Cognitive Learning Agents for Autonomous Mobility on Demand Systems

Ömer Ibrahim Erduran

Department of Computer Science, Goethe University, Frankfurt am Main, Germany

Abstract

In my PhD thesis, the concept of Cognitive Agents with extended Learning capabilities for Autonomous Mobility on Demand (AMoD) scenarios is investigated. Specifically, the focus is set on the Ride-hailing concept with a fleet of autonomously driving vehicles. The Agent-based approach provides the possibility to consider cognitive agent architectures for different types of agents in the given scenario. In this regard, the vehicle agents are built up based on the Belief-Desire-Intention (BDI) architecture. My dissertation combines two paradigms that are considered significant research areas, namely Machine learning (ML) and Agent-oriented Programming (AOP). Therefore, a structured overview of the research made so far is provided pointing out significant areas in the AMoD application scenario which are worthwhile to work on. For each of the areas, I describe why the setting and combination of MAS and ML are relevant and interesting for in-depth investigation.

Keywords

BDI Agent, Machine Learning, Agent-Oriented Programming, Mobility on Demand, Multi-Agent System, Neuro-Symbolic AI

1. Introduction

The basis of my dissertation relates to work in Multi-Agent system (MAS) research, which has been studied over the last decades [1]. Advances in Deep Learning have made remarkable progress in data-driven learning approaches. Using extensive computational processing for Reinforcement Learning, it is likewise possible to make achievements in different applications. Due to the proliferation of ML techniques in multi-agent scenarios, which have appeared especially recently [2], the concept of MAS is particularly prominent in ML research, where multiple agents are interacting with each other to accomplish a common goal. In this work, a MAS consisting of autonomous *Software Agents* is considered with explicitly programmed capabilities. A Software Agent is an entity that interacts with its environment by sensing information and producing actions. In the field of Agent-Oriented Programming, the most predominant agent architecture is the *Belief-Desire-Intention* architecture, in short BDI, which is based on the *Sense-Think-Act* cycle of cognitive Agents. The main characteristic of cognitive Software Agents is that the capabilities which the human calls intelligent behavior have to be hard coded into the agent architecture. In ML, the intelligent behavior of the considered

Doctoral Consortium of the 45th German Conference on Artificial Intelligence (KI2022), September 20, 2022, Trier (Virtual), Germany
erduran@cs.uni-frankfurt.de (Ö. I. Erduran)
0000-0002-1586-0228 (Ö. I. Erduran)

learning agent emerges during the learning phase. In the underlying scenario of AMoD, certain challenges are investigated and extended with ML methods. This dissertation aims to address the issue of extending cognitive agent models with ML capabilities and provides initial results considering the mentioned application scenario. Integrating ML techniques is a recent and open issue in Agent-Oriented Programming [3, 1, 4]. For example, one of the key limitations of the BDI architecture is the lack of generating new plans during processing [5]. In a traffic simulation environment, the learning procedure of the agent is then evaluated. The results of this dissertation will provide insight into different approaches to integrating learning capabilities into the BDI agents interacting in a MAS. Moreover, a novel approach to learning in a MAS with BDI agents is presented investigating the potential and limitations of the framework.

2. Problem formulation and Related work

The underlying problem formulation is based on the *Dynamic Pickup and Delivery problem* (DPDP) [6], which is a variation of the *Vehicle Routing Problem* (VRP). Given is a *trip request* d_i with a passenger, a potential set of trip requests R and a complete and directed graph $G = (V, A)$ with a node set $V = \{0\} \cup \{i^+ \mid i \in R\} \cup \{i^- \mid i \in R\}$, which means that V contains the origin and destination of all trip requests and moreover, the vertex 0 which represents the *depot*. Furthermore, an edge set $A = \{(i, j) : i, j \in V, i \neq j\}$ is considered, where each edge $(i, j) \in A$ has a non-negative length or cost c_{ij} and a non-negative travel time t_{ij}. At each time step t, each vehicle $k \in K$ is either serving a node, waiting at a node, or moving towards a node. The vehicles can decide to *accept* or *reject* a trip request [6]. Solving VRP and DPDP is already tackled with ML approaches [7, 8]. In [8], the problem of fleet management is formulated as a Multi-agent Reinforcement Learning problem. Here, the MAS is learning as a whole to dispatch and reposition vehicles. In the work of [9], different assignment strategies with flexible parameters for a variation of vehicle allocation are investigated and compared with each other. In [10], an overview of fleet management problems and approaches is presented containing optimization approaches for different scenarios, but neglecting learning approaches. Finally, [11] presents a mixed-integer program for optimizing vehicle-sharing systems. The approach in my work differs from the works mentioned since I consider the BDI agent architecture and integrate ML into the agent reasoning cycle.

3. Current state

So far, I investigated the representation of the considered application scenario as well as concrete research questions. The addressed research questions in my work are focused on the interaction between the cognitive BDI architecture which is enhanced with learning capabilities and its effects to the AMoD simulation environment. The primary question of my research is: *To what extent can cognitive BDI agents benefit from learned behaviors?* This question addresses the issue described in the previous chapter. More specifically, I consider Deep Learning algorithms to influence the decision-making of the BDI agents. Thus, I analyze the effects of learning by investigating the BDI reasoning steps, like *goal selection* as well as the effects in the simulation environment.

- *How can Neural Networks enhance the decision-making of BDI agents?*
- *Which tasks are suitable for ML in Mobility on Demand?*
- *What are the effects of multiple cognitive BDI agents learning in an MAS?*

In my first work, I tackled the fleet positioning of vehicle agents with cluster analysis using generated GPS coordinates and open-source customer trip requests from a bike-sharing fleet [12]. In the second work, I worked on the application scenario representation considering the BDI agent architecture and a trip request negotiation process [13]. Furthermore, I covered the relevant literature on the considered research question of integrating ML into BDI Agents as a survey [14]. Due to the mentioned requirements, extensive research has been done to investigate development platforms for Software Agents as well as Mobility on-demand simulation platforms. Here, I picked out *JadeX* [15] as an agent development platform and MATSim [16] as the traffic simulation environment. The main characteristic of JadeX is that Goals and Plans are formulated explicitly. Therefore, I currently investigate the learning behavior on a goal level for BDI agents which addresses the decision-making step. MATSim and JadeX are implemented in Java. The integration of ML algorithms is therefore also considered in Java. The libraries *DL4J* [1] and *DJL* [2] provide Deep Learning and Reinforcement Learning algorithms that will be applied to the BDI agent cycle.

4. Planned contributions

The core contribution of my thesis is the investigation of the BDI architecture connected with learning capabilities. As an application scenario, the ride-hailing application is considered. The subject of my research is therefore an agent architecture, where different capabilities and concepts are investigated with a focus on learning algorithms as extensions. The research environment is the development of a fleet management system with high-level learning strategies for cooperating autonomous vehicles in Mobility on Demand scenarios. In the following, the specific contributions are described.

4.1. Fleet Management

The first task considers the whole vehicle fleet and its utilization starting with the positioning and rebalancing of the vehicle agents. In general, a fleet coordination challenge is addressed. This problem is tackled using spatiotemporal data and self-organizing and communicating agents. One research direction is focusing on challenges that arise during the processing of the fleet. Since the vehicle agents contain a cognitive thinking phase, a travel time prediction component will be integrated into the agent's thinking phase leading to more informed decisions and actions. Finally, the battery charging behavior of a fleet is investigated which also influences the decision-making of the vehicle agents. Here, the avoidance of running out of battery power is one central question. This challenge will be tackled with Reinforcement Learning by training the BDI agents to learn battery management.

[1] https://deeplearning4j.konduit.ai/
[2] https://djl.ai/

4.2. Decision making

The second task contains multiple contributions and represents the main part of my thesis concerning the cognitive part, where learning capability is employed in the BDI cycle to decide about committing to customer trip requests. Similar to the work of [3], I consider the question of integrating Reinforcement Learning methods in the BDI architecture for the decision-making of vehicle agents. Considering different agent capabilities, I compare different agent types including learning agents. Starting with a single vehicle agent and its decision-making, the question is extended to the whole fleet representing a MAS with cognitive learning agents. The certain task is the trip assignment step. In this case, a utility-based negotiation is considered as well as a learned utility function with Deep Reinforcement Learning (DRL), where the decision is based on the corresponding reward function.

The communication of agents in MAS is significant for coordination. In the BDI architecture, a common method to realize communication is using standardized predefined speech acts. The messaging process requires an extensive engineering process, where each messaging type and direction has to be considered in order to provide communication to the MAS. Learning when to communicate with other agents is crucial for efficient problem-solving. Furthermore, learning on a fleet level is a novel approach for cognitive BDI agents which has to be investigated since nearly all of the works published in this intersection, focus on the single-agent setting.

Since the decision-making steps inside the typical BDI agent is predefined, some steps are suitable for specific ML algorithms. During *Goal selection*, the agent decides, which *Goal(s)* it should pursue and thus which *plans* it should process. Learning which goal to pursue is a novel approach in BDI and ML integration. This approach will be investigated with decision trees and neural networks as well as evaluated in the AMoD simulation environment.

5. Conclusion

This dissertation focuses on a variety of integration methods for cognitive agent architecture and ML methods in a Mobility on Demand application. The presented areas are tackled in current research with a variety of approaches. Fundamentally, my approach considers AOP for ML as well as the given application scenario of ride-hailing and therefore differs from the majority of current work in this area. The contribution of my dissertation delivers cognitive Software Agents enhanced with ML techniques processing in a BDI manner. The mentioned open issues in the previous section lead to a need for a thorough investigation of cognitive agents that are capable to learn. Since they are used for industrial applications, this research intersection of considering cognitive agents for AMoD enables the investigation for novel research insights with respect to autonomy in MAS and fleet applications.

References

[1] A. Guerra-Hernández, A. El Fallah-Seghrouchni, H. Soldano, Learning in BDI Multi-agent Systems, in: Computational Logic in Multi-Agent Systems, volume 3259, Springer Berlin Heidelberg, 2004, pp. 218–233.

[2] J. Foerster, I. A. Assael, N. De Freitas, S. Whiteson, Learning to Communicate with Deep Multi-Agent Reinforcement Learning, Advances in NeurIPS 29 (2016).
[3] M. Bosello, A. Ricci, From Programming Agents to Educating Agents – A Jason-Based Framework for Integrating Learning in the Development of Cognitive Agents, in: EMAS, Springer International Publishing, 2020, pp. 175–194.
[4] L. Padgham, S. Sardina, S. Sen, Incorporating learning in bdi agents (2008).
[5] R. H. Bordini, A. El Fallah Seghrouchni, K. Hindriks, B. Logan, A. Ricci, Agent programming in the cognitive era, AAMAS 34 (2020).
[6] G. Berbeglia, J.-F. Cordeau, G. Laporte, Dynamic pickup and delivery problems, European journal of operational research 202 (2010) 8–15.
[7] M. Nazari, A. Oroojlooy, L. V. Snyder, M. Takác, Deep reinforcement learning for solving the vehicle routing problem, arXiv preprint arXiv:1802.04240 (2018).
[8] J. Jin, M. Zhou, W. Zhang, M. Li, Z. Guo, Z. Qin, Y. Jiao, X. Tang, C. Wang, J. Wang, Coride: joint order dispatching and fleet management for multi-scale ride-hailing platforms, in: Proceedings of the 28th CIKM, 2019.
[9] A. C. Regan, H. S. Mahmassani, P. Jaillet, Evaluation of dynamic fleet management systems: Simulation framework, Transportation research record 1645 (1998) 176–184.
[10] M. Bielli, A. Bielli, R. Rossi, Trends in models and algorithms for fleet management, Procedia-Social and Behavioral Sciences 20 (2011) 4–18.
[11] R. Nair, E. Miller-Hooks, Fleet management for vehicle sharing operations, Transportation Science 45 (2011) 524–540.
[12] Ö. I. Erduran, M. Minor, L. Hedrich et al., Multi-agent Learning for Energy-Aware Placement of Autonomous Vehicles, in: the proceedings of ICMLA 2019, IEEE, Boca Raton, FL, USA, 2019, pp. 1671–1678.
[13] Ö. I. Erduran, M. Mauri, M. Minor, Negotiation in ride-hailing between cooperating bdi agents, in: Proceedings of the 14th ICAART Volume 1, INSTICC, SciTePress, 2022, pp. 425–432.
[14] Ö. I. Erduran, Machine Learning Algorithms for Cognitive and Autonomous BDI Agents, in: P. Reuss, V. Eisenstadt, J. M. Schönborn (Eds.), Proceedings of the LWDA 2022 Workshops: FGWM, FGKD, and FGDB, Hildesheim (Germany), Oktober 5-7th, 2022, volume 3341 of *CEUR Workshop Proceedings*, CEUR-WS.org, 2022, pp. 112–123.
[15] A. Pokahr, L. Braubach, W. Lamersdorf, Jadex: A BDI Reasoning Engine, in: Multi-Agent Programming, volume 15, Springer US, Boston, MA, 2005, pp. 149–174.
[16] D. Singh, L. Padgham, K. Nagel, Using MATSim as a Component in Dynamic Agent-Based Micro-Simulations, in: Engineering Multi-Agent Systems, Springer International Publishing, 2019, pp. 85–105.