

Multi-tasking resource-constrained agents reach higher accuracy when tasks overlap*

Andreas Kalaitzakis and Jérôme Euzenat

Univ. Grenoble Alpes, Inria, CNRS, Grenoble INP, LIG, F-38000 Grenoble, France
{Andreas.Kalaitzakis, Jerome.Euzenat}@inria.fr

Abstract. Agents have been previously shown to evolve their ontologies while interacting over a single task. However, little is known about how interacting over several tasks affects the accuracy of agent ontologies. Is knowledge learned by tackling one task beneficial for another task? We hypothesize that multi-tasking agents tackling tasks that rely on the same properties, are more accurate than multi-tasking agents tackling tasks that rely on different properties. We test this hypothesis by varying two parameters. The first parameter is the number of tasks assigned to the agents. The second parameter is the number of common properties among these tasks. Results show that when deciding for different tasks relies on the same properties, multi-tasking agents reach higher accuracy. This suggests that when agents tackle several tasks, it is possible to transfer knowledge from one task to another.

Keywords: Cultural knowledge evolution · Knowledge transfer · Multi-tasking.

1 Introduction

Agents have been previously shown to improve their accuracy as a result of cultural knowledge evolution. The latter studies agents that evolve their knowledge representations, based on their perception and the feedback they receive from other agents. Recent work on cultural knowledge evolution focuses on agents tackling a single task: taking an abstract decision within an abstract domain. In [3], agents are forced to take identical decisions regarding a set of environment objects. Eventually, agents learn to agree over a single decision task, yet not necessarily on the same basis. For example, two agents may both decide to visit *Barcelona*. Agent α may base its decision on the *temperature* property, while agent β may base its decision on the *ticket_price* property.

However, several tasks may exist. We build on previous works by introducing agents capable of taking abstract decisions within several domains. To do so, agents classify objects into ontology classes and associate these classes with different decisions for different tasks. We consider that realistic agents should not be able to develop ontologies containing all class descriptions. Thus, we limit the number of classes to be maintained

* Author version of the paper appeared in: Vadim Malvone, Aniello Murano (eds), Proc. 20th European conference on multi-agents systems (EUMAS), Napoli (IT) *Lecture notes in computer science* 14282:425-434, 2023

within an agent’s ontology. When this limit is reached, agents will try to forget knowledge that is not relevant to the tasks they favor. Deciding for different tasks may rely on a set of common properties. For example, the property *temperature* may be used in order to choose a destination (task 1). The same property may also be used to decide whether to wear a T-shirt (task 2). However, the property *temperature* may be completely irrelevant to choosing a movie (task 3). We assume that when this set is not empty, agents carrying several tasks may develop multi-purpose knowledge, i.e., knowledge that can be transferred among different tasks. Based on this, we formulate the following hypothesis: multi-tasking agents tackling tasks that rely on the same properties, are more accurate than multi-tasking agents tackling tasks that rely on different properties. We test this hypothesis by varying two parameters. The first parameter is the number of tasks assigned to each agent. The second parameter the number of common properties shared among the different tasks. Two variations of the second parameter are examined. Tasks either rely on the same properties, or rely on different ones. We then evaluate agent ontologies based on their contribution to promote successful interactions and provide accurate decisions. Based on this evaluation, the following is shown: when agents tackle tasks based on common properties, knowledge built by an agent while tackling one task, improves its accuracy on another task. We thus conclude that it is possible to transfer knowledge from one task to another.

After discussing related work in §2, preliminaries regarding the entities that constitute the environment as well as the notation that describes it, are introduced in §3. In §4, an outline of the experiment is provided, including how agents learn their initial ontologies, interact with each other and adapt when they disagree. Section 5 presents our hypothesis and the protocol used to test it. Results are presented in §6 and conclusions are provided in §7.

2 Related work

It has been shown that referential games[9] facilitate the establishment of communication protocols between communicating agents. [11] argues that a communication protocol emerges when agents attempt to minimize the computational complexity of semantic interpretation. [7] studies a framework where two agents develop a language in order to succeed in a referential game. [6] shows that implicit cultural transmission leads to greater language compositionality. While our work relies on successfully communicating agents, our focus is on how this successful communication allows for better task completion.

Different examples of multi-tasking agents exist in literature. Indicatively, multi-task learning has been shown to significantly improve classification in a variety of areas, e.g., adversary robustness [10], visual interconceptual similarity [4] and phenotype learning [5]. Agents have also been used to study the impact of multi-task learning on emerging communication protocols. In [12], cooperative multi-agent reinforcement learning is considered. Our work is related to these works, since they consider agents that perform several tasks. However our focus is not on agents that improve their accuracy individually. Here we study agents that improve their accuracy through social transmission.

Social transmission among agents has been studied in [3] and [13]. In [13], the authors examine how concepts are organized and how their collective behavior can be established autonomously. In [3], a two-stage experiment is used, where agents first learn a classifier and then interact in pairs. Through an adaptation mechanism, it is shown that the agents achieve better knowledge, without adopting identical ontologies. We differentiate from these, by introducing memory-limited agents that tackle several tasks.

3 Experimental framework

3.1 Environment

Agents evolve in an environment populated by objects described by a set \mathcal{P} of boolean properties. Objects are therefore described by the presence or absence of a property $p \in \mathcal{P}$, denoted by p and $\neg p$ respectively. Hence, there are $2^{|\mathcal{P}|}$ object types, that are gathered in a set \mathcal{I} .

3.2 Tasks

The term task refers to a piece of work, carried out by an agent. Here, we will concentrate on a set of decision tasks: making a decision about an object. There may be different tasks $t \in \mathcal{T}$ associated to a different set of possible decisions \mathcal{D}_t . Each object o can be considered with respect to any task $t \in \mathcal{T}$. A function $h^*(o, t) \rightarrow \mathcal{D}_t$ provides the correct, unknown to agents, decision for an object o with respect to a task t . For example, $h^*(tomato, coloring)$ will provide the decision *red*.

3.3 Agents

Agents are autonomous, co-existing entities, able to perceive and distinguish objects based on their properties. In this context, a population of multi-tasking agents \mathcal{A} is assigned different subsets of \mathcal{T} . To tackle these tasks, agents build and evolve private ontologies, expressed in \mathcal{ALC} [2]. Each agent α uses its ontology to compute a function $h^\alpha(o, t) \rightarrow \mathcal{D}_t$ which, given an object o and a task t , provides a decision $h^\alpha(o, t)$. The right part of Figure 1 shows an example of a multi-task ontology constructed by an agent α . The bottom part represents the private ontology \mathcal{O}^α of agent α , allowing it to classify objects of the environment. The top part shows a set of decision ontologies, each one containing the valid decisions for a respective task t . An agent α learns at most one decision for an object o and a task t . Thus, each leaf of \mathcal{O}^α cannot be aligned more than once with the same decision ontology.

4 Experiment outline

In this paper, we examine if knowledge can be transferred from one task to another. To this end, a two-stage experiment is used. In the first stage, agents induce private ontologies based on randomly selected labeled examples. In the second stage, agents

go through a fixed number of interactions. For each interaction, two randomly selected agents will have to decide with respect to an object o and a task t . When agents disagree, one of the two agents adapts its ontology. More details about how agents learn, are assigned tasks, interact, release resources and adapt their ontologies are presented in subsections 4.1, 4.2, 4.3, 4.4 and 4.5 respectively.

4.1 Initial ontology induction

We approach multi-task learning as a problem of inducing an ontology capable of providing a decision for any task $t \in \mathcal{T}$. Different algorithms may be used, affecting the final accuracy of agents. This paper does not examine how different learning algorithms impact the achieved accuracy. This paper examines how cultural evolution improves the accuracy of multi-tasking agents. Thus, details about the learning algorithm are omitted. A learning example can be seen in figure 1. By the end of its initial ontology induction phase, the agent α is able to classify an object described by $p_1 \sqcap p_2$ but unable to decide about the task t_1 .




Fig. (1) Given a set of labeled examples, agents will induce a decision tree. The latter is subsequently transformed into an ontology. Each color represents a different decision.

4.2 Task assignment

Agents are assigned with different subsets of \mathcal{T} of the same size. The latter varies from 1 to $|\mathcal{T}|$ and remains constant for the duration of an experiment. Based on it, all possible task permutations of the same size are initially produced. Each permutation corresponding to a different subset of \mathcal{T} , is then assigned to an even number of agents. Thus, the number of agents is always a multiple of the number of the different subsets of \mathcal{T} .

4.3 Interaction

For each interaction, two randomly selected agents α and β are asked to provide a decision for an object o with respect to a task t . The agents provide their decisions based on the respective functions $h^\alpha(o, t)$ and $h^\beta(o, t)$. If an agent is unable to provide a decision, then one decision is randomly selected. The agents will then disclose their decisions to each other. If $h^\alpha(o, t) = h^\beta(o, t)$, the agents agree and their interaction is

considered as successful. On the contrary, their interaction ends as a failure. In this case, one of the two agents may adapt its ontology. In order to decide which agent will adapt, an evaluation set is randomly selected. It contains samples labeled with respect to the task t . The agents are evaluated against this set and a score is assigned to each one of them. The agent with the lowest score may adapt its ontology.

4.4 Resources release

When an agent's resources are exhausted, it tries to forget knowledge as follows (Figure 2). Leaf nodes that satisfy the following criteria are removed: (a) they have the same immediate parent node (b) they are associated with the same decision regarding all tasks assigned to the agent. The process is repeated recursively, as long as leaf nodes satisfying (a) and (b) exist.




Fig. (2) Let an agent α assigned the task t_2 , with t_2 relying on the property set \mathcal{P}_{t_2} . The property $p_1 \notin \mathcal{P}_{t_2}$, thus p_1 does not allow for distinguishing different decisions for the task t_2 . In this example, the agent has associated the same decision (in red), to both $p_4 \sqcap \neg p_1$ and $p_4 \sqcap p_1$. These two classes can be removed without any loss of accuracy with respect to t_2 . For the task t_2 , the parent node will now be associated with the decision d_2 (red). For the task t_1 , the parent node will now be associated with one of two decisions previously associated with its former descendent nodes. Here, the decision d_3 (gray) was randomly selected.

4.5 Adaptation

Our adaptation mechanism extends the one presented in [3]. Based on it, an agent can either replace an existing decision or split a class into two sub-classes (Figure 3). The agent does this on the basis of a property that distinguishes the current object from the objects classified by the class to be split. Only the decisions concerning the current task are affected.

5 Experimental setting

5.1 Hypothesis

- Multi-tasking agents tackling tasks that rely on the same properties, are more accurate than multi-tasking agents tackling tasks that rely on different properties.

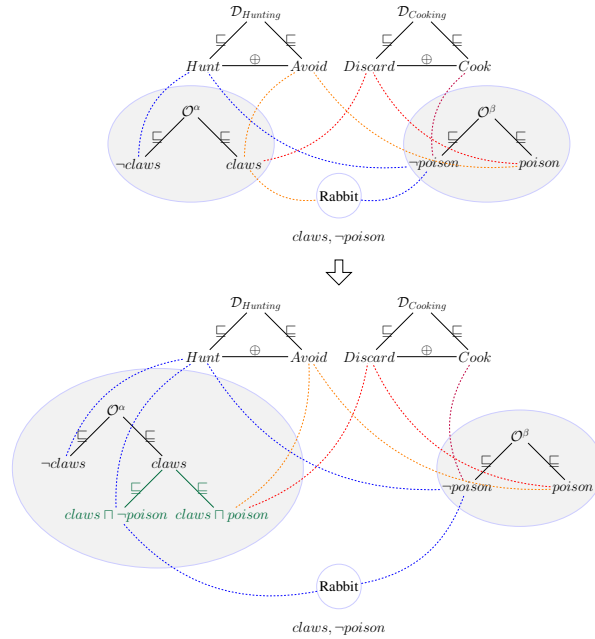


Fig. (3) The agent α will split the class *claws* into two sub-classes using the property *poison*. The first ($claws \sqcap \neg poison$) will be associated with the decision of the agent β . The second ($claws \sqcap poison$) will be associated with all decisions previously associated with the class *claws*.

5.2 Parameters

The experiment is executed under 6 setups. Each setup is run 20 times and its results are averaged. One run consists of 80000 interactions with each interaction taking place among two randomly selected agents. The total population of agents is 18. Their environment contains 64 different object types, each one perceivable through 6 different binary properties. The agents are initially trained with respect to 3 tasks. Taking 1 out of 4 decisions with respect to each task relies on 2 properties. These properties are either the same for all tasks, or different for each task. Agents induce an initial ontology based on a random 10 % of all existing labeled examples. The agents are assigned 1 to 3 tasks. Agent evaluation is based on 60% of all samples.

5.3 Measures

Success rate, as introduced in [3] is defined as the proportion of successful interactions, over all performed interactions until the n^{th} interaction. Task accuracy adapts the accuracy measure introduced in [3] to different tasks. It is defined as the proportion of object types for which a correct decision would be taken with respect to a task t , by an agent α on the n^{th} iteration of the experiment.

$$tacc(\alpha, n, t) = \frac{|\{o \in \mathcal{I} : h_n^\alpha(o, t) = h^*(o, t)\}|}{|\mathcal{I}|}$$

6 Results and discussion

We hypothesize that when tasks rely on common properties, it is possible for agents to build multi-purpose knowledge. To test this hypothesis, the accuracy of the following two populations was compared. The first consists of agents assigned up to 3 tasks for which all properties are shared. The second consists of agents assigned up to 3 tasks for which no properties are shared. Figure 4 depicts the evolution of the agents (a) average accuracy, (b) accuracy on their best task and (c) success rate, for different number of tasks and common properties. Figure 4a shows that assigning more tasks to agents,




Fig. (4) (a) average accuracy, (b) accuracy on best task and (3) success rate for different number of assigned tasks and common properties.

significantly improves their average accuracy. This improvement is higher when agents tackle tasks that rely on the same properties. On the one hand, when tasks rely on different properties, agents tackling 3 tasks are 9% more accurate than agents tackling 1 task. On the other hand, when tasks rely on common properties, agents tackling 3 tasks are 55% more accurate than agents tackling 1 task. This shows that agents tackling tasks relying on a common set of properties, may improve their accuracy on one task by carrying out another task. Results thus support our hypothesis.

Figure 4b shows two things. First, agents tackling tasks that rely on the same properties achieve a higher accuracy on their best task, compared to agents tackling tasks that rely on different properties. This indicates that while the agents may abstain from some tasks, their ontologies contain multi-purpose knowledge, acquired during the initial ontology induction phase. This further supports our hypothesis. Second, when tasks rely on different properties, the effect of the number of tasks assigned to each agent on the accuracy for its best task, is statistically insignificant ($p > 0.05$). This indicates that when tasks rely on different properties, learning to decide with respect to one task is not related to learning to decide with respect to a different task.

Figure 4c shows that tackling less tasks or having tasks that rely on common properties improves the success rate. This is due to two reasons. The first is that the fewer the assigned tasks, the fewer are the decisions over which agents need to agree. The second is that the more tasks rely on common properties, the less non relevant knowledge may be present to an agent’s initially induced ontology. Furthermore, while success rate improves over the course of the experiment, it does not converge to 1. This indicates that the final ontologies do not allow agents to reach consensus. This can be explained by the limitation of resources: agents may lack the resources required to learn to decide accurately for all assigned tasks and objects. As a result, they are able to decide accurately for different subsets of the existing object types at a given time. Thus, unless the different subsets coincide for all agents, consensus cannot be achieved. The latter is true even when agents interact over the same single task.

6.1 Statistical analysis

Analysis of variance shows that the number of common properties among different tasks, has a statistically significant impact ($p < 0.05$) on all measures. The number of assigned tasks has a statistically significant impact on (1) the success rate and (2) the average accuracy. When tasks rely on common properties, the latter has a statistically significant impact on the agents accuracy on their best task.

7 Conclusion

We hypothesize that agents tackling tasks that rely on common properties, benefit from the formation of multi-purpose knowledge. We test this hypothesis by introducing agents that learn and evolve ontologies with respect to several tasks. The experimental results support this hypothesis. On the one hand, it is shown that when agents tackle tasks that rely on common properties, knowledge is transferred from one task to another. On the other hand, when these tasks rely on different properties, tackling additional tasks does not affect the agents accuracy on their best task. Thus, deciding between tackling one or several tasks depends on the agents objective and the environment setup. The agent objective corresponds to whether they seek to optimize their accuracy on average or on their best task. The environment setup corresponds to whether the tasks depend on common properties or not. The experiments rely on minimal hypotheses about the environment, hence the results apply to a wide range of environment. These may serve as an insight on how agents evolve their knowledge within more complex environments. For example, one may consider environments where some tasks share properties, while other tasks are completely independent.

Acknowledgements

This work has been partially supported by MIAI @ Grenoble Alpes (ANR-19-P3IA-0003).

Data availability

The cultural evolution simulator used for our experiments can be found in [1]. Settings, results and the data analysis notebook are available in [8].

References

1. Lazy lavender. <https://gitlab.inria.fr/moex/lazylav>. (2023)
2. Baader, F., Calvanese, D., McGuinness, D., Nardi, D., Patel-Schneider, P.F. (eds.): The Description Logic Handbook: Theory, Implementation, and Applications. Cambridge University Press (2003)
3. Bourahla, Y., Atencia, M., Euzenat, J.: Knowledge improvement and diversity under interaction-driven adaptation of learned ontologies. In: Proc. 20th ACM international conference on Autonomous Agents and Multi-Agent Systems (AAMAS), London, United Kingdom. pp. 242–250 (2021)
4. Fan, J., Gao, Y., Luo, H.: Integrating concept ontology and multitask learning to achieve more effective classifier training for multilevel image annotation. IEEE Transactions on Image Processing **17**(3), 407–426 (2008)
5. Ghalwash, M., Yao, Z., Chakraporty, P., Codella, J., Sow, D.: Phenotypical ontology driven framework for multi-task learning. In: Proceedings of the Conference on Health, Inference, and Learning. p. 183–192. CHIL '21, New York, USA (2021)
6. Harding Graesser, L., Cho, K., Kiela, D.: Emergent linguistic phenomena in multi-agent communication games. In: Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP). pp. 3700–3710. Hong Kong, China (2019)
7. Havrylov, S., Titov, I.: Emergence of language with multi-agent games: Learning to communicate with sequence of symbols. In: 5th International Conference on Learning Representations (ICLR 17, workshop track). Toulon, France (2017)
8. Kalaitzakis, A.: 20230505-MTOA experiment description (2023), <https://sake.re/20230505-MTOA>
9. Lewis, D.K.: Convention: A Philosophical Study. Cambridge, MA, USA: Wiley-Blackwell (1969)
10. Mao, C., Gupta, A., Nitin, V., Ray, B., Song, S., Yang, J., Vondrick, C.: Multitask learning strengthens adversarial robustness. In: Vedaldi, A., Bischof, H., Brox, T., Frahm, J.M. (eds.) Computer Vision – ECCV 2020. pp. 158–174. Cham (2020)
11. Steels, L.: What triggers the emergence of grammar? In: AISB'05: Proceedings of the Second International Symposium on the Emergence and Evolution of Linguistic Communication (EELC'05). pp. 143–150. Hatfield, United Kingdom (2005)
12. Thomas, J., Santos-Rodriguez, R., Anca, M., Piechocki, R.: Multi-lingual agents through multi-headed neural networks. vol. 4. Tromsø, Norway (2023)
13. Wang, J., Gasser, L.: Kmutual online ontology alignment. In: Proc. 1st ACM international conference on Autonomous Agents and Multi-Agent Systems (AAMAS), Bologna, Italy (2002)