Vol.4,Issue 01,May 2024
Journal of Artificial Intelligence General Science JAIGS

Home page http://jaigs.org

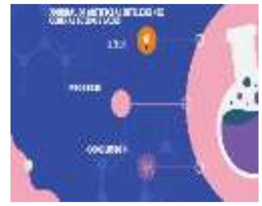

Ethical Considerations in AI Simulations for Designing Assistive Technologies

Evin Miser[1], Orcun Sarioguz[2]

[1]Ankara University

[2]Anadolu University

ARTICLEINFO
Article History:
Received: 01.05.2024
Accepted: 15.05.2024
Online: 20.05.2024
Keyword: intelligent assistive technology, ethical reflection, simulation, dementia care, conceptual approach

ABSTRACT

Current ethical debates on the use of artificial intelligence (AI) in healthcare approach AI technology in three primary ways. First, they assess the risks and potential benefits of current AI-enabled products using ethical checklists. Second, they propose ex ante lists of ethical values relevant to the design and development of assistive technologies. Third, they advocate for incorporating moral reasoning into AI's automation processes. These three perspectives dominate the discourse, as evidenced by a brief literature summary. We propose a fourth approach: viewing AI as a methodological tool to aid ethical reflection. This involves an AI simulation concept informed by three elements: 1) stochastic human behavior models based on behavioral data for simulating realistic scenarios, 2) qualitative empirical data on value statements regarding internal policy, and 3) visualization components to illustrate the impact of variable changes. This approach aims to inform an interdisciplinary field about anticipated ethical challenges or trade-offs in specific settings, prompting a re-evaluation of design and implementation plans. This is particularly useful for applications involving complex values and behaviors or limited communication resources, such as dementia care or care for individuals with cognitive impairments. While simulation does not replace ethical reflection, it allows for detailed, context-sensitive analysis during the design process and before implementation.Finally, we discuss the quantitative analysis methods enabled by stochastic simulations and the potential for these simulations to enhance traditional thought experiments and future-oriented technology assessments.

Introduction

In the science fiction movie Dark Star (1974, directed by John Carpenter), a starship captain argues with an AI-controlled bomb about whether it should detonate. The crux of the dispute is whether the bomb's decision to detonate is based on accurate data—specifically, whether it correctly interpreted a detonation order through its sensory input. The captain's philosophical arguments about the limits of the bomb's self-awareness inadvertently turn the bomb into a nihilist, leading it to detonate and kill the ship's human crew. Such scenes act like thought experiments, a method commonly used in philosophy and science to anticipate implications or test the ethical and epistemic coherence of ideas (Walsh, 2011). In the case of Dark Star, the risk presented is that future AI could develop its own morality, potentially harmful to humans. However, it is uncertain whether the insights from thought experiments apply to real-world AI systems, such as those under development for dementia care (Schweda et al., 2019).

This paper aims to expand the thought-experiment approach by exploring the use of computational simulations as a tool for ethical reflection on human-AI interaction. We propose that algorithmic simulations can enhance ethical reflection by providing empirical, simulated data during the design phase, thereby improving the anticipation of ethical issues in various AI applications.

According to the High Level Expert Group on Artificial Intelligence (2019), AI systems are defined as software systems that analyze their environment and take actions to achieve specific goals independently. This broad definition does not specify the type of algorithms used—whether symbolic rule-based or sub-symbolic neural networks—nor does it define the level of automation (Shneiderman, 2021, p. 48).

Traditionally, ethics in technology development is guided by general principles used in thought experiments to test which principles have desirable consequences or are likely to gain public acceptance. Recently, the use of thought experiments, such as trolley-dilemma scenarios for automated vehicles, has been criticized for being too narrow or abstract (Goodall, 2019; De Freitas et al., 2021). These experiments often lack empirical validity and may be biased by the perspectives of ethicists or technology developers.

Mainstream ethical evaluation approaches for new technologies, such as biotech, nanotech, or AI, often treat technology as an object of ethical reflection—this is termed the "ethics of AI." While this approach is suitable for biosciences or nanotechnology, it need not be the sole method for reflecting on AI. Empirically informed ethical reasoning, which collects qualitative data from practitioners, stakeholders, and affected individuals, has recently become a standard. This approach reduces bias, including expert bias, and enhances the generalizability to real-world situations (Schicktanz et al., 2012; Mertz et al., 2014). Despite its advantages, this method also has epistemic limitations.

In applied ethics, AI is often considered a feature of specific products, either for environmental analysis and goal-oriented actions or for aiding moral decisions. The literature review (see Section 2) shows this approach's dominance.

Due to the complex and sometimes opaque outcomes of AI technology, "explainability" and "trust" have become essential criteria for ethical evaluation (High Level Expert Group on Artificial Intelligence, 2019; Amann et al., 2020; Coeckelbergh, 2020; Markus et al., 2021; Border and Sarder, 2022). However, these criteria assume a human-AI interaction framework where users can understand and trust the AI system, which is not always possible in healthcare and disability settings. Here, users may not be able to monitor, interact with, or comprehend an AI system's outputs, making trust unfeasible.

Our "AI-Assisted Ethics" approach aims to anticipate ethical trade-offs and social implications in complex, contextualized settings where criteria like trust or explainability may not be appropriate. Complexity is particularly significant in scenarios where the direct anticipation of outcomes is limited due to the heterogeneous nature of involved individuals. These diverse interactions with AI can limit the generalizability of empirical observations.

Our article results from interdisciplinary cooperation between ethicists, social scientists, engineers, and machine learning specialists, combining insights from various sub-studies focused on intelligent assistive technologies (IAT) in healthcare, particularly dementia care. The following proposal emerges from this comprehensive exchange. First, we summarize the general discussion on AI ethics, machine ethics, and human-AI interaction ethics in dementia care (Section 2). Given the challenges in assessing AI technology impacts in such cases, we developed a conceptual approach to using AI as a simulation tool to anticipate ethical and social issues in IAT implementation (Section 3). We focus on dementia care (Section 3.1), outlining an in silico simulation concept (Section 3.2) involving multiple agents interacting with different ethical values. Through 'Ethical Compliance Quantification,' we can quantitatively compare design alternatives and inform stakeholder discussions. We present results from a simulation model in technology-assisted dementia care to illustrate this approach, discussing its advantages and challenges. This simulation is informed by value statements from interviews and employs stochastic human behavior models encompassing behavioral data.

AI as a Product of Technology and an Object of Ethical Reflection

In this section, we summarize the main discussions around AI ethics and ethical machines, focusing on the ethics of human-AI interaction in the context of "intelligent assistive technology" (IAT) for the care of older adults and persons with dementia. Various authors have developed catalogs of values and ethical principles to guide these ethical assessments (Currie et al., 2020; Spiekermann, 2016; Umbrello and van de Poel, 2021; van Wynsberghe, 2013; Schicktanz and Schweda, 2021). Key ethical criteria include self-determination, non-maleficence, beneficence, privacy, and sustainability (Hofmann, 2013; Novitzky et al., 2015; Ienca et al., 2018; Vandemeulebroucke et al., 2018). These principles are often prioritized differently, and while sometimes proposed as design guidelines, they are more frequently used to evaluate existing technologies.

This emphasis on post-development assessment might have led Ienca and colleagues to advocate for "a coordinated effort to proactively incorporate ethical considerations early in the design and development of new products" (Ienca et al., 2018, p. 1035). Similarly, McLennan and colleagues, in a recent paper on "embedded ethics," stress the importance of an "ongoing practice of integrating ethics into the entire development process" through a "truly collaborative, interdisciplinary enterprise" (McLennan et al., 2022, p. 3). This approach reflects earlier efforts from the 1990s, when engineers and philosophers developed strategies to integrate ethical issues and values into human-machine interaction design, known as computer ethics, social informatics, participatory design, and value-in-design. Friedemann and Kahn (2007) distinguish three main approaches regarding values and ethical principles in technology development:

1. Embodied Approach: Designers incorporate values directly into the technology.

2. Exogenous Approach: Users determine and impose values after the technology is implemented.

3. Interactional Approach: Values emerge from the interaction between designers and users, including approaches like value-in-design and participatory design.

All three approaches are present in current AI technology design.

Key questions for assessing such technologies include

Whether AI/IAT technologies should be used, if they pose ethically acceptable risks, if the opportunities they present outweigh these risks, and whether their use might conflict with fundamental human rights and ethical principles such as human dignity, self-determination, or justice. In fields like elder care, this assessment often occurs post-prototype development rather than during the design phase.

Conversely, machine ethics focuses on whether AI technologies, which can operate with varying degrees of autonomy, can and should be designed to function in morally acceptable ways. This involves ethical considerations around appropriate concepts and standards and questions about morality itself, such as moral agency and responsibility, as well as technical challenges in implementing these concepts through algorithms and "training" (Anderson and Anderson, 2007). This debate includes:

Top-Down Approaches: These attempt to specify moral precepts deductively by applying general moral norms. Examples include operationalizing fundamental philosophical principles like utilitarianism (maximizing utility) or Kantian ethics (universalizability of maxims) into algorithms that guide the AI's moral behavior. Van Wynsberghe (2013) argues for embedding moral values in care robots during their development.

Bottom-Up Approaches: These focus on developing moral competencies inductively through pertinent moral experiences. An example is MIT's Moral Machine experiment (Awad et al., 2018), which collects large datasets of human responses to online moral dilemmas, thereby teaching AI systems about morality through learned experiences.

Each approach has its own set of challenges and considerations, reflecting the complexity and nuance required in developing morally and ethically sound AI systems.

Methodology: Model Conceptualization

Premises for AI-Assisted Ethics in Supporting IAT Development

Ethical reflection on modern technologies, including AI, involves several critical steps: identifying problems (not just dilemmas), considering relevant facts, applying various ethical approaches and principles to explore alternative conclusions, testing for consistency with established norms, evaluating the applicability of abstract rules to concrete situations, justifying specific decisions (an aspect of explainability), and ensuring the societal legitimacy of the entire reflective process. This process is inherently complex and cannot be reduced to a fixed set of values. Most ethical approaches begin with a priori moral intuitions and theoretical generalizations, such as "values."

When ethical considerations are applied during or before the development of a technology, they often rely on principles that may be too broad for concrete design decisions—a limitation of top-down models. To make ethical issues more relevant to specific design decisions, analogies from previous situations are often extrapolated to new contexts. This extrapolation is prone to error, and not all errors are apparent before product implementation. Bridging the gap between overly general and overly specific (but extrapolated) recommendations for ethical design is crucial, especially when dealing with vulnerable populations like persons with dementia.

Adapting IAT systems to complex settings—characterized by multiple agents with different goals, varying moral intuitions, and different cognitive states and communication skills—during the design phase requires a different approach. Ethical reflection should accompany the design process, enabling experiments with different designs to identify practical moral problems and potential value conflicts. However, in situ experimentation can be unethical for vulnerable people, such as those with dementia, and it can prompt fear and anger, as seen with the COACH prompting system for handwashing (Mihailidis et al., 2018). Moreover, a review by Alkadri and Jutai (2016) found that many technologies for this target group are weak regarding safety and efficacy.

In our field of study, technology-assisted dementia care, another significant challenge is communication between humans and AI. Machines "explaining" their decision criteria to humans is not feasible for persons with dementia, who have varying and limited capacities for effective communication. This group also cannot provide detailed feedback to designers or scholars, limiting the effectiveness of participatory design approaches.

To address these challenges, we propose rethinking AI as an integral tool in the ethical design process, not just a product of technology. We suggest using in silico simulation, a computational simulation of the technology in its environment, as a proxy for in situ experimentation. Ideally, these simulations should include multiple human agents, represent their goals, model their internal decision-making processes (from deterministic to stochastic models), and incorporate their environment, including the device or procedure under development. Simulations can be run repeatedly at little cost and without exploiting vulnerable populations. They allow for assessing the effects of a product on agents in various settings while varying inputs, enabling reflection on the model-building process (Chandrasekharan et al., 2013).

Unlike in situ experiments or thought experiments, which focus on outcomes concerning ethical acceptability, inefficiency, and safety, in silico simulations provide a flexible and ethical way to explore the complex interactions between AI and human agents in a controlled, replicable environment.

Concept for an AI-Assisted Simulation

Our AI-assisted ethics simulation (Figure 1) consists of several key elements, detailed below.

Multiple Agents

A multi-agent simulation environment creates a virtual world where simulated agents interact. In this context, "simulation" means representing the state of the world and its agents through a set of variables in a programming language. These variables (referred to as the variable "score") reflect the state of the simulated world at any given time. For example, a variable called "location" might contain two scores indicating an agent's position in a two-dimensional simulation world.

The simulation progresses in steps, with each step manipulating the variables according to rules that define the temporal evolution of this simulated world. For instance, a "move" rule might dictate that an agent with a "destination" variable different from its "location" variable will update its location by a "step length" towards the "destination." These rules are implemented as pieces of program code.

Agent interactions are modeled by rules that depend on and modify the variables representing the state of two or more agents. Simulation environments typically allow for stochastic rules, where outcomes depend on a random process. For example, the step length in a "move" rule might be determined by sampling from a normal distribution defined by a mean step length and a certain standard deviation. Stochastic rules are particularly important in human-agent simulations to capture non-deterministic behavior.

A simulation run involves initializing the state variables (e.g., the location and current disorientation state of a simulated patient, locations of simulated caregivers) with predefined scores (such as location coordinates) and then advancing the simulation step-by-step until a specified termination condition is met (such as reaching a certain time point or simulation state). When using stochastic rules, different runs of the simulation may result from the same initial conditions. By conducting multiple runs, it is possible to analyze the statistical properties and temporal development of state variables in the simulation by aggregating data from the run protocols.

Ethical Compliance Quantification (ECQ)

The objective of the ECQ evaluation is to provide a quantitative measure of how well an IAT adheres to a value model, based on a set of run protocols generated by the simulation. From an ethicist’s perspective, the value model is crucial because it translates the sequence of events in the simulation runs, particularly the actions of the IAT in specific situations, into an ethical assessment.

To illustrate how to define such a value model, consider the care facility floor planning example discussed earlier. The first step is to identify the relevant values or value dimensions. Since we will assign scores to these values, a set of scores—one for each value—defines a point in a multi-dimensional space where each dimension corresponds to a value. A simple ethical value system might include "efficiency" and "fairness."

The next step is to create formulas that calculate quantitative scores for the value dimensions of "efficiency" and "fairness" based on data from a simulation run. In our hypothetical scenario, where stakeholders move between locations, efficiency might be calculated as the ratio of straight-line distance to distance traveled, while fairness might be represented by the ratio of efficiency scores between different stakeholders (with the value "1" indicating optimal fairness when all stakeholders experience equal efficiency). This ECQ setup can then be used to compare different floor plans based on their ratings across various value dimensions.

This simple thought experiment highlights the core challenge of defining a value model: ensuring it accurately reflects how values relate to the real world. For example, the simplistic definition of fairness used here might prompt questions about its adequacy, as it ignores factors like the physical fitness of stakeholders, which could affect their ability to walk longer distances. Stakeholders might argue for a correction factor in the fairness computation to account for physical fitness.

This example demonstrates the multilateral nature of the value-sensitive design process required for defining a value model, as it makes value judgments explicit and reveals the flexibility available in mapping event sequences to value ratings. Simulation-based ECQ also enables the assignment of numbers to qualitative value statements by, for instance, counting how often a certain qualitative requirement is met or violated across multiple simulation runs.

The ECQ concept offers something impossible in the real world: the ability to evaluate the quantification across different design alternatives for all involved models. By varying the IAT policy, one can assess the impact of different design alternatives on value compliance (policy feedback), potentially leading to an optimal IAT policy. Adjusting the value model allows for evaluating the plausibility of the resulting value quantification and the value model itself (value operationalization feedback). Finally, varying the stakeholder and IAT models enables the assessment of the ECQ results' sensitivity to the ecological validity of the simulation model.

An Example

SimDem (Shaukat et al., 2021) is a simulation system we developed to analyze a smart-watch-based IAT in a nursing home for dementia patients. The smart watch assists residents by detecting deviations from their intended routes and providing directional prompts to help them reach their destination. A fundamental design question is determining the number of guidance interventions that should trigger the system to assume the wearer is permanently disoriented, prompting an alert to a caregiver. Based on previous expert interviews, our reasoning, and literature research, we identified values such as "safety" and "fairness" for assessment.

The first step in performing ECQ is creating the simulation model. In this case, it is a 2D virtual nursing home based on the floor plan of a real nursing home. The wayfinding behavior for the patient simulation (see Figure 2) has been assessed through observations of real subjects.

The simulated IAT—the smart watch—has a certain probability of detecting disorientation and a certain probability that an intervention will help the wearer regain orientation. Both probabilities are design parameters of the simulation model. With this setup, multiple simulation runs can be performed to analyze the quantitative effect of different assistive strategies on the values of interest. Figure 3 shows aggregated results from 1809 runs using different value models, illustrating the importance of value model definition.

First, we operationalized "safety." There are multiple ways to do this, such as considering the relative amount of time a patient is disoriented as "unsafe" time. This approach produced the reddish-colored box plots in Figure 3, labeled "Safety (Original)." However, this operationalization proved implausible. The plot showed that the strategy of immediately calling a nurse (Nhelp = 0) resulted in higher non-compliance (i.e., longer time in an unsafe state) than the strategy of waiting for five failed interventions (Nhelp = 5). Clearly, the more failed interventions we wait for, the longer the disoriented patient wanders unguided. This mismatch with common sense indicates a flawed operationalization.

The ECQ approach allows the discovery of such mistakes in value operationalization by visualizing the value scores across different strategies. In this example, once visualized, it became clear that the value operationalization was incorrect: a nurse accompanying a patient should be considered safe, regardless of the patient's disorientation state.

Providing a more plausible value operationalization is straightforward: we should only consider the time a patient is disoriented without nurse guidance as "unsafe" time. Using this refined operationalization of "safety," we now see that immediately calling a nurse is indeed safer than waiting for multiple interventions (see Figure 3, purple box plot labeled "Safety (Refined)"). Additionally, this plot reveals that having no IAT at all ("Nurse Only") is the least safe strategy, aside from leaving the patient completely unattended ("No Help"). Without a smart watch to detect disorientation, nurses must actively discover disoriented patients, highlighting that the smart watch always enhances safety in this simulation setting.

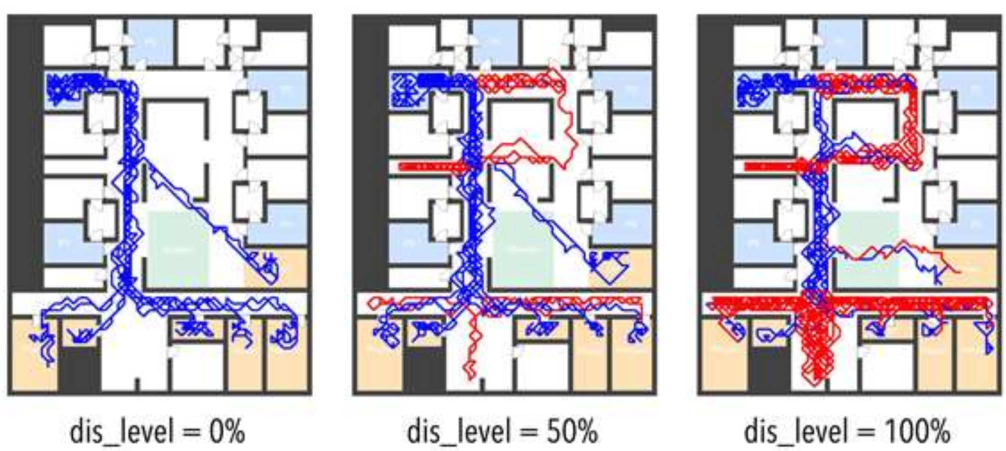

FIGURE 2
Visualization of Trajectories for Simulated Patient Agents with Various Levels of Disorientation Probability (dis_level). Blue = patient state oriented; Red = patient state disoriented.

Discussion

Simulation-based ECQ is a method for exploring the ethics design space, developing "ethics awareness" in designers, and informing ethicists about the outcomes of different scenarios and how various variables influence the process. It allows for the simulation or anticipation of complex ethical trade-offs, not just hypothetically or generally (as in thought experiments), but as visualized trade-offs concerning human-AI and human-human interactions that those involved cannot easily explain or rationalize. In this discussion, we will focus on three main challenges:

1. How qualitative values can (or must be) operationalized for computational simulations and what this requires.

2. In which contexts and for what purposes the advantages of AI-assisted simulations outweigh their disadvantages and limitations.

3. Why AI-assisted ethics simulations can be compared to thought experiments but offer innovative epistemic dimensions for ethical reasoning.

First, the methodology discussed in this paper does not focus on a specific value—such as autonomy or safety—but on improving the process of ethical reflection for IAT development by considering a diversity of values. Aliman and Kester (2019) argue for a consequentialist approach that predicts the overall utility of a future outcome for a given population. While their argument supports AI-assisted simulation concerning utility, our methodology proposes a strategy to understand the impact of an IAT on various moral values that can be operationalized. In this sense, our proposal is agnostic regarding the specific values considered during design, but it is not agnostic regarding the requirement for a participatory and pluralistic approach. Our methodology acknowledges the central challenge of value operationalization. Unless a value is operationalized, it cannot be analyzed by ECQ. While this may seem a drawback, we view it as an advantage. ECQ challenges value experts to operationalize their value concepts, providing the opportunity to utilize such operationalization. A claim that a value cannot be quantified can now be contested by providing an operationalization, challenging the opponent to show where it violates the value system. One core benefit of using ECQ is to expose situations where an operationalization cannot be found or agreed upon. By forcing stakeholders to define their value concepts explicitly, ECQ reveals conflicts that are independent of "machine ethics" but arise from our own inconsistent or ambiguous opinions on moral behavior. Behavioral economics has shown that individuals may make contradictory assessments of situations depending on whether they are experienced or remembered (Aliman and Kester, 2019), highlighting this problem.

Conclusion and Outlook

AI-assisted simulations offer a promising avenue for addressing the limitations of current empirically informed ethical reasoning, as well as traditional approaches like thought experiments and forecasting methods. They enable the exploration of numerous complex scenarios with flexibility and provide objective observations that can be visualized and analyzed systematically. Visualization is especially relevant as it highlights trade-offs and observations.

Our contribution focuses on how empirical data about stakeholders' value preferences and potential behaviors can inform a "supertool" to simulate various ethically relevant parameters and outcomes. Ethically motivated empirical research and AI-assisted simulation strategies complement each other, offering a new methodology for empirically informed ethical reflection—AI-assisted ethics. However, interdisciplinary collaboration and learning are essential to effectively combine theoretical and methodological perspectives.

While our focus has been on AI technology, AI-assisted ethics is not limited to this domain. It can be applied to various fields, leveraging its potential for technology development, particularly in areas with low practical implementation and empirical experience. In the future, AI model simulations could incorporate the inner states and behaviors of simulated agents, enhancing the understanding of interventions and their effects.

References List:

[1]. Prakash, S., Malaiyappan, J. N. A., Thirunavukkarasu, K., & Devan, M. (2024). Achieving Regulatory Compliance in Cloud Computing through ML. AIJMR-Advanced International Journal of Multidisciplinary Research, 2(2).

[2]. Malaiyappan, J. N. A., Prakash, S., Bayani, S. V., & Devan, M. (2024). Enhancing Cloud Compliance: A Machine Learning Approach. AIJMR-Advanced International Journal of Multidisciplinary Research, 2(2).

[3]. Devan, M., Prakash, S., & Jangoan, S. (2023). Predictive Maintenance in Banking: Leveraging AI for Real-Time Data Analytics. Journal of Knowledge Learning and Science Technology ISSN: 2959-6386 (online), 2(2), 483-490.

[4]. Eswaran, P. K., Prakash, S., Ferguson, D. D., & Naasz, K. (2003). Leveraging Ip For Business Success. International Journal of Information Technology & Decision Making, 2(04), 641-650.

[5]. Prakash, S., Malaiyappan, J. N. A., Thirunavukkarasu, K., & Devan, M. (2024). Achieving Regulatory Compliance in Cloud Computing through ML. AIJMR-Advanced International Journal of Multidisciplinary Research, 2(2).

[6]. Malaiyappan, J. N. A., Prakash, S., Bayani, S. V., & Devan, M. (2024). Enhancing Cloud Compliance: A Machine Learning Approach. AIJMR-Advanced International Journal of Multidisciplinary Research, 2(2).

[7]. Biswas, A. (2019). Media Insights Engine for Advanced Media Analysis: A Case Study of a Computer Vision Innovation for Pet Health Diagnosis. International Journal of Applied Health Care Analytics, 4(8), 1-10.

[8] Chopra, B., & Raja, V. (2024). Toward Enhanced Privacy in Digital Marketing: An Integrated Approach to User Modeling Utilizing Deep Learning on a Data Monetization Platform. Journal of Artificial Intelligence General science (JAIGS) ISSN: 3006-4023, 1(1), 91-105.

[9]. Raja, V. (2024). Fostering Privacy in Collaborative Data Sharing via Auto-encoder Latent Space Embedding. Journal of Artificial Intelligence General science (JAIGS) ISSN: 3006-4023, 4(1), 152-162.

[10]. Raja, V. ., & chopra, B. . (2024). Exploring Challenges and Solutions in Cloud Computing: A Review of Data Security and Privacy Concerns. Journal of Artificial Intelligence General Science (JAIGS) ISSN:3006-4023, 4(1), 121–144. https://doi.org/10.60087/jaigs.v4i1.86

[11]. SARIOGUZ, O., & MISER, E. (2024). Data-Driven Decision-Making: Revolutionizing Management in the Information Era. Journal of Artificial Intelligence General science (JAIGS) ISSN: 3006-4023, 4(1), 179-194. DOI: https://doi.org/10.60087/jaigs.v4i1.131

[12]. Raja, V. (2024). Exploring Challenges and Solutions in Cloud Computing: A Review of Data Security and Privacy Concerns. Journal of Artificial Intelligence General science (JAIGS) ISSN: 3006-4023, 4(1), 121-144.

[13]. Nadella, G. S. (2023). Validating the Overall Impact of IS on Educators in US High Schools Using IS-Impact Model–A Quantitative PLS-SEM Approach (Doctoral dissertation, ProQuest University (Demo)).

[14]. Nadella, G. S., Gonaygunta, H., Kumar, D., & Pawar, P. P. (2024). Exploring the impact of AI-driven solutions on cybersecurity adoption in small and medium enterprises. World Journal of Advanced Research and Reviews, 22(1), 1190-1197.

[15]. Nadella, G. S., Meduri, K., Gonaygunta, H., Addula, S. R., Satish, S., Harish, M., & Maturi, S. K. S. P. (2024). Advancing Edge Computing with Federated Deep Learning: Strategies and Challenges. International Journal for Research in Applied Science and Engineering Technology, 12(4), 3422-3434.