open up large new domains of problems for AI-assisted solutions. As a direct implication, mission-critical systems and foundations will need to become rapidly adaptable to ever faster-moving models and capabilities on the reference task of AI-enhanced AI. In addition, we will increasingly rely on a partnership approach to integration between human experts and AI assistants to ensure that domain experts can take internal ownership of machine assistance pipelines. Such a partnership can use increased investments in semi-supervised methods and first-tech proper generalization of domain-specific tasks based on few clarification queries.

10.2. Regulatory and Ethical Considerations

As the use of AI technologies grows, crucial questions increasingly arise related to regulatory compliance or ethical alignment. Important work is required to understand how to ensure that AI models and tools are in alignment with legal regulations, such as data privacy laws, which are becoming increasingly prevalent and that require careful attention. Beyond regulatory considerations, important guidelines also exist for the ethical use of AI in mission-critical domains so that they can serve their intended purpose without inflicting harm on the people or communities they purport to help. Ethical considerations are not just about how the systems are used or how their outputs are applied; they also include the question of how trusted and fair the models and tools themselves are in their design and fair evaluation. For example, there have been calls for the responsible use of face recognition systems, stating that face recognition is not infallible. Human reviews are essential before any law enforcement action is taken based on the results of a face recognition scan. This is a recognition of the fact that face recognition systems, as built and trained today, are well known to have disparate performance across demographic groups.

The ethical considerations expand even further for categories of AI applications, such as emotion recognition, that could be considered highly intrusive. Such AI applications have little real benefit while posing great costs to societal trust in the technology as well as personal concerns regarding how the data is being used, who it is being shared with, and how it is being protected. These concerns have led to organizations discontinuing their offerings or entire lines of products based on emotion recognition technology. The reasons for their withdrawal vary, with some companies discontinuing products due to a lack of accuracy to the standard that the companies hold themselves to, while others note the ethical concerns of the technology, specifically the use of the technology without user knowledge.

11. Conclusion

Using artificial intelligence (AI) systems that automatically interpret sensory input to produce actions at least as good as the best human actors or counselors is attractive for mission-critical applications. This chapter outlines key functions of such systems as: reliable reasoning from minimal knowledge; continual self-improvement through skipping minor consequences of decisions based on reasoning; collaborating with other actors, including human, to increase payoffs; ensuring security through correct predictions of key properties of the world to avoid malicious agents; and minimizing exposure to risk through insurance against worst-case results of important decisions and selecting risks appropriate for their specialties. The concept of intelligent agents is briefly summarized for readers unfamiliar with this field, followed by examples of intelligent agents in action. Constraints on optimal actions are categorized; the role of AI in mission-critical applications is discussed; and the prospects of AI technical capability and technological conversion, as well as the risks posed by intelligent agent superpowers, are addressed.

In summary, only intelligent agents producing actions within boundaries set by human values have the prospect of potentially achieving better performance than other and human actors on mission-critical applications. Computer scientists and engineers are potentially in a good position to contribute significantly to a solution of AI risk by advancing trustworthiness technology that enables on-demand robotic or advisory agents available to augment and support good action selection by their users. In this scenario, intelligent agents assisting humans with decision-making at the highest levels are, at best, analogs of advanced missiles and troops guided by national desire and religious and cultural values. Some have speculated that values are too varied for there to be any guarantee of such congruence among nations for AI-capable actors. Continued emphasis on the production of trustworthy assistants, of whatever level of sophistication, may lessen that likelihood and should also support research into AI systems capable of collaborating effectively with sufficient levels of human optimization expertise.

References

[1] Holmes J, Sacchi L, Bellazzi R. Artificial intelligence in medicine. Ann R Coll Surg Engl. 2004;86(86):334-8.

[2] Varghese C, Harrison EM, O'Grady G, Topol EJ. Artificial intelligence in surgery. Nature medicine. 2024 May;30(5):1257-68.

[3] Neisser U. General, academic, and artificial intelligence. InThe nature of intelligence 2024 Mar 8 (pp. 135-144). Routledge.

[4] Klamma R, de Lange P, Neumann AT, Hensen B, Kravcik M, Wang X, Kuzilek J. Scaling mentoring support with distributed artificial intelligence. InInternational Conference on Intelligent Tutoring Systems 2020 Jun 3 (pp. 38-44). Cham: Springer International Publishing.

[5] Otaigbe I. Scaling up artificial intelligence to curb infectious diseases in Africa. Frontiers in Digital Health. 2022 Oct 21;4:1030427.

[6] Dasawat SS, Sharma S. Cyber security integration with smart new age sustainable startup business, risk management, automation and scaling system for entrepreneurs: An artificial intelligence approach. In2023 7th international conference on intelligent computing and control systems (ICICCS) 2023 May 17 (pp. 1357-1363). IEEE.

[7] Peteiro-Barral D, Guijarro-Berdiñas B. A study on the scalability of artificial neural networks training algorithms using multiple-criteria decision-making methods.

InInternational Conference on Artificial Intelligence and Soft Computing 2013 Jun 9 (pp. 162-173). Berlin, Heidelberg: Springer Berlin Heidelberg.

[8] Kuguoglu BK, van der Voort H, Janssen M. The giant leap for smart cities: Scaling up smart city artificial intelligence of things (AIoT) initiatives. Sustainability. 2021 Nov 7;13(21):12295.

[9] Gowda D, Chaithra SM, Gujar SS, Shaikh SF, Ingole BS, Reddy NS. Scalable ai solutions for iot-based healthcare systems using cloud platforms. In2024 8th International Conference on I-SMAC (IoT in Social, Mobile, Analytics and Cloud)(I-SMAC) 2024 Oct 3 (pp. 156-162). IEEE.

[10] Awan MZ, Jadoon KK, Masood A. Scalable and effective artificial intelligence for multivariate radar environment. Engineering Applications of Artificial Intelligence. 2023 Oct 1;125:106680.

[11] Landin M. Artificial intelligence tools for scaling up of high shear wet granulation process. Journal of Pharmaceutical Sciences. 2017 Jan 1;106(1):273-7.

[12] Panda SP. Securing 5G Critical Interfaces: A Zero Trust Approach for Next-Generation Network Resilience. In2025 12th International Conference on Information Technology (ICIT) 2025 May 27 (pp. 141-146). IEEE.

[13] Mocanu DC, Mocanu E, Stone P, Nguyen PH, Gibescu M, Liotta A. Scalable training of artificial neural networks with adaptive sparse connectivity inspired by network science. Nature communications. 2018 Jun 19;9(1):2383.

[14] Blanco L, Kukliński S, Zeydan E, Rezazadeh F, Chawla A, Zanzi L, Devoti F, Kolakowski R, Vlahodimitropoulou V, Chochliouros I, Bosneag AM. Ai-driven framework for scalable management of network slices. IEEE Communications Magazine. 2023 Nov 23;61(11):216-22.

[15] Sadek AH, Mostafa MK. Preparation of nano zero-valent aluminum for one-step removal of methylene blue from aqueous solutions: cost analysis for scaling-up and artificial intelligence. Applied Water Science. 2023 Feb;13(2):34.

[16] Cohen RY, Kovacheva VP. A methodology for a scalable, collaborative, and resource-efficient platform, MERLIN, to facilitate healthcare AI research. IEEE journal of biomedical and health informatics. 2023 Mar 20;27(6):3014-25.

[17] Adelodun AB, Ogundokun RO, Yekini AO, Awotunde JB, Timothy CC. Explainable artificial intelligence with scaling techniques to classify breast cancer images. InExplainable Machine Learning for Multimedia Based Healthcare Applications 2023 Sep 9 (pp. 99-137). Cham: Springer International Publishing.

[18] Sanz JL, Zhu Y. Toward scalable artificial intelligence in finance. In2021 IEEE International Conference on Services Computing (SCC) 2021 Sep 5 (pp. 460-469). IEEE.

[19] Haefner N, Parida V, Gassmann O, Wincent J. Implementing and scaling artificial intelligence: A review, framework, and research agenda. Technological Forecasting and Social Change. 2023 Dec 1;197:122878.

[20] Sai S, Chamola V, Choo KK, Sikdar B, Rodrigues JJ. Confluence of blockchain and artificial intelligence technologies for secure and scalable healthcare solutions: A review. IEEE Internet of Things Journal. 2022 Dec 29;10(7):5873-97.

[21] Moro-Visconti R. Artificial Intelligence-Driven Digital Scalability and Growth Options. InArtificial Intelligence Valuation: The Impact on Automation, BioTech, ChatBots,

FinTech, B2B2C, and Other Industries 2024 Jun 2 (pp. 131-204). Cham: Springer Nature Switzerland.

[22] Oikonomou EK, Khera R. Designing medical artificial intelligence systems for global use: focus on interoperability, scalability, and accessibility. Hellenic Journal of Cardiology. 2025 Jan 1;81:9-17.

[23] Sayed-Mouchaweh M, Sayed-Mouchaweh, James. Artificial Intelligence Techniques for a Scalable Energy Transition. Springer International Publishing; 2020.

[24] Govea J, Ocampo Edye E, Revelo-Tapia S, Villegas-Ch W. Optimization and scalability of educational platforms: Integration of artificial intelligence and cloud computing. Computers. 2023 Nov 1;12(11):223.

[25] Hammad A, Abu-Zaid R. Applications of AI in decentralized computing systems: harnessing artificial intelligence for enhanced scalability, efficiency, and autonomous decision-making in distributed architectures. Applied Research in Artificial Intelligence and Cloud Computing. 2024;7(6):161-87.

[26] Pazho AD, Neff C, Noghre GA, Ardabili BR, Yao S, Baharani M, Tabkhi H. Ancilia: Scalable intelligent video surveillance for the artificial intelligence of things. IEEE Internet of Things Journal. 2023 Mar 31;10(17):14940-51.

[27] Sakly H, Guetari R, Kraiem N, editors. Scalable Artificial Intelligence for Healthcare: Advancing AI Solutions for Global Health Challenges. CRC Press; 2025 May 6.

[28] Shivadekar S, Halem M, Yeah Y, Vibhute S. Edge AI cosmos blockchain distributed network for precise ablh detection. Multimedia tools and applications. 2024 Aug;83(27):69083-109.

[29] Bano S, Tonellotto N, Cassarà P, Gotta A. Artificial intelligence of things at the edge: Scalable and efficient distributed learning for massive scenarios. Computer Communications. 2023 May 1;205:45-57.

[30] Mishra A. Scalable AI and Design Patterns: Design, Develop, and Deploy Scalable AI Solutions. Springer Nature; 2024 Mar 11.

[31] Panda SP. Augmented and Virtual Reality in Intelligent Systems. Available at SSRN. 2021 Apr 16.

[32] Rane N, Choudhary S, Rane J. Artificial intelligence for enhancing resilience. Journal of Applied Artificial Intelligence. 2024 Sep 9;5(2):1-33.

[33] Nath PC, Mishra AK, Sharma R, Bhunia B, Mishra B, Tiwari A, Nayak PK, Sharma M, Bhuyan T, Kaushal S, Mohanta YK. Recent advances in artificial intelligence towards the sustainable future of agri-food industry. Food Chemistry. 2024 Jul 30;447:138945.

[34] Abisoye A, Akerele JI. A scalable and impactful model for harnessing artificial intelligence and cybersecurity to revolutionize workforce development and empower marginalized youth. International Journal of Multidisciplinary Research and Growth Evaluation. 2022 Jan;3(1):714-9.

[35] Raman R, Buddhi D, Lakhera G, Gupta Z, Joshi A, Saini D. An investigation on the role of artificial intelligence in scalable visual data analytics. In2023 International Conference on Artificial Intelligence and Smart Communication (AISC) 2023 Jan 27 (pp. 666-670). IEEE.

[36] Panda SP. The Evolution and Defense Against Social Engineering and Phishing Attacks. International Journal of Science and Research (IJSR). 2025 Jan 1.

[37] Newton C, Singleton J, Copland C, Kitchen S, Hudack J. Scalability in modeling and simulation systems for multi-agent, AI, and machine learning applications. InArtificial Intelligence and Machine Learning for Multi-Domain Operations Applications III 2021 Apr 12 (Vol. 11746, pp. 534-552). SPIE.

[38] Bestelmeyer BT, Marcillo G, McCord SE, Mirsky S, Moglen G, Neven LG, Peters D, Sohoulande C, Wakie T. Scaling up agricultural research with artificial intelligence. IT Professional. 2020 May 21;22(3):33-8.

[39] Meir Y, Sardi S, Hodassman S, Kisos K, Ben-Noam I, Goldental A, Kanter I. Power-law scaling to assist with key challenges in artificial intelligence. Scientific reports. 2020 Nov 12;10(1):19628.

[40] Shivadekar S, Kataria DB, Hundekar S, Wanjale K, Balpande VP, Suryawanshi R. Deep learning based image classification of lungs radiography for detecting covid-19 using a deep cnn and resnet 50. International Journal of Intelligent Systems and Applications in Engineering. 2023;11:241-50.

[41] Panda SP. Relational, NoSQL, and Artificial Intelligence-Integrated Database Architectures: Foundations, Cloud Platforms, and Regulatory-Compliant Systems. Deep Science Publishing; 2025 Jun 22.

[42] Shlezinger N, Ma M, Lavi O, Nguyen NT, Eldar YC, Juntti M. Artificial intelligence-empowered hybrid multiple-input/multiple-output beamforming: Learning to optimize for high-throughput scalable MIMO. IEEE Vehicular Technology Magazine. 2024 May 20;19(3):58-67.

[43] Samuel O, Javaid N, Alghamdi TA, Kumar N. Towards sustainable smart cities: A secure and scalable trading system for residential homes using blockchain and artificial intelligence. Sustainable Cities and Society. 2022 Jan 1;76:103371.

[44] Villegas-Ch W, Govea J, Gurierrez R, Mera-Navarrete A. Optimizing security in IoT ecosystems using hybrid artificial intelligence and blockchain models: a scalable and efficient approach for threat detection. IEEE Access. 2025 Jan 22.

[45] Mungoli N. Scalable, distributed AI frameworks: leveraging cloud computing for enhanced deep learning performance and efficiency. arXiv preprint arXiv:2304.13738. 2023 Apr 26.

[46] Panda SP. Artificial Intelligence Across Borders: Transforming Industries Through Intelligent Innovation. Deep Science Publishing; 2025 Jun 6.

[47] Cheetham AK, Seshadri R. Artificial intelligence driving materials discovery? perspective on the article: Scaling deep learning for materials discovery. Chemistry of Materials. 2024 Apr 8;36(8):3490-5.

[48] Panda SP, Muppala M, Koneti SB. The Contribution of AI in Climate Modeling and Sustainable Decision-Making. Available at SSRN 5283619. 2025 Jun 1.

[49] Shivadekar S. Artificial Intelligence for Cognitive Systems: Deep Learning, Neuro-symbolic Integration, and Human-Centric Intelligence. Deep Science Publishing; 2025 Jun 30.

[50] DeCost BL, Hattrick-Simpers JR, Trautt Z, Kusne AG, Campo E, Green ML. Scientific AI in materials science: a path to a sustainable and scalable paradigm. Machine learning: science and technology. 2020 Jul 14;1(3):033001.