ISSN: 3006-4023 (Online), Vol. 2, Issue 1
Journal of Artificial Intelligence General Science (JAIGS)

Journal Homepage: https://jaigs.org/index.php/JAIGS

Exploring the Latest Trends in Artificial Intelligence Technology: A Comprehensive Review

Jeff Shuford[1], Md.Mafiqul Islam[2]

[1]Nationally Syndicated Business & Technology Columnist,USA
[2]Department of Information Science and Library Management, University of Rajshahi, Bangladesh

Abstract

Artificial intelligence (AI) has become increasingly pervasive across various domains, including smartphones, social media platforms, search engines, and autonomous vehicles, among others. This study undertakes a scoping review of the current landscape of AI technologies, following the PRISMA framework, with the aim of identifying the most advanced technologies utilized in different domains of AI research. Three reputable journals within the artificial intelligence and machine learning domain, namely the Journal of Artificial Intelligence Research, the Journal of Machine Learning Research, and Machine Learning, were selected for this review. Articles published in 2022 were scrutinized against certain criteria: the technology must be tested against comparable solutions, employ commonly approved or well-justified datasets, and demonstrate improvements over comparable solutions. A crucial aspect of technology development identified in this review is the processing and exploitation of data collected from diverse sources. Given the highly unstructured nature of data, technological solutions should minimize the need for manual intervention by humans. The review indicates that creating labeled datasets is a labor-intensive process, leading to increased research focus on solutions leveraging unsupervised or semi-supervised learning technologies. Efficient updating of learning algorithms and the interpretability of predictions emerge as key considerations in the development of AI technologies. Moreover, in real-world applications, ensuring safety and providing explainable predictions are imperative before widespread adoption can be achieved. Thus, this review underscores the importance of addressing these factors to facilitate the responsible and effective integration of AI technologies into various domains.

Keywords: *Artificial Intelligence, Deep Learning, Machine Learning, Natural Language Processing, Reinforcement Learning, Scoping Review*

Article Information:
Article history: *Received:* 01/01/2024 *Accepted:* 10/01/2024 *Online:* 07/02/2024 *Published*: 07/02/2024
Corresponding author: Jeff Shuford Email: jshuford024@gmail.com

© The Author(s) 2024. Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permitsuse, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the originalauthor(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other thirdparty material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the mate-rial. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation orexceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit http://creativecommons.org/licenses/by/4.0

Introduction

Artificial intelligence (AI) technologies are employed across diverse domains to tackle various tasks, categorized based on the nature of the task and the data involved. This study examines several such domains:

1. Natural Language Processing (NLP): Focuses on processing human language and encompasses tasks such as text classification, summarization, translation, and distinguishing between fake and real news.

2. Computer Vision: Involves processing and utilizing images and videos, including tasks such as identifying artificial from real images and predicting events like disease outbreaks.

3. Reinforcement Learning: In this domain, an AI agent interacts with an unknown environment, receiving rewards based on its actions. It finds applications in tasks like imperfect-information games and simulating animal or human cognition.

4. Motion Planning: Concerned with modelling the motions of autonomously operating systems, such as self-driving cars, and is crucial in the development of autonomous vehicles.

These domains are organized hierarchically into branches, each containing various tasks and problems, along with corresponding technological solutions. In this scoping review, 21 different tasks are identified and grouped into five branches: Natural Language Processing, Computer Vision, Robotics and Motion, Reinforcement Learning, and Others. Each branch encompasses tasks and problems tackled with state-of-the-art technological solutions tested against comparable alternatives.

While numerous scoping reviews exist, many focus on specific fields such as healthcare, construction, education, or automotive industries. However, there's a notable gap in reviews mapping technological solutions across different tasks and problems.

Given the rapid advancement of AI technologies, this scoping review aims to systematically identify the most recent and advanced solutions developed and used in 2022 across various tasks and domains. The paper's structure includes a methodology section, results for each branch, a discussion summarizing key findings, and concluding remarks with recommendations for future research directions.

Method

This study employed a scoping review approach, following the guidelines outlined by Arksey and O'Malley (2005) and the extension provided by the Preferred Reporting Items for Systematic Reviews and Meta-Analyses (PRISMA). PRISMA offers a standardized checklist of reporting items essential for conducting scoping reviews. For this review, scientific publications from the domains of artificial intelligence and machine learning were scrutinized. Specifically, articles published in the year 2022 were selected for analysis.

Three prominent journals—Journal of Artificial Intelligence Research, Journal of Machine Learning Research, and Machine Learning—were chosen based on their high-level classification by the Finnish Publication Forum (pub) and their relevance to the research domain. Only journals with a classification level of 3, denoting the highest level, and a primary focus on artificial intelligence or machine learning technology, were considered. All volumes and issues published in 2022 in these selected journals were included in the identification process.

Each article underwent individual inspection, with strict criteria applied to qualify technological solutions for inclusion: the selected technology must be compared against comparable solutions, utilize commonly approved or well-justified datasets during application, and demonstrate improvements over existing solutions in specific areas.

A total of 248 articles were initially screened based on title and abstract to identify potentially relevant publications. Subsequently, 76 articles were evaluated based on full-text examination to identify pertinent publications. From these, 27 articles were ultimately selected for data charting. The flowchart depicting the selection process is illustrated in

Figure 1, adhering to the PRISMA framework.

To facilitate data extraction, a structured data charting form was devised. Key factors searched for included the type of task and problem addressed by the solution, the proposed solution itself, the underlying technology, datasets used for testing, solutions compared against during evaluation, and testing results. Articles were categorized based on the task and problem that the technology aimed to address.

The outcomes of the scoping review are as follows:

Based on the findings of the scoping review, artificial intelligence technologies are categorized into five distinct branches: Natural Language Processing, Computer Vision, Robotics and Motion, Reinforcement Learning, and Others. These branches were established to categorize technological solutions based on the specific tasks and problems they aim to address.

In this study, Natural Language Processing encompasses tasks centered around the processing of textual data, while Computer Vision focuses on tasks involving the processing of visual data. Reinforcement Learning pertains to tasks that utilize reinforcement learning agents, and Robotics and Motion encompass tasks related to the motion of autonomous systems, including trajectory prediction.

The category labeled as others includes tasks that do not fit neatly into any of the aforementioned branches but may incorporate elements from multiple branches. Within each branch, various tasks and problems are introduced along with the most recently proposed solutions. This categorization was informed by the observations gleaned from the scoping review, highlighting differences in tasks and problems based on the nature of the processed and generated data. Additionally, reinforcement learning agents are distinguished as a separate branch due to their unique characteristics.

In the subsequent sections, each branch is presented along with the tasks and problems it encompasses. The branches are introduced in the following sequence: Natural Language Processing, Computer Vision, Robotics and Motion, Reinforcement Learning, and Others. Furthermore, each branch is accompanied by a diagram for illustrative purposes

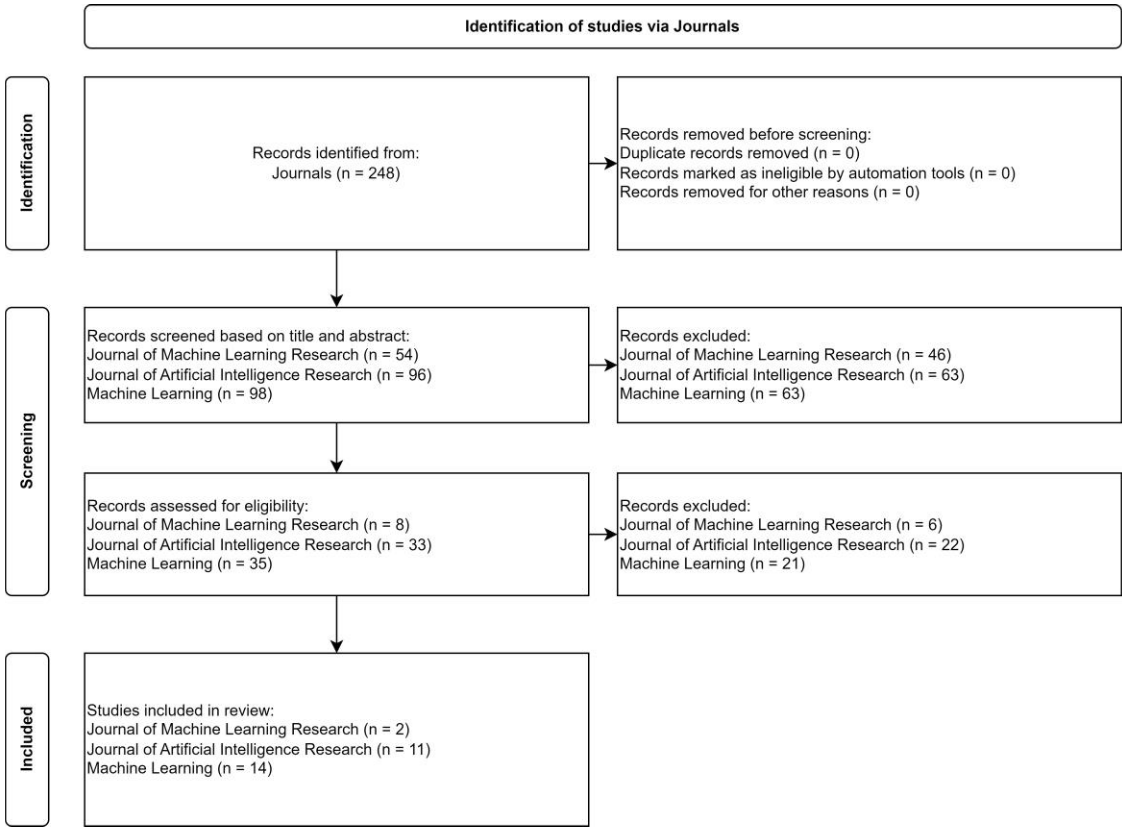

Natural Language Processing

Natural language processing (NLP) constitutes a pivotal branch of artificial intelligence dedicated to the comprehension and processing of human language. This field underpins numerous real-world applications, including text translation and summarization, offering solutions to tasks that would otherwise be impractical or excessively laborious when handled by human workers. NLP extensively leverages both traditional machine learning techniques and neural network solutions. Notably, the integration of deep learning methodologies with NLP has yielded state-of-the-art results (Campesato, 2021). This branch encompasses tasks and problems associated with the processing of textual data. The structure of this branch is depicted in Figure 2.

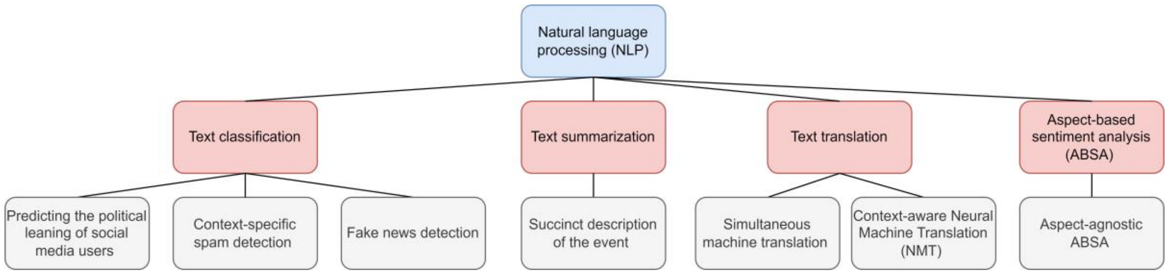

Text Classification

Text classification serves a critical role in various applications such as fake news detection, spam identification, and predicting political leanings (de Souza et al., 2022; Kawintiranon et al., 2022; Fagni and Cresci, 2022). Typically, supervised learning algorithms are employed in text classification tasks, albeit facing challenges due to the necessity of labeled datasets (de Souza et al., 2022). To address this limitation, extensive research is being conducted on unsupervised and semi-supervised learning approaches (de Souza et al., 2022; Fagni and Cresci, 2022).

Fake News Detection

Detecting fake news often relies on binary classification and requires a sizable and balanced set of labeled news articles, although such datasets may not reflect real-world scenarios adequately (de Souza et al., 2022). One-Class Learning (OCL) algorithms, such as k-Means and k-Nearest Neighbors Density-based (k-NND), have shown promising results but are subject to limitations (de Souza et al., 2022). Positive and Unlabeled learning (PUL) algorithms, like Rocchio Support Vector Machine (RC-SVM) and Positive and Unlabeled learning by Label Propagation (PU-LP), mitigate the need for extensive labeling efforts and have demonstrated superior performance in various scenarios (de Souza et al., 2022).

Context-Specific Spam Detection

Identifying context-specific spam entails distinguishing irrelevant content from traditional spam, which poses a significant challenge on social media platforms (Kawintiranon et al., 2022). State-of-the-art solutions often rely on models based on Random Forest, leveraging both content and user information for accurate detection (Kawintiranon et al., 2022).

Predicting the Political Leaning of Social Media Users

With the increasing consumption of political content on social media, predicting the political leanings of users has become crucial for various analyses (Fagni and Cresci, 2022). Content-based solutions, utilizing unsupervised learning approaches like Parties enriched + clustering, offer effective means of predicting political preferences, albeit requiring labeled data for optimal performance (Fagni and Cresci, 2022).

Text Summarization

In light of information overload from diverse sources, text summarization plays a vital role in providing concise event descriptions (Chen et al., 2022). Multi-document Event Summarization (MES) frameworks, such as Event-Pg, aim to capture the essence of events using various techniques, including encoder-decoder neural network models and graph-based methods (Chen et al., 2022).

Text Translation

Conventional machine translation and real-time translation techniques, including Simultaneous Machine Translation, are continuously evolving to improve translation quality and efficiency (Haralampieva et al., 2022). Context-aware Neural Machine Translation (NMT), employing pre-trained language models like BERT, enhances translation accuracy by integrating contextual information effectively (Wu et al., 2022).

Aspect-Based Sentiment Analysis (ABSA)

Aspect-based sentiment analysis focuses on discerning sentiment in sentences, especially when analyzing comments online (Xing and Tsang, 2022). Aspect-Agnostic ABSA methods, incorporating Long Short-Term Memory networks (LSTM), Graph Convolutional Networks (GCN), and pre-trained BERT models, improve sentiment analysis accuracy by preserving aspect-specific information effectively (Xing and Tsang, 2022).

Computer Vision

Computer vision encompasses diverse applications in our daily lives, facilitated by advancements in deep learning techniques (Dadhich, 2018). Tasks and problems involving the processing of visual data are included in this branch, addressing various challenges and applications (Dadhich, 2018)..

Image Recognition

With the advancement of generative models, distinguishing artificial images from real ones has become imperative. Utilizing wavelets in image processing presents a promising solution to this challenge (Wolter et al., 2022). Another key issue in image recognition algorithms is ensuring continuous learning, with recent solutions focusing on updating specific modules rather than retraining the entire model (Skantze and Willemsen, 2022).

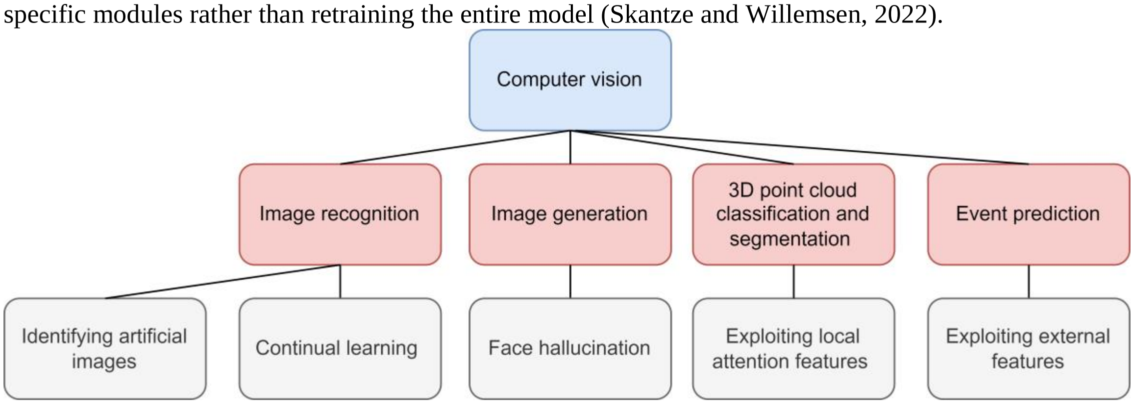

Identifying Artificial Images

Generative adversarial networks (GANs) have dual applications, from generating images to creating deceptive deepfakes, emphasizing the need for technologies to discern real images from fake ones (Wolter et al., 2022). Deepfake detection methods operate in either the frequency or spatial domain, employing techniques like discrete cosine transform (DCT) or convolutional neural networks (CNNs) for effective classification (Wolter et al., 2022). Wavelets, traditionally used in applied mathematics, are now being integrated into neural networks to enhance image processing capabilities (Wolter et al., 2022).

Continual Learning

Multimodal representation learning algorithms, such as Contrastive Language-Image Pre-Training (CLIP), combine language and vision to ground meaning effectively (Skantze and Willemsen, 2022). However, ensuring continual learning remains a challenge, addressed by CoLLIE, which updates a transformation model without retraining the entire CLIP model (Skantze and Willemsen, 2022).

Image Generation

Face recognition algorithms often struggle with low-resolution images, prompting the use of super-resolution techniques to enhance image quality (Han et al., 2022). Face hallucination methods, like the C-Face network, aim to reconstruct high-resolution images from low-resolution inputs while preserving identity features (Han et al., 2022).

3D Point Cloud Classification and Segmentation

3D point clouds offer valuable geometric information but pose challenges due to their unstructured nature (Gao et al., 2022). The Spatial Depth Attention (SDA) network integrates global and local attention features, improving classification and segmentation tasks (Gao et al., 2022).

Event Prediction

Spatio-temporal event data, such as disease outbreaks, require predictive modeling to anticipate future occurrences (Okawa et al., 2022). Considering external features is essential for accurate event prediction and subsequent risk assessment (Okawa et al., 2022).

Exploiting External Features

The Hawkes process, a mathematical framework commonly used to model events like infectious diseases and earthquakes, typically overlooks external factors such as weather and population distribution, which can significantly influence event triggers (Okawa et al., 2022). Integrating external information, readily available from platforms like GIS, into the Hawkes process has been explored in several studies, albeit with methods requiring handcrafted features that are unable to leverage unstructured data, such as images (Okawa et al., 2022).

Okawa et al. (2022) introduce Convolutional Hawkes Process (ConvHawkes), a novel architecture that enhances Hawkes processes by incorporating georeferenced images using a Convolutional Neural Network (CNN) with continuous kernel convolution. This approach, consisting of external effect and spatio-temporal decay components, leverages CNNs to transform images into latent feature maps, expanding them onto the continuous spatio-temporal space to capture information from external factors. By integrating neural network models into the Hawkes process formulation, ConvHawkes surpasses alternative solutions like Spatio-temporal homogeneous Poisson Process (SPP) and Recurrent Marked Temporal Point Process (RMTPP), demonstrating superior performance across various metrics (Okawa et al., 2022).

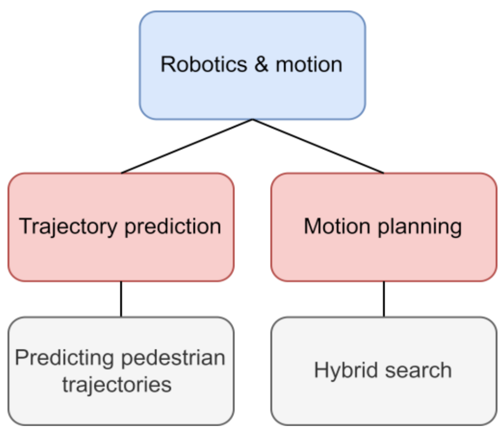

Robotics and Motion

In the development of autonomous systems, trajectory prediction and motion planning are pivotal (Prédhumeau et al., 2022; Strawser and Williams, 2022). For instance, self-driven cars need to predict the trajectories of external factors like pedestrians to prevent accidents (Prédhumeau et al., 2022). Motion planning ensures that self-driven cars are prepared for diverse scenarios, including uncertain factors (Strawser and Williams, 2022). This section encompasses tasks and challenges related to motion planning and trajectory prediction in autonomous systems.

Trajectory Prediction

Trajectory prediction is essential for the development of autonomous vehicles, particularly in forecasting pedestrian trajectories, given the evolving urban landscape (Prédhumeau et al., 2022).

Predicting Pedestrian Trajectories

Pedestrian safety is paramount in autonomous vehicle development, especially in shared spaces where vehicles and pedestrians interact. Predicting pedestrian trajectories in such scenarios requires a combination of expert and data-driven models. While expert models like the Social Force Model (SFM) are widely used, data-driven models offer higher accuracy but demand substantial data and struggle with real-time predictions in crowded spaces. Prédhumeau et al. (2022) propose an Agent-Based Modeling (ABM) approach, which simulates pedestrian behavior in shared spaces more realistically, capturing the diversity of pedestrian actions and improving collision prediction compared to traditional SFM-based methods.

Motion Planning

Motion planning faces challenges due to uncertainties inherent in autonomous environments, necessitating quick decision-making by agents. Solutions like hybrid search algorithms have emerged to address these challenges effectively (Strawser and Williams, 2022).

Hybrid Search

Motion planning, especially for autonomous underwater vehicles, requires handling numerous scenarios efficiently while meeting predefined requirements. Hybrid search algorithms, such as the one proposed by Strawser and Williams (2022), leverage a combination of region and trajectory planners to produce feasible solutions faster than traditional methods like Chance-Constrained Rapidly-Exploring Random Tree (CC-RRT) and Disjunctive Linear Programming (DLP). Their solution, integrating region and trajectory planners with risk models, demonstrates superior performance in generating feasible trajectories within specified constraints.

Reinforcement Learning

Reinforcement learning (RL) stands as a cornerstone of artificial intelligence, where agents interact with unknown environments to optimize actions based on rewards (Sewak, 2019). This section encompasses tasks and challenges involving RL agents, illustrating the evolving landscape of advanced artificial intelligence agents.

Idle education

Applications of imitation learning (IL) in fields like robotics and natural language processing (NLP) have yielded encouraging results. IL aligns expert data's behaviors and learning policies. Policy is retrieved using limited expert knowledge that has been acquired. One method for IL that has shown promise is inverse reinforcement learning, or IRL. Using generative-discriminative framework, IRL has demonstrated some significant performance breakthroughs. One of the main issues with these frameworks has been the imbalance between the discriminator and the generator. Various approaches have been presented to address this problem (Zhang et al., 2022).

Steer clear of gauche painting

Due to the differing rates of learning between the discriminator and the generator, imbalance between them arises particularly with Generative Adversarial Imitation Learning (GAIL) and poses a serious issue during training. The generator generates state-action pairings and is a Reinforcement Learning (RL) agent. Using supervised learning, the discriminator picks up new information far more quickly than the RL generator. This results in a situation where the generator performs poorly even while a well-trained discriminator is obtained (Zhang et al., 2022).

Zhang et al. (2022) present a method that uses a discriminator—more akin to a teacher—to try and produce a generator that is well-trained. To prevent gradient vanishing, the discriminator provides the generator with appropriate fictitious rewards. Generative Adversarial Imitation Learning with Variance Regularization (GAIL-VR) is the name of the suggested remedy. The fundamental problem with GAIL is that, because of the disparity in learning speeds, each state-action combination generated by the generator may be identified by the discriminator. resulting in a scenario where the generator receives meager rewards and little diversity. The generator must be able to create state-action pairings from the expert data that the discriminator is unable to distinguish in order to receive larger rewards. Attempts have been made to eliminate the gradient vanishing phenomenon by refining the discriminator algorithm in order to overcome this issue. GAIL-VR was contrasted with other methods, including Wasserstein Adversarial Imitation Learning (W-GAIL) and regular GAIL. According to the findings, GAIL-VR achieves the highest training speed and average reward in practically all MuJoCO environments. It was determined that both low and high dimensions state-action spaces exhibit good performance from GAIL-VR. Additionally, the suggested method was quite stable and the same parameters worked in a variety of settings.

The game of imperfect information

Reinforcement learning algorithms can benefit from imperfect-information games like HUNL Poker and StarCraft II (Liu et al., 2022). (Bertsimas and Paskov, 2022). There are issues where a great number of decisions must be taken in a short amount of time due to the vast amount of states and action spaces (Liu et al., 2022). Research has been done on various approaches, and efforts have been made to find more effective solutions to these issues (Liu et al., 2022; Bertsimas and Paskov, 2022). One significant obstacle to reinforcement learning

One drawback of algorithms has been their interpretability. Prior to being widely used, the interpretability issue must be resolved (Bertsimas and Paskov, 2022).

The Whole-Length Starcraft Ii Game

Profound promise has been demonstrated by Deep Q-Network (DQN) and Deep Reinforcement Learning (DRL) in addressing problems like Atari games, Go, and self-driving cars. Large-scale issues, however, still seem to be difficult for RL algorithms to solve. StarCraft has gained popularity as a setting for investigating the potential of reinforcement learning algorithms. StarCraft offers a sizable map with enormous state and action zones as an environment. StarCraft offers imprecise information, and thousands of choices must be made in intervals of ten to thirty minutes. Since it's a multi-agent game, working with other players may be necessary. Based on StarCraft II (SC2), the learning platform used with RL algorithms is known as the StarCraft II Learning Environment (SC2LE) (Liu et al., 2022).

The most encouraging outcomes with regard to RL in SC2 gameplay have been attained by AlphaStar. There have been some claims that AlphaStar uses too much human knowledge. AlphaStar only employs human-envisioned techniques; it does not develop any original ones. It has been suggested that AlphaStar has not fully resolved the SC2 bug for these and other reasons. An architecture like AlphaStar is used by two open-source projects: DI-Star and SC2IL (Liu et al., 2022).

The difficult problem is divided into multiple smaller problems using the Hierarchical Reinforcement Learning (HRL) technique, and each smaller problem is solved separately. A scenario known as "the curse of dimensionality" occurs when there is an exponential increase in the number of explorable states and state-space has a large dimension. Because the problem is divided into smaller subproblems, HRL can be used to solve the curse of dimensionality. Option, MaxQ, and ALISP are a few examples of conventional HRL algorithms. Recently, other HRL algorithms have been presented, including FeUdalNetwork, Option-Critic, and Meta Learning Shared Hierarchies (MLSH). According to Liu et al. (2022), MLSH, which is based on meta-learning, has outperformed the PPO algorithm in some tasks.

Liu et al. (2022) suggested the HRL approach as a remedy to the SC2 issue. Two distinct timelines and two different kinds of policies are included in the answer. The controller obtains a global observation over an extended period of time and selects a sub-policy on the basis of that observation. The controller-selected sub-policy chooses a macro action in a brief amount of time. Each sub-policy gets its own local observation as well as its own action-space and reward goals. The architecture is two layers deep, but it can be made three layers deep by adding more layers beneath the sub-policy layer, for example. Large state and action spaces can be divided into smaller ones using HRL. The algorithm is easier to train when sub-policies have their own action spaces. During the training phase, curriculum learning was also implemented. The curriculum was created with a range of difficulty levels in mind, from easy to challenging, to assist agents in their training. The controller was configured to choose a sub-policy every eight seconds, and the sub-policy would execute a macro-action once every second. A non-hierarchical design with the controller and sub-policies turned off was used to test the HRL algorithm. The findings show that when the problem gets harder, the HRL algorithm will perform noticeably better than the non-hierarchical approach. The suggested remedy was contrasted with TStarBots and mini-AlphaStar (mAS), which comprise nearly all

an element of AlphaStar. The findings showed that while TStarBots needed a lot more human expertise and computer power, they outperformed the HRL solution in terms of performance. The HRL approach outperforms comparable methods in terms of resource efficiency and performs better even with constrained computational resources.

Accurate Framework

One major disadvantage of reinforcement learning (RL) agents is their lack of interpretability. This is a serious issue that will hinder the technology's widespread adoption if it is not resolved. One of the theories for why RL agents become difficult to interpret is the use of neural networks (Bertsimas and Paskov, 2022).

Bertsimas and Paskov are studying interpretable RL-agent (2022). This study compares three distinct learning algorithms: Feedforward Neural Network, Extreme Gradient Boosted Trees (XGBoost), and Optimal Classification Trees (OCT) against Slumbot in the HUNL Poker game as well as against each other. A new method for converting the game state into a vector—a mix of two elements—is presented. Utilizing the Counterfactual Regret Minimization (CFR) self-play algorithm, the average strategy in HUNL Poker is determined. The three algorithms discussed above attempt to learn this CFR-discovered tactic. It can be assumed that OCT is the most interpretable neural network while Feedforward Neural Network is the least interpretable. While XGBoost is not as non-interpretable as a feedforward neural network, it is still not very clear why. The goal of the project was to develop a framework that would make the entire process—from feature generation to model training—as interpretable as feasible. Every algorithm performs better than Slumbot, according to the results. The top performer is an agent built on a feedforward neural network, closely followed by XGBoost. While using considerably fewer parameters, OCT-based agents are not far behind Neural Network-based agents. It was determined that using an interpretable architecture, it is possible to build a very potent HUNL Poker agent. Human-readable printouts can be generated by using OCT as a learning algorithm to examine the approach the system takes.

Cognitive thinking that is physical

When applied to challenging games, deep reinforcement learning (DRL) algorithms have demonstrated encouraging outcomes. However, these algorithms typically lack the ability to mimic fundamental human cognitive abilities like spatial reasoning. In artificial intelligence contests, Animal-AI (AAI) has served as a testbed for assessing physical cognitive reasoning (Mitchener et al., 2022).

Critical Perception

The best DRL systems, according to the results, were unable to complete tasks like spatial elimination, which calls for common sense physical reasoning. DRL systems have weak generalization abilities with unknown samples and are opaque. Neural networks, which are employed as function approximators in DRL systems, carry these inherent defects. Numerous studies have been done that look into the combination of symbolic and neurological systems. Positive outcomes have been attained as symbolic systems typically improve DRL approaches' performance and interpretability (Mitchener et al., 2022).

Mitchener et al. (2022) introduced the Detect, Understand, Act (DUA) strategy. The DUA employed two distinct operational levels: macro-level and micro-level. The mapping of environmental observation for discrete actions and temporal abstraction are done at the micro-level. Macro-level is a timescale that associates options with symbolic states and is made up of a vast number of environment timesteps. The three parts of the suggested remedy are Detect, Understand, and Act. The Detect module transforms data from the environment into a meaningful representation at each timestep. The Understand module processes this symbolic representation, which is also referred to as meaningful representation. The Learned Meta-Policy is used by the Understand module to start the appropriate option. The Act module has these options, which are DRL agents who have already undergone training. Instructions on how to filter the input data delivered to these DRL agents are provided by the Understand module. Raw photos are filtered by the Detect module to produce more usable features. A set of bounding boxes is created by parsing images. These bounding boxes provide information that is converted into an Answer Set Programming (ASP) application. Both the ASP program and the Inductive Logic Answer Set Programming (ILASP) learner are sub-modules of the Understand module. The symbolic meta-policies utilized in the Understand module are learned by Inductive Meta-Policy learning (IMP). Proximal Policy Optimization (PPO) is the DRL algorithm used by the Act module. The top competitors in the 2019 AAI tournament were compared to DUA. According to the outcomes, the suggested solution had cutting-edge functionality and would have placed third in the AAI competition for 2019. It was determined that DUA may generalize to more difficult tasks and produce outcomes that are easy to understand.

Rewarding learning that is safe

A great deal of interaction with the environment is needed for training Reinforcement Learning (RL) agents in simulated environments. When attempting to train physical systems in real-world settings, this might provide a significant challenge. Safety concerns need to be taken into account more carefully in real-world settings, and it would be possible to include safety aspects into system behavior. For example, while flying in an area where people are present, autonomous drone routes may be restricted. The safety and sample efficiency aspects of a framework are necessary in real-world scenarios. When it comes to safety, there are two distinct methods. The safety of the learning process or the governing agent/policy may be prerequisites (Cowen-Rivers et al., 2022).

Value-At-Risk Conditionality as a Safety Limit

Cowen-Rivers et al. (2022) introduce SAfe Model-Based and Active reinforcement learning (SAMBA) as a potential approach. SAMBA is suitable for usage in continuous state and action space for learning control, and it employs Conditional Value-at-Risk (CVaR) as a safety constraint. The PILCO serves as a model for the SAMBA structure, which is modified using two cutting-edge techniques that enable both active exploration and safety. The three main goals of SAMBA are to adhere to safety regulations, maximize active exploration, and reduce costs. Samba was evaluated in comparison to unconstrained model-free algorithms like PPO, unconstrained model-based algorithms like PlaNet, and expectation-constrained model-free algorithms like Safety-constrained

PPO (SPPO). In light of the findings, SAMBA lowers sample and overall costs.

Utilized in training as (TC). In comparison to competing systems, SAMBA's safety performance was on level with or better. Additionally, it was determined that sample efficacy did not compromise testing-phase safety.

A visible reward system

An agent that uses reinforcement learning (RL) must interact with its surroundings in order to learn from its experiences. It's a popular belief that an RL agent is unaware of its reward function. This implies that while the reward function might contain certain high-level concepts that the agent could use, the agent only asks the function questions based on the circumstances at hand. A better and quicker learning outcome could be obtained by having access to the reward function's structure (Icarte et al., 2022).

Developing Internal Structure Reward Functions

Icarte et al. (2022) introduce the reward machine approach. Several reward functions make up the reward machine. An agent behaves the same way inside the reward machine as it does when it moves between states within an environment. Based on the state, the reward machine's output determines which reward function the agent should utilize. The agent is aware of the number of states that are accessible and can make use of this knowledge as they learn. Two structures were used with the reward machine: Hierarchical RL for Reward Machines (HRM) and Counterfactual Experiences for Reward Machines (CRM). CRM makes advantage of its prior experiences to produce the appropriate behavior at various reward machine states. When an agent is assigned a mission that requires them to pick up a coffee before reaching their location, they will not receive a reward if they arrive at their destination before picking up the coffee. As soon as CRM locates the coffee, it will use this information of how to get there. The difficulty with HRM is broken down into smaller issues known as choices. For instance, HRM would learn numerous regulations, including picking up the mail before the coffee and the coffee before the mail, if the task involved picking up the coffee and mail before arriving at the location. Furthermore, a new technique named Automated Reward Shaping (RS) is presented, the main idea of which is to give the agent interim rewards as it learns the task. HRM and CRM are used with RS. In various test settings, Q-learning (QL), Double Deep Q-Network (DDQN), and Deep Deterministic Policy Gradient (DDPG) were employed as the main off-policy learning techniques. Comparisons were made between CRM and HRM as well as between vanilla versions of the main off-policy learners. The findings show that CRM works better than alternative alternatives in the majority of the test situations. HRM performed best in a setting where both the state space and the action were continuous. Performance was enhanced in certain test conditions by the proposed RS. It was determined that while the suggested solutions performed better than the baseline solutions in each experiment, the computing cost of CRM and HRM was higher.

References

[1]. Khan, R. A. (2023). Meta-Analysis of Cyber Dominance in Modern Warfare: Attacks and Mitigation Strategies. *Turkish Journal of Computer and Mathematics Education (TURCOMAT)*, *14*(03), 1051-1061. Retrieved from https://www.turcomat.org/index.php/turkbilmat/article/view/14288

[2].Ray, R. K., Linkon, A. A., Bhuiyan, M. S., Jewel, R. M., Anjum, N., Ghosh, B. P., ... & Shaima, M. (2024). Transforming Breast Cancer Identification: An In-Depth Examination of Advanced Machine Learning Models Applied to Histopathological Images. *Journal of Computer Science and Technology Studies*, *6*(1), 155-161. https://www.doi.org/10.32996/jcsts.2024.6.1.16

[3] Islam, M., & Shuford , J. . (2024). A Survey of Ethical Considerations in AI: Navigating the Landscape of Bias and Fairness. *Journal of Artificial Intelligence General Science (JAIGS) ISSN:3006-4023*, *1*(1). https://doi.org/10.60087/jaigs.v1i1.27

[4] Akter, most. S. (2024). Interdisciplinary Insights: Integrating Artificial Intelligence with Environmental Science for Sustainable Solutions. *Journal of Artificial Intelligence General Science (JAIGS) ISSN:3006-4023*, *1*(1). https://doi.org/10.60087/jaigs.v1i1.28

[5] khan , M. R. . (2024). Advancements in Deep Learning Architectures: A Comprehensive Review of Current Trends. *Journal of Artificial Intelligence General Science (JAIGS) ISSN:3006-4023*, *1*(1). https://doi.org/10.60087/jaigs.v1i1.29

[6] Rana , M. S. ., & Shuford , J. . (2024). AI in Healthcare: Transforming Patient Care through Predictive Analytics and Decision Support Systems. *Journal of Artificial Intelligence General Science (JAIGS) ISSN:3006-4023*, *1*(1). https://doi.org/10.60087/jaigs.v1i1.30

[7] Mia , M. R. ., & Shuford , J. . (2024). Exploring the Synergy of Artificial Intelligence and Robotics in Industry 4.0 Applications . *Journal of Artificial Intelligence General Science (JAIGS) ISSN:3006-4023*, *1*(1). https://doi.org/10.60087/jaigs.v1i1.31

[8] Klinkenberg , D. . (2024). The Gnostic Code . *Journal of Artificial Intelligence General Science (JAIGS) ISSN:3006-4023*, *1*(1). https://doi.org/10.60087/jaigs.v1i1.32

[9] Carrasco Ramírez., D. J. G. ., Islam, M. ., & Even, A. I. H. . (2024). Machine Learning Applications in Healthcare: Current Trends and Future Prospects. *Journal of Artificial Intelligence General Science (JAIGS) ISSN:3006-4023*, *1*(1). https://doi.org/10.60087/jaigs.v1i1.33

[10] Islam, M. (2024). Applications of MachineLearning(ML): The real situation of the Nigeria Fintech Market. *Journal of Artificial Intelligence General Science (JAIGS) ISSN:3006-4023*, *1*(1). https://doi.org/10.60087/jaigs.v1i1.34

[11] Shuford, J. . (2024). Quantum Computing and Artificial Intelligence: Synergies and Challenges. *Journal of Artificial Intelligence General Science (JAIGS) ISSN:3006-4023*, *1*(1). https://doi.org/10.60087/jaigs.v1i1.35

[12] Shuford, J. (2024). Deep Reinforcement Learning Unleashing the Power of AI in Decision-Making. *Journal of Artificial Intelligence General science (JAIGS) ISSN: 3006-4023*, *1*(1). https://doi.org/10.60087/jaigs.v1i1.36

[13] Islam, M. M. . (2024). The Impact of Transfer Learning on AI Performance Across Domains . *Journal of Artificial Intelligence General Science (JAIGS) ISSN:3006-4023*, *1*(1). https://doi.org/10.60087/jaigs.v1i1.37