THÈSE

Pour obtenir le grade de

DOCTEUR DE L'UNIVERSITÉ GRENOBLE ALPES

École doctorale : MSTII - Mathématiques, Sciences et technologies de l'information, Informatique
Spécialité : Mathématiques et Informatique
Unité de recherche : Laboratoire d'Informatique de Grenoble

L'évolution culturelle des connaissances dans les agents cognitivement limités : précision, spécialisation et avantages

Cultural evolution of knowledge within cognitively-restricted agents: accuracy, specialization and benefits

Présentée par :

Andreas KALAITZAKIS

Direction de thèse :

Jérôme EUZENAT — Directeur de thèse
DIRECTEUR DE RECHERCHE, INRIA CENTRE GRENOBLE-RHONE-ALPES

Rapporteurs :

OLIVIER BOISSIER
PROFESSEUR, ECOLE NAT SUP MINES ST ETIENNE
ZAHIA GUESSOUM
MAITRESSE DE CONFERENCES HDR, UNIVERSITE DE REIMS - CHAMPAGNE ARDENNE

Thèse soutenue publiquement le **29 novembre 2024**, devant le jury composé de :

JULIE DUGDALE, — Présidente
PROFESSEURE DES UNIVERSITES, UNIVERSITE GRENOBLE ALPES
JERÔME EUZENAT, — Directeur de thèse
DIRECTEUR DE RECHERCHE, CENTRE INRIA UNIVERSITE GRENOBLE ALPES
OLIVIER BOISSIER, — Rapporteur
PROFESSEUR, ECOLE NAT SUP MINES ST ETIENNE
ZAHIA GUESSOUM, — Rapporteure
MAITRESSE DE CONFERENCES HDR, UNIVERSITE DE REIMS - CHAMPAGNE ARDENNE
MAXIME MORGE, — Examinateur
MAITRE DE CONFERENCES HDR, UNIVERSITE LILLE 1 - SCIENCES & TECHNOLOGIES

Acknowledgement

With these lines, a journey comes to an end, a journey that has outlasted the duration of this PhD. During this journey, I was rarely alone. For the most part, I was accompanied, encouraged, and supported by key people during pivotal moments, many of whom shaped my path in ways that I only fully appreciated later.

Among these, my director, Jérôme Euzenat, stands out as a central figure. His profound erudition, his integrity as a researcher, and his inherently uncompromising nature were a driving force in my growth, as I strived to convince him and justify my results. I am particularly grateful because, regardless of all the ups and downs that came with this thesis, Jérôme's door always remained open, ready for discussion or to review unnecessarily long and complex sentences. I complete this journey having gained a better understanding in the field but, most importantly, research methodology and ethos.

Equally instrumental in shaping this thesis were the manuscript's reviewers, Olivier Boissier and Zahia Guessoum. Their expertise, insight, and constructive feedback were pivotal, not only in the completion of this thesis but also in helping me gain perspective. I am equally grateful to the defense committee members, Julie Dugdale and Maxime Morge, for the interesting and engaging discussion that followed the PhD presentation, as well as for the opportunity to present my work before scholars who are experts in their respective domains.

I would like to thank all current and previous mOeX members—some present in body, others in spirit. Adriana, Alban, Helga, Linda, Line, Khadija, Yasser: we shared not only an office but also our struggles, reliefs, and excitement. Our numerous discussions on the cultural evolution of knowledge representations and prior work in this emerging field significantly influenced the early stages of my PhD.

As mentioned before, the extent of someone's contribution is often invisible at the moment. Such is the contribution of Professor Paraskevi Fragopoulou. Without her early research stimuli and the opportunities she provided, it is questionable whether I would have followed this academic path.

One of the most visible contributions in everyday life comes from friends. They are the ones who offer acceptance, without any judgment. Julia, Despoina, Milto, Thoma, Niko, Mimi, Theodore, your presence, whether in person or from a distance, made this journey significantly easier. I would also like to extend my gratitude to my friends and students, who honored me with their presence at my PhD defense.

In the end, the only people whose contribution is even greater, are those who raised you and the person with whom you share your life. Parents' unconditional love, encouragement, and support for their child's achievements,

regardless of how trivial, are what shape the future. They are the ones who draw the bow so that the arrow may later fly straight. This journey is completed today, but in truth, it began decades ago. I will never forget their commitment to providing me with an education that would allow me to lead a better life than the one they led. Still etched in my memory is the first office they purchased for me at the age of six, when they could barely afford more crucial necessities. I cherish those years and hope that I have flown at least as high as you expected and hoped for. Ultimately, there is only one person who stands by you when you flourish and is also there to pick up the pieces when you fail miserably. The person who believes in you even when you don't believe in yourself. I am infinitely grateful not only for sharing my life with Sofia (Dr. Kleisarchaki) but also for sharing this PhD journey with her. I dedicate this PhD to Sofia, truly believing that it is as much hers as it is mine. After all, it was her steps I tried (in vain) to follow academically since our first days in undergraduate computer studies. Her drive for excellence and her enthusiasm for our discipline were an example that was impossible to ignore.

".. you can't connect the dots looking forward; you can only connect them looking backward. So you have to trust that the dots will somehow connect in your future. You have to trust in something — your gut, destiny, life, karma, whatever." (Steve Jobs, 2005) With the passage of time, I can now see how, based on luck, effort, support, and encouragement, these dots were connected. May the end of this journey mark the beginning of an even more exciting one.

Contents

1 Introduction **1**
1.1 Context 1
1.2 Thesis synopsis 2
1.3 Key contributions 5
1.4 Experiments reproducibility 6
1.5 Manuscript outline 6

2 Specialization of cultural traits in artificial populations **7**
2.1 Social learning and cultural evolution 7
2.2 Agents and multi-task learning 10
2.3 Experimental cultural evolution 13
2.4 Cultural evolution of languages 20
2.5 Cultural evolution of unipotent knowledge 22
2.6 Specialization of multipotent traits 25
2.7 Conclusion 28

3 A framework for adaptive resource-restricted multipotent agents **29**
3.1 Introduction 29
3.2 Experimental framework 30
3.3 Experiment outline 31
3.4 Conclusion 36

4 Knowledge transferability among different tasks **39**
4.1 Experimental setting 40
4.2 Results and discussion 43
4.3 Statistical Analysis 44
4.4 Conclusion 44

5 Knowledge specialization in resource-limited agents **49**
5.1 Experimental setting 50

5.2 Results and discussion 55
5.3 Statistical Analysis 60
5.4 Conclusion 64

6 Influence of knowledge forgetting on long-term correctness 65
6.1 Experimental setting 66
6.2 Results and discussion 67
6.3 Statistical Analysis 70
6.4 Conclusion 71

7 Task exploration in resource-limited agent populations 77
7.1 Experimental setting 78
7.2 Results and discussion 81
7.3 Statistical Analysis 86
7.4 Conclusion 89

8 Conclusion 95
8.1 Summary 95
8.2 Contributions 96
8.3 Perspectives 96

Bibliography 99

— 1 —
Introduction

1.1 Context

Artificial agents are often entrusted with tasks that, until recently, were carried out by humans. These tasks range from content generation to law enforcement, healthcare applications and more. In order to accomplish them, agents take different decisions based on how they perceive objects and what they have learned about them so far. For example, a healthcare agent may have learned that an elevated heart rate is a sign of acute hypoxic exposure, while a law enforcement agent may have learned that the same symptom corresponds to a usage of illegal substances. Given their individual knowledge, some agents will thus be more suitable for accomplishing some tasks, while other agents will be more suitable for accomplishing different ones.

Yet, this may change over time. Like humans, agents may revise their individual knowledge based on the feedback they receive from the environment and social transmission. For example, interacting agents may exchange knowledge by teaching one another. Alternatively, an agent may revise its knowledge by observing the consequences of other agents' decisions. Consequently, an agent may unconsciously adopt behavioral traits of agents accomplishing different tasks, pushing this agent to reconvert to a different task.

How are knowledge traits distributed, when different tasks compete over agent resources? At one end of the spectrum, one may conceive agent societies consisting of agents accomplishing several tasks such as farming, fishing, hunting and craftsmanship. It is expected that by tackling several tasks, these agents may benefit from developing multi-purpose knowledge. This would in turn increase their facility to address different, potentially unforeseen tasks. Skills developed for hunting, may be reused for fishing and vice-versa. On the other end of the spectrum, societies consisting of highly specialized agents with complementary skill-sets can be conceived. In such societies, some agents become hunters, some fishermen and some cooks. It is expected that highly specialized agents will excel in a few tasks, relying on other agents for the remaining tasks. Based on these, mixed societies in which highly specialized agents co-exist alongside multi-purpose agents may also be conceived. The main objective of our thesis is to examine whether agents specialize, and if so, whether this is a natural tendency or the result of imposed economic or social conventions. We decompose this objective into three main questions. First, do agents transfer knowledge from one task to another? Second, does assigning different tasks to different agents lead to agent specialization, and if so, is this beneficial for the societies they belong to? Third, is there any agent reproduction strategy allowing for equitable task exploration, i.e., balanced collective advancement in all tasks?

One way to answer these questions is through cultural evolution, specifically through cultural evolution of knowledge. Cultural evolution is the theory of evolution, applied to cultural traits. A cultural trait is a distinct feature that is shared and evolved within a population. It can be tangible, like a tool, a weapon, a piece of clothe or intangible, like a language, a norm, a belief. Intangible cultural traits are essentially learned behaviors passed from one generation to the next one, allowing individuals within a population to interact, communicate and understand one another. These traits can be furtherly distinguished based on whether they are unipotent or pluripotent. As unipotent, we consider a cultural trait that serves a single purpose, e.g., a norm that is valid only on a special occasion. As pluripotent, we consider a cultural trait that can serve several purposes, e.g., knowledge applying on different tasks.

Cultural evolution of knowledge uses agents to study how knowledge representations evolve, when pressure is exerted on agents for achieving a common goal. This goal may be survival, maximizing diversity or other. Here we focus on agents whose goal is to agree with each other.

In this context, knowledge traits can be considered being equivalent to genotypes and agent behaviors being equivalent to phenotypes [32]. As such, knowledge traits not only affect how agents decide, but also constitute artefacts that can be horizontally (intra-generational transmission) or vertically (inter-generational transmission) transmitted. Essentially, knowledge exchange leads to knowledge revision, which introduces knowledge variations upon which natural selection can act. Ultimately, the cultural evolution of knowledge allows us to identify which knowledge traits persist over consecutive generations and how these traits are distributed within agent populations.

1.2 Thesis synopsis

It has been previously shown that, by trying to agree, agents not only reach consensus, but also improve the quality of their knowledge [11, 12, 92]. These works examine agents that tackle the same task and have unlimited memory. However, realistic agents have limited memory and may accomplish several tasks. Our goal is to understand how accomplishing several tasks affects the knowledge acquired by a population of agents. To that end, we represent knowledge using ontologies and study their cultural evolution over three dimensions: (1) the number of tasks assigned to agents, (2) the size of their memory and (3) the overlap among the different tasks.

In Chapter 4 we consider agents sharing a growing set of decision tasks. These tasks consist of reasoning over different object properties in order to take abstract decisions within different abstract domains. Our hypothesis is that multi-tasking agents tackling tasks that rely on the same properties (100% overlap) are more accurate than multi-tasking agents tackling tasks that rely on different properties (0% overlap). We tested this hypothesis by varying the number of tasks assigned to the agents and whether these tasks overlap or not. Results show that when deciding for different tasks relies on the same properties, multi-tasking agents are significantly more accurate. This suggests that it is possible to transfer knowledge from one task to another, or put differently that by farming agents can become better at cooking.

In Chapter 5 we consider agent populations assigning different tasks to different agents. The agents are initially trained using labeled examples corresponding to all tasks. Subsequently, each agent receives individual compensation for accomplishing a limited subset of tasks. Our assumption was that by trying to specialize on fewer tasks, agents will improve their maximum accuracy. Hence this should increase the compensation collectively acquired by a population of agents. We have shown this to be partially true. Contrary to our assumption, results show that the number of tasks assigned to agents has no effect on their maximum accuracy. However, it is shown to affect not only the total compensation acquired by agent populations, but also the way this compensation is distributed to the individual agents. Given a set of non overlapping tasks, the fewer the tasks assigned to each agent, the higher and more equitably distributed the compensation acquired by them. Further examination of agent ontologies showed findings in [57] to be a consequence of parasitic knowledge. As parasitic, we define the knowledge that occupies an agent's memory while being unrelated to the tasks this agent accomplishes. Depending on their remaining memory, some agents will succeed in learning accurate ontologies for the task they accomplish and some other will not. As a result, the former agents cannot converge on their decisions with the latter ones. When agents have limited resources, learning to agree and specializing on fewer tasks are mutually exclusive. Agents can either improve their communication or their individual maximum accuracy. This was found to be true even for mono-tasking agents that share the same task and their memory is enough for accommodating the knowledge required for 3 tasks.

In Chapter 6, we deal with parasitic knowledge by comparing two knowledge removal approaches. The first approach, used in [56, 57], consists of removing leaf classes that (a) share the same parent and (b) are associated to the same decisions for the tasks the agents undertake. This method removes some parasitic knowledge, without affecting agents accuracy on the tasks they undertake. However, it does not remove parasitic knowledge that was acquired early, e.g., during the agents initial training. The second approach consists of removing randomly selected ontology classes. This approach is agnostic about whether the removed knowledge is parasitic or not and thus may affect the agents' accuracy on any task, carried out or not. When agents accomplish one task, this approach improves their maximum accuracy as it allows for not only relearning previously removed knowledge, but also acquiring new knowledge for which there was not enough memory. When agents accomplish several tasks, this approach is detrimental to both their average and maximum accuracy. Results show that when agents remove randomly selected classes, consensus will be reached regardless of the number of tasks they carry out. However, when agents accomplish several tasks, this approach leads to a gradual deterioration of the agents' accuracy. This deterioration will eventually stop, with the agents ultimately agreeing based on low accuracy knowledge.

We showed that memory-limited agents will specialize regardless of how many tasks they accomplish. Additionally, we showed that fixed size agent populations do not benefit from specialization, unless each task is assigned to an equal number of agents. When each task is not assigned to an equal number of agents, a population will collectively become more accurate on a single task. This will thus render agent populations very efficient at some tasks and very inefficient at others.

This is due to two reasons. First, the fewer the agents confronting their knowledge for a task, the fewer the adaptations that take place. This limits the exploration of more successful ways of dealing with this task. Second, the less agents contribute on a task, the less knowledge variation about this task exists. As a consequence, the less competition will take place among different decisions and the less successful will the selection be in promoting accurate decisions. For example, let an agent population relying on cooking and farming but having less cooks than farmers. Given this task distribution, this population will be very efficient at growing crops but very inefficient in transforming harvest into food. Thus, in order to satisfy its needs in food, this population will need to either resort in exchanging goods with other populations, or make its cooks work more.

In [56, 57], this is achieved by a priori assigning each task to an equal number of agents. However, tasks cannot always be equitably distributed, as new tasks may be introduced at any moment. Can a population advance its knowledge on all tasks, without a priori assigning each task to an equal number of agents? We have previously established that knowledge improvement continues while agents confront their knowledge with that of other agents. When disagreements no longer occur, the agents' accuracy stabilizes. This signifies that what could be learned through social transmission, has already been learned. As a consequence, agents with a higher than average success rate have lower potential in promoting new decisions than agents with a lower than average success rate. The presence of agents with consistently lower than average success rates, may arise from the fact that most agents tackle over-explored tasks, while different few agents tackle the under-explored ones. Our hypothesis is that the agents that have the highest impact in advancing their populations knowledge, are the ones that have a lower than average success rate. These are the agents that possess rare knowledge. Our assumption is that its rarity reflects that it is shaped for an under-represented task, i.e., a task on which the majority of agents is not yet competent. This task could thus benefit from additional knowledge confrontation. It has been argued that while researchers are conservative when they choose which areas to explore, focusing on under-explored areas is more likely to achieve high impact [37, 80].

In chapter 7, different reproduction strategies are considered. Our hypothesis was that the reproduction of agents with low success rate, i.e., agents possessing rare knowledge, will allow their societies to restore and maintain the balance in task distribution. This hypothesis was based on the assumption that at each generation, over-explored tasks will be assigned to less agents and under-explored tasks will be assigned to more agents. Eventually, this will lead to an equilibrium where the probability of producing off-springs will be equal for all agents, i.e., there will be no marginal agent sub-population. Results reject this hypothesis. They show that the vast majority of agent populations will inevitably converge into a single specialization. Favoring the reproduction of less successful agents was shown to delay and at times halt this convergence. However, this leads to a significant deterioration of knowledge correctness. Agents do not improve their accuracy on the tasks on which they were previously inefficient. On the contrary, they deteriorate their accuracy on the tasks they previously excelled.

1.3 Key contributions

First, we show that when an agent carries out a task, this affects its capacity to carry out a different task. The higher the number of tasks assigned to individual agents, the higher is their average accuracy. The more these tasks overlap in terms of required capacities, the more the agents tackling them improve their average accuracy.

Second, we show that specialization emerges when agent resources are restricted and not when less tasks are assigned to individual agents. Specialization was not found to increase the agents' maximum accuracy. Yet, tackling less tasks proved beneficial for agent populations. The less tasks assigned to individual agents, the higher and more equitably distributed is the wealth collectively collected by agent populations. However, the benefits identified during the second contribution were found to be valid only when (1) the agent populations and tasks remained constant for the duration of the experiment and (2) each task was assigned to an equal number of agents.

Last, we show that the natural tendency of agent populations is to specialize in the same task, i.e., the majority of agents will excel on one task. Depending on which agents and thus which knowledge is selected for vertical transmission, convergence on a single task may be delayed but not avoided.

1.4 Experiments reproducibility

All the experiments in this thesis are reproducible. Source code is stored in an open repository [61] and experiment results have been fully logged and are publicly available [49–55].

1.5 Manuscript outline

Chapter 2 has a dual objective. First, it provides fundamental preliminaries regarding adaptive artificial populations. Second, it examines how multi-potent populations acquire, collectively evolve and specialize individual knowledge representations. To that end, several examples from artificial and living populations are presented.

Chapter 3 describes in detail the experimental framework used throughout this thesis. It is divided into two parts. The first part consists of all the definitions regarding the entities that constitute the environment, alongside with the notation describing it. The second part consists of the outline of the general experiment, used throughout all our experiments. It describes how agents are trained, are assigned tasks, interact, remove knowledge and adapt their ontologies.

The remaining of the this manuscript is organized over 5 chapters. Four of these chapters correspond to the experiments carried out. Chapter 4 demonstrates how knowledge is transferred from one task to another, when these tasks overlap. It presents a set of hypotheses, the parameters used to test them, as well as the results showing that multi-tasking is beneficial for artificial populations. Chapter 5 shows how memory limitations and not selective task accomplishment, pushes agents to specialize their knowledge. Likewise Chapter 4, different hypotheses are presented and tested, showing collective but not individual benefits arising from specialization. In Chapter 6, the problem of parasitic knowledge is tackled. In particular, Chapter 6 presents an experiment that compares different methods for removing parasitic knowledge. Its results highlight the long-term benefits that arise from the short-term removal of pertinent knowledge. Chapter 7 demonstrates that realistic artificial populations of multi-potent agents fail to cover all their needs unless social and economic conventions are imposed.

Finally, in Chapter 8 a conclusion of this thesis is provided, consisting of its summary, key contributions and arised perspectives.

— 2 —

Specialization of cultural traits in artificial populations

This thesis relies on experimental cultural evolution to study how agents specialize cultural traits over consecutive generations. In this chapter, we review literature dealing with: (1) how the transmission of cultural traits leads to behavioral change, (2) when this is considered as cultural evolution, (3) how cultural evolution is simulated, (4) how languages and (unipotent) knowledge evolve and (5) how artificial and natural populations benefit from the formation and specialization of (multipotent) traits.

Section 2.1 is divided into two subsections. Subsection 2.1.1 concentrates on what social learning is, how it affects natural populations and why it constitutes an alternative for multi-agent systems learning. Subsection 2.1.2 provides a definition for cultural evolution and provides empirical evidence for the presence of such processes in nature. In Section 2.2, a brief introduction to adaptive agents and multi-task learning is provided. Section 2.3 introduces experimental cultural evolution, presenting several experiment design choices, population topologies and interaction modalities. Section 2.4 and Section 2.5 focus on the cultural evolution of conceptual artefacts. The first one focuses on the cultural evolution of languages and the second one focuses on the cultural evolution of uni-potent knowledge. In Section 2.6, literature concerning specialization is reviewed. It is divided into two subsections. Subsection 2.6.1 discusses empirical studies on how humans and animals specialize. Subsection 2.6.2 discusses literature regarding model-based agent specialization. This chapter's conclusion is provided in section 2.7.

2.1 Social learning and cultural evolution

Knowledge specialization within the cultural evolution of knowledge can be studied in two ways: either by using agents interacting and learning from one another, or through model-based simulations. Since our focus is on knowledge, we chose to base our work on agents coordinating while carrying out tasks. To that end, literature about social learning and its enduring effects through cultural evolution is reviewed here.

2.1.1 Social learning

With the exception of solitary species, breathing organisms tend to either take actions collectively, or alter their behavior based on observing others. Their decisions are thus a result of individually acquired knowledge, negotiation, teaching and imitation to name a few [13, 47, 78, 87]. Wolfs may coordinate on how to take down a bear and a lizard may learn that a praying mantis is a force to be reckoned. This can, for example, be based on observing the fatal outcome of another lizard trying to devour one. While few works concentrate on what triggers social learning, extensive literature exists on how teaching and imitation [88, 93] affect learners [14, 59].

In [47], Galef and Heyes review the most impacting works from the 80s and 90s within the field of study of animal social learning. Through this review, not only they assess the maturity of this rapidly growing field, but also predict the direction research will follow within this area. They argue that what was until then of interest to a limited niche of experimental psychologists and biologists, will attract scientists from a multitude of other disciplines. These include anthropology, behavior genetics, economics and robotics to name a few.

[78] provides a survey on different experimental approaches followed by researchers studying animal social learning, both in the wild and controlled (laboratory) environments. Through this work, Reader and Biro highlight that these environments should be considered as complementary. That is because studies conducted in the wild often lack sufficient control over the learning process, while studies conducted in captivity often fail to replicate empirical data due to not capturing vital environmental conditions. While our work is not directly related to animal social learning, the argument regarding the validity of "in vitro" social experiments is something that should be of our concern as well. Conclusions drawn from agent-based experimentation, while valid under the examined conditions, may not be generalizable to the agents' living counterparts. Simply put, what is valid for agents is not necessarily valid for humans and other living beings.

[87], through the study of meerkats, supports that social learning is equally present in non-primate mammals. Thornton and Clutton-Brock show that the collective adoption of socially available information depends on the context and associated costs. In particular, that social learning leads to the formation of group-wide, persisting traditions. These are cultural traits that persist, i.e., that can be persistently traced over time. They support that this happens when developing a skill is intellectually demanding, or the deviation from a norm is associated with a high cost/risk (e.g., high presence of predators). Lastly, they argue that persisting

traditions may modify the selection pressures and thus have an effect on how genes evolve. We consider their findings interesting and relevant to the cultural evolution of knowledge. First, they demonstrate the reciprocal influence between cultural evolution and genetic evolution. Second, they associate the adoption of a cultural trait with the complexity of a task.

Likewise nature, (agents') environment conditions are subject to change both spatially and temporally. Thus, individual agents cannot a priori learn behaviors for all possible situations. This is especially true for static agent environments, i.e., environments that do not change over time. To that end, [2] argues that artificial populations (agents) may use different ways to learn how to perform a task. [70] shows experimentally that multi-agent systems can improve the performance of their agents, by extending beyond individual learning. On the one hand, logic-based agents can take rational decisions based on incorporating domain knowledge. On the other hand, nature inspired agents may learn how to accomplish tasks based on observing other agents or interacting with them. For example, an agent may learn to avoid an action after observing a penalty previously attributed to another agent. Our framework relies on both methods, by implementing logic-based agents capable of altering their behavior, based on feedback they receive from other agents.

2.1.2 Cultural evolution

While different species have demonstrated adaptation triggered by social learning [62, 87, 94, 95], opinions differ as to whether social learning necessarily implies cultural evolution.

Often this disagreement is due to the lack of explicit cumulative cultural evolution definition. [68] highlights this ambiguity, by reviewing how this term is used by researchers and provides a core definition of it. According to Mesoudi : "the minimum requirements for a population to exchibit CCE (cumulative cultural evolution) are (i) a change in behaviour (or product of behaviour, such as an artefact), typically due to asocial learning, followed by (ii) the transfer via social learning of that novel or modified behaviour to other individuals or groups, where (iii) the learned behavior causes an improvement in performance, which is a proxy of genetic and/or cultural fitness, with (iv) the previous threesteps repeated in a manner that generates sequential improvement over time".

This sequential improvement over time is discussed in [8]. The latter argues that an emerged collective behavior or performance, is only a snapshot within an enduring evolutionary process. In particular, it supports that each such snapshot depends on the collective's previous memory and history of actions. [36] examines the variation of cultural traits, specifically how complex human creativity leads to the creation of new cultural traits. Drawing insights from various disciplines (cognitive science, psychology, archeology to name a few), Fogarty proposes a mathematical model for deliberate-creativity. The latter takes as input a number of problems seeking solutions and the existing cultural traits within a population. Given these, it calculates the expected number of new ideas and traits, generated by deliberate-cognitive creativity.

These traits can be transmitted horizontally [23, 42], vertically or obliquely and the effects of this transmission may be observed several generations later. [79] discusses the role of social learning in cultural evolution and human adaptability, allowing humans to develop and evolve complex skills and behavioral traits they perfect over several generations.

Sometimes cultural evolution is explained using the key principles of the genetic evolution. Yet, while in genetic evolution genetic information is only transmitted during reproduction, cultural information may be adopted, dropped and re-adopted at any time. [85] champions that cultural evolutionists must abandon the idea that the success and proliferation of a cultural trait may be predicted by a measure similar to that of biological fitness. Additionally, contrary to cultural traits, genes cannot be accumulated by an individual. Yet, cultural accumulation is what probably allowed humans, to tackle tasks of increasing complexity.

[86] proposes a model of cultural accumulation, taking into account an individual's contribution in its population's total amount of cultural traits. This model relies on three assumptions. First, that culture is transmitted in a single learning event per individual. Second, that cultural evolution does not affect genetic evolution and vice versa. Third, that acquiring a cultural trait makes nor less, not more propable to acquire another cultural trait. Results show two things. First, that a population's structure does not affect the total accumulated culture. The latter is shown to be proportional to the population's size. Second, that cultural accumulation within an individual is not affected by the population's size.

Researchers working on cultural evolution often use mathematical models or experimental simulations to test their hypotheses. [71] argues that while social learning in nature can be extremely complex, models may still be adequate for explaining the behavior of individuals. Using such models, [30] compares the adaptiveness of individuals that use individual learning and individuals that rely on social learning.

2.2 Agents and multi-task learning

2.2.1 Agents

Agents are autonomous software entities that accomplish tasks, often designed in the image of human reasoning. Demazeau defines agents as a software or physical entities, capable of accomplishing a certain attributed mission in an autonomous manner or in cooperation with other agents [17]. They can be classified into two main categories: deliberative agents and reactive agents [97].

Reactive agents are agents with simple internal representations, whose actions depend solely on the reflexive reaction to an external stimuli. Their domain of application spans from temperature control systems to security systems and robotic vacuum cleaners. For example, in the case of robotic vacuum cleaners, the stimuli may be an impact and a reflexive reaction a backward motion followed by a turn of 90 degrees. Among the simplest ways to implement such agents is through predefined condition action rules, i.e., rules in the form: if condition then action. When the condition becomes true, the appropriate predefined action is taken. Another way to implement a reactive agent is through finite state machines (FSM).

Contrary to reactive agents, the internal representations of deliberate agents are complex. According to Wooldridge [96], a deliberative agent is an intelligent entity that possesses an explicitely represented symbolic model of the world, and in which decisions (for example about what actions to perform) are made via symbolic reasoning. Deliberative agents are thus agents who possess knowledge about their surrounding environment and are able to plan their actions accordingly. This makes this class of agents ideal for multi-agent system simulation experiments. Deliberative behaviours are often implemented using the belief–desire–intention software model [16, 76, 77]. According to BDI, agents' functions are organized into 3 components: beliefs, desires and intentions. Beliefs encapsulate what an agent believes about its environment, alongside what it can perceive using its sensors. Desires describe what this agent's objectives are, intentions describe strategies or sequences of actions in order to achieve these objectives.

A common critisism for BDI agents is the absence of mechanisms allowing agents to learn from past behaviour or other agents and thus adapt to new situations [41, 73]. One extension of the classic BDI model allowing for agents learning new behaviours is JaCaMo [9, 10].

This thesis focuses on the formation, evolution and specialization of multipurpose knowledge. We are thus naturally interested in a subclass of deliberative agents, adaptive agents, i.e., agents with the capacity to learn and evolve their reasoning. One agent type within this subclass is the reinforcement-learning agent, which through trial and error revises its behavior in order to maximize a reward function. Subsection 2.2.2 reviews literature using reinforcement-learning agents tackling several tasks.

However, since this thesis focuses on cultural evolution and agent-to-agent interaction, reinforcement learning is not required. Our definition of adaptive multi-tasking agents, based on symbolic reasoning and ontologies is provided in Chapter 3. Unless stated otherwise, the term agents will hereafter correspond to adaptive agents.

2.2.2 Multi-task learning

Multitask learning has been shown to significantly improve classification in a variety of areas, e.g., adversarial robustness [64], visual interconceptual similarity [33], phenotype learning [39].

Machine learning models achieve high accuracy on a range of computer vision benchmarks. Yet, they remain vulnerable to adversarial attacks, when imperceptible input perturbations are purposely introduced to the model. Adversarial robustness refers to the ability of a learning model to remain unafected by such decoys. With artificial intelligence being increasingly present in time critical applications, e.g., autonomous vehicles, improving the stability and robustness of machine learning models is of paramount importance. As a result, it remains an ongoing challenge in machine learning research. To that end, [64] empirically and theoretically supports that the number of classification tasks a model is trained on, is related to its robustness. In particular, experiments on two real-world datasets show that the higher the number of classification tasks a network is trained on, the more complex it becomes to successfully attack it. Results thus suggest (1) that increasing the number of labeled examples does not necessarily guarantee adversarial robustness and (2) that current deep networks remain vulnerable because they train on too few classification tasks.

[33] tries to minimize the number of classifiers required. The proliferation of digital cameras and the role social media play in human expression, have led to an exponential growth of the number of pictures available online. The latter makes searching for an image within such large datasets difficult. Without efficient and intuitive keyword search, this task becomes increasingly difficult, if not impossible. It is thus crucial to improve multilevel image annotation. This is important as an image may contain several levels of semantics. Annotations should be able to exploit all local and global image features, along with any existing semantic connection among different concepts. A greedy solution to the problem is to train a high number of independent classifiers. In this case, each classifier will be trained to detect a specific (local or global) feature of the image in question. The objectives in [33] are three.

First, to exploit the semantic connections that exist within different concepts. Second, to detect the semantic connection between images, when these images dispose very diverse local or global features. Third, to distinct between images that are semantically un-related, yet sharing local or global features. To meet these objectives, the authors proposed a novel, ontology-based scheme for achieving automatic hierarchical classification. Results show that the proposed scheme achieves not only to improve on image classification, but also to successfully bridge the semantic gap.

Features and concepts are often under-represented within large datasets. For example, despite the large number of patients present within electronic health records, little data exists concerning patients that suffer from Mitral valve stenosis. When data is noisy, sparse or non-uniformly distributed, predicting a medical outcome can be difficult. In order to improve the efficiency of prediction tasks under such conditions, [39] follows an approach similar to that followed by [33]. In order to overcome the scarsity of specific concepts and features, multitask learning is enhanced using an ontology. This ontology, essentially a relationship graph, relates different phenotypes (concepts) and medical outcomes (predictions). Exploiting these relationships, allows for predicting medical outcomes based on phenotypes that have a higher presence within a dataset. Results obtained using the proposed ontology augmented framework, showed a significant improvement in learning performance. Using a publicly available medical record database, authors demonstrated the efficiency of their approach, achieving an increased efficiency on several real medical outcome predictions over different MTL schemes.

In [33,39,64], it has been demonstrated that solving tasks jointly using a single neural network, can improve prediction accuracy, data efficiency and training time. However, due to a phenomenon called "negative transfer", this is not always true. In this case, the performance of multitask classifiers can be significantly inferior to that of smaller independent networks. This can for example be because one task dominates the other during training, or because the learned features have high distinguishing capacity in one task but the exact opposite on a different one. Thus, whether multitask learning is beneficial or not, relies on the relationship between the jointly trained tasks.

[83] studies how tasks cooperate and compete during classifiers training. To that end, a framework is proposed for assigning different tasks to a small number of neural networks. Cooperating tasks are in this case computed by the same neural networks, while competing tasks are computed by different ones. Although the proposed framework offers a trade-off between time and accuracy, its training computation cost is often less appealing. In response to this, the authors present two distinct strategies. Results show that employing these strategies outperforms single-task networks as well as multitask networks with all tasks trained jointly.

2.3 Experimental cultural evolution

To understand cultural evolution, experimental methods have been employed to test hypotheses and predictions of cultural evolution models, providing insight into real world patterns. These experiments make the liaison between earlier theoretical models and real-world data (e.g., archeological, historical). Unlike theoretical models, experiments avoid the bias introduced by the researchers' intuitions and assumptions, on how individuals should behave within the cognitive mechanism of cultural evolution.

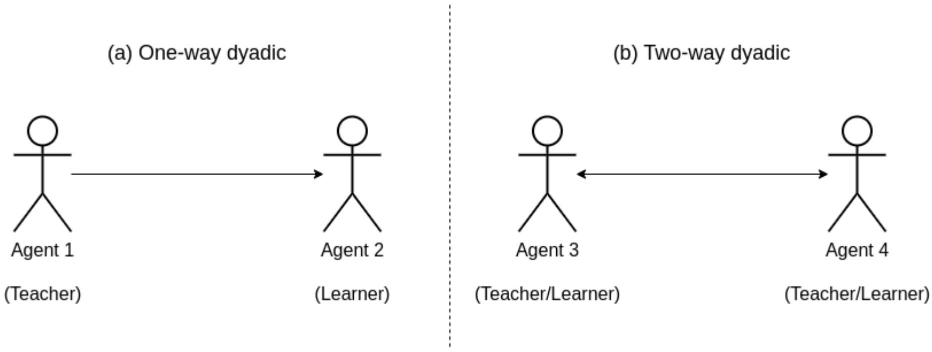

Figure 2.1 Dyadic experiments in multi-agent systems, illustrating interactions between two agents. The figure demonstrates both one-way and two-way communication between agents.

2.3.1 Experiment types

Different types of cultural evolution experiments have been proposed for different fields and applications. [67] offers a classification of these experiments, based on how cultural traits are transmitted from one individual to another.

Dyadic

Dyadic interactions : participants are placed in pairs, where the first individual is the learner and the second individual is the teacher. This kind of experiment is for example recurrent in the experimental study of languages and communication protocols. For example, [34] uses dyadic interactions to compare two types of language evolution. To that end, individuals participate into a graphical communication task. Who these participants interact with, depends on the transmission protocol. The transmission protocol corresponds to either iterative learning, or social collaboration. When participants evolve languages through iterative learning, linguistic traits are transmitted vertically within locked participants dyads. Simply put, one parent will transmit cultural information to its offspring. At each generation, a previously considered offspring, i.e., learner, becomes a parent. Thus, iterative learning slowly evolves languages in isolation, generation after generation. When participants evolve languages based on social collaboration, linguistic traits are also transmitted horizontally. Any participant can transmit cultural information to any other participant, as long as they are directly connected. As expected, this leads to participants rapidly converging to a single, publicly accepted language. Experiments in the cultural evolution of knowledge naturally fall within dyadic experiments, with knowledge being transmitted either horizontally between dyads of the same generation or vertically between dyads of consecutive generations.

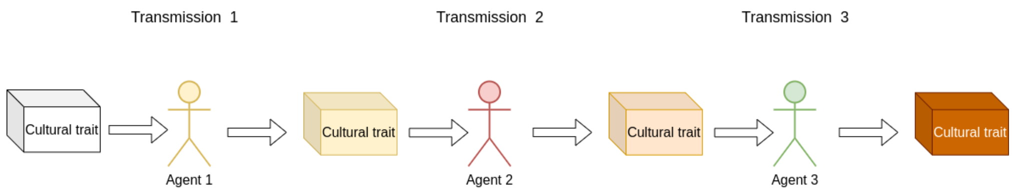

Figure 2.2 Transmission chain experiment illustrating the sequential transmission of cultural traits from one generation or agent to the next.

Transmission chains

In transmission chain experiments, individuals are positioned linearly one after the other, as if they were links within a chain. The first individual is given an artefact to complete or process. After this process, the altered artefact is passed to the second individual and so on. Several parallel transmission chains may exist. This method has been used to study content biases in the repeated transmission of written texts [5, 31, 48].

[5] uses 92 transmission chains, each one consisting of 4 randomly selected and assigned to their in-chain positions psychology students. The objective of this work is to directly compare the survival of negatively and positevely evaluated information, across several transmission episodes. Results suggest that a content transmission bias exists, when an original message is transformed before being re-transmitted. In particular, humans are shown to either interpret ambiguous information as negative, or favor the transmission of non-ambiguous negative information.

[48] follows a similar approach, conducting a recall-based transmission experiment. It uses 30 chains with 4 generations of participants that read, recall and pass descriptions of dominant, prestigious and social football players to the next generations. Results show a content bias towards prestigious and dominant profiles over highly social ones.

Closed groups

In closed group experiments, the population is organized into sub-populations and each individual is assigned to one. Likewise transmission chains, individuals are asked to complete tasks, e.g., create an artefact or take an action given a situation. This type of experiment is often used for (1) comparing the performance between agents that apply social learning and agents that learn individually and (2) identify when individuals choose to employ social learning instead of asocial learning. Examples of closed group experiments are [15, 29, 40, 65, 66, 89, 90].

[15] provides, through an online social experiment, empirical evidence on whether and when humans follow prestige biased social learning. It employs 225 participants organized into three groups, playing a two-round quiz game. When

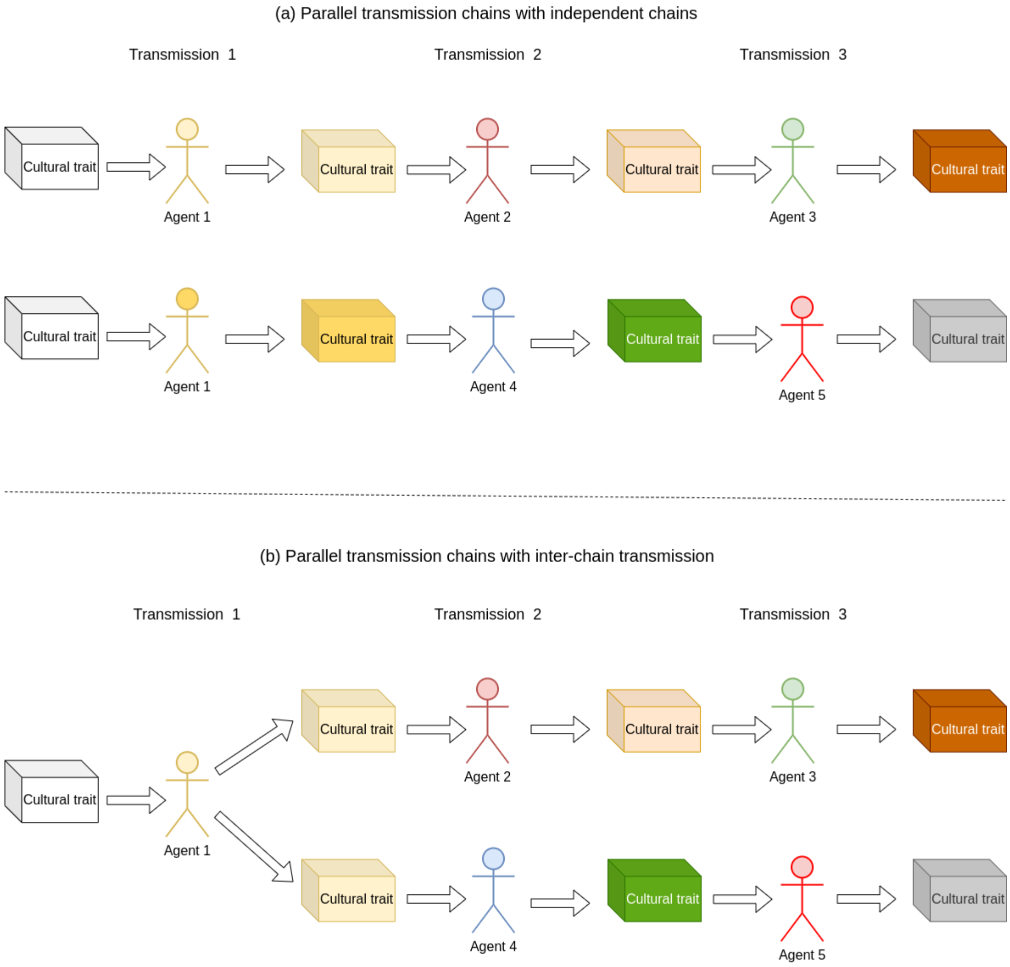

Figure 2.3 A parallel transmission chains experiment, illustrating multiple chains of sequential transmission of information or cultural traits within a population. These chains may be independent or dependent through the transmission of a cultural trait to more than one individuals belonging to parallel chains.

a participant does not know an answer, this participant may choose to copy the answer of another person of the same group. Depending on its group, selecting from whom to copy the answer may reside on different cues. These cues may be unrelated (control group), i.e., cues that are not systematically linked to performance or success, un-reliable (prestige group), i.e., the number of times someone has been copied or reliable (success group), i.e., the number of times someone has been successful. Results show that prestige biased learning may emerge as a social learning strategy, when direct success information is not available or un-reliable. This highlights the flexibility of the human species in adopting social learning strategies, leading to the acceleration of cumulative cultural evolution.

[29] discusses a type of transmission bias, conformity. The latter is an exag-

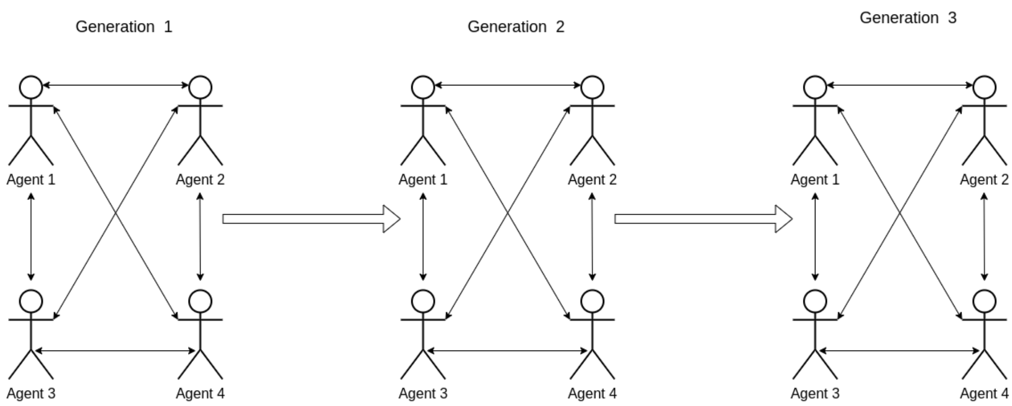

Figure 2.4 Illustration of a closed group experiment in multi-agent systems, illustrating the behavior and interactions of a fixed group of agents within a controlled environment. Such experiments allow researchers to study dynamics of interaction, cooperation, and decision-making within a stable group over time.

gerated tendency to follow the majority, sufficiently strong to increase the size of this majority over time. Results show that conformist populations may outperform non-conformist ones, as long as some basis for effective individual learning is present.

In [90], 56 participants perform a rest-less two-armed bandit task on a computer screen. The objective of this experiment is to investigate whether the exploitation/exploration tendency of humans in isolated settings, is able to predict the use of socially available information in group settings. To that end, participants carry a first task in isolation, a second task where the choices of another participant are available and a third task, (once more) in isolation. Results showed that no significant individual association exists, between the exploration frequency in the asocial task and the use of the social information in the social task. This suggests that how humans act in isolated armed-bandit scenarios, is independent to how they act in group armed-bandit scenarios.

Similar questions are examined in [66]. The latter examines how people choose among different learning strategies and how these strategies reflect on their acquired payoff. To that end, the task of designing an arrowhead was assigned to two groups of people, one group of demonstrators and one group of learners. Demonstrators were given the chance to improve their arrowhead design in solitude (asocial learning), based on the payoff they acquired over 30 consecutive hunting episodes. Contrariwise, learners had one more option, copy an existing demonstrator design based on a series of different social learning strategies : payoff-bias, conformity, averaging and random copying. While social learning offers the advantage of avoiding arrowhead designs that are locally optimal, results are in line with [90]. Overall, the vast majority (77.5%) of the learning performed was asocial. Payoff-bias (copying the most successful demonstrator) accounted for another

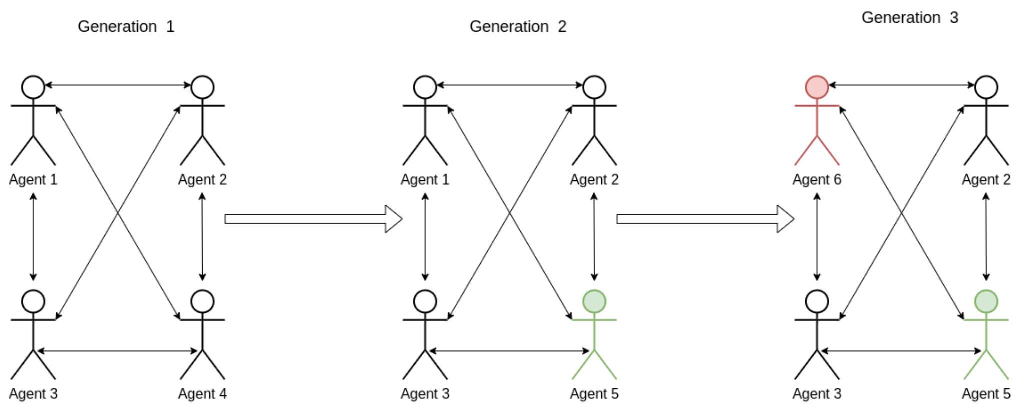

Figure 2.5 Illustration of a migration multi-agent systems experiment. It illustrates the movement and interaction of agents within a simulated environment. The experiment is used to study patterns of migration, spatial distribution, and interaction dynamics.

19%, leaving a mere 11% to other social learning strategies.

Migration methods

Migration methods experiments are experiments consisting of groups that are not fixed. Individuals are either replaced by new, naive individuals [4,21,22,25,81], or individuals are swapped among different groups [27]. These works aim to examine whether new group members, i.e., immigrants, adapt to an already established culture.

[4] builds upon previous research, examining how traditions (cultures) appear and evolve when group individuals are replaced by new, naive ones. To that end, microsocieties of 4 individuals unanimously choose among 2 different types of anagrams and receive payoff for solving them. Blue anagrams lead to higher short-term payoff, accompanied by a time-out penalty. After solving each anagram, players choosing blue are obliged to wait between 1 and 3 minutes. Since each generation spans over 12 minutes, the longer the time-out, the more payoff is potentially lost in the long run. Red anagrams induce no time-out penalty, yet their individual short-term payoff is low. Thus, choosing to solve a red anagram is a strategic choice that leads to a higher long-term payoff. Getting the newly added individuals inline with the established tradition, resides in the previously existing senior individuals. To that end, senior individuals used arguments that were factual, untrue or simply directly imposed their point of view. Results showed evidence of two evolving traditions : (1) regarding which type of anagrams to solve and (2) how to get the new individuals inline. While in the majority of cases groups initially favored solving blue anagrams, these groups eventually adopted traditions that maximized higher long-term (red anagrams) payoffs. The stronger

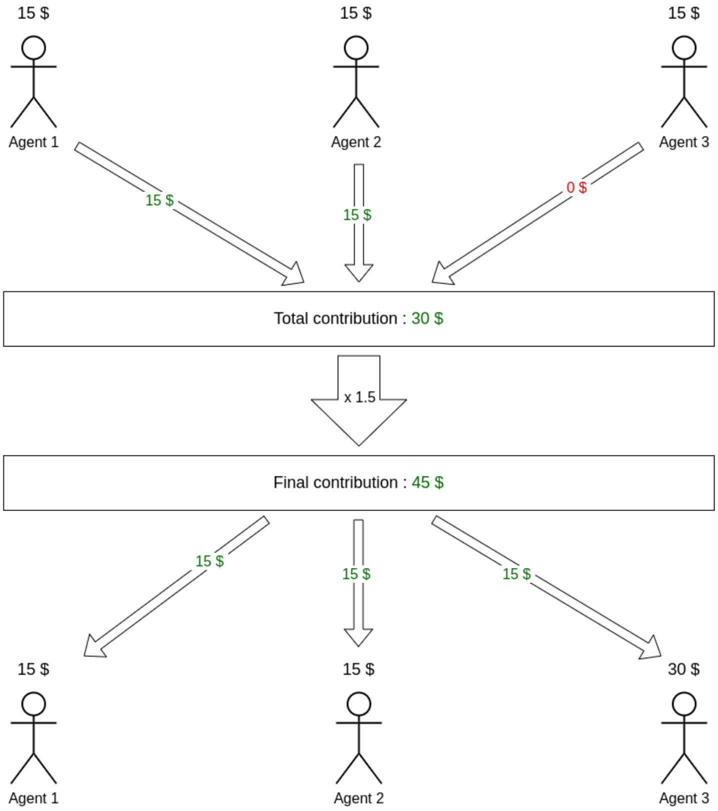

Figure 2.6 Illustration of a public goods game, where participants decide how much of a given sum of money to contribute to a public good. The total contribution is multiplied and equally distributed among all participants, irrespective of individual contributions.

the dependence between penalty (time-out duration) and payoff, the strongest the traditions adopted by the microsocieties were.

[25] examines how cultural selection occurs in non-human species. In particular : (1) whether more successful foraging traditions are culturally selected and diffused and (2) whether this is affected by replacing population individual with new, naive ones. To that end, an experiment is performed based on 18 wild-caught bird populations of the Parus major family. Result showed that these demographic replacement processes increases the efficiency of foraging within bird populations. This was not found to be a result of higher innovation rates, but rather due to higher adoption rates of novel efficient behaviors introduced by immigrant birds.

Economic games

Inspired by economic theory, these games usually focus on cooperative dilemmas e.g., public goods game [20, 45, 46, 60]. It has been previously hypothesized that social learning renders societies more cooperative. [20] tests this hypothesis, using a public-goods game organized over several sessions and groups. A public-goods game is an experimental tool often used in economic and social sciences, allowing for studying cooperation and collective action dilemmas [7, 35]. Participants are given an initial endownment of money, part of which they can attribute to a public good. The endownment benefits all participants equally, regardless of their individual contribution. The dilemma arises because, from an individual point of view, it is rational for each participant to contribute as little as possible, hoping that others will contribute enough to provide the public good. However, if everyone adopts this strategy, the public good may not be adequately funded. Providing these groups with varying sources of social information, allows to test whether this changes their cooperation levels over time. Results showed that the access in social information, i.e., how others contribute and what they get in return, did not increase cooperation, eventually leading to more "selfish" individual behaviors. This was found to be especially true when information on how successful other individuals are, was available. In that case, individuals leveraged this information to acquire higher individual gains, in detriment to the overall cooperation levels.

[60] examines whether cooperative strategies become traditions that are socially transmitted and as a result, subject to cultural evolution. To that end, individuals from different villages of a broader Indian forager-horticulturist society are recruited. Results suggest that social learning in cooperative environments differs from case to case, depending on the characteristics of each individual, as well as these of its environment.

2.4 Cultural evolution of languages

[44] studies how languages evolve over time. To that end, it proposes an experiment where two neural network-modeled agents, without common prior language, engage in a referential game [63]. Referential games are used among others in NLP, computational linguistics and experimental cultural evolution. Generally, the term refers to a task used to study reference resolution, i.e., understanding references that others make to objects or concepts. The goal of such games is to investigate whether and how well, participants can achieve a state of conceptual interoperability, i.e., understand each other. Typically, a participant (the speaker) uses its own language to convey information about an object to another participant (the listener). The listener's role is to match the conveyed information to an object based on its understanding about the environment and the speaker's language.

In the case of [44], the game consists of matching a sequence of symbols to a target image. At each round, a target image is randomly selected and presented to the sender. The sender then comes up with a message (finite sequence of symbols) representing this target image. It then sends this message to the listener, along with a set of images. The set contains decoy images as well as the target one. The goal of the listener is, based on the received message, to distinguish the target image from the decoy images. Ultimately, the goal of [44] is to examine whether the messages exhanged between the interacting agents, share characteristics with natural languages. These characteristics often considered a prerequisite for developing a general-purpose AI. Counterintuitively, results show that training convergence decreases with the length of the transmitted message. Put differently, the longer the sequences used by the senders (to describe target images), the faster the shared languages emerge. Results also suggest the following about the resulting language. First, there is redundancy in message encoding, meaning that the same semantic content can be conveyed through multiple paraphrases. Second, there is a higher encoding and decoding robustness, indicating that the semantic content of the symbols in each message has similar discriminative power. As a result, even if a message is corrupted, the remaining correctly conveyed symbols are sufficient for the encoding and decoding of the original message.

[43] uses similar agent to agent interaction games to study contact linguistics [69]. Contact linguistics examine how languages evolve, when communities speaking one language are suddenly exposed to communities that speak a different one. In the context of our thesis, languages can be seen as equivalent to knowledge representations shaped for different purposes and isolated communities as agents specialized on different tasks. Thus, results acquired within contact linguistics can provide insight into what happens to knowledge when specialized agents interact.

Do these languages evolve to co-exist, or does one language fall to oblivion while the other one evolves to dominate over all communications? If the former is true, do languages gradually form a linguistic continuum, allowing for neighboring languages for being more inter-interpretable? To answer these questions, Myers-Scotton introduced a multi-agent framework where neural network-modeled agents exchange messages about their perceptual input. Results lead to several interesting conclusions on how languages evolve. First, under the examined framework, two agents (language users) are sufficient for a shared language to emerge. Second, whether two languages will converge depends on the size, density, and topology of the examined communities. When the population ratios of the communities are imbalanced, a new majority language emerges as a result of convergence. When this is not the case, both languages are subjected to gradual simplification, yet continue to co-exist. The term that is often used for describing such languages is "creole". Lastly, when a linguistic continuum emerges, agents of neighboring communities are perfectly capable of communicating with each other. Their communication is as easy as if they did belong to the same community. The latter is less and less true as the distance between two communities increases.

[26] studies a specific emerging characteristic of natural languages, compositionality. A language is considered compositional when a concept can be expressed by combining words, each one with its own meaning. Compositionality is what allows to express an infinite number of concepts, using a finite set of words. If

it was to be achieved by general purpose A.I, this would greatly improve the communication among different artificial entities, or among artificial entities and humans. In the context of knowledge, compositionality may allow an agent to generalize its individual knowledge representation. This can be particularly interesting when agents are trained based on limited sets of training examples. For example, an agent that has learned to deal with maintaining car engines and motorcycle brakes, will be able to understand how car brakes work althouth it has never seen car brakes before.

Work on evolutionary linguistics [69] has previously shown that transmitting a language from one generation of humans to the next one, significantly enchances its compositionality. [91] tests this concept on artificial populations that are modeled using neural networks. Ultimately, it bridges research on language compositionality using several human generations and research on how languages emerge when artificial agents interact. To that end, an experiment based on a cooperative dialog-based reference game is conducted. Each interaction takes place between randomly selected pairs of agents, communicating using discrete symbols. Periodically, agents are replaced with new, naive ones. Naive agents are initially provided with a randomly generated language, while the previously existing agents communicate using an already grounded one. Pressure is thus put on naive agents to converge to the grounded languages, leading to a setting that is equivalent to that of iterative learning. Results show the following: First, replacing agents induces cultural transmission, which significantly increases the observed language compositionality. Second, the selection of agents for replacement has no statistically significant effect on the observed compositionality. [84] argues that a communication protocol emerges when agents attempt to minimize the computational complexity of semantic interpretation.

2.5 Cultural evolution of unipotent knowledge

The cultural evolution of knowledge has been studied in [11, 12, 92].

In order for agents to successfully coordinate, agents exchange representational objects that can locally be interepreted in ways that make collective sense. [92] examines how can collective behavior be stabilized, while each agent organizes its knowledge, reasons with it and acts autonomously. Put differently, [92] investigates how agents that initially think and to that end act differently, can gradually agree on a way of doing things, without necessarily adopting one agent's way. To that end, authors propose an experiment based on an algorithm they call "mutual ontology alignment". This experiment has two stages.

During the first stage, a population of agents is initialized with each agent assigned a random mapping of the domain of natural numbers. This mapping is organized over several categories (concepts), each one consisting of randomly selected numbers. For example, a mapping may consist of two categories, one for numbers that are even and one for numbers that are odd.

During the second stage, the agents undergo a finite number of interactions, each one taking place between two randomly selected agents. For each interaction, one of these two agents is designated as the teacher, while the other agent is designated as the learner. The teacher will initiate the interaction, by sending to the learner one number. This number will be randomly selected within a randomly selected category of the teacher's mapping. The learner will receive this number and find the category to which it belongs, according to its own mapping. The learner will then select a different number within this category and send it back to the teacher. If these two numbers fall in the same category within the teacher's mapping, the current interaction ends and the experiment proceeds with the next interaction. If these two numbers do not fall in the same category within the teacher's mapping, then one or both agents may adapt their mapping.

Pressure is thus applied to reduce mapping mistakes over time. On an abstract level, a population collaborativelly designs an ontology that satisfies mutual circumstances based on non ideal individual ontologies. This work provides a first formalized approach to a dynamic process of mutual continuous co-alignment of a collective concept. While it does not identify all the conditions under which convergence can be achieved, it experimentally shows that the latter can indeed be achieved. In [92], when the experiment ends, while each agent possesses its own individual mapping, all mappings are identical. Agents always act in the same way, because their knowledge is organized in the same way as well. However, this is not what we necessarilly observe in humans. Furthermore, no notion of whether the constructed mapping is accurate or not is provided.

When agents and humans alike tackle a task, they end up building individual knowledge that may be diverse, incomplete or inaccurate. Yet, unless consensus is achieved, they cannot successfully cooperate. [75] reviews different algorithms used for achieving consensus in artificial populations.

Inspired by research on cultural evolution of languages, [11] studies how knowledge evolves culturally, as a result of social transmission. In particular, [11] answers three main questions. First, whether the cultural evolution of knowledge representations allows individuals to agree on their decisions. Second, whether this process improves knowledge correctness. Third, whether individuals agree because their knowledge is identical.

To answer these questions, an experiment sharing similarities with that of [92] is proposed. It is based on populations of artificial agents, each one representing its knowledge using an individual ontology. This ontology allows for classifying objects into classes based on boolean properties and aligning these classes with learned decisions. The experiment is organized over two stages.

During the first stage, each agent induces an initial ontology, based on a set of labeled examples. Each example consists of a set of properties, constituting an object type and a label, i.e., the correct, uknown to agents decision. For example, one labeled example could be the following. For an object that has nails, is orange and has black stripes, one should *run*. Since this set is different and incomplete for each agent, different agents are initially able to accurately decide with respect to different objects (or situations).

During the second stage, pairs of agents will go over a finite number of interactions. At each interaction, a randomly selected pair of agents will be presented with a randomly selected object. The agents will reason based on their knowledge and the object's properties. It should be noted that in this work, all properties are observable. Based on its reasoning, each agent will then decide what action to take with respect to the presented object, subsequently disclosing its decision to the other agent. If the agents agree, their interaction ends successfully. If the agents disagree, then one of the two agents will adapt its ontology. The latter can be done by substituting the existing decision, (if all properties are already present in the ontology), or by extending its ontology (if properties not yet evaluated exist). Which of the two agents will adapt its knowledge, is decided based on implicit environmental feedback, received before each interaction. Semantically, the latter corresponds to the prestige or success bias often found in human interactions.

Knowledge is selected given the pressure applied on agents agreement. The adaptation leads to variation, which in turn allows for subsequent selection and potentially further variation. Since the experiment is based on a closed world with predefined object types and properties, knowledge selection and adaptation stops when disagreements no longer occur.

Results show that all three paper's hypotheses are supported. First, agents improve the ratio of successful over all interactions, until it converges to 1. Authors call this measure success rate. Second, while the pressure is applied on agents agreeing on their decisions and not becoming more accurate, knowledge accuracy significantly improves over time. What makes this interesting is that mean accuracy improves without explicit environment feedback about decisions correctness. Third, unlike [92], agents' ontologies do not become identical. The agents thus learn to agree, yet not based on identical reasoning.

In [11] the population of agents remains unchanged for the duration of the experiment. Thus, the observed variation is solely due to the initial ontology training and horizontal social transmission. [12] extends [11] by adding inter-generational (vertical) transmission to the existing intra-generational (horizontal) one. Previous results on inter-generational transmission were obtained through drastic selection of who becomes a parent and which knowledge this parent transmits to its descendants. The objective of [12] is to study how this drastic selection of agents affects the quality of the transmitted knowledge. It is assumed that allowing for oblique social transmission, will increase ontology variation. Exceeding the previously achieved variation will allow for agents putting their knowledge in question for longer, eventually leading to building more accurate knowledge. Results showed the following: First, combining inter-generational and intra-generational transmission improves the accuracy of agents' ontologies. Second, when the parent selection constraints are relaxed, the agents' lifespan significantly affects their knowledge correctness. If the lifespan of agents is short, agents may transmit to their descendents, knowledge that has not been sufficiently vetted, and thus is of low quality (accuracy). If the lifespan of agents is sufficiently long, even if they vertically acquire low quality knowledge due to relaxed parent selection, agents have a chance in improving it through knowledge selection and horizontal transmission.

The objective of [1] is to study cultural transmission between and within generations. To that end, it follows an approach similar to [12], periodically replacing agents by new, naive ones. The naive agents are initially assigned either a random ontology, or an ontology induced using decisions their parents would take for a randomly selected set of objects. Results suggest that intra-generational transmission enhances the variability within cultural evolution. This can be particularly useful when the examined populations live in environments that are subject to constant change. Despite this, the importance of having both types of transmission for population evolution is highlighted.

2.6 Specialization of multipotent traits

2.6.1 Specialization in empirical studies

Negative transfer or insufficient resources can be expected by individuals trying to evolve knowledge with respect to several tasks. Empirical data show that intentionally or not, humans deal with this problem by specializing their knowledge and skills. Recent works attribute the appearance of specialization to the formation of comparative advantage [99]. As an example, an individual is better at making farming tools and another individual is better in using these tools in agriculture. Specialization ensures that these individuals exploit their skills optimally, by allowing them to allocate more time and resources to what they do best. This has been previously shown to lead to higher exploration of novel techniques and thus more efficient production [98].

While intuitive, the aforementioned example comes in contradiction with empirical studies in hunter-gatherer populations [18]. In these, a good hunter may give away some of his pray, not necessarily as a direct exchange for a service or good. This may be inherent to the risk of relying on someone else [82]. If the blacksmith has no clue on how to hunt and if the hunter was to disappear unexpectedly, the blacksmith would starve to death. This risk could be reduced if several individuals tackled each task. This could bring competition and assure survival in case of sudden disappearance of some individuals.

Regardless of the lack of experimental results regarding the origin and impact of specialization, increased efficiency-productivity is something axiomatically accepted in the vast majority if not in all theoretic papers regarding specialization. This is mainly based on the assumption that the more time an individual spends on a task, the better this individual becomes at it. While this intuition might hold for human beings, it does not necessarily extend to other species or machines, nor is it necessarily connected to the process of specialization itself.

2.6.2 Model-based agent specialization

The opposite side of empirical specialization studies are theoretical models. While these have proven to be the cornerstone of analysis by allowing testing hypotheses and assumptions, most of the existing theory regarding the division of labor consists of verbal reasoning. Since specialization is crucial for markets, economists, keen on mathematical models were among the first to provide models bridging specialization, the division of labor and the market [99].

Others have approached specialization using controlled experiments. [38] studies how artificial entities specialize when several rewarding [72] tasks can be performed. The objective of [38] is to examine how the number of agents and the number of times each task is performed before another task, encourage (agent) specialization. To that end, an experiment based on reinforcement learning modeled agents is proposed. These agents live in a gridworld environment and receive reward when succeeding in carrying out tasks. Topographically, the environment is divided into three areas of equal surface (number of squares), within which agents with partial observability and resources are randomly placed at each episode. These tasks are sequential and consist of moving these resources from one area to the next one. When the agents succeed in this, they receive a reward, based on a reward function that favors quick learning and disfavors useless actions or idling (waiting behaviors). In the end of the experiment, agents are individually evaluated based on the absolute difference between the times they played each task divided by all played games regardless of the task. According to this, an agent that does one task independently of its respawn location, is considered as specialized. Results show that specialization is on average higher with 6 agents than it is with 2 agents. Furthermore, the same results also show that not all agents specialize to the same degree. These agree with [74], where an advantage of sacrificing cognitive resources in favor of specific tasks, improves efficiency in ant colonies. However, [74] disagrees with [28], where ants focusing on one task do not necessarily get more efficient at it. Further results also correlate the number of times each task is played before the next one, with the fairness and the stability of the agents learning process. The authors show that the lower this number is, the more difficult it is for agents to exploit the environment in order to obtain their reward.

The specialization literature presented so far in Subsection 2.6.2 focuses on two aspects. The first aspect concentrates on the origins of specialization, i.e., the in vitro and in vivo conditions that push artificial and natural populations towards specialization. The second aspect contentrates on the benefits these populations may acquire from it. An important aspect concerning this thesis is however overlooked, that of cultural evolution. None of these works discusses (1) how cultural traits become specialized and (2) how these specialized cultural traits are accumulated over time.

How specialized cultural traits are accumulated over time has been recently discussed in [6]. In [6], it is assumed that the effects of specialization may be ambivalent. On the one hand specialization may allow for collectively undertaking complex tasks, i.e., tasks that would otherwise require more cognitive resources. On the other hand, specialization may lead to populations that are less robust when faced with disruptive events, i.e., sudden demographic and environmental changes.

While losing rare knowledge can hinder the efficiency and survivability of any population type, its effects are assumed to be amplified in specialized populations. To test this assumption, Ben-Oren et al. propose a framework inspired by [58]. This framework is based on the concept of cultural innovations, for which the term tools is used. Tools, i.e., cultural traits, occur in specialized and non specialized groups of individuals alike. The more a population is divided into groups that use non-overlapping sets of tools, the more this population is considered to be specialized. The proposed framework does not explicitely model any form of cultural transmission. Instead, the tools are periodically created, independently of the population's division and the existence of other tools. The innovation rate, i.e., the frequency with which new tools are created, is linearly dependent on the population size. Instead of relying on variation and selection, all tools are associated with a randomly selected positive selection coefficient s. Each unestablished tool can be stochastically lost with a probability of $1-s$. Based on this framework, the authors examine the relationship between cultural specialization, cultural repertoire size (how many distinct tools exist) and the robustness of this repertoire to disruptive events. Results show that the relationship between the population size and the number of distinct tools is nonlinear and differs depending on the degree of specialization.

In summary, specialization was found to increase the number of distinct tools in large populations and decrease it for small ones. This can be explained by the fact that in non-specialized populations, the population's total repertoire of cultural traits is limited by the capacity of each individual to acquire and accumulate cultural traits. In specialized populations, this is not a requirement. In this case, each individual is required to know only a small fraction of all existing cultural traits.

2.7 Conclusion

Researchers have previously examined the implications of social learning and social transmission on the evolution of cultural traits (§2.1). Inspired by this, significant work has been done in the areas of multi-agent systems and multi-task learning (§2.2), experimental cultural evolution (§2.3), specialization (§2.6), and pairwise intersections of these, such as cultural evolution of languages (§2.4) and cultural evolution of knowledge (§2.5). However, there has been little to no consideration of multi-tasking agents, cultural evolution of knowledge and specialization altogether. The objective of this thesis is to bridge these areas by examining how cognitive restrictions and inter-agent interactions push multi-tasking agents to specialize their knowledge. To that end, an experimental framework for evolving multi-potent knowledge using agents is proposed in Chapter 3, allowing for performing experiments on the formation of general-purpose knowledge (Chapter 4), agent specialization (Chapter 5), knowledge forgetting (Chapter 6) and task exploration (Chapter 7).

— 3 —

A framework for adaptive resource-restricted multipotent agents

3.1 Introduction

Experimental cultural evolution of knowledge provides valuable insights into how knowledge representations, such as ontologies, evolve as a result of inter-agent interaction. Recent works are based on frameworks dealing with agent populations that are limited to one task and whose resources are unlimited. Using these frameworks, it has been demonstrated that agents will continue to improve their knowledge correctness, until they agree on all decisions.

Our objective is to study how the cultural evolution of knowledge is affected, when several tasks compete over agent resources. To that end, we successively consider (1) whether agents undertaking semantically overlapping tasks develop general purpose knowledge, (2) what pushes agents to specialize and how agent specialization affects them and their populations and (3) whether a population of agents can equitably improve its knowledge on all tasks. We refer to these considerations as (a) knowledge transferability, (b) agent specialization and (c) task exploration respectively. These considerations, combined, will answer how knowledge regarding different tasks is distributed within a population of agents.

None of the three considerations can be taken into account by the existing frameworks. The framework closest to our objective is the one used in [11, 12]. However, it does not address: (a) tackling several tasks, (b) task assignment, (c) task overlap and (d) resource limitations.

To address these, we enhance this framework as follows. Unlike previous frameworks where agents are limited to a single task, the proposed framework allows agents to undertake several tasks. Deciding for these tasks relies on specific properties, which can be shared or distinct. Tasks relying on the same properties are considered overlapping, while tasks relying on different properties are considered independent. This distinction allows for examining two main things. First, the transferability of knowledge among different tasks. Second, the formation and evolution of specialized knowledge representations. Each agent can be assigned any number of tasks. Agents tackle these tasks by operating within restricted memory, reflecting real-world resource limitations. This allows for studying the effects of tasks competing for memory on the cultural evolution of multitasking knowledge. Through these, a comprehensive framework for adaptive resource-restricted agents is provided.

This chapter is divided into two sections. §3.2 provides all the preliminary definitions regarding the agents, the different tasks, the objects and the environment that surrounds them. In §3.3, an outline of our experiments is provided. This outline describes in detail how agents are assigned tasks, acquire their initial ontologies, interact with other agents, adapt their ontologies and deal with the exhaustion of their resources.

3.2 Experimental framework

Our experimental framework extends the framework used in [11, 12]. To provide clarity and continuity, we present definitions and notations of the original framework, as they are fundamental to understanding our extensions. Here, we introduce agents capable of tackling several tasks that may overlap to varying degrees. This is discussed in Subsections 3.2.2 and 3.2.3.

3.2.1 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 only $2^{|\mathcal{P}|}$ object types, that are gathered in a set $\mathcal{I}$.

3.2.2 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 $\mathcal{D}_t$. In this context, each object o can be considered in regard 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 in regard to a task t.

3.2.3 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

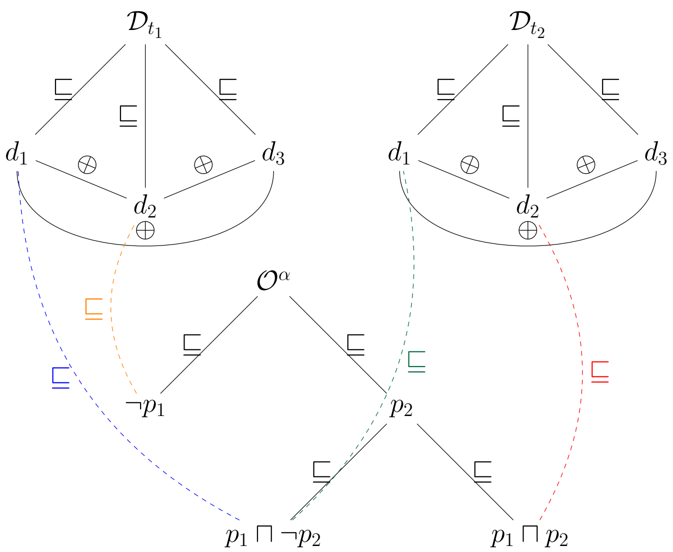

Figure 3.1 Example of a multi-task ontology for an agent α. Each color represents a different decision.

agents $\mathcal{A}$, carries out multiple tasks. To this end, agents build and evolve knowledge in the form of ontologies, private to each agent, expressed in $\mathcal{ALC}$ [3]. Each agent k uses its knowledge to compute a function $h^k(o,t) \rightarrow \mathcal{D}_t$ which, given an object o and a task t, provides a decision $h^k(o,t)$. Figure 3.1 shows an example of multi-task knowledge constructed by an agent k. The bottom part represents the private ontology $\mathcal{O}^\alpha$ of agent k, 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. Given that an agent k learns at most one decision for an object o and a task t, each leaf of $\mathcal{O}^k$ cannot be aligned more than once with the same decision ontology.

3.3 Experiment outline

Experiments are initialized with a learning phase. At the end of this phase, each agent has learnt a private ontology. Once their ontologies learnt, agents will go through a fixed number of interactions. Depending on an interaction's outcome, one agent will adapt its ontology. Figure 3.2 illustrates the flow of the experiment from the perspective of agent α. More details about how agents learn, define their scope, interact, forget knowledge and adapt their ontologies are presented in subsections 3.3.1, 3.3.2, 3.3.3, 3.3.4 and 3.3.5 respectively.

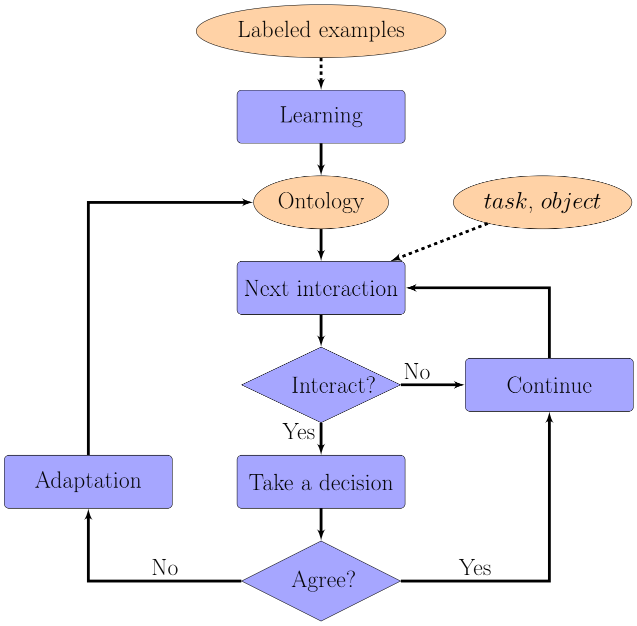

Figure 3.2 Experiment outline. Elements in bisque color constitute inputs for different activities (e.g., labeled examples). Elements in violet constitute either activities (rectangles), or decisions (diamonds).

3.3.1 Agent training

[92] and [11] have already demonstrated the capacity to effectively learn an ontology in order to carry out a single type of task. We approach multi-task learning as a problem of building a multi-purpose ontology. This can be achieved either through a training set containing labeled samples regarding different tasks, or through a merging of ontologies learnt from single tasks [24], ultimately leading to an ontology able to provide a decision for any task $t \in \mathcal{T}$. Here we propose a multi-purpose ontology induction algorithm. We refer to this algorithm as interleaved learning. It is based on a set of labeled samples covering a subset of $\mathcal{I}$, and thus does not represent all object types. The algorithm is given a set of samples $\langle o, t, h^*(o,t) \rangle$ containing:

- an object o, characterized by a set $\mathcal{P}(o) = \{p_1(o), p_2(o) \dots p_n(o)\}$ of boolean properties
- a task $t \in \mathcal{T}$
- a decision $h^*(o,t) : \mathcal{I} \rightarrow \mathcal{D}_t$

Based on these samples, a decision tree is induced. It consists of (intermediate) parent and leaf nodes. Each parent node n_{parent} contains the following attributes:

- the property evaluated $evp(n_{parent})$

- children nodes $child_{true}$ and $child_{false}$, corresponding to $evp(n_{parent})$ being true or false respectively

Each leaf node n_{leaf}, contains the following attributes:

- a set of objects $obj(n_{leaf})$ classified by n_{leaf}
- a set of decisions $des(n_{leaf})$ attached at a node n_{leaf}, with at most one decision per task

o	p_1	p_2	p_3	p_4	t	$h^*(o,t)$
o_1	1	1	0	0	t_1	d_1
o_2	1	0	0	1	t_2	d_1
o_3	0	1	0	1	t_1	d_2
o_4	0	0	1	0	t_2	d_2
o_5	1	1	0	1	t_1	d_1
o_6	1	1	0	0	t_2	d_1
o_7	0	1	1	0	t_1	d_2
o_8	1	0	1	1	t_2	d_2

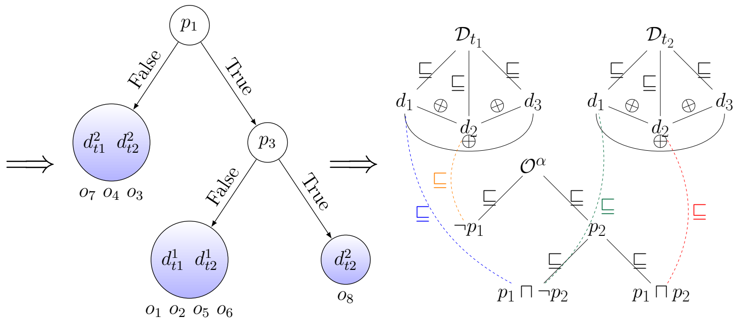

Figure 3.3 Given a set of labeled examples, agents will induce a decision tree. The latter is subsequently transformed into an ontology. Each color represents a different decision.

A learnt tree can be observed in Figure 3.3 center. This decision tree is then transformed to an ontology in $\mathcal{ALC}$. Such an ontology can be observed in Figure 3.3 right. Interleaved learning comes with a side effect. It is possible for an agent to be able to classify an object but be unable to provide a decision regarding some tasks. For example, given the ontology in Figure 3.1 and a given object o, described by $p_1 \sqcap p_2$, it is impossible to decide with respect to task t_1.

3.3.2 Scope definition

Here we implement agents that undertake different subsets of tasks based on prioritizing some tasks over others. To that end, we introduce the function $prior^\alpha(t) \to \mathcal{N}$ returning a unique integer assigned to the task t by an agent α. This integer ranges from 1 to $|\mathcal{T}|$ and represents the priority for the task t among all existing tasks. Provided with their priority rankings, agents define their corresponding individual scope, i.e., the set of tasks they undertake during the experiment. In order to decide whether a task should be included in the task scope, agents compare the priority of each task t with a variable m, i.e., the maximum accepted priority rank. If $prior^\alpha(t) \leq m$, then the task is inside the agent's scope. If $prior^\alpha(t) > m$, then the task is outside the agent's scope. For example, if m equals 1, then the agent α's scope will contain only one task, while if m equals $|\mathcal{T}|$, the agent α's scope will contain all existing tasks.

3.3.3 Interaction

For each interaction (Algorithm 1), two agents α and β are randomly selected from a population of agents. α and β are then presented with an object o and a task t. They will then provide decisions based on their respective functions $h^\alpha(o,t)$ and $h^\beta(o,t)$. If an agent is unable to provide a decision, then one 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 ends successfully. On the contrary, their interaction ends as a communication failure. In this case, one of the two agents will adapt its knowledge.

Algorithm 1 Interaction

Two agents α and β are randomly selected
A task $t \in \mathcal{T}$ and an object $o \in \mathcal{I}$ are randomly selected
t and o are assigned to agents α and β
if α or β contains no decision **then**
    A random decision $d \in \mathcal{D}_t$ is assigned to the agent(s) without a decision
end if
The agents compare their respective a decisions
if $h^\alpha(o,t) = h^\beta(o,t)$ **then**
    The communication ends successfully
else
    A subset $\mathcal{I}^* \subseteq \mathcal{I}$ is randomly selected
    $\mathcal{O}^\alpha$ and $\mathcal{O}^\beta$ are evaluated with respect to $\mathcal{I}^* \subseteq \mathcal{I}$ and task t
    if $\mathcal{O}^\alpha$ is evaluated higher than $\mathcal{O}^\beta$ **then**
        β will adapt its ontology
    else
        α will adapt its ontology
    end if
end if

3.3.4 Memory management

Realistic agents do not have unlimited memory. We model this by limiting the number of classes per agent's ontology. When this limit is reached, then agents will try to forget a part of their knowledge in order to re-gain memory resources in favor of the tasks they are assigned. We assume that if restricted enough, an ontology will be able to contain only the properties required to take the correct decision with respect to a single task. For example, if deciding with respect to the task t_1 depends on $\mathcal{P}_{t_1} = \{p_1, p_2\}$ and deciding with respect to the task t_2 depends on $\mathcal{P}_{t_2} = \{p_3, p_4\}$, then if the ontology evaluates only p_1 and p_2, then if a decision regarding the task t_2 is returned, the probability of this decision to be correct is on average $\frac{1}{|\mathcal{D}_{t_2}|}$. For example, let deciding whether the agent should take an umbrella (decision 1) or not (decision 2), relying solely on the property *dewpoint*. If the agent's ontology does not contain the property *dewpoint*, then the probability of an associated decision to be correct is on average 50%. In this subsection, two mechanisms to forget knowledge are defined: (1) knowledge generalization and (2) random knowledge removal.

Knowledge generalization

Knowledge generalization is achieved by removing parent nodes that satisfy the following criteria: (a) their children are leaf nodes, and (b) their children are associated with the same decision regarding all tasks within the agent's scope. Criteria (a) guarantees that no unnecessary forgetting will take place. Criteria (b) guarantees that no knowledge regarding the assigned tasks will be lost. The process is repeated recursively, as long as parent nodes satisfying (a) and (b) exist.

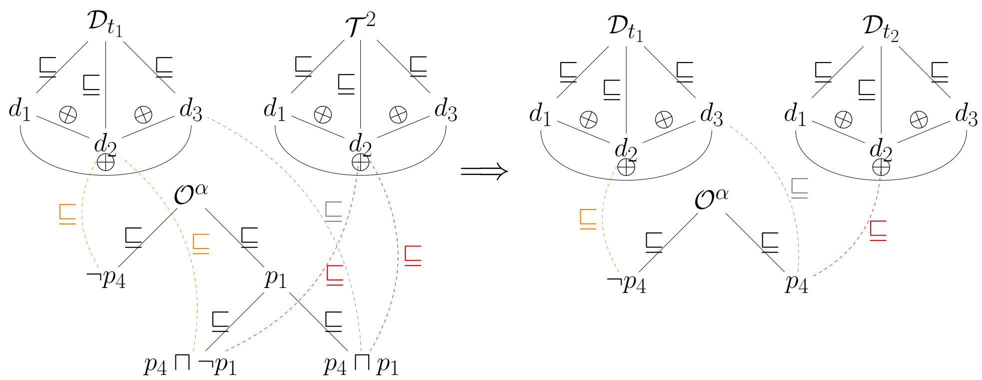

Figure 3.4 Example of knowledge forgetting while interacting with respect to the task t_2. Let the task t_2 relying on the property set $\mathcal{P}_{t_2}$. The property $p_1 \notin \mathcal{P}_{t_2}$, thus p_1 does not allow for distinguishing different decisions for the task t_2. In this example, the 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 merged without any loss of knowledge with respect to t_2.

Random knowledge removal

Random knowledge removal consists of removing all the descendant classes of a randomly selected class. The randomly selected class will then become a leaf class and will be associated with decisions previously associated to its leaf descendant classes. When several decisions exist for the same task, one decision per task is randomly selected.

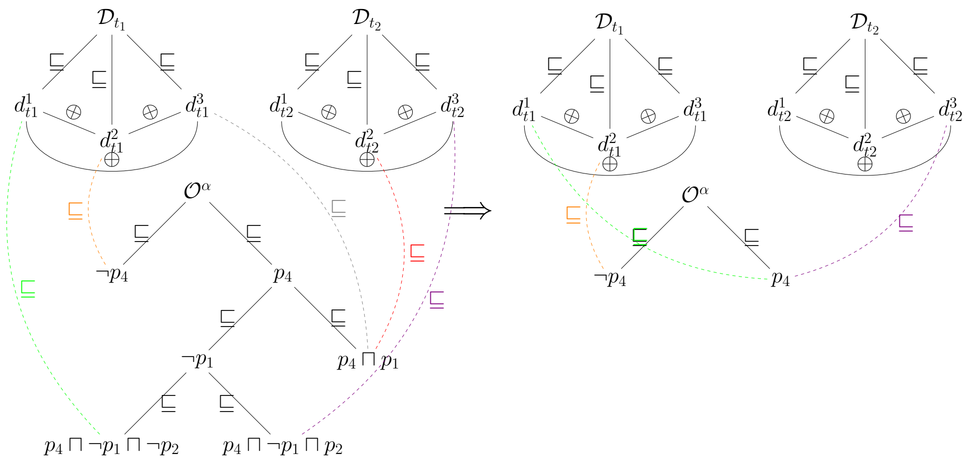

Figure 3.5 Let an agent α assigned the task t_2, with t_2 relying on the property set $\mathcal{P}_{t_2}$. The agent α will randomly select the class p_4 and remove its three descendant classes. The class p_4 will now be associated with decisions associated with the recently removed classes $p_4 \sqcap \neg p_1 \sqcap \neg p_2$, $p_4 \sqcap \neg p_1 \sqcap p_2$ and $p_4 \sqcap p_1$. When more than one decision exist, one is randomly selected. For example, the class $p_4 \sqcap \neg p_1 \sqcap \neg p_2$ is associated with the decision d^1_{t1} and the class $p_4 \sqcap \neg p_1 \sqcap p_2$ is associated with the decision d^3_{t1}. After the knowledge removal, the class p_4 will be associated with the randomly selected d^1_{t1} (green).

3.3.5 Adaptation

We extend the adaptation mechanism presented in [11], by taking into account the presence of several tasks. As depicted in Figure 3.6, when a communication failure occurs, one of the two agents will adapt its ontology. This agent can either change an existing decision, or divide a class into two subclasses as described in the algorithm 2. In this case, a property is chosen based on the class description provided by the non-adapting agent. This property will allow the adapting agent for distinguishing the current object from all the other objects currently classified in the class to be divided. Only decisions concerning the current task are affected.

3.4 Conclusion

In this chapter, we presented an extended framework for simulating how knowledge culturally evolves and specializes, when agents interact. We initially introduced: (1) the objectives of our thesis, (2) the limitations that render existing framework unsuitable for our objectives and (3) the enhancements required for achieving them. Based on these, an extension of [61] allowing for the formation of both multi-purpose and specialized knowledge was proposed. First, rather than sharing the same single task, agents can now tackle several tasks. These tasks may be the

Algorithm 2 Ontology adaptation

α asks β for the definition of the class classifying the object o
β answers with $\mathcal{C}^\beta$
if $\mathcal{C}^\alpha \not\sqsubseteq \mathcal{C}^\beta$ **then**
    α asks β for its decision regarding the object o and the task t
    β answers with $h^\beta(o,t)$
    a property $p : p \in \mathcal{C}^\beta, p \notin \mathcal{C}^\alpha$ is selected.
    α splits $\mathcal{C}^\alpha$ into $\mathcal{C}_1^\alpha \equiv \mathcal{C}^\alpha \sqcap \neg p$ and $\mathcal{C}_2^\alpha \equiv \mathcal{C}^\alpha \sqcap p$
    α associates $\mathcal{C}_1^\alpha$ with all the decisions previously associated with $\mathcal{C}^\alpha$
    if No decision regarding t is associated with $\mathcal{C}^\alpha$ **then**
        α associates $\mathcal{C}^\alpha$ with the decision $h^\beta(o,t)$
    end if
    α associates $\mathcal{C}_2^\alpha$ with the decision $h^\beta(o,t)$
else
    $\mathcal{C}^\alpha = \mathcal{C}^\alpha \sqcap p$
    Associate $\mathcal{C}^\alpha$ with all the decisions previously associated with $\mathcal{C}^\beta$
end if

same for all agents, or different agents may be assigned different tasks. Second, constraints on agents memory resources were introduced, thereby simulating more realistic cognitive limitations. Third, the dependency on different object features when deciding for different tasks was introduced. Given agents' limited memory resources, this promotes knowledge diversity and thus potentially specialization among agents. Fourth, different parent selection criteria for agent reproduction were introduced, allowing for studying how different evolutionary strategies, impact knowledge specialization. The framework presented here will allow us to investigate the effects of cultural evolution on how artificial populations that have limited resources and span over several generations, evolve and specialize multipotent knowledge.

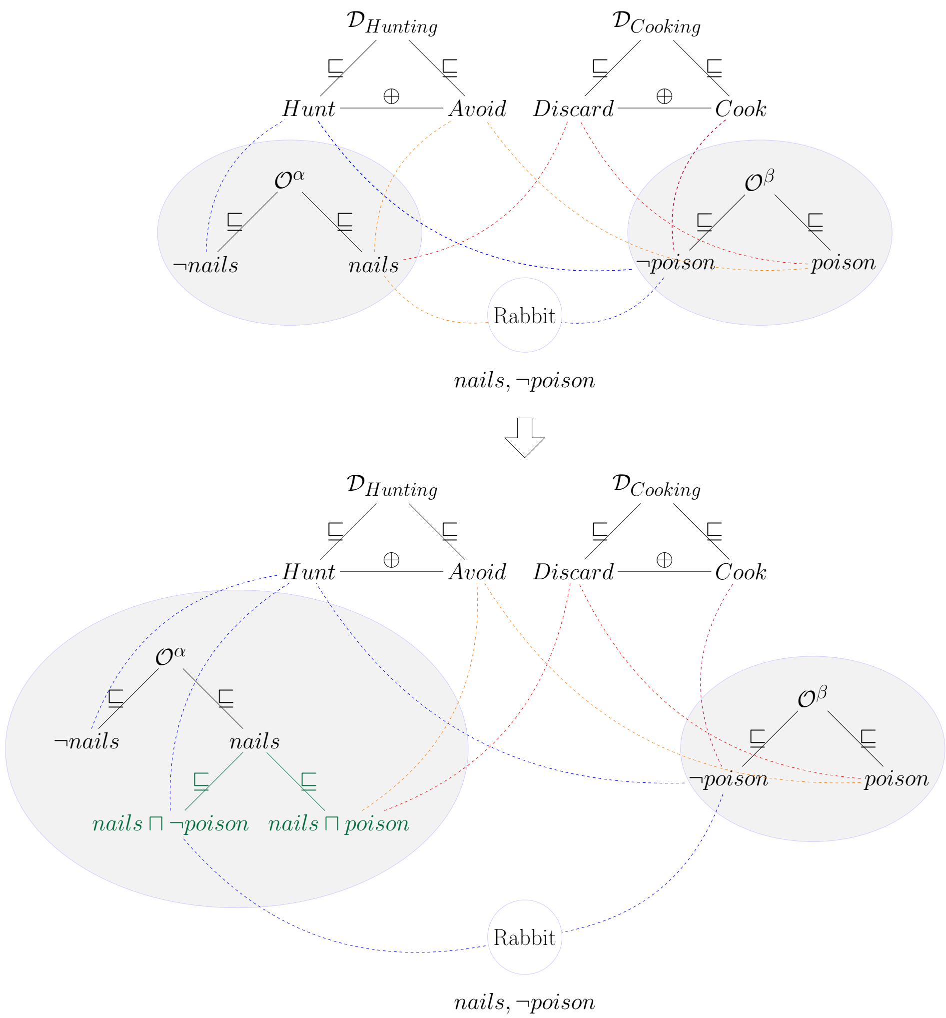

Figure 3.6 Subsequently to a communication failure, the agent α will split the class *nails* into two sub-classes by additionally evaluating the property *poison*. The first sub-class, $nails \sqcap poison$ will be now associated with the same decision as agent β's ($Hunt$). The second decision will be associated with the decision previously associated to the class *nails* ($Avoid$).

— 4 —

Knowledge transferability among different tasks

The cultural evolution of knowledge has been previously shown to enhance the accuracy of interacting agents [11,12]. In these, agents share the same single task: taking an abstract decision within an abstract domain. In [11], the agents' goal is 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 [11, 12] by introducing agents capable of taking abstract decisions within several domains.

Deciding for these 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.

4.1 Experimental setting

4.1.1 Hypothesis

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

4.1.2 Parameters

The experiment is executed under six setups, corresponding to agents tackling one to three tasks, that either fully overlap (100%) or do not overlap at all (0%). Each setup is run 20 times and its results are averaged. One run consists of 80000 interactions with each interaction taking place among two agents. These agents are randomly selected out of a total population of 18 agents. When their memory is exhausted, these agents free memory by generalizing their knowledge (Subsection 3.3.4). Their environment contains 64 different object types, each one perceivable through 6 different binary properties. The agents are initially trained with respect to all $|\mathcal{T}| = 3$ tasks. For each task, 4 different decisions exist. Deciding with respect to different tasks always relies on 2 out of the 6 binary properties. When tasks overlap, these 2 properties are the same for all tasks. When the tasks do not overlap, deciding with respect to different tasks relies on a different pair of properties. To induce its initial ontology, each agent takes into account a randomly selected 10% of all existing labeled examples. As a consequence, there may be overlap between the training of two agents. The agents subsequently undertake tasks within their scope. To decide which agent adapts its knowledge in case of disagreement, agents compare their scores. Their individual scores are attributed to them by the environment between two consecutive interactions. They are calculated based on a randomly selected corpus consisting of 60% of all existing training samples.

4.1.3 Measures

Success rate

To evaluate the outcome of agents' interactions and conclude on whether agents converge on their decisions, we use the success rate, used in [11]. Each interaction k is considered successful if the decision of the first randomly selected agent α_k about an object o_k and a task t_k, $h_k^{\alpha_k}(o_k, t_k)$, is the same with that of the second randomly selected agent β_k about the same object o_k and task t_k, $h_k^{\beta_k}(o_k, t_k)$. The success rate

Table 4.1 Summary table of parameters (independent variables)

Parameter	**Value**
Runs	20
Interactions	80000
Generations	1
Agents $\|\mathcal{A}\|$	18
Memory capacity	12 classes
Memory release technique	Generalization
Tasks $\|\mathcal{T}\|$	3
Assigned tasks (scope)	{1,2,3}
Object types $\|\mathcal{I}\|$	64
Properties	6
Properties per task $\|\mathcal{D}_t\|$	2
Task overlap	{0%, 100%}
Decisions per task	4
Training ratio	0.1
Score ratio	0.6

measure can then be defined as the proportion of successful interactions over all performed interactions until the n^{th} interaction. It is calculated using the following formula:

$$\mathrm{SR}(n) = \frac{1}{n} \sum_{k=1}^{n} \mathbb{I}[h_k^{\alpha_k}(o_k, t_k) = h_k^{\beta_k}(o_k, t_k)]$$

Where:

- α_k and β_k are agents randomly selected for the k^{th} interaction.
- o_k is the object and t_k is the task involved in the k^{th} interaction.
- $\mathbb{I}$ is the indicator function that returns 1 if the condition inside is true, and 0 otherwise.
- n is the total number of interactions.

Task accuracy & task accuracy derived measures

To evaluate the agents' performance on specific tasks, we adapt the accuracy measure introduced in [11]. We refer to this as task accuracy $TAC(\alpha, n, t)$. We define it as the proportion of object types for which a correct decision is made, with respect to a given task t, by an agent α on the n^{th} iteration of the experiment. It is calculated using the following formula:

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

Where:

- α is the agent whose accuracy is being evaluated.
- n is the iteration number.
- t is the specific task being considered.
- $\mathcal{I}$ is the set of all object types in the environment.
- o is an object type within the set $\mathcal{I}$.
- $h_n^\alpha(o, t)$ is the decision made by the agent α for the object o on the task t at the n^{th} iteration.
- $h^*(o, t)$ is the correct decision for the object o on the task t.

Given $TAC(\alpha, n, t)$ the tasks in which an agent α performs the best BPT and the worst WPT can be identified. BPT and WPT are defined as follows:

$$\mathrm{BPT}(\alpha, n) = \arg\max_{\forall t \in T} \mathrm{TAC}(\alpha, n, t)$$

$$\mathrm{WPT}(\alpha, n) = \arg\min_{\forall t \in T} \mathrm{TAC}(\alpha, n, t)$$

They allow for calculating an agent α's accuracy on its best and worst performing task respectively:

$$\mathrm{BTAC}(\alpha, n) = TA(\alpha, n, \mathrm{BPT}(\alpha, n))$$

$$\mathrm{WTAC}(\alpha, n) = TA(\alpha, n, \mathrm{WPT}(\alpha, n))$$

Throughout this thesis, the terms best task accuracy (BTAC) and worst task accuracy(WTAC) are interchangeable with maximum and minimum accuracy.

Lastly, an evaluation of an agent's mean performance over all tasks in $\mathcal{T}$, can be obtained by calculating its mean accuracy:

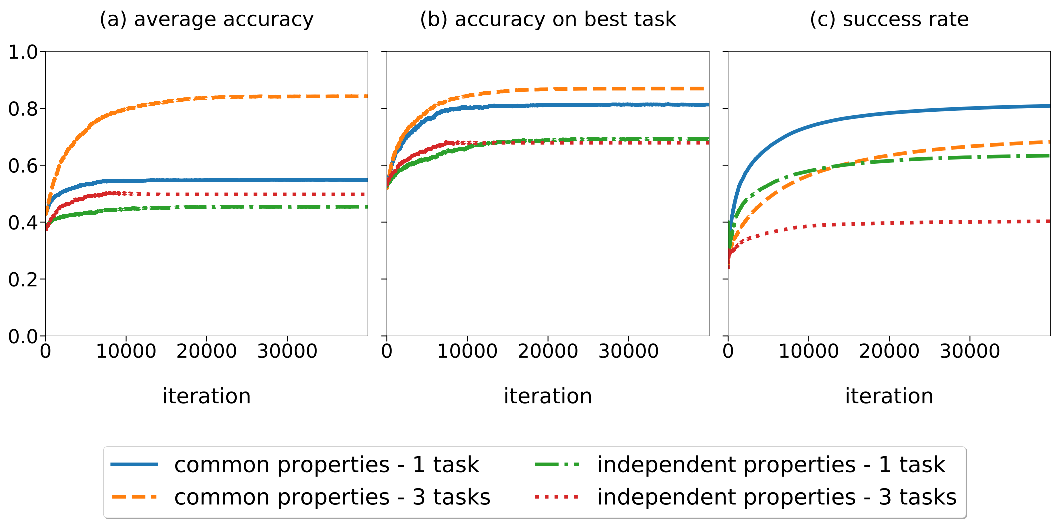

Figure 4.1 (a) average accuracy, (b) accuracy on best task and (c) success rate for different number of assigned tasks and common properties.

$$\text{AVAC}(\alpha, n) = \frac{1}{|T|} \sum_{t \in T} \text{TAC}(\alpha, n, t)$$

4.2 Results and discussion

Figure 4.1a shows that assigning more tasks to agents, 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 4.1b 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 ($p > 0.01$). 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 4.1c 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 one is that the fewer the assigned tasks, the fewer are the decisions over which agents need to agree. The second one 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. As shown by the blue curve in Figure 4.1c, the latter was shown true even when agents interact over the same single task.

4.3 Statistical Analysis

Analysis of variance shows that the number of common properties among different tasks, has a statistically significant impact ($p < 0.01$) 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 (maximum accuracy). The boxplots 4.2, 4.3 and 4.4 portray how each respective measure varies with the tasks overlap and the number of tasks and generations.

Table 4.2 Summary of p-values from one-way ANOVA for different numbers of tasks. The table is organized by measure and percentage of overlapping properties. It examines the effect of the number of tackled tasks for a given overlap percentage.

Measure	**0% Overlap**	**100% Overlap**
Average accuracy	1.60e-06	2.37e-18
Accuracy on best task	0.064	0.044
Success rate	1.30e-17	0.004

4.4 Conclusion

In this chapter we considered knowledge transferability, i.e., how knowledge acquired while doing one task can benefit a different task. We hypothesized that agents tackling tasks that rely on common properties, benefit from the formation

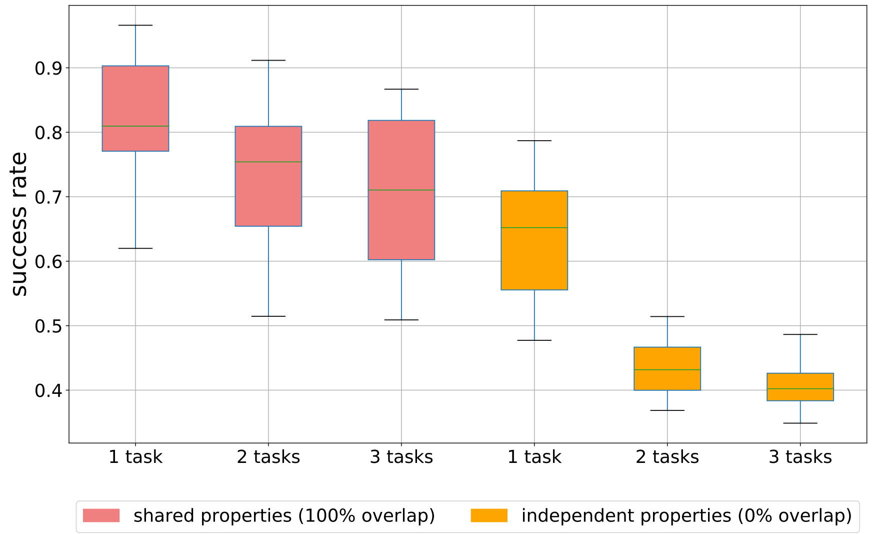

Figure 4.2 Success rate grouped by the number of tasks and generations. Boxes in red correspond to setups where the carried tasks depend on the same properties (100% overlap). Boxes in yellow correspond to setups where the carried tasks depend on different properties (0% overlap).

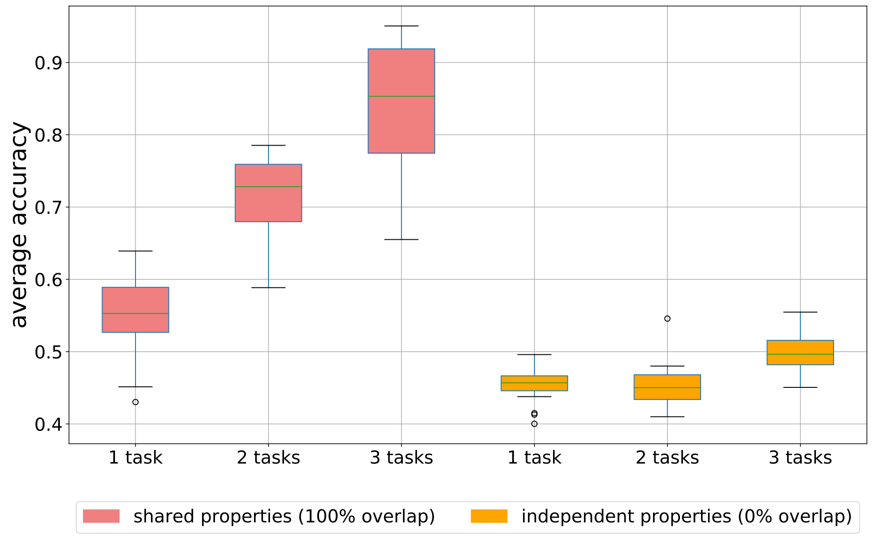

Figure 4.3 Average accuracy grouped by the number of tasks and generations. Boxes in yellow correspond to setups where the carried tasks depend on different properties (0% overlap).

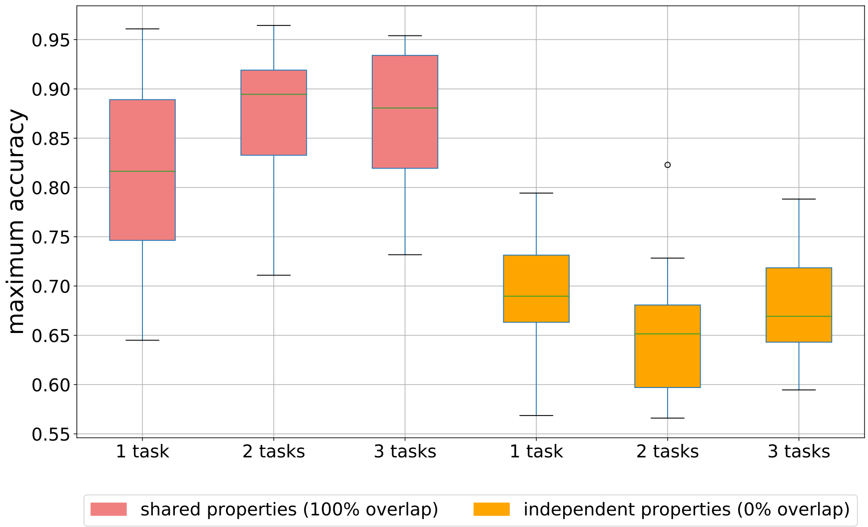

Figure 4.4 Maximum accuracy grouped by the number of tasks and generations. Boxes in yellow correspond to setups where the carried tasks depend on different properties (0% overlap).

Table 4.3 Summary of p-values from one-way ANOVA analysis for different numbers of tasks. The table is organized by measure and percentage of overlapping properties. It examines the effect of task overlap percentage for a given number of tasks.

Measure	1 Task	2 Tasks	3 Tasks
Average accuracy	5.72e-08	2.22e-21	4.32e-06
Accuracy on best task	5.52e-05	1.68e-12	5.92e-12
Success rate	8.08e-07	2.30e-13	2.22e-13

of multi-purpose knowledge. We test this hypothesis by introducing agents that tackle an increasing number of increasingly overlapping 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 setup of their environment. The agents 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 environments. These may serve as an insight on how agents evolve their knowledge within more complex environments. Such environments are examined in the following chapters.

— 5 —

Knowledge specialization in resource-limited agents

Chapter 4 focuses on how knowledge is transferred from one task to another, when taking accurate decisions for different tasks relies on shared properties. Among others, results showed that the more tasks overlap, the more beneficial multitasking is. However, little is known about agent specialization. Do agents specialize their knowledge when they focus on fewer tasks? Do any benefits arise when different tasks are assigned to different agents? We approach these questions from two perspectives. From an agent society's perspective, the interests of an agent may come in conflict with these of its society. Simply put, what is beneficial for the individual prosperity of an agent may not necessarily be equally beneficial for agent societies and vice versa. Our first hypothesis is that agents that undertake fewer tasks will achieve a higher accuracy on their best task than agents that undertake all available tasks. We assume that by interacting on a limited set of tasks, agents will refrain from polluting their ontologies with irrelevant knowledge. These agents are expected to be more accurate on some tasks than others, to the detriment of their average accuracy. Our second hypothesis is that the more accurate agents become on their best task, the more prosperous their societies are. Our assumption is that by analogy to our world, the agents' world will thrive through specialization.

We organize this chapter in two parts. In the first one, agents with unlimited memory undertake a limited set of tasks, henceforward referred to as scope. Taking correct decisions for each task relies on different properties. For example, deciding on whether to take an umbrella may rely on the property *relative_humidity* while deciding on whether to wear a t-shirt may rely on the property *temperature*. In the second part, agents undertake the same tasks, while facing memory limitations. Likewise Chapter 4, we implement these limitations by setting a maximum number of classes to be maintained in an agent's ontology. When the agents reach these limitations, they remove parts of their knowledge that are not relevant to their scope. To that end, agents employ knowledge generalization, presented in Subsection 3.3.4. Two sets of measures are used to evaluate the results. The first set evaluates results from the perspective of individual agents and allow for assessing whether an agent is specialized. We consider agents specialized if the following criteria are met:

1. Agents average accuracy is lower than their accuracy on their best task.
2. Agents tackling several tasks are less accurate on their best tasks, than agents that tackle a single task.

The second set evaluates results from the perspective of agent societies. From this perspective we are interested among others in: (1) the ontologies' contribution in promoting agreement among agents, (2) the societies' total collectively accumulated points, and (3) how these points are distributed among the agents. Section 5.1 presents a set of hypotheses along with the protocol used for testing them. Test results are presented in §5.2.

5.1 Experimental setting

5.1.1 Hypotheses

Hypothesis 1: The fewer tasks the agents tackle, the higher is their accuracy on their best task.

Hypothesis 2: The more agents improve their best task accuracy, the more prosperous and equitable their populations are.

5.1.2 Parameters

Each experiment is executed under 12 setups, corresponding to agents that tackle 1 to 3 independent tasks with their memory limited to 4, 8, 12 and 64 (unlimited memory) classes respectively. Each setup is run 20 times and its results are averaged. One run consists of 80000 interactions with each interaction taking place among two agents. These agents are randomly selected out of a total population of 18 agents. Depending on the setup, each agent is assigned 1 to 3 tasks. When the agents' memory is exhausted, they free it by generalizing their knowledge (Subsection 3.3.4). Their environment contains 64 different object types, each one perceivable through 6 different binary properties. The agents are initially trained with respect to all $|\mathcal{T}| = 3$ tasks. For each task, 4 different decisions exist. Choosing one decision relies on 2 out of the 6 binary properties. Since the 3 tasks do not overlap, deciding with respect to different tasks relies on a different pair of properties. Each agent induces an initial ontology based on a randomly selected 10% of all existing labeled (training) examples. As a consequence, there may be overlap between the training of two agents. The agents subsequently undertake tasks within their scope. To decide which agent adapts its knowledge in case of disagreement, agents compare their scores. These are attributed to them by the environment between two consecutive interactions. They are calculated based on a randomly selected corpus consisting of 60% of all existing training samples.

Table 5.1 Summary table of parameters (independent variables)

Parameter	**Value**
Runs	20
Interactions	80000
Generations	1
Agents $\|\mathcal{A}\|$	18
Memory capacity	{4,8,12,64} classes
Memory release technique	Generalization
Tasks $\|\mathcal{T}\|$	3
Assigned tasks (scope)	{1,2,3}
Object types $\|\mathcal{I}\|$	64
Properties	6
Properties per task $\|\mathcal{D}_t\|$	2
Task overlap	0%
Decisions per task	4
Training ratio	0.1
Score ratio	0.6

5.1.3 Measures

Additionally to the measures defined and used in Chapter 4, here we define : (1) the scope accuracy (SCAC), (2) the correct decision rate (CDR) and (3) the P90/P10 of agents' accumulated points.

Scope accuracy

Similarly to different accuracy measures defined in Chapter 4, the scope accuracy (SCAC) evaluates the knowledge correctness of agents. In particular it evaluates the average accuracy of agents, in the tasks they undertake. Scope accuracy is based on the tasks each agent undertakes. These constitute a subset of $\mathcal{T}$, denoted as $\mathcal{S}$. Based on this, scope accuracy is defined as follows:

$$\mathrm{SCAC}(\alpha, n, S) = \frac{1}{|S|} \sum_{t \in S} \mathrm{TAC}(\alpha, n, t)$$

Where:

- α is the agent whose accuracy is being evaluated.
- n is the iteration number.
- t is the specific task being considered.
- S is the set of tasks each agent undertakes.
- $\mathrm{TAC}(a, n, t)$ is the accuracy of agent α, on task t, for the iteration n (defined in Chapter 4).

Correct decision rate

In order to evaluate the efficiency of an agent population in taking decisions, the measure of correct decision rate (CDR) is introduced. It is defined as the proportion of correct decisions taken, over all decisions taken until the n^{th} interaction. When the two interacting agents agree, this decision (h_k^{final}) is the one taken by both agents. When the two interacting agents disagree, the decision taken into account is that of the agent exhibiting the highest accuracy in the currently involved task. We formulate correct decision rate as follows:

$$\mathrm{CDR}(n) = \frac{1}{n} \sum_{k=1}^{n} \left[\mathbb{I}(h_k^{\text{final}} = h^*(o_k, t_k)) \right]$$

Where:

- n is the total number of interactions or iterations until now.
- k is the current interaction.
- o_k is the object involved in the k-th interaction.
- t_k is the task involved in the k-th interaction.
- α is the first randomly selected agent involved in the k-th interaction.

- β is the second randomly selected agent involved in the k-th interaction.
- $h^{\alpha}(o_k, t_k)$ is the decision made by agent α for object o_k and task t_k.
- $h^{\beta}(o_k, t_k)$ is the decision made by agent β for object o_k and task t_k.
- $h^{*}(o_k, t_k)$ is the correct decision for object o_k and task t_k.
- $\mathrm{TAC}(\alpha, k, t_k)$ is the Accuracy of agent α at iteration k for task t_k.
- $\mathrm{TAC}(\beta, k, t_k)$ is the Accuracy of agent β at iteration k for task t_k.
- h_k^{final} is the final decision made in the k-th interaction.

And h_k^{final} is (defined as follows:

$$h_k^{\text{final}} = \begin{cases} h^{\alpha}(o_k, t_k) & \text{if } h^{\alpha}(o_k, t_k) = h^{\beta}(o_k, t_k) \\ h^{\alpha}(o_k, t_k) & \text{if } h^{\alpha}(o_k, t_k) \neq h^{\beta}(o_k, t_k) \text{ and } \mathrm{TAC}(\alpha, k, t_k) > \mathrm{TAC}(\beta, k, t_k) \\ h^{\beta}(o_k, t_k) & \text{if } h^{\alpha}(o_k, t_k) \neq h^{\beta}(o_k, t_k) \text{ and } \mathrm{TAC}(\beta, k, t_k) > \mathrm{TAC}(\alpha, k, t_k) \end{cases}$$

Decile dispersion ratio (P90/P10)

In order to evaluate the inequality in points acquired within agent populations, an adaptation of the decile dispersion ratio (P90/P10) [19] is used. We define it as the ratio of the sum of points collectively acquired by the poorest 90% of the agent population (sum_{P90}), over the sum of points collectively acquired by the richest 10% of the agent population (sum_{P10}). This measure thus expresses the prosperity of the richest agents, as multiples of that of the poorest ones. It is calculated as follows:

- Let $A = \{a_1, a_2, \ldots, a_m\}$ be the set of m agents participating in the experiment.
- Let P_i^n be the total points accumulated by agent a_i after n iterations of the experiment, where i ranges from 1 to m.
- Let $\mathcal{P}' = \{P_1', P_2', \ldots, P_m'\}$ be the sorted set of points such that $P_1' \leq P_2' \leq \ldots \leq P_m'$.

The 90th Percentile is the value within which 90% of the points in $\mathcal{P}$ fall. It is calculated as follows:

1. Compute the position $k_{90} = 0.9 \times (m + 1)$.
2. Separate k_{90} into its integer part $\lfloor k_{90} \rfloor$ and fractional part f_{90}, such that $k_{90} = \lfloor k_{90} \rfloor + f_{90}$.
3. Compute the P90:

$$P90 = \begin{cases} P'_{\lfloor k_{90} \rfloor} & \text{if } k_{90} \text{ is an integer} \\ P'_{\lfloor k_{90} \rfloor} + f_{90} \times (P'_{\lfloor k_{90} \rfloor + 1} - P'_{\lfloor k_{90} \rfloor}) & \text{otherwise} \end{cases}$$

The 10th Percentile is the value below which 10% of the points in $\mathcal{P}$ fall. It is calculated as follows:

1. Compute the position $k_{10} = 0.1 \times (m + 1)$.
2. Separate k_{10} into its integer part $\lfloor k_{10} \rfloor$ and fractional part f_{10}, such that $k_{10} = \lfloor k_{10} \rfloor + f_{10}$.
3. Compute the P10:

$$P10 = \begin{cases} P'_{\lfloor k_{10} \rfloor} & \text{if } k_{10} \text{ is an integer} \\ P'_{\lfloor k_{10} \rfloor} + f_{10} \times (P'_{\lfloor k_{10} \rfloor + 1} - P'_{\lfloor k_{10} \rfloor}) & \text{otherwise} \end{cases}$$

The respective sums of points within the 90th and the 10th percentiles are calculated as follows:

$$\text{sum}_{P90} = \sum_{i=1}^{\lfloor k_{90} \rfloor} P'_i + f_{90} \times P'_{\lfloor k_{90} \rfloor + 1}$$

$$\text{sum}_{P10} = \sum_{i=1}^{\lfloor k_{10} \rfloor} P'_i + f_{10} \times P'_{\lfloor k_{10} \rfloor + 1}$$

Finally:

$$\text{DDR} = \frac{\text{sum}_{P90}}{\text{sum}_{P10}}$$

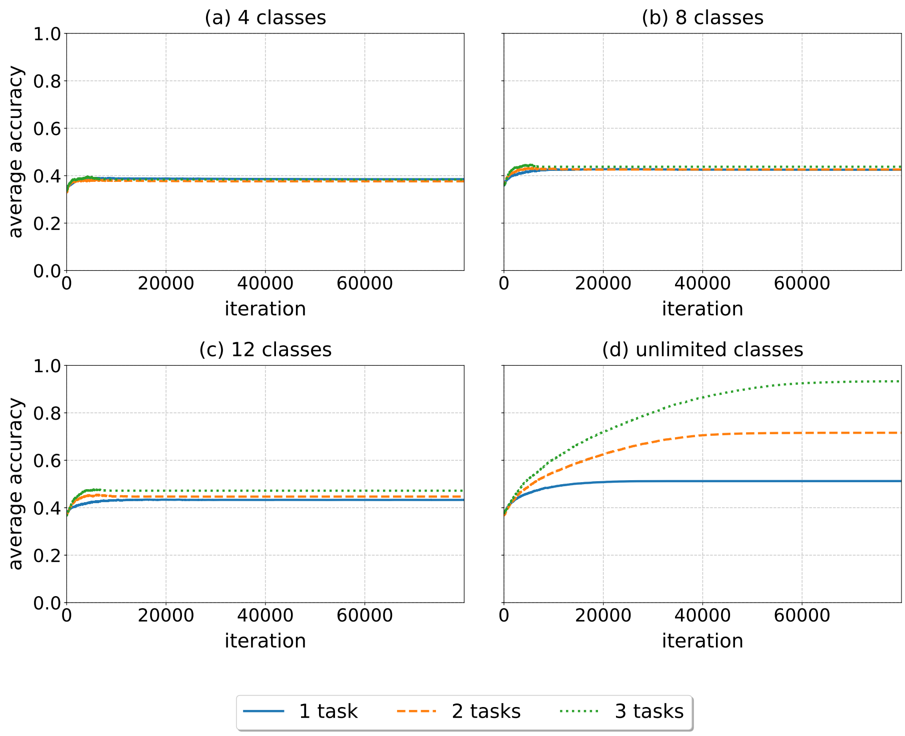

Figure 5.1 Average accuracy for different memory and scope sizes. Different subplots correspond to different memory sizes (4, 8, 12 and unlimited classes). Different curves correspond to different scope sizes $|scope| = \{1, 2, 3\}$.

5.2 Results and discussion

Here we present and discuss results acquired by varying the agents' scope and memory size. Each figure consists of 4 subfigures, each one corresponding to a different memory size (4, 8, 12 and unlimited classes). Within such subfigure, different scope sizes are represented by curves of different color and pattern.

5.2.1 From an agent's perspective

Figure 5.1 portrays the evolution of the average accuracy for different scope and memory sizes. It shows that the agents average accuracy is not affected by the size of their scope. This is due to the imposed agent limitations. Our environment consists of 3 tasks, each one relying on 2 out of 6 total properties. The agents' ontologies are restricted in such a way, that on average they contain 2 different properties. Two cases can be distinguished. In the first case, an agent learns to take highly accurate decisions for one task and inaccurate decisions for the

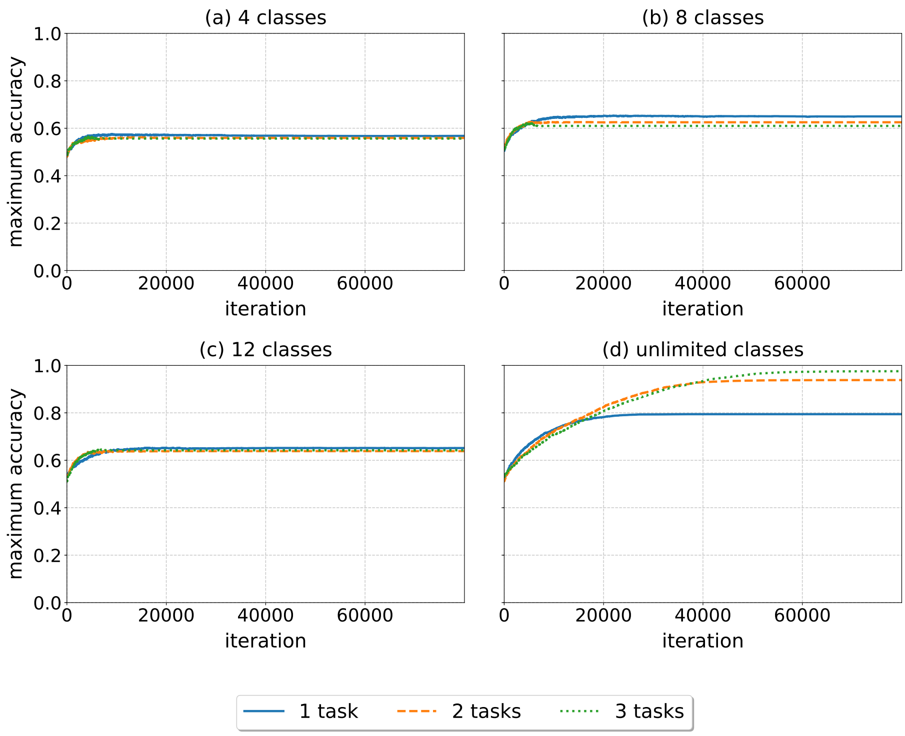

Figure 5.2 Maximum accuracy for different memory and scope sizes. Different subplots correspond to different memory sizes (4, 8, 12 and unlimited classes). Different curves correspond to different scope sizes $|scope| = \{1, 2, 3\}$.

remaining tasks. In the second case, an agent learns an ontology that allows it to take averagely accurate decisions for several tasks. This translates into agents that demonstrate average accuracies that are statistically indistinguishable to each other.

Figure 5.2 portrays the evolution of agents' accuracy on their best task (maximum accuracy), for different scope and memory sizes. The sub-figure 5.2(d), corresponds to unrestricted agents. The sub-figures 5.2(a), 5.2(b) and 5.2(c), correspond to agents with ontologies limited to 4, 8 and 12 leaf classes respectively. Examining the sub-figure 5.2(d), the following observation may be drawn. Agents that have unlimited memory do not benefit from limiting their scope. More specifically, Figure 5.2(d) shows that agents tackling fewer tasks, demonstrate a lower best task accuracy compared to agents that tackle all tasks. In conjunction with Figure 5.1, the sub-figures (a), (b) and (c) show that the best accuracy agents whose memories are limited to 4, 8 and 12 classes is significantly higher than their average accuracy. Hence, these agents specialize on specific tasks. However, this specialization is not related to the size of their scope, but rather to their memory limitations. These agents can become accurate at most on one task, regardless of the size of their scope. In other words, the fact that agents may abstain from all interactions but those that regard the single task they undertake does not allow them to improve their average accuracy on their best task (maximum accuracy). On the one hand, the best task accuracy of restricted agents is not affected by the size of their scope. On the other hand, the best task accuracy of agents with unlimited memory is higher when agents tackle all tasks. Based on Figures 5.1 and 5.2, hypothesis 1 is rejected.

Additionally, results show that when agents have limited resources, the goals of maximizing accuracy in specific tasks and reaching consensus are mutually exclusive. Put differently, agents can either specialize, or learn to agree on their decisions. When an agent population favors the former, the vast majority of agents do not become 100% accurate on the task they specialize. We found that this is due to the way agents free their memory after its exhaustion. Examining different agent ontologies showed the following: the closer to the root of an ontology the irrelevant knowledge is, the more difficult it is to remove it. We refer to this type of knowledge as parasitic and discuss how to eliminate it in Chapter 6.

5.2.2 From a society's perspective

Figure 5.3 displays the evolution of the average success rate for different scope sizes and agent limitations. It shows that the success rate of restricted agents stabilizes, yet does not converge to 1. This indicates that when the memory of agents is limited to 4, 8 and 12 classes, their final ontologies do not allow for agreeing on all decisions. We assume that if restricted enough, an ontology will be able to contain only the properties required to take the correct decision with respect to a single task. As a result, it is to be expected that agents interacting over a number of tasks that requires more memory, will not be able to converge on all decisions. Yet, the latter is shown to be true even when agents interact over a single task. Due to the initial ontology induction algorithm, the ontology of an agent that undertakes the task t may contain properties that do not belong to $\mathcal{P}_t$. At the same time, no guarantee is offered that the agents will replace these properties with properties that belong in $\mathcal{P}_t$. Two distinct cases can be envisioned. In the first case, an agent will gradually discard all properties that are not related to the task it undertakes. This agent will learn an ontology that will allow for taking correct decisions with respect to all encountered object types. In the second case, an agent will fail to discard all the properties that are not related to the task it undertakes. Agents that fall under the second case will repeatedly replace properties that belong in $\mathcal{P}_t$ with different properties that belong in $\mathcal{P}_t$ as well. As a result, these agents are able to take correct decisions for different subsets of the existing object types at a given time instance.

Figure 5.4 displays the evolution of the total accumulated points for different

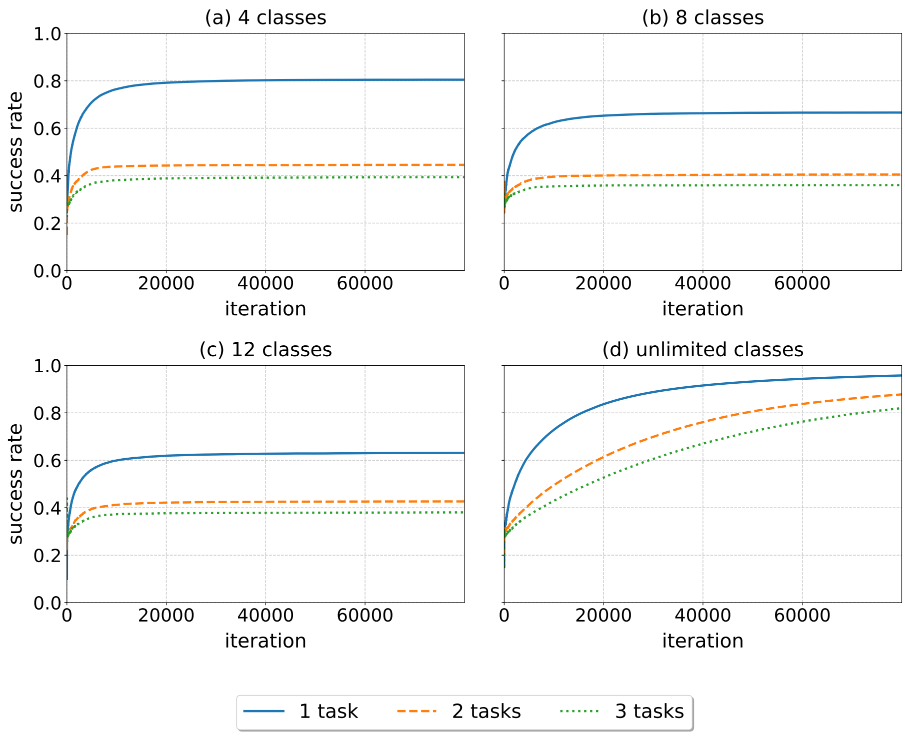

Figure 5.3 Average success rate for different memory and scope sizes. Different subplots correspond to different memory sizes (4, 8, 12 and unlimited classes). Different curves correspond to different scope sizes $|scope| = \{1, 2, 3\}$.

scope and memory sizes. It shows that the smaller is the agents scope, the higher is the number of total points accumulated by their population. Results thus support the hypothesis 2.

Figure 5.5 portrays the evolution of the decile dispersion ratio (P90/P10) of these points, for different scope sizes. Results show that for most examined setups the higher the number of tasks assigned to agents, the less equitable points distribution is. This can be explained as follows. The more agents coordinate over fewer shared tasks, the more probable it is that their ontologies will share properties as well. Based on this, these agents will be comparably accurate on these tasks and agree more often. As a result, these agents will also more often fail or succeed in pairs, sharing any acquired compensation. Let us consider agents whose scope consist of a single task. In the examined setup, deciding for this task relies on 2 properties and the agent ontologies are limited to 4 classes. Given their initial training, these agents may find themselves unable to discard all the properties that are not related to the tasks within their scope. This is especially true when these unrelated properties are found closer to their ontology's root. The more the

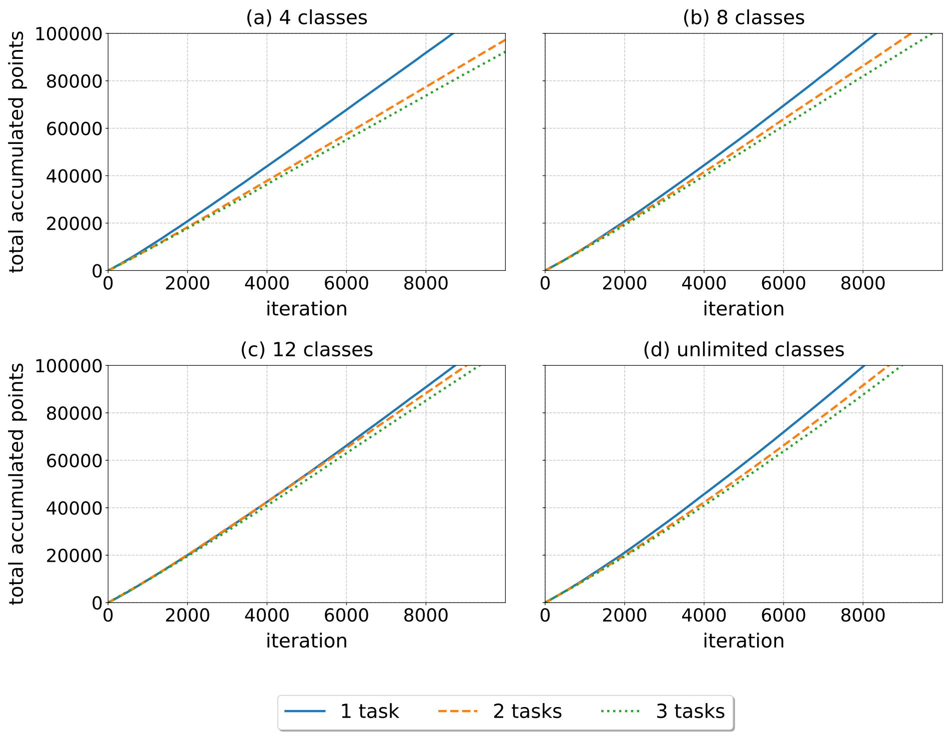

Figure 5.4 Total accumulated points for different memory and scope sizes. Different subplots correspond to different memory sizes (4, 8, 12 and unlimited classes). Different curves correspond to different scope sizes $|scope| = \{1, 2, 3\}$.

agents tackle the same task, the more probable it is that the remaining properties are shared among them. On the contrary, when agents tackle all tasks, the number of different properties arrangements is higher. As seen in Figure 5.3, this leads to agents that disagree, even when deciding with respect to a common task. Given our protocol, this means that even if the taken decision is accurate, the respective points will not be shared among the two interacting agents.

Closing, it is important to highlight that specializing in fewer tasks was found to be beneficial only when each task was assigned to an equal number of agents. When this was not the case, agents collectively specialized in a small subset of tasks. As a result, these agent populations became highly accurate in these tasks, but also highly inaccurate in the remaining ones. Whether and how agents can become equally accurate on all tasks is examined in Chapter 7.

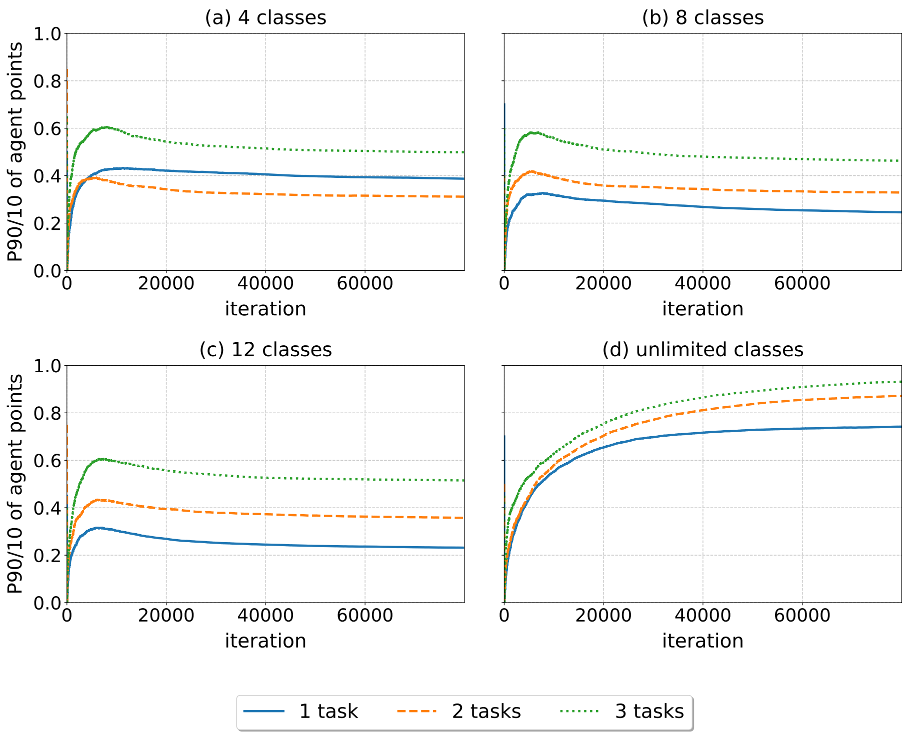

Figure 5.5 Average points dispersion ratio for different memory and scope sizes. Different subplots correspond to different memory sizes (4, 8, 12 and unlimited classes). Different curves correspond to different scope sizes $|scope| = \{1, 2, 3\}$.

5.3 Statistical Analysis

Our analysis of variance (one-way ANOVA) shows that the agents scope size, has a statistically significant impact ($p \leq 0.01$) on (1) the success rate, (2) the total accumulated points and (3) the points distribution. The size of the scope was not shown to have a statistically significant impact ($p > 0.01$) on the agents (1) average accuracy and (2) best task accuracy. This was found to be also true for agents whose ontologies are limited to 8 and 12 classes. The boxplots 5.6, 5.7, 5.8, 5.9, 5.10 and 5.11 portray how each respective measure varies with the memory and the scope size of agents. They are grouped by memory and scope size with different colors corresponding to different memory sizes (4, 8, 12 and unlimited classes) and bar colors corresponding to different number of tasks.

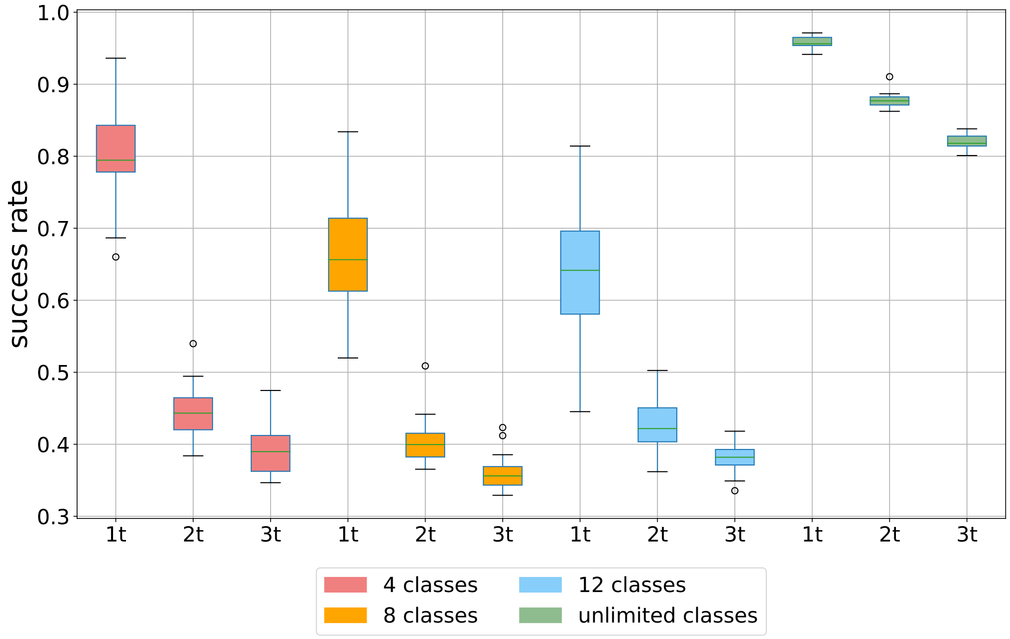

Figure 5.6 Success rate.

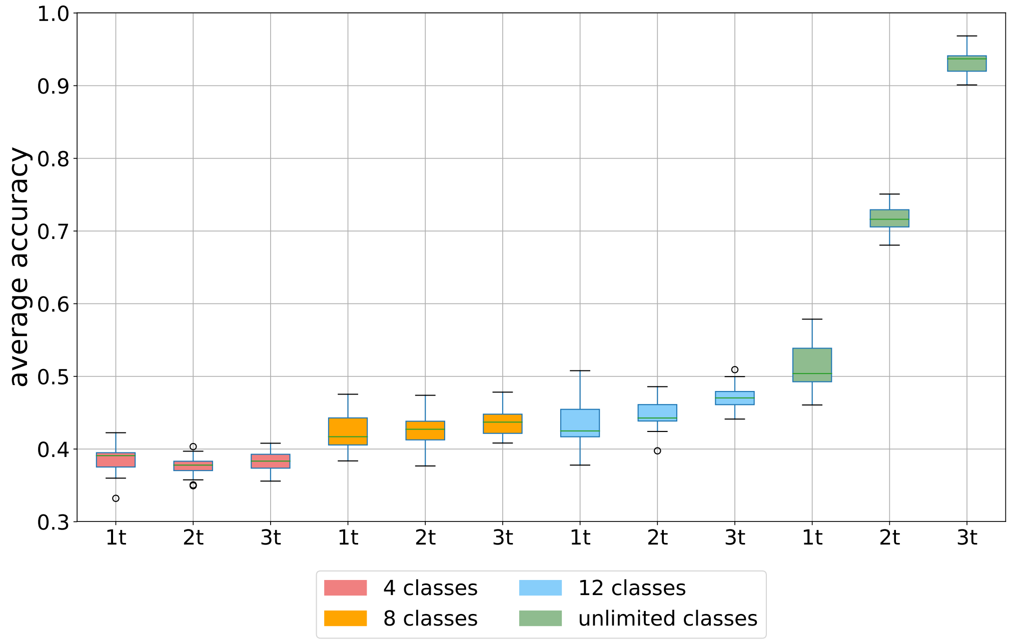

Figure 5.7 Average accuracy.

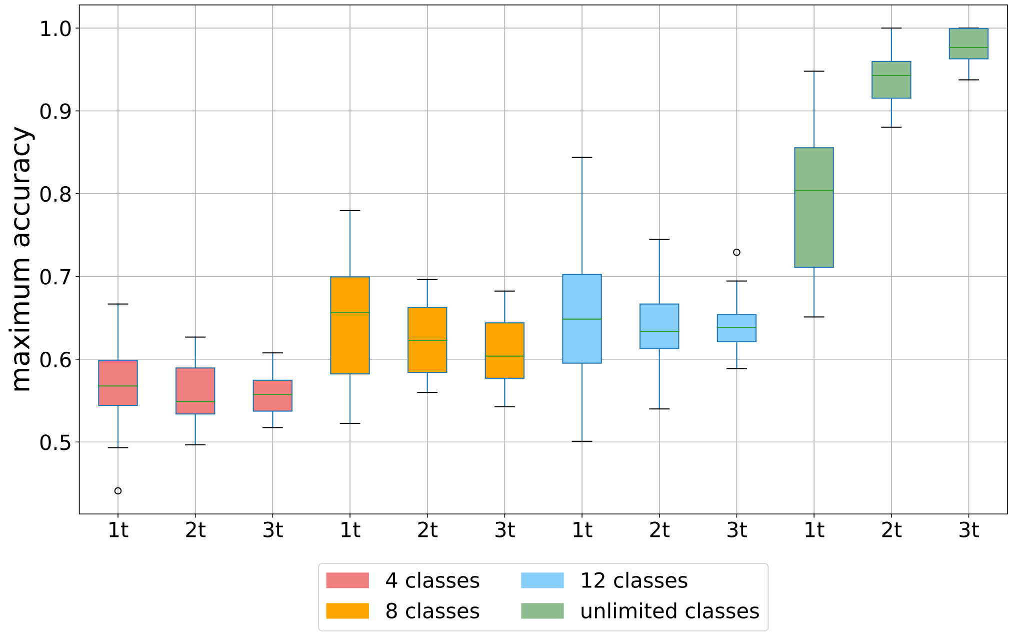

Figure 5.8 Maximum accuracy.

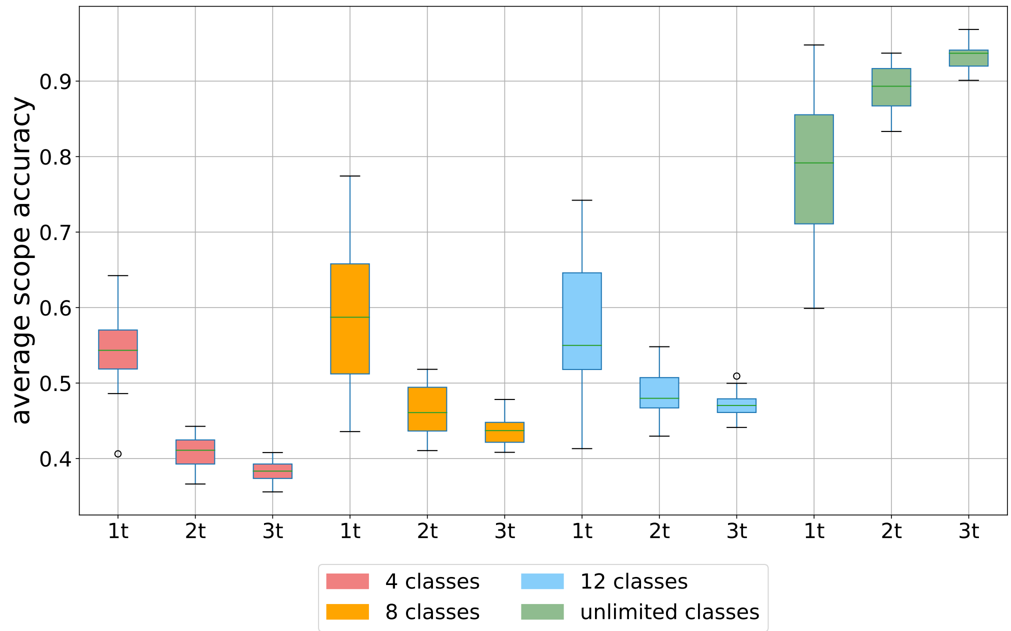

Figure 5.9 Scope accuracy.

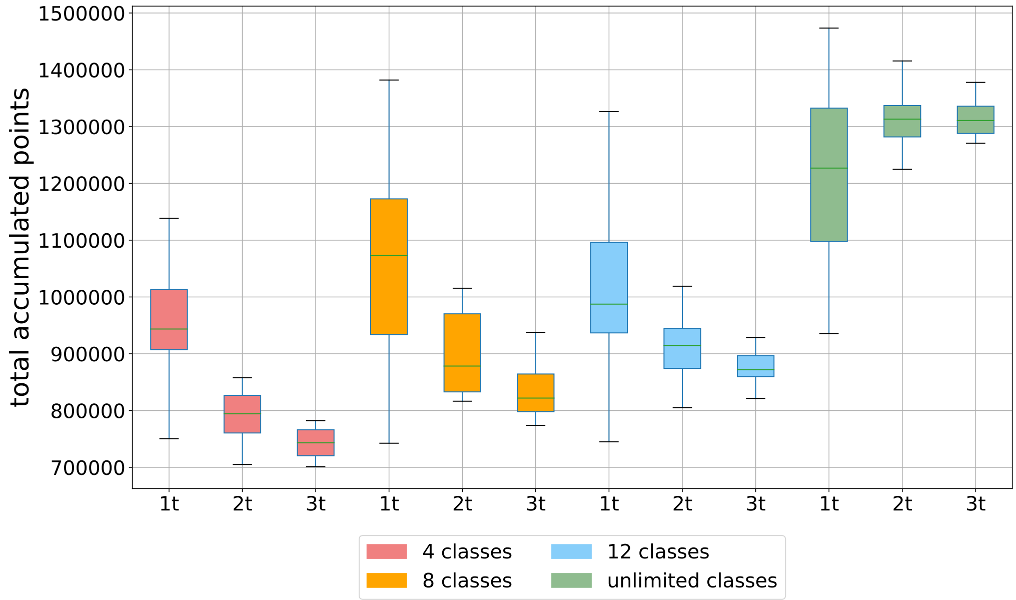

Figure 5.10 Total accumulated points.

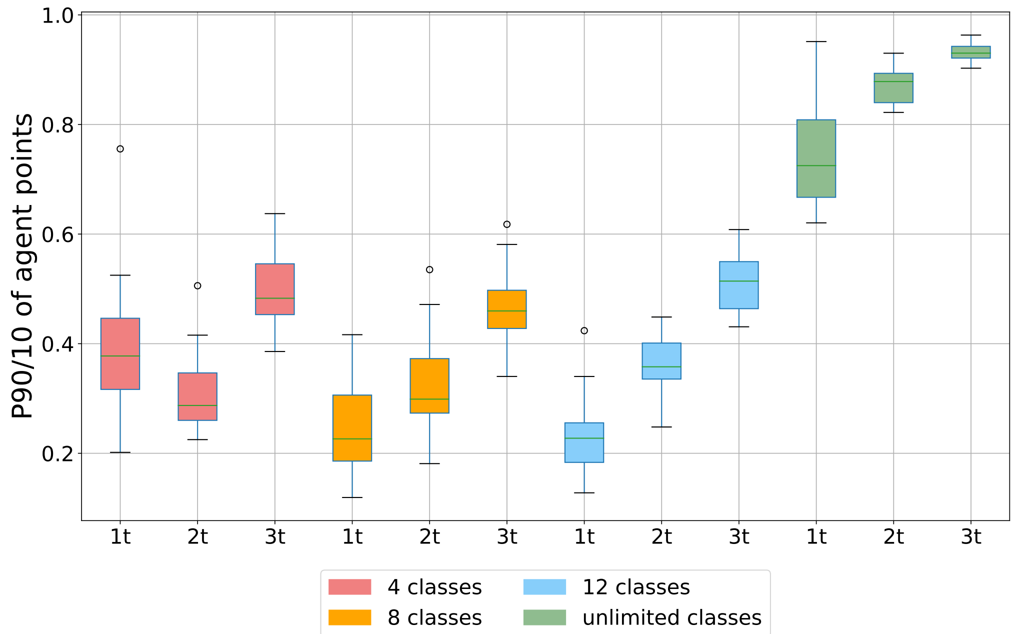

Figure 5.11 Total points decile ratio.

Table 5.2 Summary of p-values from one-way ANOVA for different numbers of tasks. The table is organized by measure and memory size. It examines the effect of agents' scope for a given memory size.

Measure	**4 Classes**	**8 Classes**	**12 Classes**	**Unlimited Classes**
Average accuracy	0.180	0.160	8.23e-06	1.17e-49
Accuracy on best task	0.640	0.090	0.750	1.06e-14
Success rate	2.64e-33	2.91e-25	1.46e-19	5.34e-46
Total compensation	2.56e-15	1.86e-08	2.56e-05	0.004
Compensation P90-10	3.57e-07	2.01e-10	2.16e-20	7.99e-15

5.4 Conclusion

In this chapter, we considered how assigning different tasks to various agents affects both individual agents and agent populations. From an individual agent perspective, we examined whether focusing on a limited scope of tasks improves each agent's accuracy in their best task. Results show that restricting the agents' memory leads to knowledge specialization, regardless of the size of their scope. From a population perspective, we explored whether assigning different tasks to different agents leads to more prosperous and equitable populations. Results show that by assigning fewer tasks to agents, agent populations not only collectively accumulate more points but also distribute these points more evenly among their agents.

— 6 —

Influence of knowledge forgetting on long-term correctness

Chapter 5 focuses on how specializing on a limited subset of tasks called scope, affects the prosperity of individual agents and agent societies. Among others, the results acquired from it revealed the presence of parasitic knowledge. Can agents eliminate this parasitic knowledge without sacrificing the benefits of knowledge specialization? To answer this question, we extend the specialization experiment by introducing two modifications.

The first modification affects how agents remove knowledge. In Chapter 5, agents will avoid to remove knowledge if this has a negative impact on their scope accuracy. However, results acquired and discussed in Chapter 5, show that this hinders agent specialization. Given the way our agents train and evolve their ontologies, their knowledge is not necessarily mutually complementary. In other words, more than one agents may possess the same knowledge. We assume that if agents do not remove the same knowledge at the same time, even if an agent loses useful knowledge, this will be eventually relearned from agents who still retain it. Our main hypothesis is that discarding randomly selected knowledge will allow for eliminating unrelated properties and relearning useful knowledge will allow agents to achieve consensus and maximize their maximum accuracy. We expect this to translate into a higher and more optimized specialization. We implement this modification by allowing agents to remove randomly selected classes from their ontologies. Results show that while removing randomly selected knowledge leads to a short-term accuracy deterioration, it allows agents to eliminate parasitic knowledge. This practically means that in order for an agent to become highly accurate on a task, it must temporarily discard accurate and relevant knowledge. It is important to emphasize the following. The above benefits apply to populations that specialize on a single task. Otherwise, this knowledge removal mechanism leads to a gradual degeneration of what agents collectively know for each task. Therefore, if agents specialize on more than one task, freeing memory using the generalization mechanism described in Subsection 3.3.4 leads to more accurate agents.

The second modification consists of periodically replacing agents with new, naive ones. Previous work [11] has shown that when experiments span over several generations, agents improve both their average accuracy and communication. We assume that the introduction of generations will enhance the elimination of parasitic knowledge.

6.1 Experimental setting

6.1.1 Hypothesis

Agents that free memory by removing randomly selected knowledge, will reach consensus after a finite number of interactions and maximize their maximum accuracy.

6.1.2 Parameters

The experiment is executed under 12 setups, acquired by varying: (1) the number of tasks assigned to each agent, (2) the number of agent generations and (3) the mechanism used for freeing memory. Each setup is run 10 times and its results are averaged. One run consists of 80000 interactions with each interaction taking place among two agents. These agents are randomly selected out of a total population of 36 agents. Depending on the setup, each agent is assigned 1 to 3 tasks and has an ontology (memory) capable of accommodating a maximum of 8 classes. Each run spans over either 1 or 4 generations. When a run spans over 4 generations, one generation occurs every 20000 interactions. When a generation occurs, the 18 longest existing agents are replaced by 18 new agents. When the agents' memory is exhausted, they free it either by generalizing their knowledge (Subsection 3.3.4), or by removing randomly selected classes (Subsection 3.3.4). The agents' environment contains 64 different object types, each one perceivable through 6 different binary properties. The agents are initially trained with respect to all $|\mathcal{T}| = 3$ tasks. For each task, 4 different decisions exist. Deciding with respect to different tasks always relies on 2 out of the 6 binary properties. The 3 tasks do not overlap, thus deciding with respect to each one relies on a different pair of properties. To induce its initial ontology, each agent takes into account a randomly selected 10% of all existing labeled examples. As a consequence, there may be overlap between the training of two agents. The agents subsequently undertake tasks within their scope. To decide which agent adapts its knowledge in case of disagreement, agents compare their scores. These are attributed to them by the environment between two consecutive interactions. They are calculated based on a randomly selected corpus consisting of 60% of all existing training samples.

Table 6.1 Summary table of parameters (independent variables)

Parameter	**Value**
Runs	10
Interactions	80000
Generations	{1,4}
Agents $\|\mathcal{A}\|$	36
Replaced agents	18
Memory capacity	8 classes
Memory release technique	{Generalization, Random class removal}
Tasks $\|\mathcal{T}\|$	3
Assigned tasks (scope)	{1,2,3}
Object types $\|\mathcal{I}\|$	64
Properties	6
Properties per task $\|\mathcal{D}_t\|$	2
Task overlap	0%
Decisions per task	4
Training ratio	0.1
Score ratio	0.6

6.2 Results and discussion

Here, we present and discuss results acquired by varying how agents remove parts of their knowledge in response to memory exhaustion.

Figure 6.1 displays the evolution of the average success rate for different scope sizes, number of generations and knowledge removal mechanisms. It shows two things. First, generalizing agents do not reach consensus, regardless of whether the experiment spans over one or several generations. This confirms the results acquired and discussed in Chapter 5. Second, agents removing randomly selected knowledge will reach consensus. However, results show that this is only possible when agents accomplish a single task. Hypothesis is thus partially supported. Additionally, results show that when the experiment spans over several generations, the success rate does not increase monotonically. This was found to be true for both of the examined mechanisms.

Figure 6.2 displays the evolution of agents' average accuracy for different scope sizes, number of generations and knowledge removal mechanisms.

Figure 6.3 displays the evolution of best task (maximum) accuracy for different

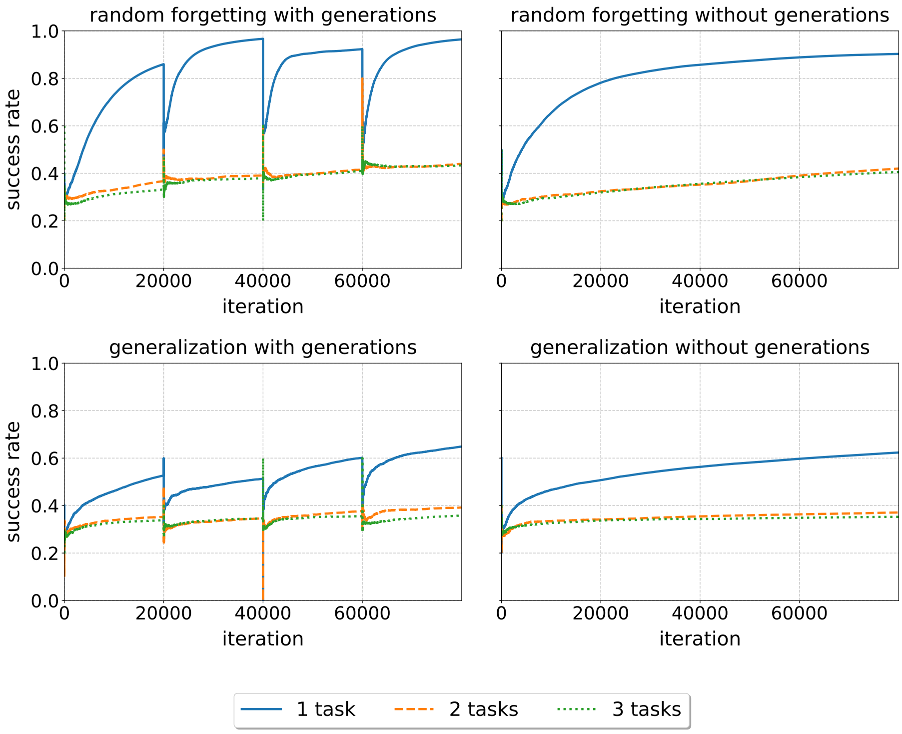

Figure 6.1 Average success rate for different memory and scope sizes. Different subplots correspond to different combinations of number of generations and knowledge removal mechanisms. Different curves correspond to different scope sizes $|scope| = \{1, 2, 3\}$.

scope sizes, number of generations and knowledge removal mechanisms.

Concerning agents forgetting randomly selected knowledge, both Figure 6.2 and Figure 6.3 lead to the same conclusion: forgetting randomly selected knowledge is beneficial for agents coordinating over a single shared task and detrimental to agents coordinating over several tasks. When agents coordinate over a single task, they achieve the highest possible maximum accuracy by maintaining more properties required for this task, until the end of the experiment. This is because the presence of more properties required for accomplishing this task, allows for higher discrimination among the available decisions and thus better selection. When agents coordinate over several tasks, forgetting randomly selected knowledge significantly deteriorates their maximum accuracy. This is because they disagree more often. Given their ontologies' limited size, agents are found dropping and adopting properties that are important for different tasks at each iteration. Given the pressure on agent agreement, this forces agents' ontologies to gradually converge on the same properties, until the agents agree on low accuracy knowledge.

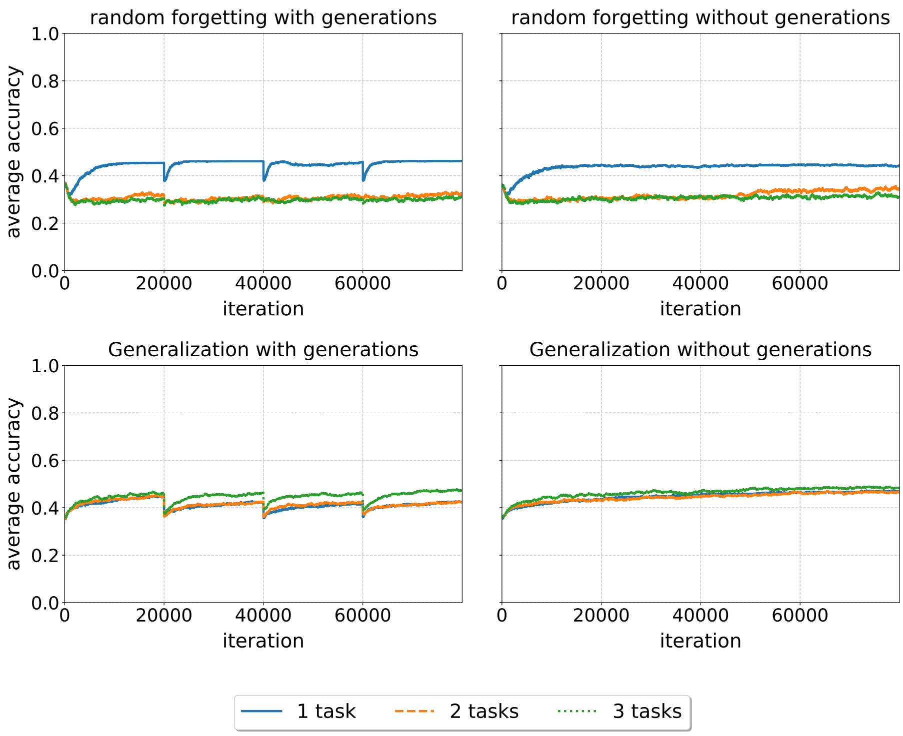

Figure 6.2 Average accuracy by number of generations, knowledge forgetting mechanism and scope sizes. Different subplots correspond to different combinations of number of generations and knowledge removal mechanisms. Different curves correspond to different scope sizes $|scope| = \{1, 2, 3\}$.

The same tendencies are observed with respect to the agents' average accuracy. However, the benefit from randomly forgetting knowledge when coordinating over a single shared task is limited. This can be explained by the fact that the average accuracy of single-tasking agents is mainly affected by their maximum accuracy, in this case their accuracy on the single assigned task. The accuracy of the remaining tasks remains low in most cases, as the properties present in the agents' ontologies do not allow for accurate discrimination.

Concerning agents that forget through generalization, both Figure 6.2 and Figure 6.3 confirm the results acquired in Chapter 5: the number of tasks each agent is assigned, has a statistically insignificant effect on the observed maximum and average accuracy.

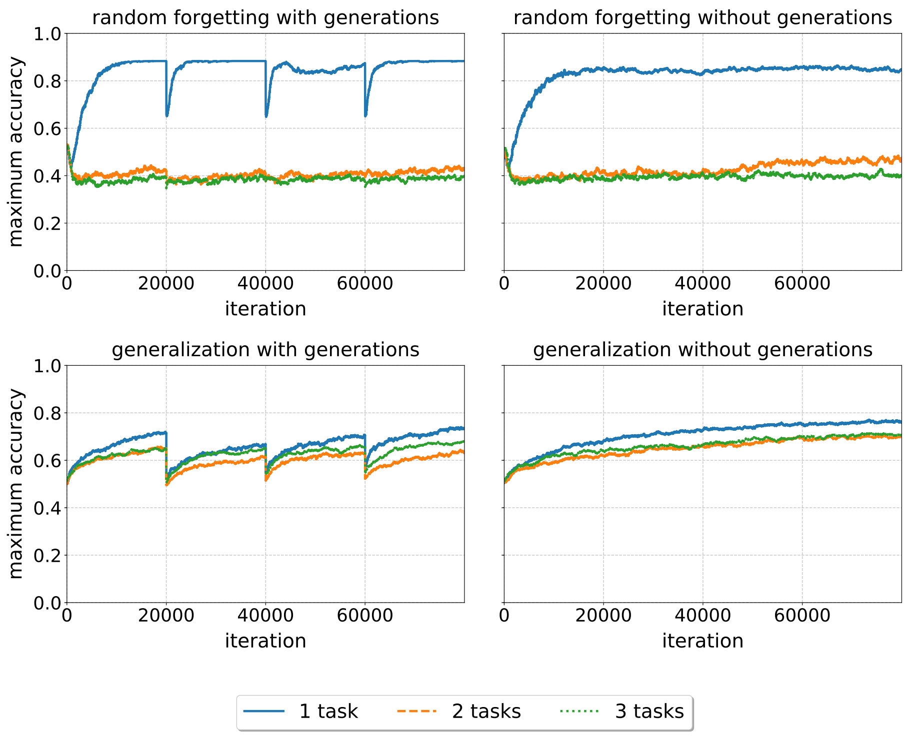

Figure 6.3 Maximum accuracy for different number of generations, forgetting mechanism and scope sizes. Different subplots correspond to different combinations of number of generations and knowledge removal mechanisms. Different curves correspond to different scope sizes $|scope| = \{1, 2, 3\}$.

6.3 Statistical Analysis

Analysis of variance shows the following three. First, when the knowledge removal mechanism and scope are fixed, the number of generations is insignificant for the final success rate, maximum accuracy and scope accuracy values. Second, when the scope and the generation interval are fixed, the knowledge removal mechanism is significant for the final success rate values. For the final maximum and scope accuracy values, the knowledge removal mechanism is significant unless the experiment spans over 4 generations and agents are assigned a single task. Third, when the knowledge removal mechanism and the generation interval are fixed, the scope size is significant for the final success rate and scope accuracy values. The scope size is also significant when the agents remove randomly selected knowledge. When this is not the case, the scope size is significant only when the experiment spans over several generations. Figures 6.5, 6.6, 6.7 and 6.8 portray how each respective measure varies with the number of tasks, the number of generations

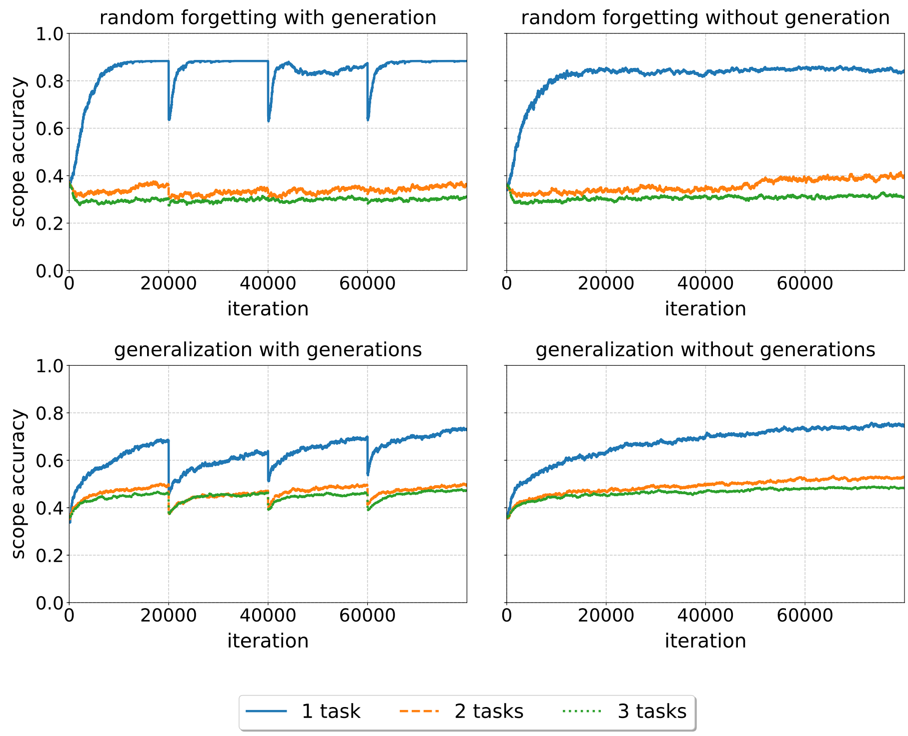

Figure 6.4 Scope accuracy for different number of generations, knowledge forgetting mechanism and scope sizes. Different subplots correspond to different combinations of number of generations and knowledge removal mechanisms. Different curves correspond to different scope sizes $|scope| = \{1, 2, 3\}$.

and the knowledge removal mechanism. Different colors correspond to different combinations of knowledge removal mechanisms and number of generations. Bars of the same color correspond to different numbers of tasks for the same number of generations and the same knowledge removal mechanism.

6.4 Conclusion

In Chapter 5, the presence of parasitic knowledge was identified. Under the examined conditions, it prevented agents from acquiring all the knowledge that was collectively available for a task. As a result agents (1) did not achieve consensus and (2) did not improve their knowledge when narrowing their scope. In this chapter, we examined how freeing memory through the removal of randomly selected knowledge, reflects on agents' knowledge and behavior. Our assumption was that short-term drawbacks, may be out-weighted by long-term benefits. To that end, the following hypothesis was formulated: Agents removing randomly

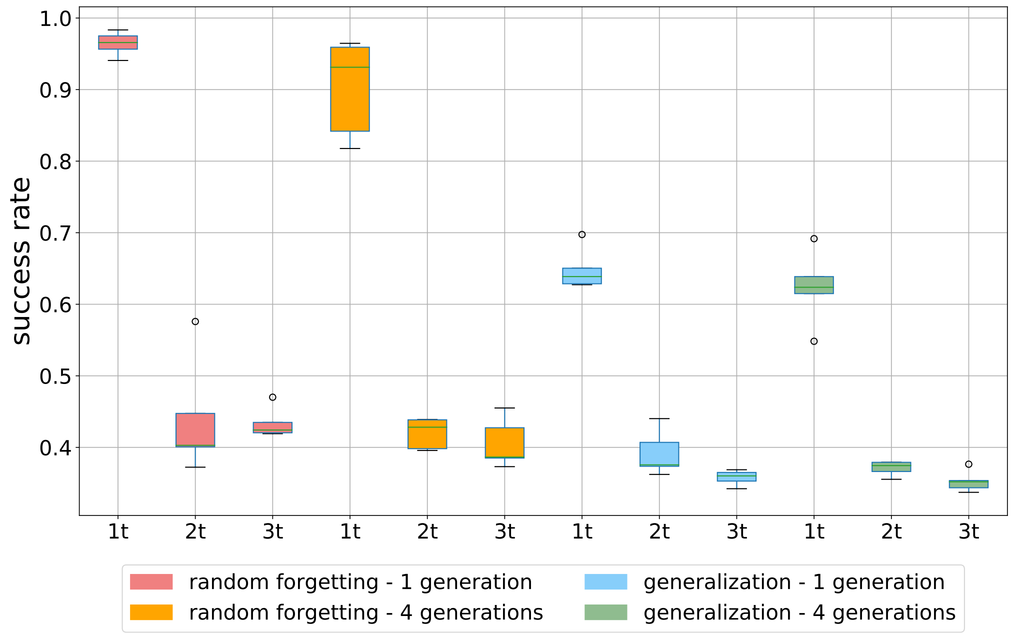

Figure 6.5 Success rate grouped by the number of tackled tasks, number of generations, and knowledge removal mechanism.

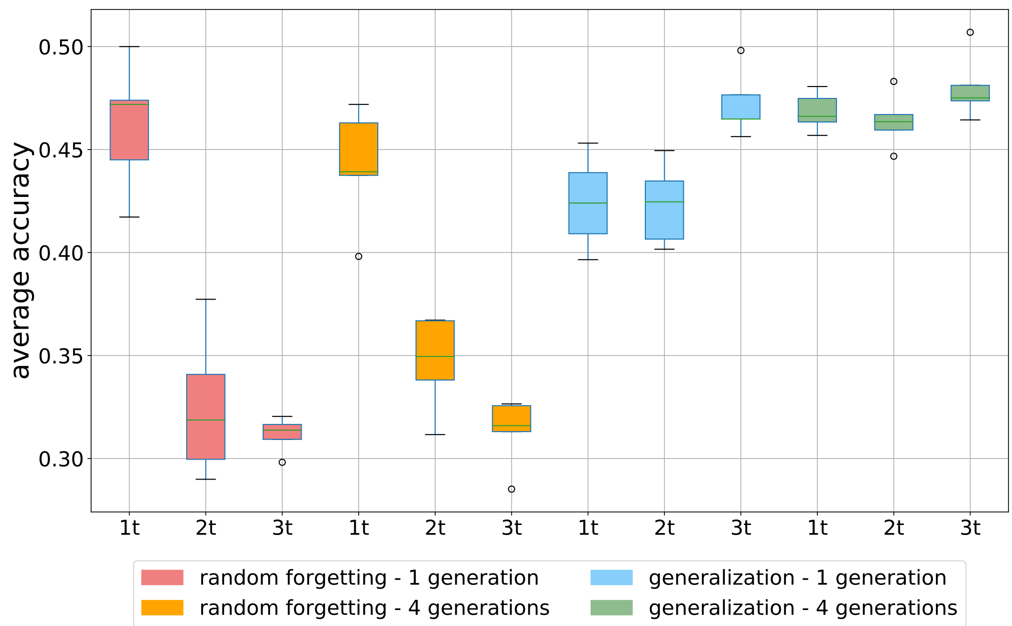

Figure 6.6 Average accuracy grouped by the number of tackled tasks, number of generations, and knowledge removal mechanism.

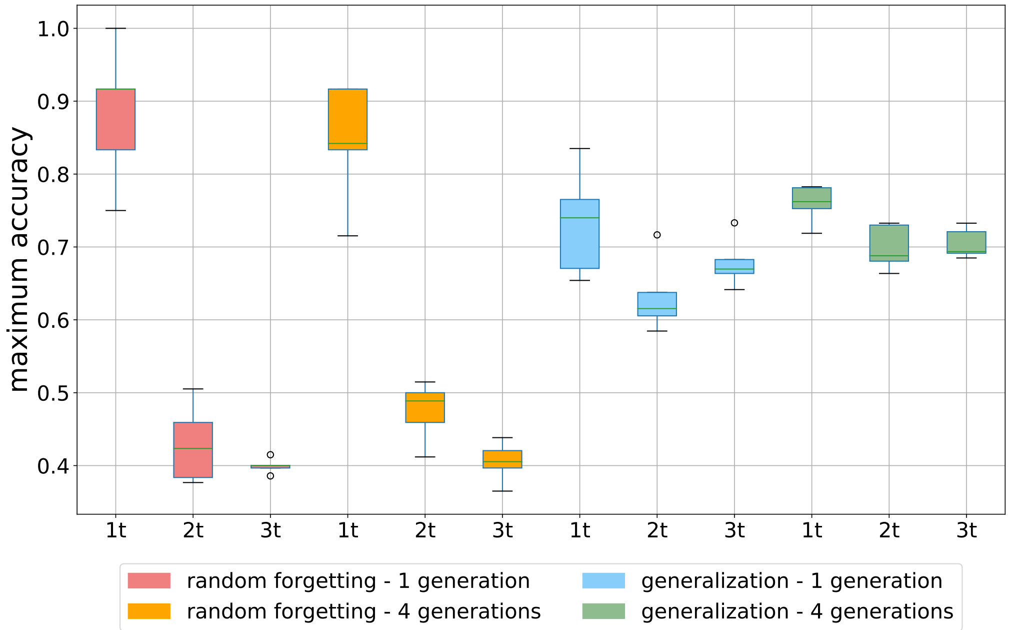

Figure 6.7 Maximum accuracy grouped by the number of tackled tasks, number of generations, and knowledge removal mechanism.

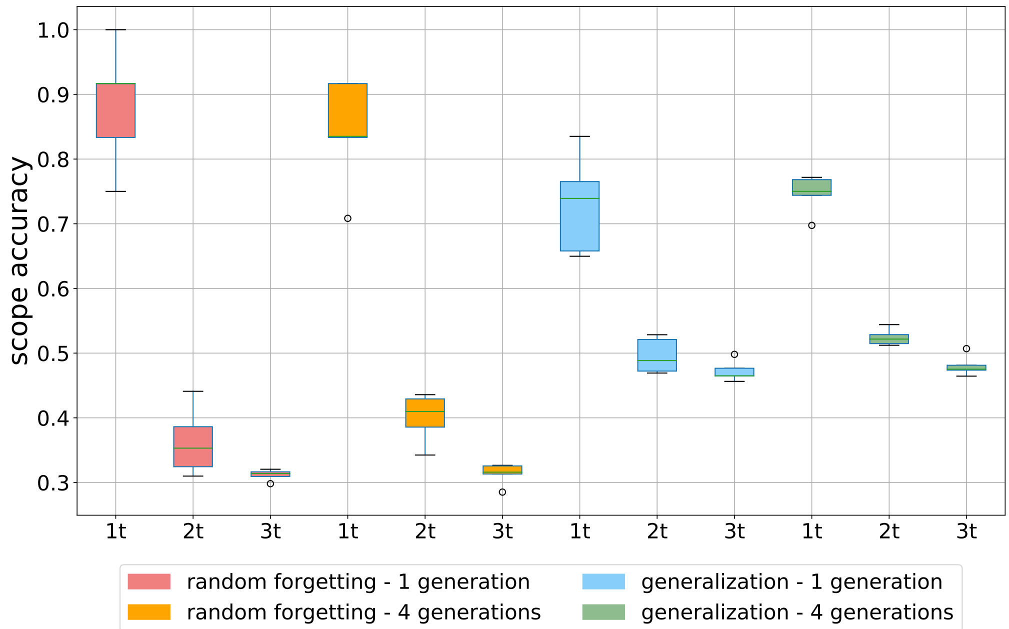

Figure 6.8 Scope accuracy grouped by the number of tackled tasks, number of generations, and knowledge removal mechanism.

Table 6.2 Summary of p-values from one-way ANOVA for full and reduced models of different parent selection methods. The full model includes all used parent selection methods (min/max success rate, min/max compensation, random selection). The reduced model only includes (1) min success rate and (2) random parent selection. The p-values concern multitasking agents and are organized by measure and memory size (maximum number of ontology classes). It examines the effect of different parent selection methods for a given agent memory size.

Measure	**Generalization**		**Random Forgetting**	
	1 Gen	**4 Gens**	**1 Gen**	**4 Gens**
Success rate	7.37e-10	1.45e-08	7.65e-10	8.34e-10
Average accuracy	0.0027	0.170	2.65e-06	4.42e-06
Maximum accuracy	0.041	0.005	5.37e-08	5.68e-08
Scope accuracy	3.27e-06	2.45e-10	8.60e-09	6.23e-09

Table 6.3 Summary of p-values from one-way ANOVA analysis for full and reduced models of different parent selection methods. The full model includes all used parent selection methods (min/max success rate, min/max compensation, random selection). The reduced model only includes (1) min success rate and (2) random parent selection. The p-values concern multitasking agents and are organized by measure and memory size (maximum number of ontology classes). It examines the effect of different parent selection methods for a given agent memory size.

Measure	**1 Task**		**2 Tasks**		**3 Tasks**	
	1 Gen	**4 Gens**	**1 Gen**	**4 Gens**	**1 Gen**	**4 Gens**
Success rate	2.58e-08	8.45e-05	0.240	0.001	9.54e-05	0.013
Average accuracy	0.065	0.085	0.0006	9.42e-06	4.83e-08	2.26e-07
Maximum accuracy	0.023	0.058	0.0002	9.80e-06	1.18e-07	5.14e-08
Scope accuracy	0.220	0.045	0.001	0.0001	4.83e-08	2.26e-07

selected knowledge, will reach consensus after a finite number of interactions. Results support this hypothesis. Agents reach consensus, regardless of the size of their scope. However, unless each agent tackles a single task, removing randomly selected ontology classes, significantly degrades its knowledge. Results show that when agents focus on a single task, temporarily removing relevant and accurate knowledge allows for more accurate knowledge representations in the long term. This was found to be true for all knowledge correctness measures and independent from the number of generations. We can thus conclude that forgetting randomly selected knowledge may prove beneficial, yet under a very narrow set of conditions. This suggests that selecting one memory approach from another, should always be examined in conjunction with task assignment.

— 7 —

Task exploration in resource-limited agent populations

We have previously shown that when agents possess limited resources, they specialize. In chapter 5, we examined the repercussions of this specialization. In particular, we examined how assigning different tasks to different agents, reflects on individual agents and agent populations. The tasks examined in chapter 5 are independent. Accurately deciding for different tasks thus rely on different properties. Each of these tasks was assigned to an equal number of agents. Results showed that agent populations equitably explore and become accurate in all existing tasks. Becoming collectively accurate in all tasks, did not allow these agent populations for improving their maximum accuracy in the tackled tasks. Yet, it allowed these agent populations to collectively obtain more income (points) and distribute it in a more equitable manner to the individual agents.

However, the examined setup is based on an artificially imposed equidistribution of tasks to the agents. Little is known on whether the observed collective benefits can be obtained without a priori guaranteeing that each task is constantly tackled by an equal number of agents. Our assumption is that when two agents repeatedly disagree, this is because their knowledge is shaped for serving different purposes, i.e., tasks. Hence, favoring the reproduction of agents that disagree the most with other agents, should restore and maintain balance in task exploration. Based on this, the following hypothesis was formulated: 'Favoring the reproduction of agents with the lowest success rate will allow agent populations to equally explore all tasks'. We tested this hypothesis by conducting a replacement experiment, i.e., a portion of the agent population is periodically replaced by new, naive agents.

The replacement takes place in two steps. During the first step, new agents are introduced. Each new agent is trained by two previously existing agents, henceforth called parent 1 and parent 2 respectively. To that end, these parents generate and provide training examples based on what they know, i.e., their current ontologies. Those examples follow the same distribution their accuracy follows, over all existing tasks. For example, if agent *Alice* is 100% accurate in hunting but 0% accurate in the remaining tasks, *Alice* will solely provide examples for hunting. However, this is an extreme case. In most of the cases, parents will provide examples for several tasks, tasks for which they are not entirely accurate. As a result, inaccurate knowledge will be passed from one generation of agents to the next one. The latter can be seen both as an introduction of noise or variance. It may allow for agent putting their knowledge under additional scrutiny. Yet, given these agents' limited resources, this may also prove detrimental to their knowledge's overall correctness. Different parent selection criteria have been examined. According to these, parents may be selected randomly, based on their success rate and the points they have acquired so far. Acquired points may serve as an indication of knowledge correctness. During the second step, a predefined number of agents are removed. Their selection is based on the number of iterations for which they are present. The number of removed agents equals that of agents added. This allows for maintaining a population of constant size, avoiding any effect related to increasing the collective size of a population's memory resources.

7.1 Experimental setting

7.1.1 Hypothesis

Favoring the reproduction of agents with the lowest success rate will allow agent populations to equally explore all tasks.

7.1.2 Parameters

The experiment is executed under 20 setups, acquired by varying: (1) the number of tasks assigned to each agent, (2) the size of their memory and (3) the reproduction strategy applied at each generation occurrence. Each setup is run 20 times and its results are averaged. One run consists of 80000 interactions with each interaction taking place among two agents. These agents are randomly selected out of a total population of 36 agents. Each run spans over 4 generations, each one taking place every 20000 interactions. When a generation occurs, the 9 longest existing agents are replaced by 9 new agents. The 9 new agents are trained by previously existing agents. Choosing which agents will train the new agents relies on 5 different agent reproduction strategies. These are based on strategies that favor: (a) agents with the lowest success rate, (b) agents with the highest success rate, (c) agents with the lowest number of collected points, (d) agents with the highest number of collected points and (e) the random selection of parents (training agents). Depending on the setup, each agent is assigned either 1 or 3 tasks and has an ontology (memory) capable of accommodating a maximum of either 4 or 12 classes. When its memory is exhausted, it frees it by generalizing its knowledge

(Subsection 3.3.4). The agents' environment contains 64 different object types, each one perceivable through 6 different binary properties. The agents are initially trained with respect to all $|\mathcal{T}| = 3$ tasks. For each task, 4 different decisions exist. Choosing one decision relies on 2 out of the 6 binary properties. Since the 3 tasks do not overlap, deciding with respect to different tasks relies on a different pair of properties. To induce its initial ontology, each agent takes into account a randomly selected 10% of all existing labeled examples. As a consequence, there may be overlap between the training of two agents. The agents subsequently undertake tasks within their scope. To decide which agent adapts its knowledge in case of disagreement, agents compare their scores. These are attributed to them by the environment between two consecutive interactions. They are calculated based on a randomly selected corpus consisting of 60% of all existing training samples.

Table 7.1 Summary table of parameters (independent variables)

Parameter	**Value**
Runs	20
Interactions	80000
Generations	4
Agents $\lvert\mathcal{A}\rvert$	36
Replaced agents	9
Reproduction strategy	$\{min\_points$, max_points, min_srate, max_srate, $random\}$
Memory capacity	{4,12} classes
Memory release technique	Generalization
Tasks $\lvert\mathcal{T}\rvert$	3
Assigned tasks (scope)	{1,3}
Object types $\lvert\mathcal{I}\rvert$	64
Properties	6
Properties per task $\lvert\mathcal{D}_t\rvert$	2
Task overlap	0%
Decisions per task	4
Training ratio	0.1
Score ratio	0.6

7.1.3 Measures

Relative standard deviation

Additionally to the measures used in the previous chapters, here we analyze results using the relative standard deviation (RSD). RSD is a measure used to quantify the variability of a set of data relative to its mean. Here, RSD is employed to assess the consistency of accuracy across different tasks performed by a population of agents. The RSD of an agent population's accuracy is calculated as follows. First, the average accuracy for each task across all agents is determined. Then, the standard deviation of these average accuracies is calculated using the following formula:

$$s = \sqrt{\frac{\sum_{i=1}^{n}(x_i - \bar{x})^2}{n-1}}$$

Where:

- x_i is the average accuracy for task i,
- $\bar{x}$ is the mean of the average accuracies, and
- n is the total number of tasks.

The RSD is then obtained by dividing this standard deviation by the mean of the average accuracies and multiplying the result by 100. A higher RSD value indicates greater variability in the accuracy of the agent population across different tasks, while a lower value suggests more consistent performance. Put differently, the higher RSD value is, the more inequitably are tasks explored. The formula for calculating RSD is as follows:

$$RSD_{\mathrm{TAC}}(s, \bar{x}) = \frac{s}{\bar{x}} \times 100$$

We additionally use RSD with respect to how many times each task is performed, when agents do not undertake all tasks ($|\mathcal{S}_{cope}| < |\mathcal{T}|$). In this case:

- x'_i is the number of times, task i was collectively performed by an agent population,

- $\bar{x}'$ is the mean number of times a task is on average performed.

Based on these adaptations, the standard deviation of games s' is defined as follows :

$$s' = \sqrt{\frac{\sum_{i=1}^{n}(x'_i - \bar{x}')^2}{n-1}}$$

Based on s the relative standard deviation of games can be defined as follows:

$$RSD_{games}(s', \bar{x}') = \frac{s'}{\bar{x}'} \times 100$$

7.2 Results and discussion

Here we present and discuss results acquired by varying (1) which agents are selected for reproduction and (2) the size of their scope. Each figure consists of 2 subfigures, each one corresponding to a different memory size (4 and 12 classes). Within such a subfigure, different selection criterion sizes are represented by curves of different color and pattern.

7.2.1 Multitasking agents

Figure 7.1 displays the evolution of the average accuracy over all tasks (y-axis) as the number of iterations increases (x-axis), depending on (a) knowledge limit and (b) parent selection policy. It shows that favoring the reproduction of agents that are less successful or have accumulated less points, leads to a significantly lower average accuracy for a given knowledge limit. Whether agent memory is of 4 or 12 classes, there is a statistically significant difference between the random selection of parents and the selection of the parents that have the lowest success rate.

Figure 7.2 displays the evolution of the average maximum accuracy over all agents (y-axis) as the number of iterations increases (x-axis), depending on (a) knowledge limit and (b) parent selection policy. It is shown that favoring the reproduction of agents that are less successful or have accumulated less points, leads to a significantly lower average maximum accuracy for a given knowledge limit. When agent memory holds 4 classes, there is no statistically significant difference between the random selection of parents and the selection of the parents that have the lowest success rate. When the agent memories holds 12 classes, there is a statistically significant difference between the random selection of parents and

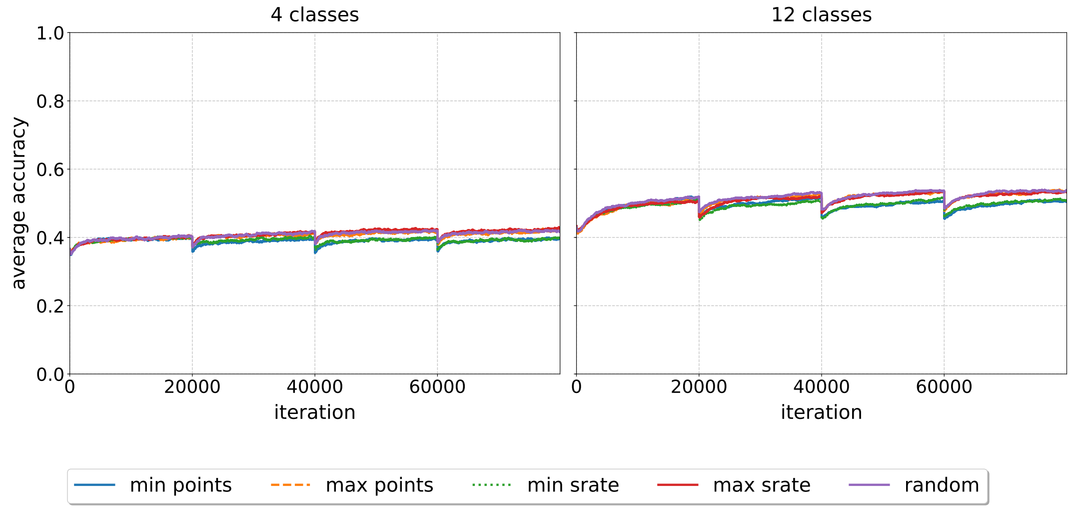

Figure 7.1 Average accuracy of multi-tasking agents for different memory and parent selection criteria. Different subplots correspond to different memory sizes (4 and 12 classes). Different curves correspond to different parent selection criteria ($minimum\_success\_rate$, $maximum\_success\_rate$, $minimum\_points$, $maximum\_points$, $random\_selection$).

the selection of parents that have the lowest success rate, i.e., that agree less with their peers.

Figure 7.3 displays the evolution of the relative standard deviation for the average accuracy of each task over all agents (y-axis) as the number of iterations increases (x-axis), depending on (a) knowledge limit and (b) parent selection policy. The following are shown. First, when the reproduction of successful agents is favored, the relative standard deviation shows a slow increase. Contrariwise, when the reproduction of the less successful agents is favored, the standard deviation reaches a local maximum. It then decreases until it stabilizes. When agent memory holds 4 classes, there is no statistically significant difference between the random selection of parents and the selection of the parents that have the lowest success rate. When the agent memory holds 12 classes, there is a statistically significant difference between the random selection of parents and the selection of parents that have the lowest success rate, i.e., agree less with their peers.

7.2.2 Single-tasking agents

Figure 7.4 displays the evolution of the average success rate (y-axis) as the number of iterations increases (x-axis), depending on (a) knowledge limit and (b) parent selection policy. Two things are shown. First, that favoring the reproduction of agents that are less successful or have accumulated less points, leads to a signifi-

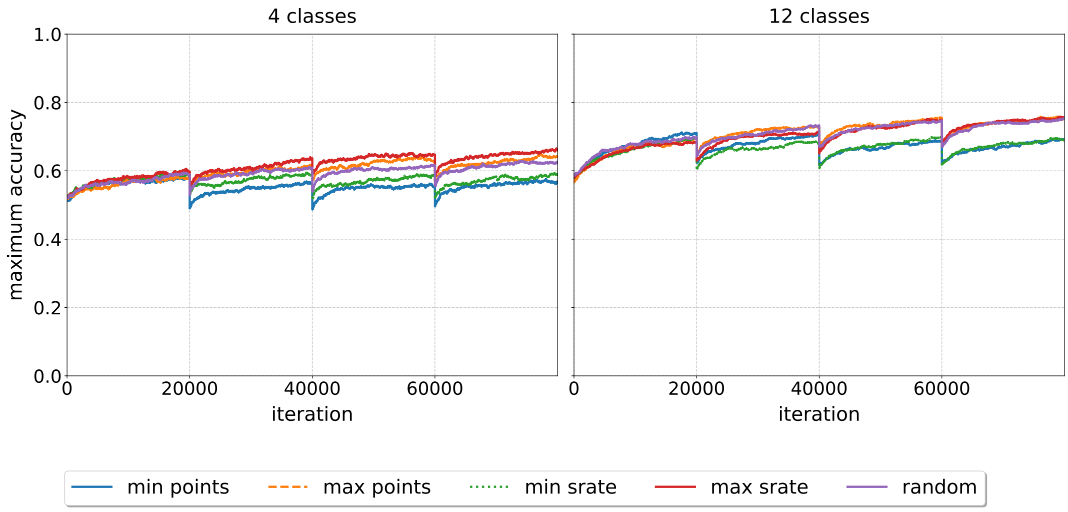

Figure 7.2 Average maximum accuracy of multi-tasking agents for different memory and parent selection criteria. Different subplots correspond to different memory sizes (4 and 12 classes). Different curves correspond to different parent selection criteria ($minimum\_success\_rate$, $maximum\_success\_rate$, $minimum\_points$, $maximum\_points$, $random\_selection$).

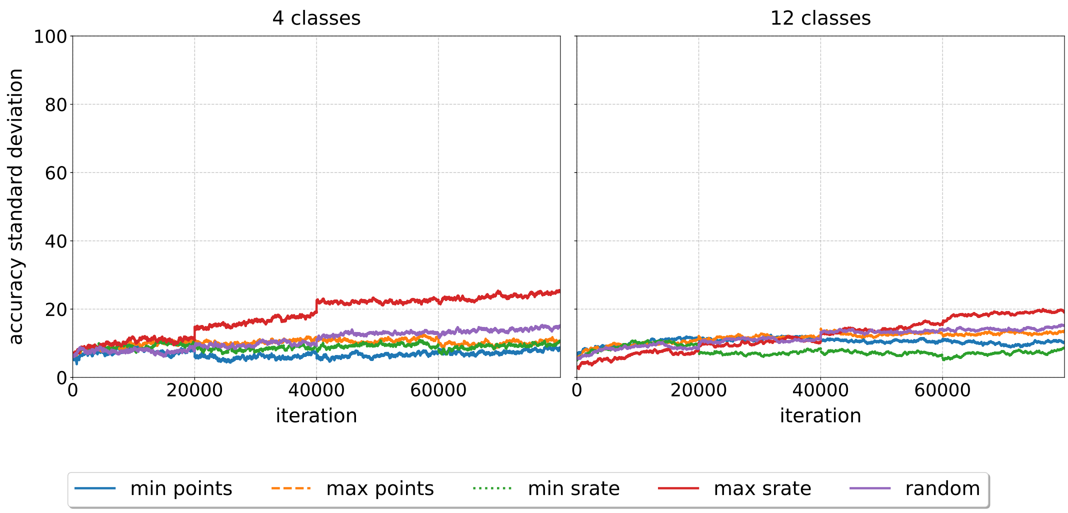

Figure 7.3 Accuracy relative standard deviation of multi-tasking agents for different memory and parent selection criteria. Different subplots correspond to different memory sizes (4 and 12 classes). Different curves correspond to different parent selection criteria ($minimum\_success\_rate$, $maximum\_success\_rate$, $minimum\_points$, $maximum\_points$, $random\_selection$).

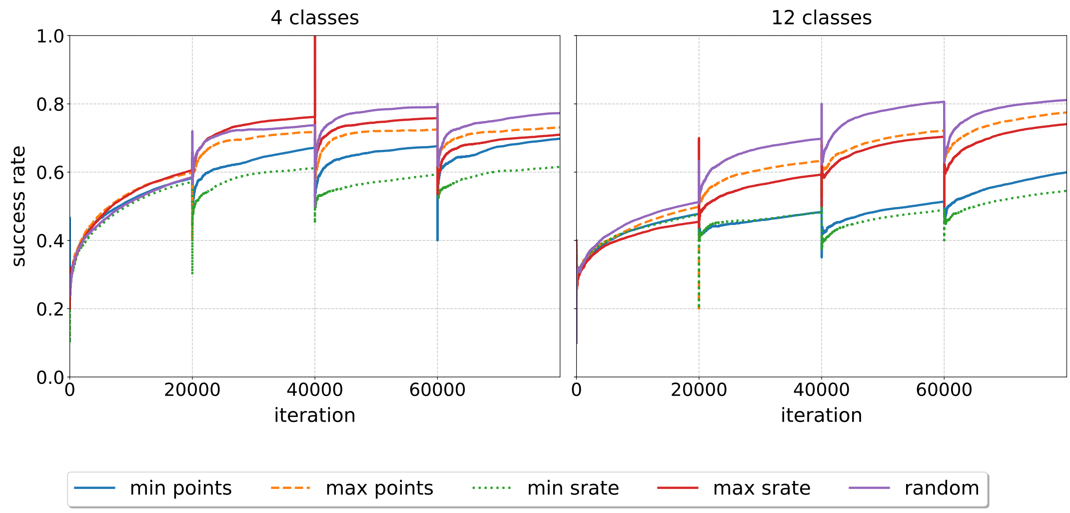

Figure 7.4 Success rate of single-tasking agents for different memory and parent selection criteria. Different subplots correspond to different memory sizes (4 and 12 classes). Different curves correspond to different parent selection criteria ($minimum\_success\_rate$, $maximum\_success\_rate$, $minimum\_points$, $maximum\_points$, $random\_selection$).

cantly lower success rate. This is to be expected according to the pressure exerted during the selection. Second, both subfigures show that the success rate stabilizes, but does not converge to 1. This indicates that the agents' final ontologies do not allow them to agree on all decisions.

Figure 7.5 displays the evolution of the average accuracy over all tasks (y-axis) as the number of iterations increases (x-axis), depending on (a) knowledge limit and (b) parent selection policy. It is shown that when agent memory consists of 12 classes, favoring the reproduction of agents that are less successful or have accumulated less points, leads to a significantly lower average accuracy for a given knowledge limit.

Figure 7.6 displays the evolution of the average maximum accuracy over all agents (y-axis) as the number of iterations increases (x-axis), depending on (a) knowledge limit and (b) parent selection policy. It is shown that when agent memory consists of 12 classes, favoring the reproduction of agents that are less successful or have accumulated less points, leads to a significantly lower average maximum accuracy for a given knowledge limit.

Figure 7.7 displays the evolution of the relative standard deviation for the average accuracy of each task over all agents (y-axis) as the number of iterations increases (x-axis), depending on (a) knowledge limit and (b) parent selection policy. The following are shown. First, when the reproduction of successful agents is favored, the relative standard deviation shows a slow increase. Contrariwise, when

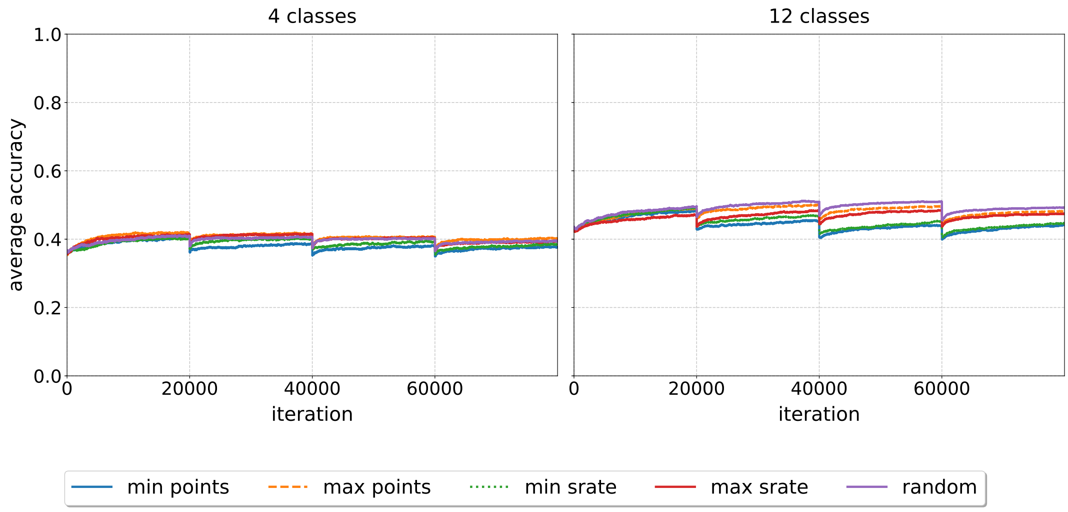

Figure 7.5 Average accuracy of single-tasking agents for different memory and parent selection criteria. Different subplots correspond to different memory sizes (4 and 12 classes). Different curves correspond to different parent selection criteria ($minimum\_success\_rate$, $maximum\_success\_rate$, $minimum\_points$, $maximum\_points$, $random\_selection$).

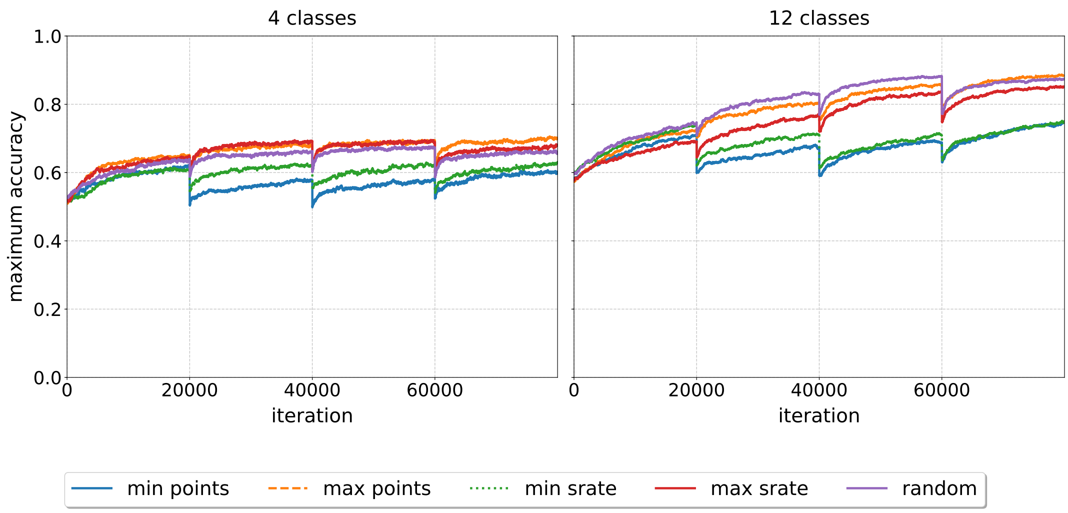

Figure 7.6 Average maximum accuracy of single-tasking agents for different memory and parent selection criteria. Different subplots correspond to different memory sizes (4 and 12 classes). Different curves correspond to different parent selection criteria ($minimum\_success\_rate$, $maximum\_success\_rate$, $minimum\_points$, $maximum\_points$, $random\_selection$).

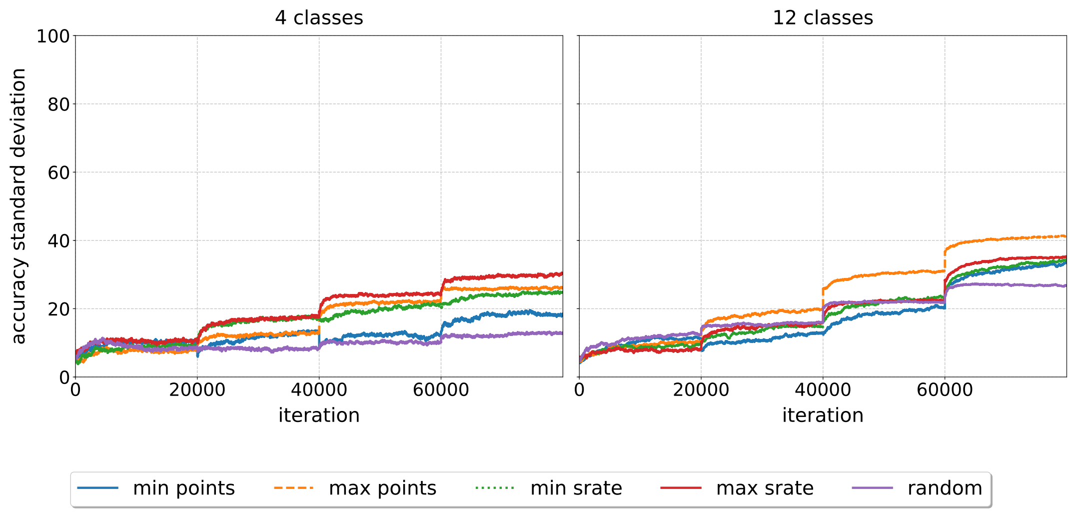

Figure 7.7 Accuracy relative standard deviation of single-tasking agents for different memory and parent selection criteria. Different subplots correspond to different memory sizes (4 and 12 classes). Different curves correspond to different parent selection criteria (*minimum_success_rate*, *maximum_success_rate*, *minimum_points*, *maximum_points*, *random_selection*).

the reproduction of the less successful agents is favored, the standard deviation reaches a local maximum. It then decreases until it stabilizes.

Figure 7.8 displays the evolution of the games relative standard deviation (y-axis) as the number of iterations increases (x-axis), depending on (a) knowledge limit and (b) parent selection policy. The following are shown. First, when the reproduction of less successful agents is favored, after each generation there is a local low. The latter is not observed when the reproduction of the most successful agents is favored. In this case, the games relative standard deviation increases more steeply after each generation. Second, for all (a), (b) and (d), after a finite number of generations, no low locals are observed. This agrees with the evolution of the accuracy relative standard deviation. The latter is monotonically increasing generation after generation for all combinations of the independent variables.

7.3 Statistical Analysis

7.3.1 Multitasking agents

Figures 7.9, 7.10 and 7.11 show the distribution of final observed values for different measures used to evaluate task exploration within agents tackling all (3) tasks. Each figure categorizes results by agent memory size (4 and 12 classes, indicated by different colors) and parent selection criteria (*minimum_success_rate*,

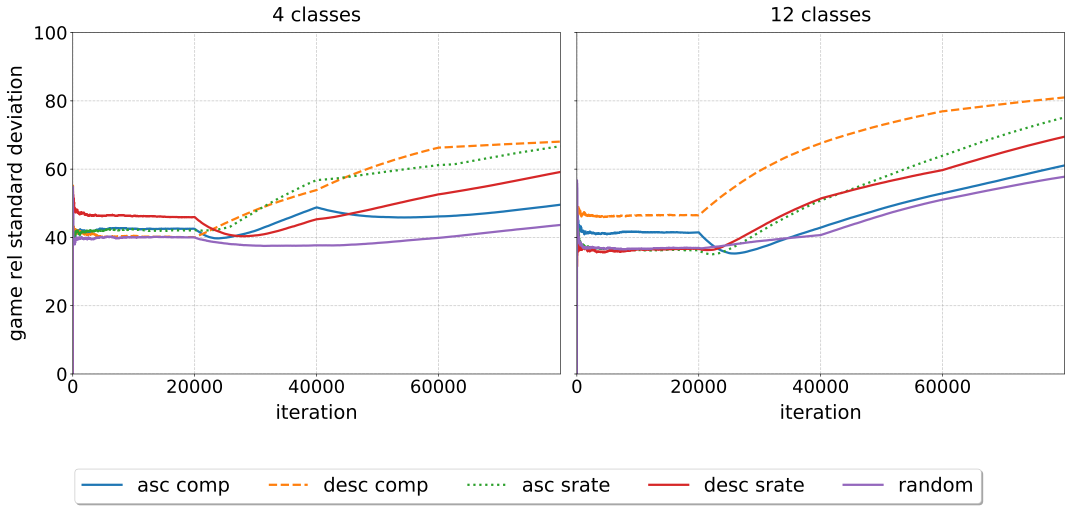

Figure 7.8 Games relative standard deviation of single-tasking agents for different memory and parent selection criteria. Different subplots correspond to different memory sizes (4 and 12 classes). Different curves correspond to different parent selection criteria ($minimum\_success\_rate$, $maximum\_success\_rate$, $minimum\_points$, $maximum\_points$, $random\_selection$).

$maximum\_success\_rate$, $minimum\_points$, $maximum\_points$, $random\_selection$, with bars of the same color representing each criterion).

Table 7.2 Summary of p-values from one-way ANOVA for full and reduced models of different parent selection methods. The full model includes all parent selection methods (min/max success rate, min/max points, random selection). The reduced model includes only (1) min success rate and (2) random parent selection. The p-values concern multitasking agents and are organized by measure and memory size (maximum number of ontology classes). The table examines the effect of different parent selection methods for a given agent memory size.

Measure	**Reduced Model**		**Full Model**	
	4 Classes	**12 Classes**	**4 Classes**	**12 Classes**
Average accuracy	0.010	7.93e-05	0.0001	6.63e-06
Maximum accuracy	0.087	0.0025	0.0001	6.38e-05
RSD_{tacc}	0.150	0.029	2.89e-05	0.021

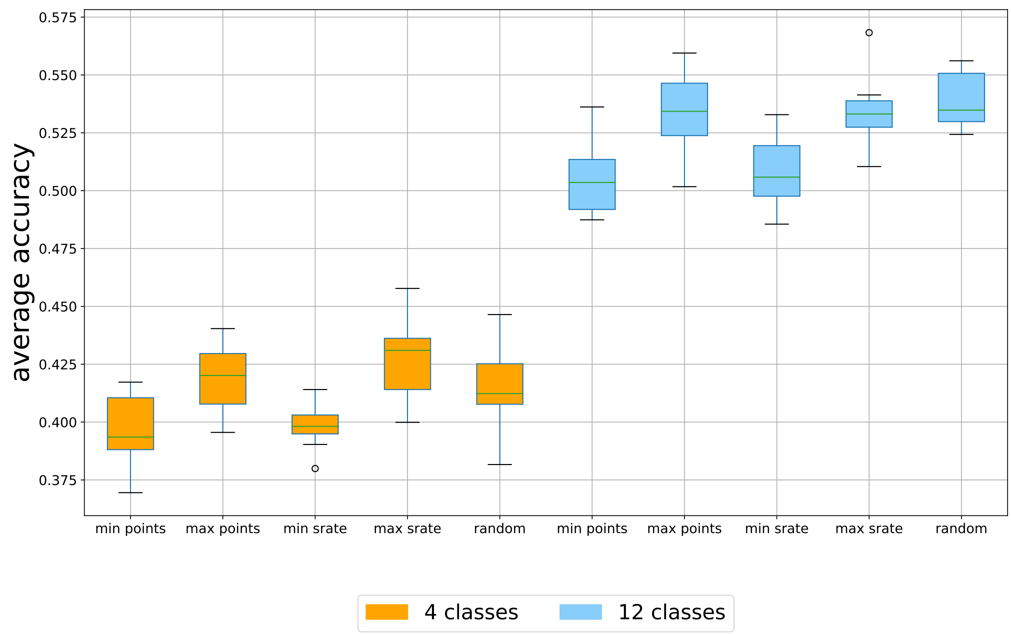

Figure 7.9 Final average accuracy of multi-tasking agents.

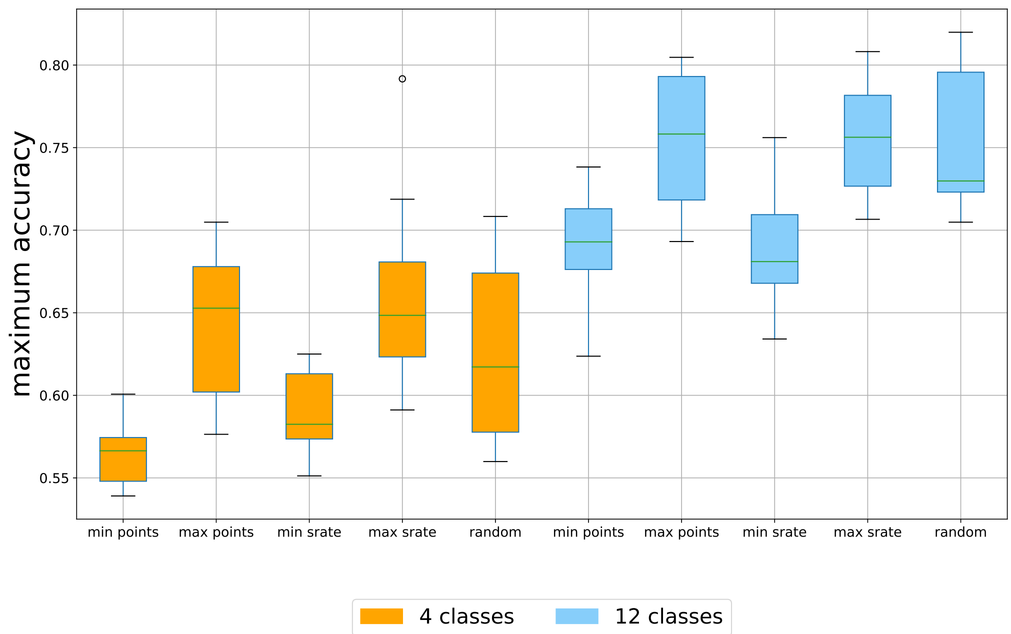

Figure 7.10 Final average maximum accuracy of multi-tasking agents.

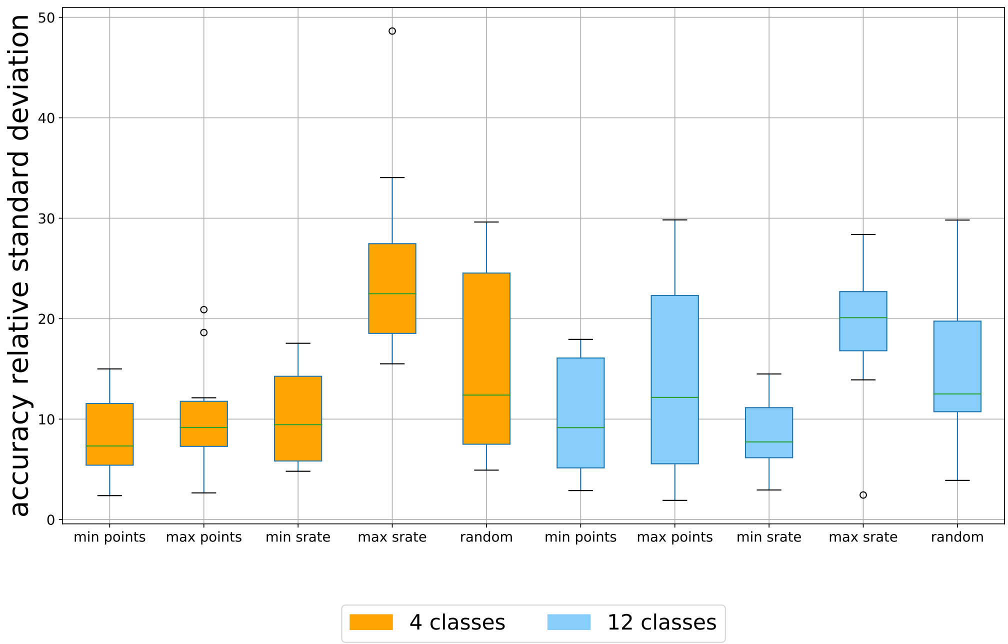

Figure 7.11 Final task accuracy relative standard deviation of single-tasking agents.

7.3.2 Single-tasking agents

Figures 7.12, 7.13, 7.14, 7.15 and 7.16 show the distribution of final observed values for different measures used to evaluate task exploration within single-tasking agents. Each figure groups results by agent memory and parent selection criteria. Different colors correspond to different memory sizes (4 and 12 classes) and bars of the same color correspond to different parent selection criteria ($minimum\_success\_rate$, $maximum\_success\_rate$, $minimum\_points$, $maximum\_points$, $random\_selection$).

7.4 Conclusion

In Chapter 5 we have established that for specialization to be beneficial for agent societies, an equitable task allocation must be artificially imposed. Without it, agent societies become collectively efficient on a small subset of tasks and inefficient on the remaining ones. In this chapter, task exploration in resource-limited populations was considered. In particular, we examined whether a population can equally improve its accuracy in all tasks, without a priori assigning each task to an equal number of agents. Our hypothesis was that favoring the reproduction of the agents with lowest success rate, agent populations will equitably explore

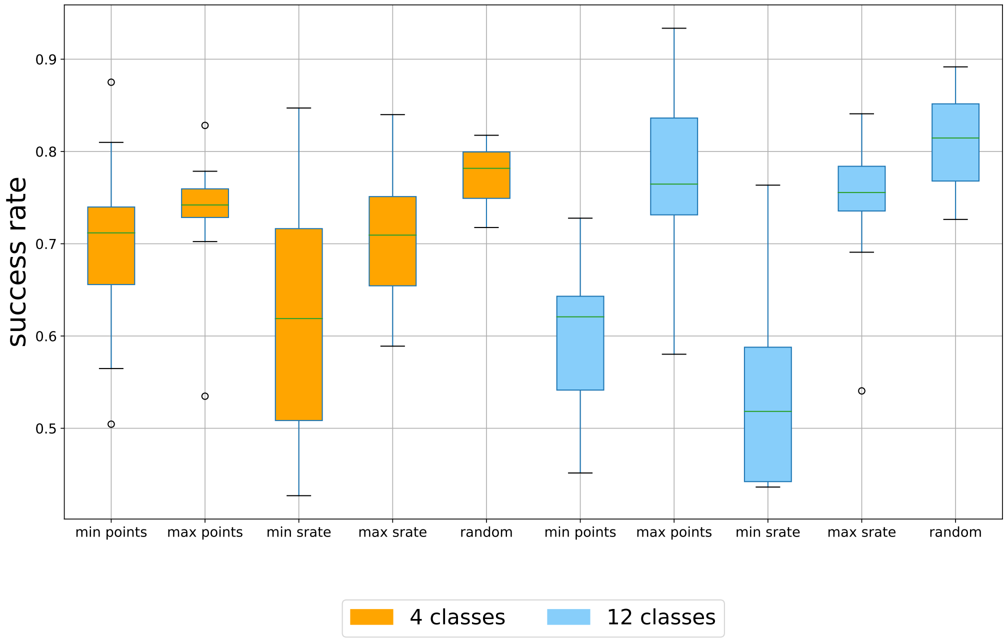

Figure 7.12 Final success rate of single-tasking agents.

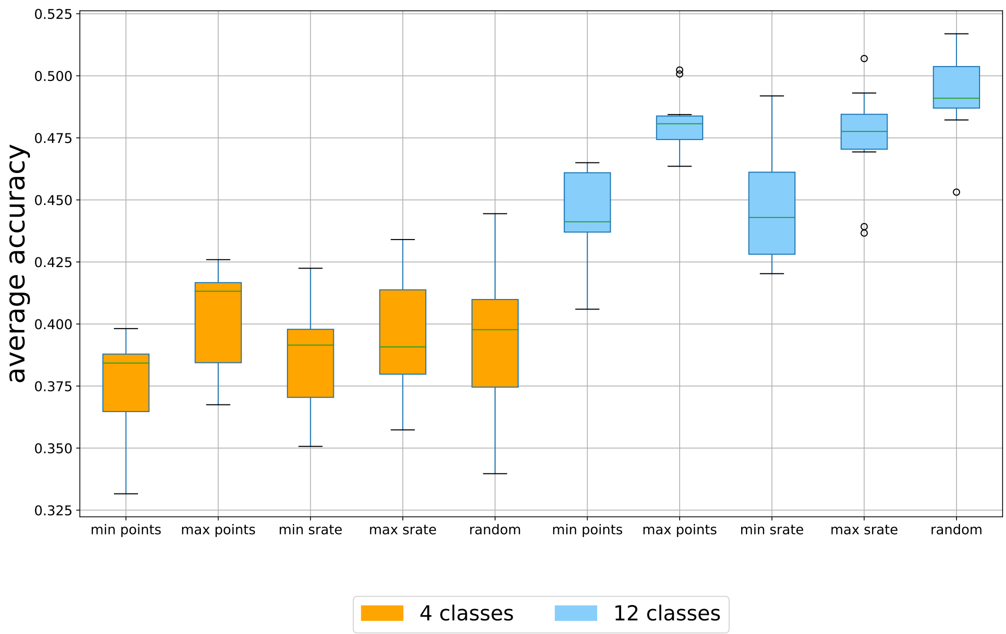

Figure 7.13 Final average accuracy of single-tasking agents.

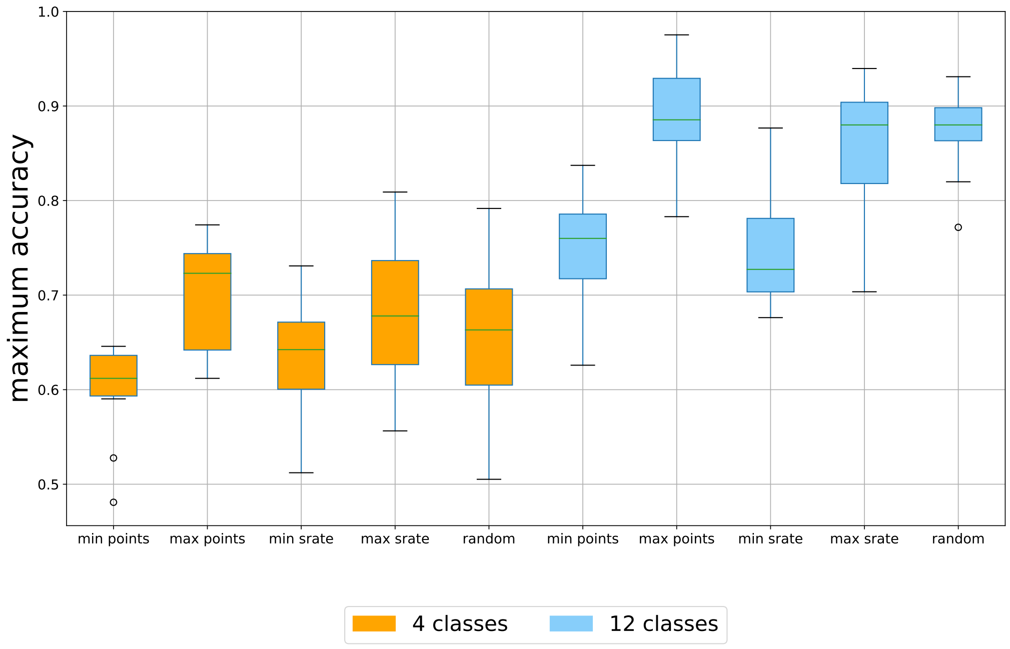

Figure 7.14 Final average maximum accuracy of single-tasking agents.

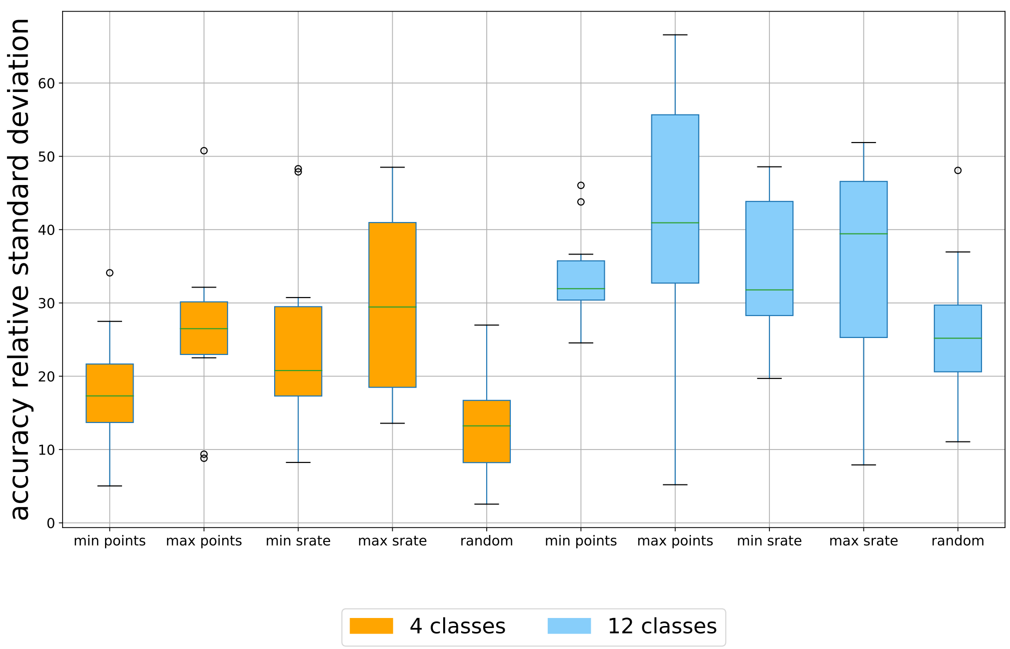

Figure 7.15 Final accuracy relative standard deviation of single-tasking agents.

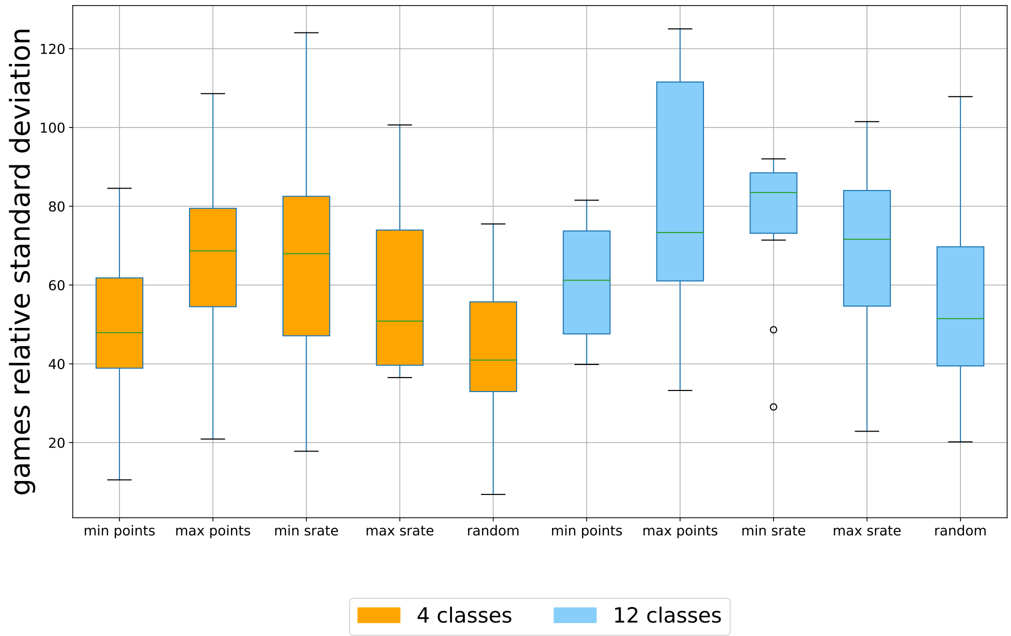

Figure 7.16 Final games relative standard deviation of single-tasking agents.

Table 7.3 Summary of p-values from one-way ANOVA for full and reduced models of different parent selection methods. The full model includes all used parent selection methods (min/max success rate, min/max points, random selection). The reduced model only includes (1) min success rate and (2) random parent selection. The p-values concern single-tasking agents and are organized by measure and memory size (maximum number of ontology classes). The table examines the effect of different parent selection methods for a given agent memory size.

Measure	**Reduced Model**		**Full Model**	
	4 Classes	**12 Classes**	**4 Classes**	**12 Classes**
Success rate	0.002	6.15e-06	0.009	2.83e-08
Average accuracy	0.530	0.0001	0.180	1.22e-06
Maximum accuracy	0.420	0.0002	0.010	7.25e-06
RSD_{tacc}	0.020	0.110	0.010	0.170
RSD_{games}	0.065	0.120	0.130	0.200

all tasks. This was based on the assumption that an agent with lower success rate, possesses rare knowledge. As rare, we define knowledge that is suitable for a task for which, the population's majority is inefficient. To test this hypothesis, we varied 2 parameters: (1) the agents' memory size and (2) the parents selection criteria. Results reject this hypothesis. They show that agent populations either collectively specialize on the same task, or degrade their accuracy on the tasks they previously excelled. This suggests that under the examined minimal assumptions, agent populations benefit more from a parents selection that is either random, or based on the highest success rate or points.

— 8 —

Conclusion

8.1 Summary

It has been previously shown that agents coordinating over a single task can evolve knowledge representations [11,12,92]. Our thesis builds upon this, by investigating how the presence of additional tasks affects this evolution. More particularly, it investigates how agent societies culturally evolve and specialize knowledge, when different tasks compete over limited agent resources. Can knowledge regarding different tasks be equitably distributed within these societies? Put differently, does the cultural evolution of knowledge guarantee that agent societies equitably improve their knowledge on all tasks? Equitable improvement of knowledge across all tasks can be achieved in various ways. On the one end of the spectrum, agents can build and evolve general-purpose knowledge, i.e., knowledge suitable for undertaking several tasks. On the other end of the spectrum, agents can build and evolve specialized knowledge, i.e., knowledge suitable for undertaking one task. To answer this question, an extended version of [61] has been proposed, allowing for controlling among others: (1) the overlap among different tasks, (2) the number of tasks assigned to agents, (3) the size of their individual resources and (4) the reproduction strategies applied from one generation of agents to the next one. Based on these controlled variables, four experiments were carried out. These experiments successively investigate: (a) whether knowledge is transferred from one task to another, (b) how task assignment and resource competition lead to agent specialization, (c) how random forgetting can benefit single-tasking agents in the long term, and (d) how different reproduction strategies affect the capacity of agent societies to efficiently tackle all tasks. Through the descriptive and statistical analysis of their results, the following can be concluded. First, agents improve their knowledge on one task by undertaking additional tasks. Second, specialization is not a result of assigning different tasks to different agents, but rather the result of limited resources. Third, in agent societies where each agent is assigned one task, agents benefit from removing randomly selected knowledge when their resources are exhausted. Finally, when agents specializing on different tasks come in contact, one task will gradually monopolize this society's memory resources. Put differently, all agents of this society will gradually specialize on the same task.

8.2 Contributions

Several contributions in the area of multi-agent systems and cultural knowledge evolution can be identified. In multi-agent systems, we contribute by proposing a modular framework allowing among others for: (1) defining different tasks and control their overlap, (2) assigning agents with one or several tasks, (3) restricting their resources, (4) defining different ways for handling their exhaustion and (5) adopting different reproduction strategies. Based on this framework, several minimal hypotheses were formed. Testing these minimal hypotheses provided a comprehensive overview of the dynamics that characterize the formation and evolution of specialized knowledge. In cultural knowledge evolution we contribute by demonstrating that:

1. Resource restricted agents do not necessarily reach consensus.
2. Agents tackling additional semantically overlapping tasks improve the correctness of their knowledge.
3. The best task accuracy of resource-restricted agents does not depend on the size of their scope.
4. Populations assigning fewer tasks to individual agents, acquire more compensation and distribute it more equitably to agents.
5. When agents carry out a single task, long-term benefits may arise from the short-term removal of pertinent knowledge.
6. Favoring the reproduction of agents with marginal knowledge will either not prevent the collective specialization of a population, or prevent it by deteriorating the knowledge correctness on the overexplored tasks.

8.3 Perspectives

The presented experiments are based on homogeneous agent populations tackling homogeneous tasks. All agents face the same memory limitations and all tasks depend on the same number of properties. Additionally, the perception of the examined agents is always correct. Based on these, three major perspectives have been identified: (1) the effects of agents properties misperception on knowledge specialization, (2) the effects of agents with different memory limitations on knowledge specialization, and (3) the effects of tasks relying on unequal number of properties on task exploration.

So far, our results have demonstrated that knowledge specialization is driven by the restriction of agent resources, rather than by the disproportionate execution of specific tasks. Delegating or executing some tasks more often than others was shown to have no effect on the maximum accuracy achieved by agents. However, all our experiments were based on the assumption that all object properties can be accurately perceived by all agents. Do agents specialize when they cannot accurately perceive all properties, and if so, do they specialize in tasks where the required object properties can be accurately perceived? Given our extended framework, it is possible to experiment with agents that make errors when perceiving object properties. These errors can be either constant, meaning an agent consistently misperceives the same properties, or random, meaning an agent misperceives different properties at different times. Constant errors may refer to sensor limitations, such as an agent being unable to perceive all shades of green, while random errors may for example be the result of inadequate attention or eye fatigue. Two outcomes can be foreseen. First, that cultural knowledge evolution will drive agents to specialize on the tasks for which they make less or no errors at all. Second, that cultural knowledge evolution will push agents to converge on inaccurate knowledge. Establishing which of the two outcomes is true, will provide an insight on how cultural knowledge evolution deals with perception errors.

While our results demonstrated that assigning agents with less tasks did not enhance their maximum accuracy, they demonstrated that this improves their societies' prosperity and equitability. However, all conducted experiments are based on the assumption that all agents possess equal memory resources. Are specialization dynamics affected when different agents possess unequal memory resources? In particular, does this lead to less equitable agent societies, or does the presence of agents with greater memory resources improve the prosperity of those with fewer resources? Agents with greater memory resources are able to accumulate more knowledge. Thus, it is likely that they benefit more from adapting their knowledge, when interacting with more successful agents. Unless the knowledge of agents with greater memory resources becomes less specialized than the knowledge of agents with fewer memory resources, our assumption is that their presence will deteriorate their society's equitability. Establishing whether agents with greater resources specialize or not, will allow us for concluding on the necessity for mechanisms to mitigate income inequalities.

Finally, as the results of Chapter 7 show, resource-restricted agent societies gradually converge on a single task. In other words, the vast majority of their agents will eventually specialize in the same task. While the effects of agent selection during reproduction on task exploration has been examined, the acquired results were based on the following assumption: deciding with respect to different tasks, relies on the same number of properties. This means that decision-wise, all tasks are of equal complexity. As a result, it is unknown whether these results are valid when agents carry out tasks of different complexity, i.e., carry out tasks that depend on different number of properties. Does the presence of less complex tasks, push agent societies to specialize on a higher number of tasks? We assume that such societies would be more resilient to abrupt condition changes. Among others, these changes may be: (1) the appearance of new tasks, (2) a change in societies' needs and (3) an abrupt change in the societies' populations. Testing whether cultural knowledge evolution pushes agent societies towards a higher number of less complex tasks, will allow us to suggest a more efficient task partitioning and organization.

Bibliography

[1] Alberto Acerbi and Domenico Parisi. Cultural transmission between and within generations. *Journal of Artificial Societies and Social Simulation*, 9(1), 2006.

[2] Eduardo Alonso, Mark D'Inverno, Daniel Kudenko, Michael Lück, and Jason Noble. Learning in multi-agent systems. *The Knowledge Engineering Review*, 16(3):277–284, September 2001.

[3] Franz Baader, Diego Calvanese, Deborah McGuinness, Daniele Nardi, and Peter F. Patel-Schneider, editors. *The Description Logic Handbook: Theory, Implementation, and Applications.* Cambridge University Press, 2003.

[4] William M. Baum, Peter J. Richerson, Charles Efferson, and Brian Paciotti. Cultural evolution in laboratory microsocieties including traditions of rule giving and rule following. *Evolution and Human Behavior*, 25(5):305–326, September 2004.

[5] Keely Bebbington, Colin MacLeod, T. Mark Ellison, and Nicolas Fay. The sky is falling: evidence of a negativity bias in the social transmission of information. *Evolution and Human Behavior*, 38(1):92–101, January 2017.

[6] Yotam Ben-Oren, Oren Kolodny, and Nicole Creanza. Cultural specialization as a double-edged sword: division into specialized guilds might promote cultural complexity at the cost of higher susceptibility to cultural loss. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 378(1872):20210418, 2023.

[7] Théodore C. Bergstrom, Lawrence E. Blume, and Hal R. Varian. On the private provision of public goods. *Journal of Public Economics*, 29(1):25–49, February 1986.

[8] Dora Biro, Takao Sasaki, and Steven J. Portugal. Bringing a time–depth perspective to collective animal behaviour. *Trends in Ecology and Evolution*, 31(7):550–562, July 2016.

[9] Olivier Boissier, Rafael H. Bordini, Jomi F. Hübner, Alessandro Ricci, and Andrea Santi. Multi-agent oriented programming with jacamo. *Science of Computer Programming*, 78(6):747–761, 2013. Special section: The Programming Languages track at the 26th ACM Symposium on Applied Computing (SAC 2011) & Special section on Agent-oriented Design Methods and Programming Techniques for Distributed Computing in Dynamic and Complex Environments.

[10] Olivier Boissier, Jomi F. Hubner, and Alessandro Ricci. The jacamo framework. In Huib Aldewereld, Olivier Boissier, Virginia Dignum, Pablo Noriega, and Julian Padget, editors, *Social Coordination Frameworks for Social Technical Systems*, pages 125–151. Springer International Publishing, Cham, 2016.

[11] Yasser Bourahla, Manuel Atencia, and Jérôme Euzenat. 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*, pages 242–250, 2021.

[12] Yasser Bourahla, Manuel Atencia, and Jérôme Euzenat. Knowledge transmission and improvement across generations do not need strong selection. In Piotr Faliszewski, Viviana Mascardi, Catherine Pelachaud, and Matthew Taylor, editors, *Proc. 21st ACM international conference on Autonomous Agents and Multi-Agent Systems (AAMAS), (Online)*, pages 163–171, 2022.

[13] Hilary O. Box and Kathleen R. Gibson. *Mammalian social learning: Comparative and Ecological Perspectives.* Cambridge University Press, October 1999.

[14] Robert Boyd and Peter J. Richerson. *Culture and the evolutionary process.* University of Chicago Press, June 1988.

[15] Charlotte Brand, Stephen. Heap, Thomas. Morgan, and Alex Mesoudi. The emergence and adaptive use of prestige in an online social learning task. *Scientific Reports*, 10, 07 2020.

[16] Michael Bratman. *Intention, Plans, and Practical Reason.* Cambridge, MA: Harvard University Press, Cambridge, 1987.

[17] Jean-Pierre Briot and Yves Demazeau. *Principes et architecture des systèmes multi-agents.* Collection IC2. Hermès-Lavoisier, 2001.

[18] Briggs Buchanan and Marcus J. Hamilton. Networks and Cultural Transmission in Hunter-Gatherer Societies. In *The Oxford Handbook of Archaeological Network Research*. Oxford University Press, 11 2023.

[19] Richard V. Burkhauser, Shuaizhang Feng, and Stephen P. Jenkins. Using the p90/p10 index to measure u.s. inequality trends with current population survey data: a view from inside the census bureau vaults. *Review of Income and Wealth*, 55(1):166–185, 2009.

[20] Maxwell N. Burton-Chellew, Claire El Mouden, and Stuart A. West. Social learning and the demise of costly cooperation in humans. *Proceedings of the Royal Society B: Biological Sciences*, 284(1853):20170067, April 2017.

[21] Christine A. Caldwell and Ailsa E. Millen. Experimental models for testing hypotheses about cumulative cultural evolution. *Evolution and Human Behavior*, 29(3):165–171, May 2008.

[22] Christine A. Caldwell and Ailsa E. Millen. Social learning mechanisms and cumulative cultural evolution. *Psychological Science*, 20(12):1478–1483, December 2009.

[23] Luigi Luca Cavalli-Sforza and Marcus W. Feldman. *Cultural transmission and evolution: A Quantitative Approach*. Princeton University Press, May 1981.

[24] Nitech V. Chawla and Kevin Bowyer. Combining decision trees learned in parallel. In *Fourth International Conference on Knowledge Discovery and Data Mining (KDD-98), New York City, New York, USA*, 1998.

[25] Michael Chimento, Gustavo Alarcón-Nieto, and Lucy M. Aplin. Population turnover facilitates cultural selection for efficiency in birds. *Research Square (Research Square)*, June 2021.

[26] Michael Cogswell, Jiasen Lu, Stefan Lee, Devi Parikh, and Dhruv Batra. Emergence of compositional language with deep generational transmission. Research Report arXiv:1904.09067, arXiv, 2019.

[27] Dominik Deffner, Vivien Kleinow, and Richard McElreath. Dynamic social learning in temporally and spatially variable environments. *Royal Society Open Science*, 7(12):200734, December 2020.

[28] Anna Dornhaus. Specialization does not predict individual efficiency in an ant. *PLOS Biology*, 6(11):1–8, 11 2008.

[29] Charles Efferson, Rafael Lalive, Peter J. Richerson, Richard McElreath, and Mark Lubell. Conformists and mavericks: the empirics of

frequency-dependent cultural transmission. *Evolution and Human Behavior*, 29(1):56–64, January 2008.

[30] Magnus Enquist, Kimmo Eriksson, and Stefano Ghirlanda. Critical social learning: a solution to rogers's paradox of nonadaptive culture. *American Anthropologist*, 109(4):727–734, December 2007.

[31] Kimmo Eriksson and Julie Coultas. Corpses, maggots, poodles and rats: Emotional selection operating in three phases of cultural transmission of urban legends. *Journal of Cognition and Culture*, 14(1–2):1–26, January 2014.

[32] Jérôme Euzenat. Replicator-interactor in experimental cultural knowledge evolution. In *WINKS 2019 - 2nd JOWO workshop on Interaction-Based Knowledge Sharing*, pages 1–11, Graz, Austria, September 2019.

[33] Jianping Fan, Yuli Gao, and Hangzai Luo. 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.

[34] Nicolas Fay, Simon Garrod, Leo Roberts, and Nik Swoboda. The interactive evolution of human communication systems. *Cognitive Science*, 34(3):351–386, 2010.

[35] Ernst Fehr and Simon Gächter. Cooperation and punishment in public goods experiments. *The American Economic Review*, 90(4):980–994, 2000.

[36] Laurel Fogarty, Nicole Creanza, and Marcus W. Feldman. Cultural evolutionary perspectives on creativity and human innovation. *Trends in Ecology and Evolution*, 30(12):736–754, December 2015.

[37] Jacob G. Foster, Andrey Rzhetsky, and James A. Evans. Tradition and innovation in scientists' research strategies. *American sociological review*, 80(5):875–908, 9 2015.

[38] Marco Jerome Gasparrini, Ricard V. Solé, and Martí Sánchez-Fibla. Individual specialization in multi-task environments with multiagent reinforcement learners. In *Artificial Intelligence Research and Development - Proceedings of the 22nd International Conference of the Catalan Association for Artificial Intelligence, CCIA 2019, Mallorca, Spain, 23-25 October 2019*, volume 319 of *Frontiers in Artificial Intelligence and Applications*, pages 339–343. IOS Press, 2019.

[39] Mohamed Ghalwash, Zijun Yao, Prithwish Chakraporty, James Codella, and Daby Sow. Phenotypical ontology driven framework for multi-task learning.

In *Proceedings of the Conference on Health, Inference, and Learning*, CHIL '21, page 183–192, New York, NY, USA, 2021. Association for Computing Machinery.

[40] Luke Glowacki and Lucas Molleman. Subsistence styles shape human social learning strategies. *Nature Human Behaviour*, 1(5), April 2017.

[41] Alejandro Guerra-Hernández, Amal El Fallah-Seghrouchni, and Henry Soldano. Learning in bdi multi-agent systems. In *Proceedings of the 4th International Conference on Computational Logic in Multi-Agent Systems*, CLIMA IV'04, page 218–233, Berlin, Heidelberg, 2004. Springer-Verlag.

[42] Michael Gurven, Raziel Davison, and Thomas S. Kraft. The optimal timing of teaching and learning across the life course. *Philosophical Transactions of the Royal Society B*, 375(1803):20190500, June 2020.

[43] Laura Harding Graesser, Kyunghyun Cho, and Douwe Kiela. 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)*, pages 3700–3710, Hong Kong, China, November 2019. Association for Computational Linguistics.

[44] Serhii Havrylov and Ivan Titov. 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.

[45] Joseph Henrich, Robert Boyd, Samuel Bowles, Colin Camerer, Ernst Fehr, Herbert Gintis, Richard McElreath, Michael Alvard, Abigail Barr, Jean Ensminger, and et al. "economic man" in cross-cultural perspective: Behavioral experiments in 15 small-scale societies. *Behavioral and Brain Sciences*, 28(6):795–815, 2005.

[46] Joseph Henrich, Richard McElreath, Abigail Barr, Jean Ensminger, Cynthia T. Barrett, Alexander H. Bolyanatz, Juan Camilo Cárdenas, Michael Gurven, Edwins Laban Moogi Gwako, Natalie Henrich, Carolyn Lesorogol, Frank W. Marlowe, David P. Tracer, and John P. Ziker. Costly punishment across human societies. *Science*, 312(5781):1767–1770, June 2006.

[47] Cecilia M. Heyes and Bennet G. Galef. Social learning in animals: The roots of culture. *Journal of Anthropological Research*, 54(1):97–99, 1998.

[48] Ángel V. Jiménez and Alex Mesoudi. The cultural transmission of prestige and dominance social rank cues: an experimental simulation. *Evolutionary Psychological Science*, 7(2):189–199, September 2020.

[49] Andreas Kalaitzakis. 20220108-MTOA (2022-01-08): Accuracy improvement from undertaking additional overlapping tasks, follows a law of diminishing returns, 2022. https://sake.re/20220108-MTOA.

[50] Andreas Kalaitzakis. 20220109-MTOA (2023-01-09): Tackling one task affects the accuracy of the same agent on a different task, 2022. https://sake.re/20220109-MTOA.

[51] Andreas Kalaitzakis. 20230110-MTOA (2023-01-10): Carrying fewer tasks is detrimental to the accuracy of agents with un-restricted resources, 2023. https://sake.re/20230110-MTOA.

[52] Andreas Kalaitzakis. 20230120-MTOA (2023-01-20): Carrying fewer tasks does not affect the accuracy of resource-restricted agents but improves the prosperity and equitability of agent societies, 2023. https://sake.re/20230120-MTOA.

[53] Andreas Kalaitzakis. 20230505-MTOA (2023-05-05): Multitasking agents improve their average and maximum accuracy when tasks overlap, 2023. https://sake.re/20230505-MTOA.

[54] Andreas Kalaitzakis. 20231120-MTOA (2023-11-20): Single-tasking agent populations cannot achieve equitable task exploration, 2023. https://sake.re/20231120-MTOA.

[55] Andreas Kalaitzakis. 20231220-MTOA (2023-12-20): Multitasking agent populations can achieve equitable task exploration to the detriment of their agents' average and maximum accuracy, 2023. https://sake.re/20231220-MTOA.

[56] Andreas Kalaitzakis and Jérôme Euzenat. Multi-tasking resource-constrained agents reach higher accuracy when tasks overlap. In Vadim Malvone and Aniello Murano, editors, *Multi-Agent Systems - 20th European Conference, EUMAS 2023, Naples, Italy, September 14-15, 2023, Proceedings*, volume 14282 of *Lecture Notes in Computer Science*, pages 425–434. Springer, 2023.

[57] Andreas Kalaitzakis and Jérôme Euzenat. À quoi sert la spécialisation en évolution culturelle de la connaissance? In Maxime Morge, editor, *Actes 31e journées francophones sur Systèmes multi-agent (JFSMA), Strasbourg (FR)*, pages 76–85, 2023.

[58] Oren Kolodny, Nicole Creanza, and Marcus W. Feldman. Evolution in leaps: The punctuated accumulation and loss of cultural innovations. *Proceedings of the National Academy of Sciences of the United States of America*, 112(49):E6762–E6769, Dec 2015. Epub 2015 Nov 23.

[59] Kevin N. Laland, Peter J. Richerson, and Robert Boyd. Developing a theory of animal social learning. In *Social learning in animals: The roots of culture (pp. 129–154). Academic Press.*, page 129–154, January 1996.

[60] Shakti Lamba. Social learning in cooperative dilemmas. *Proceedings of the Royal Society B: Biological Sciences*, 281(1787):20140417, 2014.

[61] Lazy lavender, 2020. https://gitlab.inria.fr/moex/lazylav.

[62] Ellouise Leadbeater and Lars Chittka. Social learning in insects — from miniature brains to consensus building. *Current Biology*, 17(16):R703–R713, August 2007.

[63] David Kellogg Lewis. *Convention: A Philosophical Study.* Wiley-Blackwell, Cambridge, MA, USA, 1969.

[64] Chengzhi Mao, Amogh Gupta, Vikram Nitin, Baishakhi Ray, Shuran Song, Junfeng Yang, and Carl Vondrick. Multitask learning strengthens adversarial robustness. In *Computer Vision - ECCV 2020 - 16th European Conference, Glasgow, UK, August 23-28, 2020, Proceedings, Part II*, volume 12347 of *Lecture Notes in Computer Science*, pages 158–174. Springer, 2020.

[65] Richard McElreath, Adrian Bell, Charles Efferson, Mark Lubell, Peter Richerson, and Timothy Waring. Beyond existence and aiming outside the laboratory: Estimating frequency-dependent and pay-off-biased social learning strategies. *Philosophical transactions of the Royal Society of London. Series B, Biological sciences*, 363:3515–28, 09 2008.

[66] Alex Mesoudi. An experimental comparison of human social learning strategies: payoff-biased social learning is adaptive but underused. *Evolution and Human Behavior*, 32(5):334–342, 2011.

[67] Alex Mesoudi. Experimental studies of cultural evolution. In *The Oxford Handbook of Cultural Evolution*, pages C6S1–C6S14, 05 2023.

[68] Alex Mesoudi and Alex Thornton. What is cumulative cultural evolution? *Proceedings of the Royal Society B: Biological Sciences*, 285(1880):20180712, June 2018.

[69] Carol Myers-Scotton. *Contact Linguistics: Bilingual Encounters and Grammatical Outcomes.* Oxford University Press, August 2002.

[70] Jason Noble and Daniel W. Franks. Social learning in a multi-agent system. *Computing and Informatics*, 22(6):561–574, 2004.

[71] Jason Noble and Peter M. Todd. *Imitation or something simpler? modeling simple mechanisms for social information processing*, page 423–439. MIT Press, Cambridge, MA, USA, 2002.

[72] Shayegan Omidshafiei, Jason Pazis, Christopher Amato, Jonathan P. How, and John Vian. Deep decentralized multi-task multi-agent reinforcement learning under partial observability. In *Proceedings of the 34th International Conference on Machine Learning - Volume 70*, ICML'17, page 2681–2690. JMLR.org, 2017.

[73] T. Phung, M. Winikoff, and L. Padgham. Learning within the bdi framework: An empirical analysis. In *Proceedings of the 9th International Conference on Knowledge-Based Intelligent Information and Engineering Systems - Volume Part III*, KES05, pages 282–288, 2005.

[74] Jordi Pinero and Ricard V. Solé. Statistical physics of liquid brains. *Philosophical Transactions of the Royal Society B*, 374(1774):20180376, April 2019.

[75] Anton Proskurnikov and Ming Cao. Consensus in multi-agent systems. In *Wiley Encyclopedia of Electrical and Electronics Engineering*, 11 2016.

[76] Anand S. Rao and Michael P. Georgeff. Modeling rational agents within a bdi-architecture. In *Proceedings of the Second International Conference on Principles of Knowledge Representation and Reasoning*, KR'91, page 473–484, San Francisco, CA, USA, 1991. Morgan Kaufmann Publishers Inc.

[77] Anand S Rao, Michael P Georgeff, et al. Bdi agents: from theory to practice. In *Icmas*, volume 95, pages 312–319, 1995.

[78] Simon M. Reader and Dora Biro. Experimental identification of social learning in wild animals. *Learning & Behavior*, 38(3):265–283, July 2010.

[79] Peter J. Richerson and Robert Boyd. *Not by genes alone: How Culture Transformed Human Evolution.* University of Chicago Press, June 2008.

[80] Andrey Rzhetsky, Jacob G. Foster, Ian T. Foster, and James A. Evans. Choosing experiments to accelerate collective discovery. *Proceedings of the National Academy of Sciences of the United States of America*, 112(47):14569–14574, 11 2015.

[81] Takao Sasaki and Dora Biro. Cumulative culture can emerge from collective intelligence in animal groups. *Nature Communications*, 8(1), April 2017.

[82] Paul Seabright. *The company of Strangers: A Natural History of Economic Life.* Princeton University Press, May 2010.

[83] Trevor Standley, Amir Zamir, Dawn Chen, Leonidas Guibas, Jitendra Malik, and Silvio Savarese. Which tasks should be learned together in multi-task learning? In *Proceedings of the 37th International Conference on Machine Learning*, ICML'20. JMLR.org, 2020.

[84] Luc Steels. 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)*, pages 143–150, Hatfield, 2005. University of Hertfordshire. Language.

[85] Pontus Strimling, Magnus Enquist, and Kimmo Eriksson. Repeated learning makes cultural evolution unique. *Proceedings of the National Academy of Sciences of the United States of America*, 106(33):13870–13874, August 2009.

[86] Pontus Strimling, Jonas Sjöstrand, Magnus Enquist, and Kimmo Eriksson. Accumulation of independent cultural traits. *Theoretical Population Biology*, 76(2):77–83, September 2009.

[87] Alex Thornton and Tim Clutton-Brock. Social learning and the development of individual and group behaviour in mammal societies. *Philosophical Transactions of the Royal Society B*, 366(1567):978–987, April 2011.

[88] Michael Tomasello. Chapter 15 - do apes ape? In Cecilia M. Heyes and Bennet G. Galef, editors, *Social Learning in Animals: The Roots of Culture*, pages 319–346. Academic Press, San Diego, 1996.

[89] Wataru Toyokawa, Yoshimatsu Saito, and Tatsuya Kameda. Individual differences in learning behaviours in humans: Asocial exploration tendency does not predict reliance on social learning. *Evolution and Human Behavior*, 38(3):325–333, 2017.

[90] Wataru Toyokawa, Andrew Whalen, and Kevin N. Laland. Social learning strategies regulate the wisdom and madness of interactive crowds. *Nature Human Behaviour*, 3(2):183–193, January 2019.

[91] Shresth Verma. Emergence of multilingualism in population based referential games. In *Language in Reinforcement Learning Workshop at the 37th International Conference on Machine Learning (ICML 2020)*, Vienna, Austria, July 2020.

[92] Jun Wang and Les Gasser. Mutual online ontology alignment. In *Proc. 1st ACM international conference on Autonomous Agents and Multi-Agent Systems (AAMAS), Bologna, Italy*, 2002.

[93] Andrew Whiten. Imitation of the sequential structure of actions by chimpanzees (pan troglodytes). *Journal of Comparative Psychology*, 112(3):270–281, January 1998.

[94] Andrew Whiten. Culture extends the scope of evolutionary biology in the great apes. *Proceedings of the National Academy of Sciences of the United States of America*, 114(30):7790–7797, July 2017.

[95] Andrew Whiten, Christine A. Caldwell, and Alex Mesoudi. Cultural diffusion in humans and other animals. *Current Opinion in Psychology*, 8:15–21, April 2016.

[96] Michael Wooldridge. Conceptualising and developing agents. In *Proceedings of the UNICOM Seminar on Agent Software*, pages 40–54, London, April 1995.

[97] Michael Wooldridge and Nicholas R. Jennings. Intelligent agents: theory and practice. *The Knowledge Engineering Review*, 10:115 – 152, 1995.

[98] Xiaokai Yang and Jeff Borland. A microeconomic mechanism for economic growth. *Journal of Political Economy*, 99(3):460–482, 1991.

[99] Xiaokai Yang and Siang Ng. Specialization and division of labour: A survey. In *Increasing Returns and Economic Analysis*, pages 3–63, London, UK, 1998.

Abstract

This thesis explores how the presence of several tasks affects the cultural evolution of knowledge within agent societies. Previous findings show that when agents coordinate over a single shared task, this improves their accuracy. We build upon these by investigating how coordinating over additional tasks influences the cultural evolution of knowledge, particularly when the cognitive resources of agents are limited. Our goal is to determine whether the cultural evolution of knowledge guarantees: (1) equidistribution of collective resources among different tasks, and (2) improvement of a society's knowledge across all tasks. To that end, we propose a framework allowing for examining what pushes agents to specialize and how this specialization affects both individual agents and agent societies. Our findings reveal several key insights: First, agents improve their knowledge on one task by undertaking additional tasks. Second, while assigning different tasks to different agents benefits their societies, it is not what triggers agent specialization. On the contrary, specialization emerges when agents allocate limited cognitive resources to these tasks. Third, while temporarily detrimental to their knowledge correctness, forgetting randomly selected knowledge brings long-term benefits when agents are assigned a single task. Lastly, when agents specializing in different tasks interact, one task will gradually monopolize the collectively available resources, causing the entire society to specialize in the same task. These results improve our understanding of how the formation, evolution, and prevalence of specialized knowledge affect individual agents and agent societies.

Keywords: Multi-agent simulation; Adaptive multi-agent systems; Cultural evolution

Résumé

Cette thèse explore la manière dont la présence de plusieurs tâches affecte l'évolution culturelle des connaissances au sein des sociétés d'agents. Des études antérieures ont montré que la coordination des agents autour d'une même tâche améliore leur précision. Nous nous appuyons sur ces résultats pour étudier comment la coordination de tâches supplémentaires influence l'évolution culturelle des connaissances, en particulier lorsque les ressources cognitives des agents sont limitées. Notre objectif est de déterminer si l'évolution culturelle des connaissances garantit : (1) une répartition équitable des ressources collectives entre les différentes tâches et (2) l'amélioration des connaissances d'une société dans toutes les tâches. À cette fin, nous proposons un cadre permettant d'examiner ce qui pousse les agents à se spécialiser et comment cette spécialisation affecte à la fois les agents individuels et les sociétés d'agents. Nos résultats révèlent plusieurs éléments clés. Tout d'abord, les agents améliorent leur connaissance d'une tâche en entreprenant des tâches supplémentaires. Deuxièmement, si l'attribution de tâches différentes à des agents différents bénéficie à leurs sociétés, elle n'est pas à l'origine de la spécialisation des agents. À l'inverse, la spécialisation apparaît lorsque différentes tâches sont en concurrence pour des ressources cognitives limitées. Troisièmement, bien que temporairement préjudiciable à l'exactitude de leurs connaissances, l'oubli de connaissances sélectionnées au hasard présente des avantages à long terme lorsque les agents sont affectés à une tâche unique. Enfin, lorsque des agents spécialisés dans des tâches différentes interagissent, une tâche monopolise progressivement les ressources collectivement disponibles, ce qui amène l'ensemble de la société à se spécialiser dans la même tâche. Ces résultats améliorent notre compréhension de la manière dont la formation, l'évolution et la prévalence des connaissances spécialisées affectent les agents individuels et les sociétés d'agents.

Mots clés: Simulation multi-agents; Systèmes multi-agents adaptatifs; Évolution culturelle