Applications of The Active Inference and The Free-Energy Principle Frameworks for Mimicking Social Human Behaviours on Intelligent Agents

a Systematic Literature Review

Lara Sakarya[1]
Supervisors: Chirag Raman[1]**, Ojas Shirekar**[1]
[1]EEMCS, Delft University of Technology, The Netherlands

A Thesis Submitted to EEMCS Faculty Delft University of Technology,
In Partial Fulfilment of the Requirements
For the Bachelor of Computer Science and Engineering
June 23, 2024

Name of the student: Lara Sakarya
Final project course: CSE3000 Research Project
Thesis committee: Chirag Raman, Ojas Shirekar, Willem-Paul Brinkman

Abstract

Active inference is a theory of the human brain characterising behaviour that minimises surprise. The free energy principle accounts for the adaptive behaviours of organisms through action, perception, and learning aimed at optimising reward or surprise. This study systematically reviews relevant literature to address their methodologies, relevance to mimicking social human behaviour, challenges, and limitations to guide future research by succinctly reporting previous findings and research gaps. Active inference models are extended with deep active inference, free-energy models, multimodal deep belief networks, predictive coding, and probabilistic programming. These models employ goal-directed, epistemic, reward-seeking, and decision-making behaviours and simulate cumulative culture. However, some of these models do not translate well to complex real-life applications due to their simplicity, computational demands, or the assumptions upon which they are based. Challenges with real-life applications include difficulty scaling to high-dimensional data and model simplicity. Furthermore, some experiments did not have enough data to validate or train their models.

1 Introduction

Nowadays, many programs on pattern recognition are labelled as artificial intelligence (AI), but they lack intelligence. This research paper systematically reviews the literature on the active inference framework (AIF) and the free energy principle (FEP) to gather current practices on embodied virtual agents for mimicking social human behaviour to identify the research gaps. The aim is to propose future improvements in this relatively new field to guide subsequent research.

AIF is the theory of human perception, planning, and action based on the probabilistic inference that characterises Bayes-optimal behaviour by minimising surprise in an agent's sensory observations [1]. Surprise is the model evidence's negative log representing the discrepancy between an agent's preferred sensory input and actual observations [2]. Perception constructs beliefs on probabilistic inference, and control produces the decision to minimise surprise [2]. FEP is a principle of AIF accounting for action, perception, and learning by optimising the expected reward or surprise [3] for self-organising biological agents withstanding inclination to disorder [4]. Agents resist this tendency to disorder by occupying a limited number of states and minimising their free energy through their actions. FEP provides a mathematical background for maintaining an agent's equilibrium [4] by constraining them to a limited number of states [2].

Levchuk et al. highlight that the infeasibility of distinguishing all hidden states hinders direct optimisation of surprise and model evidence [2]. Variational Free Energy (VFE) introduces an upper bound on surprise, an information-theoretic function of outcomes, their hidden cause, and the agent's internal state. A generative model depicts the internal state using a joint probability distribution over all hidden states and sensory observations that minimises VFE to infer the most likely hidden states from sensory inputs [5]. Figure 1 depicts the active inference process of minimising VFE through sensory input accumulation [6]. Predictions are created by evaluating each policy's free energy, enabling action selection for the predicted state.

AIF depicts a non-equilibrium steady-state (NESS) system that self-organises to return to a steady state after disruption [5]. Some of the current models of AIF enable agents to select policies that minimise expected free energy (EFE) [7], which functions based on prior beliefs about future outcomes and supports epistemic and exploitative behaviour [6]. Un-

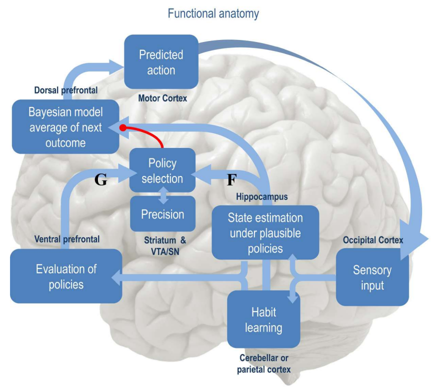

Figure 1: Belief Updating[6]

derstanding AIF relies on unveiling the adaptive behaviours of living organisms to minimise their free energy and the reason they adapt [1].

Currently, most intelligent agents can perform some but not all of the following tasks well: planning, acting on, and searching across the environment. Agents should be modelled to perform all of these tasks successfully by acquiring adaptive capabilities and fitting the requirements of their changing environment. To bridge this research gap, generalised AI systems (AI with AIF and learning) generate causes of their actions known as 'planning as inference' as opposed to generative AI systems that are restricted to content generation [8]. Intelligence can emerge when generalised AI experiences outcomes of its actions where it obtains information on the world's hidden states.

AIF agents are built with generative models to include the outcomes of their actions and intelligence of planning abilities, that is, the probabilistic depiction of sentient behaviour [1]. This model is significant for depicting human cognition since the fundamental function of the brain is controlling exchanges with the world [8]. Agents use this model to make goal-directed decisions and plan using predictive inference to learn about world states, optimise Bayesian model evidence [9], and depict sensory observations to minimise long-term surprise [10]. AIF can explain the decision and expected utility theory under embodied cognition to understand Bayesian optimal behaviour and the derivation of a sense of agency (SoA) [11].

Previous research integrates AIF with machine learning (ML) methods to endow human social intelligence and cognition. Some of these applications include deep active inference models that infer valence states [12], collaboration in heterogeneous agent teams [13], collective intelligence [14], and trust in human-computer interactions [15]. Since AIF is a theory of the human brain and behaviour, applying this framework allows agents to form adaptive and social skills that mimic human cognitive processes.

By surveying the literature, this research paper answers the question: How have active inference and the free-energy principle been applied to embodied agents and the mimicking of social human behaviours? In the context of this paper, physical robots and simulated/synthetic virtual intelligent agents are considered intelligent embodied virtual agents. The research question is split into three sub-questions to answer all the aspects of the research topic in a structured way:

- **RQ 1.** What methods and models have been used to apply the active inference and the free-energy principle to embodied agents?

- **RQ 2.** How do the active inference and the free-energy principle relate to social human behaviours?
- **RQ 3.** What are the challenges and limitations of active inference and the free-energy principle applications on intelligent agents?

2 Methodology

A systematic literature review was conducted to categorically synthesise information regarding the current practices and applications of AIF and FEP on embodied intelligent agents. This review includes a search strategy for the papers considered for the study, eligibility criteria, and synthesis of these papers to answer the sub-questions [16].

PRISMA, a 27-item checklist designed to report systematic reviews transparently, was followed for the reproducibility of this research [17]. The PRISMA flow diagram depicts the phases of the systematic review: identifying papers from databases with the designed search query, the three-step screening of these papers, and the papers included in the review after the filtering process. The book "Doing a Systematic Review: A Student's Guide" was used to structure the research paper and conduct the systematic review [18].

2.1 Search process

Cell Press, Scopus, IEEE Xplore, Web of Science, and ACM Digital Library databases were selected for the search process. Cell Press includes over 50 scientific journals with papers on life, physical, earth, and health sciences[19]. This database was used to find background information on AIF and FEP to understand the emergence of these frameworks from a biological perspective. Scopus is a comprehensive, multidisciplinary, and reliable database queried for obtaining papers on different application areas of these frameworks [20]. IEEE Xplore contains the world's most cited publications in electrical engineering, computer science, and other sciences [21]. It was chosen due to its high prestige and reliability in researching the computational aspect of AIF and FEP. Web of Science provides multidisciplinary coverage with daily updates in the biomedical, life sciences, engineering, and physics fields relevant to the background of AIF and FEP [22]. The ACM Digital Library includes journal and conference papers in scientific fields such as AI, ML, natural language processing, and human-computer interaction, which are highly relevant to this research topic [23]. Five databases were chosen to collect diverse research papers from different science fields.

The search string in Figure 2 was formulated with significant keywords and phrases extracted from the background and scoping search. This search was performed on 09/05/2024 on the selected databases to identify papers to be assessed on the eligibility criteria.

```
("active inference" OR "free energy principle") AND (cogni* OR behav* OR
intelligence) AND (optim* OR "Bayesian filtering" OR "Bayesian learning"
OR "Bayesian brain" OR "Bayesian inference" OR "homeostatic state" OR
equilibrium OR "sensory observ*" OR "variational free energy" OR surprise)
AND (agent* OR multi-agent)
```

Figure 2: Search String for Cell Press, Scopus, IEEE Xplore, Web of Science, and ACM Digital Library

2.2 Eligibility Criteria

The inclusion and exclusion criteria indicated below were used in screening papers collected through identification.

Inclusion Criteria:

- Journal articles, conference papers, and books written in English.
- Research papers focused on the computational theory of AIF and FEP.
- Research papers discussing the applications of AIF and FEP on intelligent agents.

Exclusion Criteria:

- Research papers that are not journal articles, conference papers, or books written in English.
- Research papers focused on the biological, mathematical, or physical background of AIF and FEP.

2.3 Selection Process

The selection process involved three steps: identification, screening, and inclusion. During the identification phase, the research papers were found using the search string in section 2.1. In the screening process, papers were filtered according to the criteria in section 2.2 by reading excerpts from them.

The visualisation of this process is depicted by the PRISMA flow diagram in figure 3:

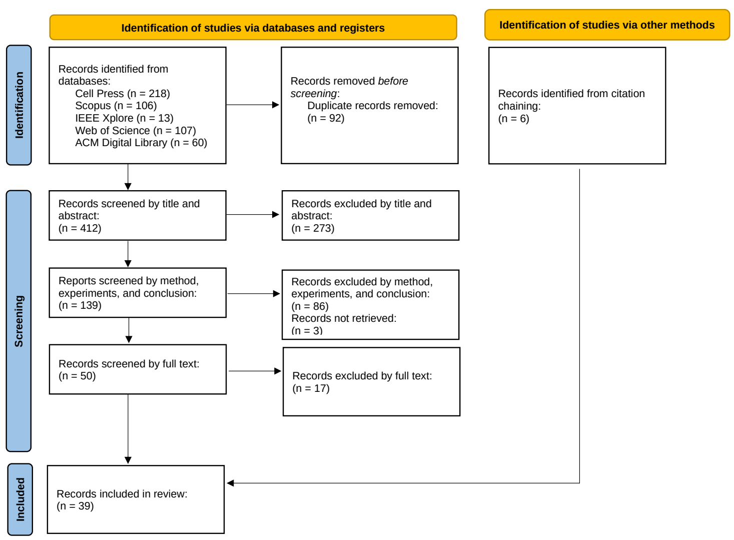

Figure 3: PRISMA flow diagram

3 Results

This section analyses and synthesises the surveyed literature to answer the sub-questions from section 1. Section 3.1 reviews the methodologies and experiments of the AIF and FEP applications. Section 3.2 connects the experiments from section 3.1 to human cognition, human social behaviour and social intelligence. Finally, section 3.3 evaluates the challenges and limitations of these applications.

3.1 RQ 1 - Models and Applications

Table 1: Summary of AIF Applications

Experiments and Simulations	References
Maze Simulations	[24], [6]
Modelling goal-directed behaviour	[25], [26], [27], [28], [29]
Modelling exploitative and exploratory behaviour	[30], [31], [32], [33], [7], [34], [35], [12]
Applications of AIF in social interactions	[36], [14], [37], [38], [39], [15], [40]
Modelling decision-making behaviour	[2], [13], [41], [10], [42]
AIF applications on robots	[9], [43], [6], [44], [45]

Maze Simulations: Chen et al. model synthetic agents in partially observable Markov Decision Processes (MDP) to test environmental volatility's effect on habit formation [24]. The embodied virtual agent seeks rewards by navigating a maze. Static environments lead to a specialist strategy where agents learn the path to previously detected rewards. On the contrary, strong habit formation is not observed in volatile environments with multiple reward locations since a generalist strategy does not adopt strong preferences. Specialised agents outperform generalists due to prioritising their decisions, while generalist agents perform slightly better than naive agents. Another maze simulation focuses on distinguishing habitual and goal-directed behaviour on the metrics of belief-based and belief-free schemes [6]. Agents avoid low entropy states lacking information and high entropy states that are ambiguous to minimise EFE.

Modelling goal-directed behaviour: AIF models of goal-directed behaviour are automated with a Probabilistic Programming (PP) toolbox to build effective, flexible, and scalable agents [25]. This dynamic approach improves on generalisable PP toolboxes with high computational cost. Bayesian thermostat and mountain car problems depicted in figure 4 show how the model leads to goal-directed behaviour. In the former experiment, the agent relocates to settle at a desired temperature in the gradient field. These experiments illustrate the scalability of AIF with automated inference methods on complex applications. Another study shows the scalability of a Deep Active Inference (DAI) approach combining FEP and AIF with deep generative models and evolution strategies using the mountain car problem [26]. The agent solves this problem while learning its environment's generative model and scales to complex real-world environments with the integration of large-scale ML. A novel approach uses Bayesian Target Modelling for Active Inference (BATMAN) to elicit goal-directed behaviour with a coupled generative model to learn future observations and present AIF as a joint framework with classical methods [27]. A cart parking simulation shows that the agent successfully positions the cart using performance feedback, validating the BATMAN approach by predicting the goal state. This novel approach for goal-directed behaviour can scale AIF agents to unpredictable environments.

Another study builds on their prior goal-directed planning model by enabling agents to observe others to infer their goals and predict future actions [28]. This T-GLean model generates a goal-directed plan in real time using stochastic gradient descent to optimise EFE.

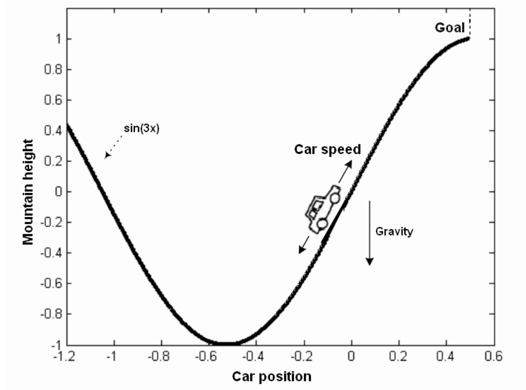

Figure 4: Mountain car problem [46]

A simulated robot navigates obstacles, and a physical robot observes and manipulates objects. Due to the real-world constraints on the physical robot, the simulated agent performs slightly better in goal-directed planning, task execution, and goal inference. A foraging geocaching task illustrates goal-directed behaviour with a generative model of epistemic (explorative) and reward-seeking (exploitative) behaviours [29]. This behaviour emerges in agents that minimise uncertainty when conducting spatial exploration. Agents explore their environment to forage for hidden objects by balancing exploration resulting from curiosity and exploitation due to their predictions.

Modelling exploitative and exploratory behaviour: Nozari et al. model self-information and exploratory behaviours with a hierarchical generative model that uses AIF and imitation learning to employ an explainable decision-making process for autonomous vehicles [30]. This hybrid approach allows agents to learn from experts and balance exploratory and exploitative behaviour according to their predictions. A lane-changing driving scenario shows this model outperforms traditional Reinforcement Learning (RL) methods. In a separate approach, a Generative Adversarial Network (GAN) builds a DAI agent to encapsulate world dynamics effectively, act on the environment, and plan actions [31]. The simulation shows that the agent adapts to volatile and non-trivial environments by balancing exploration and exploitation with policy management.

Similarly, another DAI model with Monte-Carlo (MC) sampling offers applicability in complex environments [32]. The goal is to support intelligence by scaling DAI agents for high-level tasks. In the toy environment, agents showcase epistemic exploration followed by reward-seeking when the agent is confident about its environment. In the Animal-AI environment, agents construct complex plans by avoiding obstacles. The AIF agent performs remarkably better than DQN and A2C, given the same number of training iterations, and its performance is comparable to PPO2. This comparison highlights the possibility of applying AIF to complex environments.

World and Predictive Coding (PC) models are employed for autonomous robots to actively explore their physical and social environment, obtain knowledge, and continually improve their skills [33]. AIF and FEP are significant in formulating this model for self-organising adaptive systems. This application may lead to collaboration with autonomous robots that evolve in the real world and embody artificial general intelligence (AGI).

A sophisticated AIF model also endows agents with explorative behaviour by integrating a recursive form of EFE for deep tree search on future actions and outcomes where the agent has beliefs about its beliefs [7]. Simulations show that a sophisticated agent pursues epistemic behaviour longer than the unsophisticated agent, seeking long-term goals while considering short-term challenges. This epistemic behaviour is also depicted in a reading task during foraging using a deep temporal AIF model [34]. The hierarchical generative model allows the outcomes of a level to generate hidden states at a lower level and enables agents to read and react to local and global violations.

Friston et al. model curiosity and insight with AIF through the emergence of epistemic behaviour in agents deducing their environment's statistical structure [35]. Agents' curiosity is driven by tasks involving exploring their environment to minimise uncertainty. Exploitative and explorative behaviours are also connected to emotional valence inference central to adaptive behaviour [12]. DAI allows agents to infer emotional valence through the predicted precision of their actions. Optimising the valence representation updates the Affective Charge (AC). Simulations of affective inference show the synthetic rat's adaptive behaviours in different contexts. An agent's confidence in their actions portrays a positively (exploitative behaviour) valenced state; contrarily, a lack of confidence in expected precision depicts a negatively valenced state (exploratory behaviour).

Applications of AIF in social interactions: AIF and PC are applied to social interactions to understand the role of an SoA on an agent's intention in action and outcomes [36]. In agent interactions, an agent's undesired actions may influence another agent, leading to an unsustained SoA. Effective collaboration requires minimising possible conflicts among agents with arbitrated SoAs. Agents with strong SoA are less inclined to change their intention and more likely to change the environmental state. With limited regulation, agents act more egocentrically, whereas more regulation leads agents to adapt their internal state to humans. The Predictive-coding-inspired Variational Recurrent Neural Network (PV-RNN) is advantageous over traditional deterministic RNNs since it evaluates the prediction predictability, which enriches agent interactions with others and the world.

The emergence of collective intelligence, a multi-agent system that outperforms the sum of agents' capabilities, is simulated by AIF in interactions [14]. Two agents are put in a simulated environment to sense and reach a chemical concentration (food source). The results illustrate the influence of one agent's behaviour and beliefs on the other. The collective system's ability to minimise VFE indicates social cognitive capabilities. The simplicity of the AIF agent allows for computational feasibility. In another study, Horii and Nagai propose an energy-based multimodal with active perception to enable a robot with limited resources to estimate human emotional states during social interactions using the most informative methods [37]. This approach uses multimodal signals to map human expressions to emotional states based on minimising EFE, outperforming other active perception methods.

A testable DAI multi-agent model employs agents to mimic social behaviour and simulate cumulative culture portrayed in a bi-directional interaction [38]. Minimisation of uncertainty and belief updates build cumulative culture. Agents infer each other's belief states during social interactions with generalised synchrony to converge to a state. Another research conducts a comparative analysis of neuroscience predictive models to understand how AIF connects to social cognition [39]. A new theoretical approach incorporates enactivism with predictive engagement to understand the human brain predictions' in social contexts.

Social interactions are also simulated to build a trust model that explains the sensory exchange between agents and attributes components of trust such as competence and benevolence to agents [15]. The model employs user feedback and shared behaviour and emerges trust from the need to minimise uncertainty in social interactions. Agents trust a system if they can act on that environment and reach predictable results over time. AIF accounts for affective valence, which is significant for trust mechanisms since positive feelings support trust and suggest the predictability of other agents' actions.

Hartwig and Peters model cooperative behaviours and social rules by examining social decision-making to understand how agents decide to cooperate [40]. Surprise minimisation extends classically expected utility maximisation and outperforms it due to its generality in predicting observations in various areas and maintaining the agent's exploratory behaviour. Future research can analyse more complex environments for agent interactions.

Modelling decision-making behaviour: A collection of decision-making problems assesses a team's adaptability according to its convergence to an optimal solution [2]. Alteration of the reward function's parameters simulates these problems to test the teams' capability to change, adapt, and recover. Agents acquire capabilities to function effectively in uncertain, dynamic, and complex environments. This free-energy model improves marginally over the Distributed Discrete Decision Making (D3M) heuristics but, conversely to D3M, converges much faster and maintains high convergence with the increasing objective function complexity.

In another multi-agent study comparing their model to D3M, AIF employs efficient collaboration between heterogeneous agents (human and intelligent agents) to help the decision-making process [13]. An adaptive self-organising team study explores human-machine systems' adaptation to complex scenarios. A Boltzmann distribution defines a generative model for decision-making and task-executing teams. Decision-making teams improve their organisational structure by incrementally updating their decisions' predicted probability distribution. The task execution of these teams is evaluated based on decision-making and work quality. This model converges much faster than D3M heuristics and preserves this convergence with higher objective function complexity.

Another decision-making model uses chance-constraint to employ AIF to observe the trade-off between chance-constrained violation and robustness [41]. In the experiment, drone agents try to reach a certain height in the stochastic vertical wind and showcase chance-constrained behaviour that is utilised under uncertainty. While the goal-driven agent continually interferes with corrections, the chance-driven agent avoids unnecessary interventions, reducing the cost of control.

AIF endows an agent with decision-making behaviour and homeostasis to maintain the system's internal environment [10]. The AIF agent follows a continuous action-perception cycle by updating their generative model to ensure Quality of Service (QoS) and model precision. The distributed AIF agent's high throughput in various environments showcases its decision-making and adaptive behaviour.

Metacognitive capabilities are also integrated into robot decision-making and control for brain-inspired robot controllers [42]. This application may enhance agents with self-assessment capabilities for complex cognitive tasks. A spring damper system simulation shows that the controller accomplishes the goals of given tasks and finds balance in performance, control, and confidence.

AIF applications on robots: A Multimodal Variational Autoencoder Active Inference (MVAE-AIF) models a robot arm's torque controller to scale for multimodal integration on high-dimensional inputs [9]. This controller embodies AIF with its robustness and adaptability to outperform state-of-the-art torque AIF baseline and has higher accuracy with sensory noise. Another study applies AIF to navigate a PR2 robot's 7-DoF (degrees of freedom) arm to reach a target using visual and proprioceptive sensors [44]. The robot's action control performance is accurate with correct feedback and declines marginally in the presence of sensory noise.

The iCub robot also uses visual and proprioceptive inputs with an AIF body perception and control model [45]. The robot infers its body state by minimising discrepancies between sensory predictions and observations. The reaching task simulation results indicate that the model integrates multiple sensory sources without increasing the computational demand. The robot outperforms inverse kinematics with its adaptive behaviour and accuracy.

Taniguchi et al. integrate a Multimodal Hierarchical Dirichlet Process (MHDP) to allow a robot to construct object categories by maximising information gain through action outcomes and observing multimodal information (visual, auditory, haptic) [43]. Maximising information gain is equated to minimising the expected Kullback-Leibler divergence, highly relevant to Bayesian surprise [6]. An upper-torso humanoid robot and synthetic data experiment indicate that the robot identifies object categories swiftly and accurately.

An Iterative policy selection and preference-formation approach is applied to personality formation [47]. The autonomous adaptive behaviour of FEP-AI (AIF) agents provides a way to attain Artificial Super-intelligence (ASI) with better predictive models on world states than humans. Future work will focus on simulations for prosocial personalities for homeostasis within and between individuals.

3.2 RQ 2 - Mimicking Social Human Behaviour

This section explains how AIF and FEP allow agents to adapt by controlling sensory inputs, changing future sensory predictions, and changing their internal models depicting their environment and relationships [2]. AIF simulates complex behaviours such as planning and navigation, reading, curiosity, visual foraging, the mountain car problem, and social conformity [5]. Curiosity is directly relevant to the exploratory behaviour endowed by AIF, which is shown in an experiment where agents mimic human cognitive development and learning through curiosity-driven tasks [35].

Table 2: Application Areas and Their References

Application Area	**References**
Deep active inference	[21], [32], [27], [41], [28], [16], [47], [25], [23], [48]
Collaborative and group-level applications	[39], [40], [19], [38], [25], [46], [47], [29], [23], [48]
Applications on robotic agents	[35], [17], [19], [44], [45], [25]
Indirect connections to social human behaviour	[19], [24], [41], [43], [50], [30]

DAI Applications: FEP uses a DAI model to emerge artificial general intelligence that can capture social human behaviour [26]. Another DAI model infers valence states, which can be extended to testing human or non-human animal behavioural patterns to improve on prior models [12]. Understanding these patterns could help incorporate human social behaviours into AI, leading to more advanced artificial social intelligence. DAI can be further integrated with MC methods to apply to more complex environments, advancing the understanding of human intelligence [32].

Collaborative and group-level applications: AIF and FEP enhance multi-agent social behaviour by incorporating perception and control to intelligent agents for successful collaboration and team-optimal behaviours in various situational contexts [2]. The current work focuses on a free-energy function construction to add effects of team structure and collaborative adaptation to project teams. AIF can also contribute to cooperation with the emergence of social rules to understand the dynamics of cooperative behaviour [40].

History and evolution show that forming heterogeneous groups led to the dominance of the human species by enabling adaptation, so collaboration may improve AI's adaptability [13]. These heterogeneous agents learn each agent's abilities to embody adaptive behaviours by updating their predictions on the task assignments to agents. This distributed but cooperative system with group-level adaptation allows for more flexible self-organising intelligent agents. The cooperation and flexibility of these teams resemble social organisations where agents may learn and develop social human behaviour through their exchanges with human agents. Collective intelligence supports this collaborative view and allows for interactions between agents and their physical and social environment [14]. The AIF agents perceive shared goals and align with other agents for effective coordination and better collective performance. This model builds agents with human social behaviour that can act within different social contexts, exchange beliefs, and collaborate. Trust is essential to collaboration, and AIF modelling of trust can support human-intelligent system interactions [15]. The authors also state that as AI systems evolve, they will reach a complexity where they can model humans entirely.

Peer-to-peer collaboration in decentralised distributed decision-making showcases that multi-intelligence systems can acquire purpose-driven behaviours to overcome disorder without external control [2]. PC and world state models also apply collaboration and social interactions to enhance a robot's social intelligence [33].

Human interactions and belief exchanges lead to cultural accumulation through social exchanges. The foundation of this cultural accumulation can explain the human collective intelligence by task specialisation and endow agents with a multilevel social structure and adaptive behaviour [38]. Furthermore, future work on the personality formation framework will focus on prosocial personality that plays a significant role in social human behaviour, interactions, and social coordination of personality and culture dynamics [47].

Sophisticated AIF agent behaviour may help manage social relations by considering the long-term advantages of belief exchange [7]. The SoA of these agents may be significant in these social exchanges, so the AIF and PC model provides insights into the robustness of SoA in social interactions [36].

AIF restricts social organisations to a limited number of states to persist in their survival [48]. Social organisations focus on environmental aspects that affect their process operations. Boundaries for species may naturally evolve over a long time, but social organisation may be required to define and change these boundaries hastily. Predictive errors in these boundaries pose a significant threat to the survival of social organisations, so minimising unwanted surprises is crucial. This research extends to more advanced multi-intelligence organisations that mimic human processes.

Applications on robotic agents: Horii et al. create an energy-based model for robots to recognise human emotions as depicted in figure 5 [37]. This application elevates agents' social intelligence and communication skills, leading to a more advanced portrayal of complex human social behaviours through emotional state inference. Another study incorporates the role of the body in social interactions and human cognition [39]. The combination of enactivist and cognitivist views also creates embodied self-models that operate as cybernetic controllers and predictive-memory systems [49]. Embodiment is a significant source of empirical priors, foundational to cognitive development.

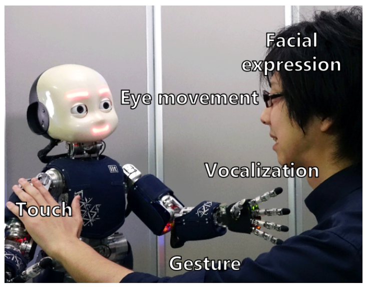

Figure 5: Multimodal human-robot interaction [37]

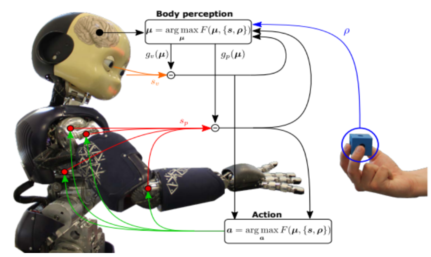

Figure 6: AIF body perception and action [45]

The significance of embodiment is further highlighted by the PR2 robot's interaction between its embodiment and the environment, showcasing social human behaviour with sensorimotor integration [44]. Oliver et al. also build a biologically plausible model of robots depicted in figure 6 to translate theoretical mathematical models to real-world applications [45]. The robot's behaviour is similar to that of a one-year-old during reaching and object-tracking tasks. Improvement in this model could produce robots resembling humans. Furthermore, AIF can be expanded to embody autonomous vehicles with imitation learning to identify when agents should mimic social human learning to adopt the expert's driving policy to succeed in a novel environment [30].

AIF to social behaviour: Some of the applications reviewed in 3.1, such as generalist and specialised agents [24], foraging geocaching task [29], drone simulations [41], maximising information gain [43], human decision-making modelling [50], and a deep temporal model for violation responses [34] may appear irrelevant to social human behaviour, but all of these applications aid in capturing social intelligence. The generalist agent adopts a flexible strategy, which can be improved to extend to social situations where the agent acts according to social contexts and cues. The evolutionary aspect of foraging is highly relevant to human decision-making skills, which can account for sentient and social behaviour. The drone simulations conducted in an unpredictable environment can also be extended to adaptive social behaviour to navigate different social contexts. Human exploratory behaviour is mimicked by maximising information gain, which could be extended to multimodal sensory input to embody human social intelligence. Lastly, violation responses indicate attentional processing and cognitive capabilities, which endow agents with human cognition that can be extended to social intelligence.

3.3 RQ 3 - Validity and Limitations

Translation to complex real-life applications: One of the main challenges in the research studies was the applicability of the models and experiments to real life. Chen et al. model agents that learn policies without cues, so agents are constructed with priors of actions leading to rewards [24]. This model is portrayed as a good approximation to actual behaviour, but the results need to be validated in a more complex environment. The exchanges in the experiments capturing cumulative culture are limited to a dyad, which does not reflect the multi-level interactions in real life [38]. Future work can improve the definition of innovation in cultural transmission and establish practical restrictions significant in real life (food, hygiene, temperature).

Furthermore, the deep learning model [26] and neural network model [28] have high complexity and computational demands that might complicate real-time applications as opposed to the controlled or simulated experimental setups. This limitation is also evident in the investigation of SoA, where the online error regression's computational demand for posterior inference is a significant bottleneck for real-time experiments, which is why the experiments were limited to pseudo-imitative interactions [36]. The GAN-based approach may also pose scaling issues since it was not tested on higher-dimensional action spaces [31]. Sophisticated AIF uses belief propagation to overcome scaling issues regarding computational efficiency, which is only tested on simple simulations [7]. Computational limitations need to be addressed for the generalisation of FEP-AI models [47].

The research of Oliver et al. also excludes real-life dynamics since they do not validate their method on real robots that may suffer from non-linear dynamics due to friction, physical constraints, uncertainties, and backlash [45]. Different robot models, such as those explored in experiments on tools constrained by linear dynamics [42] and utilising the VAE model [9], have not been tested on physical robots, making them less applicable to real-life dynamics. The increasing model complexity of valence inference and missing components of valence signals also limits the realistic applications [12].

Limitations regarding data: AI systems need to function in uncertain, noisy, and volatile environments to succeed in the future. Still, the availability of sufficient training data in these environments is low, so the successful applications are limited to a collection of tasks [13]. Due to size and type limitations, experiments are conducted on synthetic data, so unpredictable and complex real-life scenarios may not be fully captured [43].

The energy-based model [37] for human estimation performs less accurately than in recent studies due to the separate training process of the Multimodal Deep Belief Network (MDBN) and Feedforward Neural Network (FFNN). Therefore, MDBN lacks effective representation for emotion recognition in FFNN. Besides these limitations, several assumptions are made about human-agent interactions that also do not apply to real-life scenarios.

Validity of models: Future work is necessary to test the generalisability of the PR2 robot's model to more complex domains to analyse its advantages compared to other control schemes [44]. The model of Schwartenbeck et al. also suffers from validity due to an experiment conducted with a sample size of 20 people and a simple task to understand the human decision-making process [50]. This model needs to be tested on more people with more complex tasks. Spatial foraging simulations also have limitations regarding the model on prioritisation of clarity over the complexity and the assumption that the desired location is fed to the model as a prior preference [29]. Future work may improve on including the missing elements of spatial foraging.

Implementation of deep active inference [26] hardcodes the agent's expectations according to prior distributions. Still, the agent's flexibility allows it to converge swiftly to reach the set goal and build a realistic model of its environment. The generative model could participate in the optimisation process for more complex a priori beliefs for future work.

A key challenge for AIF is constructing the optimal generative model to depict all the observable data to maximise evidence due to the high sensory consequences in the real world [5]. An agent could solve this problem by creating a generative model by observing its behaviour. Another issue for this is scaling AIF due to the size of policy trees with high degrees of freedom. Fine-graining the process of preventing surprise is challenging since the current pruning strategy cannot reduce the search space sufficiently to govern inference.

A limitation of AIF with ForneyLab is that high-dimensional models could result in numerical instabilities [25]. The specific message update order may also impact the algorithm convergence, but there is insufficient information on optimal scheduling strategies. Future work can approach this scheduling with FEP using scheduling as an inference process.

4 Responsible Research and Limitations

This section explains the reproducibility and transparency of the research methodology and the consideration of possible bias factors. TU Delft Vision on Integrity [51] and Netherlands Code of Conduct for Research Integrity [52] were followed while conducting this research. As described in section 2, PRISMA guidelines were followed, which provides a clear overview of the research process. The book used as a guide in structuring the paper and reviewing is also mentioned to clarify the research further. The databases used to extract the surveyed literature were ensured to be reliable and have peer-reviewed articles to avoid misinformation.

The most significant element of bias for this research is in the literature selection. A research query was created during the background research by extracting significant keywords for AIF and FEP applications on human social behaviour. Despite using generic keywords not to limit the research scope, the query may be further optimised through more scoping research. The inclusion and exclusion criteria for filtering papers are reported 2. During the selection process, all the papers matching these criteria were selected. No positive bias was shown that could favour the results of the research question. In addition to identifying literature with the search string, citation chaining was used, which may affect reproducibility since other researchers may not select the same research papers. Still, the availability of these papers as a resource does not change.

Due to the short period of 9 weeks, ChatGPT was used as an additional resource to assist in critically analysing some of the papers. Specification of this process is depicted in the appendix A. ChatGPT was also used to improve the writing style in clarity and conciseness as specified in appendix A. Another significant constraint was the page limitation, which affected the report's coverage regarding detailed explanations of methodologies used in surveyed literature and the scarce use of visuals.

5 Discussion

This research study explored the applications of AIF and FEP on embodied virtual agents from a social human behaviour perspective to provide insights into the human brain and cognitive mechanisms. Current applications of these frameworks were reviewed and analysed to answer the sub-questions regarding the methodology used in AIF and FEP applications, endowing social human behaviour in intelligent agents with these methods, and the challenges that limit the research applications.

Various ML methods are combined with AIF to build agents that can adapt to volatile environments, balance epistemic behaviour with reward-seeking behaviour, compare long-term goals with short-term challenges, endow agents with an SoA to understand cooperative behavioural dynamics, simulate cultural accumulation, model human decision-making processes, and investigate the significance of embodiment in interacting with the environment. By mimicking these human processes in agents, new insights can be gained into human social intelligence and cognition. Besides endowing intelligent agents with human social behaviours, this research can extend AIF to biological and neuropsychological research and help uncover the unknowns of the human brain.

While not all simulations directly integrate social human behaviour, agents that minimise surprise and model basic human behaviours can be adapted to social contexts. Although adaptation may appear to only relate to physical survival, social behaviours significantly contribute to human dominance. Collective intelligence surpasses individual intelligence and is crucial for human survival [14]. By mimicking these adaptive processes, even in simple tasks, intelligent agents are endowed with adaptive capabilities to adjust to their environments, including social contexts. An agent that navigates their physical environments can also learn appropriate social behaviour according to social cues and settings. Understanding these social interactions enables seamless communication with humans, enhancing the agent's functionality and persistence in survival. Other social behaviour-mimicking applications included emotional state inference, cooperative behaviours, and cultural accumulation, all directly related to acquiring social intelligence.

Certain limitations arise as research expands, and AIF models become more complex to apply to more intricate behaviours and unpredictable situations. One of the most critical limitations is the translation of the theoretical models and simulations to real-life applications. Some models are encoded with assumptions or evaluated with simple task-related environments [9], [24], [38], [42], [44], [45], which are too restrictive to be generalised to broader applications and endow epistemic behaviour. More complex settings are required to test if agents emulate sentient behaviour. As models become more complex when applied to volatile situations, high computational demands arise, challenging real-time applications. Additionally, some robot applications did not consider the real-life dynamic constraints on robots, such as friction and physical constraints. Besides these, the lack of training data restricted the accurate analysis of the results. In theory, if applications overcome these technical limitations, AIF can depict agents with sentient behaviour, acquiring knowledge through their desires and intentions [26], [28], [36].

Despite the successful results of the applications discussed in section 3, current research needs to scale to high-dimensional inputs encountered in real life, be reproduced to ensure the generalisability of the models, be extended to incorporate real-life constraints and test with more data. However, current research provides promising results and explains the future work that can endow agents with human intelligence.

6 Conclusion

This systematic literature review concisely reports previous AIF and FEP applications to gain insights into human cognition and elevate social human behaviour and intelligence of embodied virtual agents. This study aims to guide future research by explaining the limitations of current applications, specific improvements that can be applied to current models, and a performance comparison between models. AIF proposes a promising framework for integrating social intelligence and social human behaviours into AI systems as well as understanding the human brain and cognitive processes. However, the current AIF models should be applied to real-life scenarios to ensure the accuracy and applicability of their results. This can be done by testing these models across various datasets to enhance and validate AIF applications, clarify and minimise assumptions made on the models, and evaluate agents in real environments. Future improvements should focus on integrating physical dynamics, real-life constraints, and computational aspects into simple AIF models. The complex theoretical models should be tested in real life to ensure that these applications depict sentient behaviour. With these improvements, future research can expand on these models to create highly adaptive intelligent agents that scale to complex environments 3.3.

References

[1] Thomas Parr, Giovanni Pezzulo, and Karl J. Friston. *Active Inference: The Free Energy Principle in Mind, Brain, and Behavior*. The MIT Press, March 2022.

[2] G. Levchuk, K. Pattipati, D. Serfaty, A. Fouse, and R. McCormack. Active inference in multi-agent systems: context-driven collaboration and decentralized purpose-driven team adaptation. volume 2018-March, pages 157–165, 2018.

[3] Karl Friston. The free-energy principle: a unified brain theory? *Nature Reviews Neuroscience*, 11(2):127–138, February 2010.

[4] Karl Friston. The free-energy principle: a rough guide to the brain? *Trends in Cognitive Sciences*, 13(7):293–301, July 2009.

[5] L. Da Costa, T. Parr, N. Sajid, S. Veselic, V. Neacsu, and K. Friston. Active inference on discrete state-spaces: A synthesis. *Journal of Mathematical Psychology*, 99, 2020.

[6] K. Friston, T. FitzGerald, F. Rigoli, P. Schwartenbeck, J. ODoherty, and G. Pezzulo. Active inference and learning. *Neuroscience and Biobehavioral Reviews*, 68:862–879, 2016.

[7] K. Friston, L. Da Costa, D. Hafner, C. Hesp, and T. Parr. Sophisticated inference. *Neural Computation*, 33(3):713–763, 2021.

[8] Giovanni Pezzulo, Thomas Parr, Paul Cisek, Andy Clark, and Karl Friston. Generating meaning: active inference and the scope and limits of passive ai. *Trends in Cognitive Sciences*, 28(2):97–112, 2024.

[9] C. Meo and P. Lanillos. Multimodal vae active inference controller. pages 2693–2699, 2021.

[10] Boris Sedlak, Victor Casamayor Pujol, Praveen Kumar Donta, and Schahram Dustdar. Active inference on the edge: A design study. In *2024 IEEE International Conference on Pervasive Computing and Communications Workshops and other Affiliated Events (PerCom Workshops)*, pages 550–555, 2024.

[11] K. Friston, P. Schwartenbeck, T. FitzGerald, M. Moutoussis, T. Behrens, and R.J. Dolan. The anatomy of choice: Active inference and agency. *Frontiers in Human Neuroscience*, (SEP), 2013.

[12] C. Hesp, R. Smith, T. Parr, M. Allen, K.J. Friston, and M.J.D. Ramstead. Deeply felt affect: The emergence of valence in deep active inference. *Neural Computation*, 33(2):398–446, 2021.

[13] G. Levchuk, A. Fouse, K. Pattipati, D. Serfaty, and R. McCormack. Active learning and structure adaptation in teams of heterogeneous agents: Designing organizations of the future. volume 10653, 2018.

[14] R. Kaufmann, P. Gupta, and J. Taylor. An active inference model of collective intelligence. *Entropy*, 23(7), 2021.

[15] F. Schoeller, M. Miller, R. Salomon, and K.J. Friston. Trust as extended control: Human-machine interactions as active inference. *Frontiers in Systems Neuroscience*, 15, 2021.

[16] Charles Sturt University Library. Systematic literature reviews. Library Guide, 2024. Accessed on May 24, 2024.

[17] David Moher, Alessandro Liberati, Jennifer Tetzlaff, Douglas G Altman, and The PRISMA Group. Preferred reporting items for systematic reviews and meta-analyses: the prisma statement. *PLoS Med*, 6(7):e1000097, 2009.

[18] Angela Boland, M. G. (M. Gemma) Cherry, and R. (Rumona) Dickson. *Doing a systematic review : a student's guide / edited by Angela Boland, M. Gemma Cherry Rumona Dickson.* SAGE, London, 2014.

[19] About cell. `https://www.cell.com/about`, 2024. Accessed: 2024-06-03.

[20] Scopus. `https://www.elsevier.com/products/scopus`, 2024. Accessed: 2024-06-03.

[21] About ieee xplore. `https://ieeexplore.ieee.org/Xplorehelp/overview-of-ieee-xplore/about-ieee-xplore`, 2024. Accessed: 2024-06-03.

[22] Web of science platform. `https://clarivate.com/products/scientific-and-academic-research/research-discovery-and-workflow-solutions/webofscience-platform/`, 2024. Accessed: 2024-06-03.

[23] About acm. `https://dl.acm.org/about`, 2024. Accessed: 2024-06-03.

[24] A.G. Chen, D. Benrimoh, T. Parr, and K.J. Friston. A bayesian account of generalist and specialist formation under the active inference framework. *Frontiers in Artificial Intelligence*, 3, 2020.

[25] T.W. van de Laar and B. de Vries. Simulating active inference processes by message passing. *Frontiers Robotics AI*, 6(MAR), 2019.

[26] K. Ueltzhöffer. Deep active inference. *Biological Cybernetics*, 112(6):547–573, 2018.

[27] M. T. Koudahl and B. de Vries. Batman: Bayesian target modelling for active inference. In *ICASSP 2020 - 2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 3852–3856, May 2020. journalAbbreviation: ICASSP 2020 - 2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP).

[28] T Matsumoto, W Ohata, FCY Benureau, and J Tani. Goal-directed planning and goal understanding by extended active inference: Evaluation through simulated and physical robot experiments. *ENTROPY*, 24(4), April 2022.

[29] V. Neacsu, L. Convertino, and K.J. Friston. Synthetic spatial foraging with active inference in a geocaching task. *Frontiers in Neuroscience*, 16, 2022.

[30] S. Nozari, A. Krayani, P. Marin, L. Marcenaro, D. Martin, and C. Regazzoni. Adapting exploratory behaviour in active inference for autonomous driving. 2023.

[31] F. Nikita and A. Voskoboynikov. Deep active inference agent with continuous action space. pages 215–220, 2023.

[32] Z. Fountas, N. Sajid, P.A.M. Mediano, and K. Friston. Deep active inference agents using monte-carlo methods. volume 2020-December, 2020.

[33] T Taniguchi, S Murata, M Suzuki, D Ognibene, P Lanillos, E Ugur, L Jamone, T Nakamura, A Ciria, B Lara, and G Pezzulo. World models and predictive coding for cognitive and developmental robotics: frontiers and challenges. *ADVANCED ROBOTICS*, 37(13):780–806, July 2023.

[34] Francesco Rigoli Philipp Schwartenbeck Karl Friston, Thomas FitzGerald and Giovanni Pezzulo. Deep temporal models and active inference. *Neuroscience Biobehavioral Reviews*, 77:388–402, 2017.

[35] Karl J. Friston, Marco Lin, Christopher D. Frith, Giovanni Pezzulo, J. Allan Hobson, and Sasha Ondobaka. Active inference, curiosity and insight. *Neural Computation*, 29:2633–2683, 2017.

[36] W. Ohata and J. Tani. Investigation of the sense of agency in social cognition, based on frameworks of predictive coding and active inference: A simulation study on multimodal imitative interaction. *Frontiers in Neurorobotics*, 14, 2020.

[37] Takato Horii, Yukie Nagai, and Minoru Asada. Active perception based on energy minimization in multimodal human-robot interaction. In *HAI'17*, page 103â110, New York, NY, USA, 2017. Association for Computing Machinery. journalAbbreviation: HAIâ17.

[38] N. Kastel, C. Hesp, K.R. Ridderinkhof, and K.J. Friston. Small steps for mankind: Modeling the emergence of cumulative culture from joint active inference communication. *Frontiers in Neurorobotics*, 16, 2023.

[39] Shaun Gallagher and Micah Allen. Active inference, enactivism and the hermeneutics of social cognition. *Synthese*, 195:2627–2648, 2018.

[40] M. Hartwig and A. Peters. Cooperation and social rules emerging from the principle of surprise minimization. *Frontiers in Psychology*, 11, 2021.

[41] T. Van De Laar, İ. Şenöz, A. Özçelikkale, and H. Wymeersch. Chance-constrained active inference. *Neural Computation*, 33(10):2710–2735, 2021.

[42] A. Anil Meera and P. Lanillos. Towards metacognitive robot decision making for tool selection. volume 1915 CCIS, pages 31–42, 2024.

[43] Tadahiro Taniguchi, Ryo Yoshino, and Toshiaki Takano. Multimodal hierarchical dirichlet process-based active perception by a robot. *Frontiers in Neurorobotics*, 12:22, 2018.

[44] Carlo Pezzato, Beren Millidge, Alexander Tschantz, and Christopher L. Buckley. Active inference and robot control: A case study. *Journal of Robotics and Autonomous Systems*, 130:103515, 2020.

[45] Guillermo Oliver, Pablo Lanillos, and Gordon Cheng. An empirical study of active inference on a humanoid robot. *IEEE Transactions on Cognitive and Developmental Systems*, 14(2):462–471, 2022.

[46] Florin Leon. The mountain car problem. Available at ResearchGate, 2010. Accessed: 2024-06-09.

[47] A. Safron, Z. Sheikhbahaee, N. Hay, J. Orchard, and J. Hoey. Value cores for inner and outer alignment: Simulating personality formation via iterated policy selection and preference learning with self-world modeling active inference agents. volume 1721 CCIS, pages 343–354, 2023.

[48] S Fox. Active inference: Applicability to different types of social organization explained through reference to industrial engineering and quality management. *ENTROPY*, 23(2), February 2021.

[49] A. Safron. The radically embodied conscious cybernetic bayesian brain: From free energy to free will and back again. *Entropy*, 23(6), 2021.

[50] Philipp Schwartenbeck, Thomas H B FitzGerald, Christoph Mathys, Ray Dolan, Martin Kronbichler, and Karl Friston. Evidence for surprise minimization over value maximization in choice behavior. *Scientific Reports*, 5:16575, 2015.

[51] Sabine Roeser, Merle de Kreuk, Bernard Dam, Marja Elsinga, and Rinze Benedictus. Tu delft vision on integrity 2018-2024. Policy document, Delft University of Technology, Delft, The Netherlands, September 2018.

[52] Association of Universities in the Netherlands. *Netherlands Code of Conduct for Research Integrity*, 2018. Accessed: [Insert current date].

Appendix

A LLM Usage

I used the ChatGPT version 4 with the prompt: "Please provide an overall summary as well as the summary of sections relevant to the methodology, experiments performed, results of the study, research limitations, and social human behaviour in this paper".

I used the ChatGPT version 4 with the prompt: "Can you indicate where I can improve on my writing style and make my writing more concise?".