

Knowledge Transmission and Improvement Across Generations do Not Need Strong Selection

Yasser Bourahla
Univ. Grenoble Alpes, Inria, CNRS,
Grenoble INP, LIG
F-38000 Grenoble, France
Yasser.Bourahla@inria.fr

Manuel Atencia
Universidad de Málaga
29071 Málaga, Spain
atencia@lcc.uma.es

Jérôme Euzenat
Univ. Grenoble Alpes, Inria, CNRS,
Grenoble INP, LIG
F-38000 Grenoble, France
Jerome.Euzenat@inria.fr

ABSTRACT

Agents have been used for simulating cultural evolution and cultural evolution can be used as a model for artificial agents. Previous results have shown that horizontal, or intra-generation, knowledge transmission allows agents to improve the quality of their knowledge to a certain level. Moreover, variation generated through vertical, or inter-generation, transmission allows agents to exceed that level. Such results were obtained under specific conditions such as the drastic selection of agents allowed to transmit their knowledge, seeding the process with correct knowledge or introducing artificial noise during transmission. Here, we question the necessity of such measures and study their impact on the quality of transmitted knowledge. For that purpose, we combine the settings of two previous experiments and relax these conditions (no strong selection of teachers, no fully correct seed, no introduction of artificial noise). The rationale is that if interactions lead agents to improve their overall knowledge quality, this should be sufficient to ensure correct knowledge transmission, and that transmission mechanisms are sufficiently imperfect to produce variation. In this setting, we confirm that vertical transmission improves on horizontal transmission even without drastic selection and oriented learning. We also show that horizontal transmission is able to compensate for the lack of parent selection if it is maintained for long enough. This means that it is not necessary to take the most successful agents as teachers, neither in vertical nor horizontal transmission, to cumulatively improve knowledge.

KEYWORDS

Multi-agent social simulation; Cultural evolution; Knowledge transmission; Agent generation; Cultural knowledge evolution

Reference: Yasser Bourahla, Manuel Atencia and Jérôme Euzenat. Knowledge transmission and improvement across generations do not need strong selection. In P. Faliszewski, V. Mascardi, C. Pelachaud, M.E. Taylor (eds.), Proc. of the 21st International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2022), pp163–171, Online, May 9–13, 2022.

1 INTRODUCTION

In cultural evolution [13, 22], concepts from biological evolution are applied to the culture of a society. This has been explored

This is the author’s version of the work. It is posted here for your personal use. Not for redistribution. The definitive Version of Record was published in *Proc. of the 21st International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2022)*, May 9–13, 2022

© 2022 International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). All rights reserved.

experimentally through multi-agent systems to evolve an agent population’s behaviour [6, 14], language [24] or knowledge [15] as their culture. The evolution of culture relies on its ability to display variations that can be selected and transmitted. Early work has identified different cultural transmission modes inspired by epidemiology [11, 13]: *vertical transmission* is the transmission from parents to children, *oblique transmission* is the transmission from agents of the previous generation (think about education) to those of the next generation, and *horizontal transmission* is the transmission between agents of the same generation. Here, we will also use *inter-generation transmission* for the two former and *intra-generation transmission* for the latter.

Multi-agent work in which agents adapt to interactions among themselves [24] can be considered as intra-generation transmission. This has proven to be efficient for reaching consensus and for sharing culture. Recently, it has been shown that agents that adapt their knowledge about the environment, expressed as ontologies, through social interactions in order to reach agreement are able to improve their ontology accuracy without necessarily adopting the same ontology [10]. Thus, horizontal transmission is able to improve knowledge in a society of agents.

Vertical transmission, by transmitting culture from one generation of agents to another, may be seen as the occasion to shuffle the cards. It can introduce more variation or enforce (select) a dominant culture. This calls for assessing the respective roles of inter-generation and intra-generation transmission.

Another line of work considered exactly this and showed that inter-generation transmission allows agents to improve knowledge beyond what intra-generation transmission alone does [14]. In contrast, by disabling inter-generation transmission, agents have to start improving knowledge from scratch at each generation. In a similar setting, it was also found that intra-generation transmission generates variation and that inter-generation transmission performs selection [2].

However, both works rely on strong conditions to ensure the faithful transmission of correct knowledge (using only the best part of the population as teachers, initialising agent ontologies from correct samples) or variation (adding noise in the learning process).

Here, we consider relaxing these assumptions to study the effects of horizontal and vertical culture transmission without selecting the best-performing agents and without starting with correct hints. This is achieved by offering a wider, yet differential, opportunity to breed and propagate knowledge to the whole agent society and by starting with random knowledge. This may eventually provide enough variation in the system.

To test this, we designed a new experimental framework building on these previous works. We performed experiments with inter-generation transmission through teaching [2] and intra-generation transmission through interacting [10], varying the extent of both transmission modalities and mating policies. We confirm the results of previous studies under these relaxed assumptions: inter-generation transmission increases the accuracy of knowledge obtained from intra-generation transmission. Moreover, we show that even if inter-generation transmission does not select the fittest parents, intra-generation transmission, if given enough time, can compensate for this, i.e. agents finally agree to select relevant pieces of knowledge.

The remainder of the paper is organised as follows: in Section 2, the related work is reviewed and more details on the two initial experiments are provided. Section 3 discusses the limitations of current settings and offers solutions to overcome them. Their actual implementation is precisely detailed in Section 4. Section 5 presents the performed experiments and Section 6 provides their results which are discussed in Section 7.

2 RELATED WORK

We first present work related to agents and cultural evolution generally before considering more precisely two experiments on which this work is based.

2.1 Agents and Cultural Evolution

Interactions between agents can serve as a mean of adaptation. Many approaches exploit them in multi-agent systems, for instance, to improve interoperability [5, 25] or observe emergent behaviours [19]. In some cases, by interacting with each other, agents may end up sharing and evolving a particular culture, e.g. knowledge about the environment, ways to interact, norms or conventions [16, 23]. This is typically studied by observing a population of agents that interact with each other through a well-defined protocol. The state of the system is monitored until agents reach a stable state in which their culture does not change any more. The characteristics of this culture are then assessed [24].

In this work, we are interested in the evolution of knowledge that agents use to behave. Cultural evolution has been applied to knowledge in the form of ontology alignments [3, 15]. It has also been applied to ontologies [10] in which agents adapt their concept definitions to agree with each other. They improved their knowledge about the environment without necessarily sharing the same ontology. Agents in this setting learn from others by interacting with them, which is known as social learning [18, 20].

However, when the agents reach a full agreement, they stop adapting their knowledge. In a sense, they reach a local optimum with respect to the accuracy of their knowledge. Escaping this state may be obtained by either changing the environment or the agents.

Evolving different agent generations has been considered in evolutionary multi-agent systems [12] in which agents are able to reproduce and may die. Although agents could be thought of as developing a culture, this has not been studied from that perspective. Introducing cultural mechanisms within agents has been developed in cultural algorithms [21]. However, this work has been mainly applied to optimisation problems.

It has been shown that culture transmission allows the discovery of behaviours that are inaccessible to genetic evolution [7, 17]. Multi-agent simulation of cultural evolution involving intra-generation and inter-generation transmission has been investigated in a series of experiments [2, 4, 6, 14]. Based on neural networks for embodying knowledge, they implement connection weight communication [6], supervised learning (imitation) [2, 14] or direct instruction [4] as a vertical transmission mechanism. Using the same imitation mechanism for horizontal transmission [2], it was shown that the intra-generation culture transmission provided the variation that allowed the evolutionary process to be more effective than what it is with inter-generation transmission alone. This work is however based on a drastic selection of teachers from whom agents learn, both in oblique and horizontal transmission.

Below, we discuss more lengthly the experiments developed in [2] and [10] on which we will build.

2.2 Inter-Generation Culture Inheritance

Acerbi and Parisi [2] designed an experiment in which agents use a neural network to navigate in their environment. The lifetime of agents spans 10 epochs of 50 games. At each game, the agent is presented with a situation and decides how to move. When it comes close to an edible food source, it receives a reward and when it comes close to a poisonous food source, it receives a penalty. In addition, a teacher discloses its decision to the agent, that it uses to adjust its network weights. This operation is repeated over 500 generations of 100 agents all replaced at once.

At birth, agents start with random weighted networks (W). The first 6 epochs of their lifetime are dedicated to oblique transmission: they are taught their behaviour by agents of the previous generation. The next 4 epochs implements horizontal transmission: they are taught by agents from their generation (Figure 1, left). One key point is that in both cases they are only taught by the best 5% agents in terms of accumulated rewards. Teaching is achieved by constraining the output of the neural network. Teachers may add noise in their behaviour in order to generate variation.

Experiments are reported with oblique transmission only, horizontal transmission only, or both. In particular, results show that oblique transmission alone allows agents to improve (as measured by collected rewards) over generations, though horizontal transmission does not, because agents learn from novice agents. They also show that the introduction of noise is key to speed up improvement, and that horizontal transmission can provide it.

In these experiments, the selection of the 5% best performing agents in both transmission modes (20% for [14]) is a very strong bias because it gives little opportunity to general individuals to reproduce. This does not happen in human societies.

2.3 Intra-Generation Social Learning

Bourahla and colleagues [10] designed an experiment in which agents live in an environment populated by objects identified by their features, e.g. ‘can move’ or ‘has fangs’, for which they have to make decisions, e.g. ‘hunt’ or ‘hide’. Agents’ goal is to agree with other agents on the decision to make. They follow a two-step experiment protocol (Figure 1, right). First, agents learn to make decisions about objects of their environments from a set of samples




Figure 1: Experimental frameworks of [2] (left) and [10] (right). (W =neural networks weights, S =training sample, O =ontology). In spite of structural similarities, the mechanisms for transmitting knowledge are different: both inter-generation mechanisms are supervised learning, but intra-generation transmission is different.

(S). The output of this step is, in particular, a simple ontology (O) allowing to classify objects and to determine the decision to make. Then, agents repetitively interact with each other two-by-two. They compare the decisions that they would make about objects. The interactions succeed when the agents agree on the decision; they fail otherwise. In the event of a failure, one of the agents adapts its ontology in order to adopt the decision of the other on this object.

Experiments showed that agents are able to reach a state in which they always agree. The quality of their ontologies, as measured by its accuracy, increases on average, but not in every case, and does not necessarily reach full accuracy. Finally, though the distance between agents' ontologies decreases, agents do not necessarily end up with the same ontologies.

The success criterion for agents is to agree, not necessarily to make the correct decision. The correct decision to be made is unknown to the agents. It is used only in two occasions: (1) initially, agents are trained with a sample which is correct, but incomplete; (2) the choice of which agent will adapt its knowledge is based on rewards that they have received on a (different) sample. Although (2) is legitimate, (1) is not. Hence, we relax this by using samples labelled by parents not knowing the correct decision.

3 TOWARDS AN INTEGRATED APPROACH TO KNOWLEDGE TRANSMISSION

We first discuss the limitations of the experiments presented before. Then we propose to combine them in a way addressing these limitations. This allows us to formulate our claims as hypotheses.

3.1 Limitations

The two experiments described in the last subsections have the merit to provide a model of cultural transmission in agents and to show that it has a positive effect on their efficiency. However, they do it with modalities which seem strong and unnatural:

- starting with correct samples does not correspond to what happens in real life;
- restricting the role of teachers to a very limited part of the population (5%, 20%) seems to go counter the idea of culture and its propagation.

Interestingly, these two modalities tend to enforce the faithful transmission of correct/fit culture, though the benefits of vertical transmission are supposed to come from the variation it provides. This is to the extent that [2] introduces artificial noise in the model in order to increase variation.

We conjecture that a less restricted opportunity to transmit knowledge, both across and within generations, provides enough variation to improve over horizontal transmission.

3.2 Combining the Two Experiments

Although the two experiments described above are about knowledge transmission, they differ on many notable features, such as:

- oblique and horizontal transmission vs. horizontal transmission;
- neural networks vs. ontologies and decision trees;
- selecting teachers vs. random interaction;
- positive and negative reward vs. agreement.

Some of these differences are secondary, such as the choice of the representation, though others are more important, such as those affecting transmission modalities.

To test our conjecture, we combine both approaches of Figure 1 and relax these unwanted constraints. The resulting setting, as shown in Figure 2, is such that:

intra-generation transmission is implemented by interaction exactly as in [10].

inter-generation transmission is implemented through learning. As in [2], a generation of agents is replaced by another generation. Instead of implementing oblique transmission, agents reproduce differentially with respect to a probability distribution and transmit their knowledge to their children. However, they do it with the modalities of [10].

This should allow us to determine under which conditions agents can improve generation after generation, or if they reproduce their knowledge, thus behaviour, identically.

3.3 Relaxing Elitist Assumptions

For relaxing selection, the reproduction phase is implemented by parents having children. The probability to have children will be determined either equiprobably or relative to the income of the agents (denoting their success).

Important changes concern vertical transmission. It is split into two steps: (a) the acquisition of an initial ontology by the children, and (b) the transmission of knowledge from the parents by first interacting predominantly with them and then progressively widening their interactions.

The initial ontology may either be empty, be random, be a merge of the parents' ontologies or be learned as in the first step above, from samples provided by the parents (Figure 2). Here, agents either start with random ontologies (suppressing the learning phase) or learn their ontologies under the supervision of their parents which provide the samples from which to learn.

This setting ensures a wide opportunity for agents to have children to which to transmit their knowledge. Moreover, the sample from which children learn may now be incorrect (because parents

do not have fully correct knowledge and because samples are provided by both parents) and incomplete (because they do not cover the whole object space). They thus relax the first two identified limitations.

Of course, this setting is not exempt of bias: the most successful agents, which may not have the most correct knowledge, are more likely to have children and, during horizontal transmission, success will also play a role in which agent will learn from which other. However, this bias still allows a broad fraction of the population to reproduce and transmit knowledge.

Finally, no artificial noise is added to any agent behaviour. As mentioned above, the ‘egalitarian’ recruitment of parents and the incorrectness of samples provided by parents are already sources of variation during vertical transmission. The disconnection between ontology accuracy (the measure of knowledge quality) and agreement (the goal pursued by agents) is an additional source of variation during horizontal transmission.

3.4 Hypotheses

This framework will be used to test our claims experimentally. We aim at investigating primarily two directions. First, once agents reach a global agreement, they do not adapt their knowledge anymore. As shown in [10], they may still agree on incorrect decisions and they preserve the diversity of their knowledge. The inter-generation transmission should introduce further variation allowing agents to discover new relevant pieces of knowledge. Accordingly, as in [2, 14], we hypothesise that (H_1) *vertical transmission allows new generations’ knowledge to be more accurate than that of the previous generation.*

According to [10], agents are able to improve the correctness of their decisions when they adapt their knowledge to agree with each other. This suggests that the intra-generation knowledge transmission is able to select pieces of knowledge to preserve. Thus, we hypothesise that (H_2), *interaction, used as intra-generation transmission, can compensate for the absence of parent selection.*

4 EXPERIMENTAL FRAMEWORK

In order to test the two proposed hypotheses, we have designed an experimental framework following the above principles with mechanisms for agents to mate, reproduce, die, and to transmit knowledge from parents to children.

Below we describe the general process (Section 4.1) and detail what differs from the original experiments:

- The agent life cycle (Section 4.2),
- How agents are selected to reproduce (Section 4.3),
- How agents transmit their knowledge to their children (Section 4.4).

Several questions can be raised on how these additions affect agents’ knowledge. In this paper, we focus on the impact of knowledge transmission on knowledge accuracy.

4.1 Overall Process

In these experiments, we reuse the exact same

- environments and ontologies,
- ontology learning procedure,

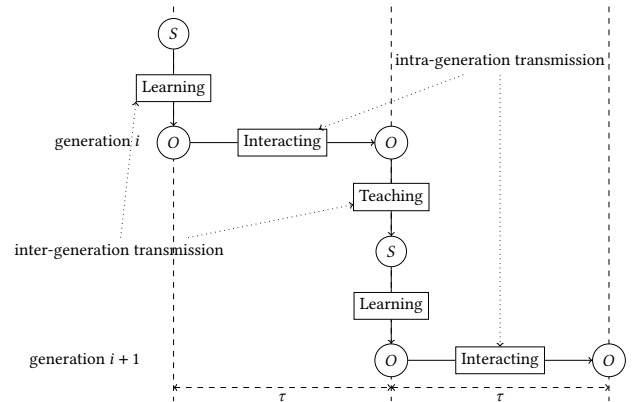


Figure 2: Combination of the mechanisms of Figure 1: horizontal transmission of knowledge is achieved through agents interacting; initial vertical transmission involves parents generating samples (S) from their ontologies (O) from which children will learn their initial ontology (O).

- social learning procedure, and
- evaluation measures,

as in [10]. The game is played with objects characterised by 4 features on which only one decision is valid among 6 possible ones.

Experiment runs are made of periods split in two parts (Figure 3):

- reproduction** in which generation $i-1$ (half of the population) is suppressed (dies) and generation $i+1$ is added (is born), and agents of the new generation acquire their initial knowledge;
- interaction** in which agents use their ontologies to agree on decisions about objects and modify them in case of disagreement.

The agents of the initial population (generations 0 and 1) are initialised with random ontologies.

This process is illustrated in Figure 2 and precisely described below.

4.2 Agent Life Cycle

An agent is born at the beginning of a period and dies at the end of the next period (it lives for 2 periods). In the first period, the agent is considered a child. During the second period, the agent is considered an adult. Between these two periods, agents may reproduce according to precise modalities explained in Section 4.3.

During the first period of their lives, (child) agents learn from their parents through inter-generation knowledge transmission detailed in Section 4.4. They gradually get detached from their parents and start interacting with other individuals of the society. During the second period of their lives, some agents have become parents and transmit their knowledge to their children. They still interact with the other individuals of the society and may adapt their knowledge.

At the end of each period, $\frac{X}{2}$ agents die and $\frac{X}{2}$ agents are born. This means that each generation is composed of $\frac{X}{2}$ agents and that the whole population is composed of X agents (X being even and greater than 4).

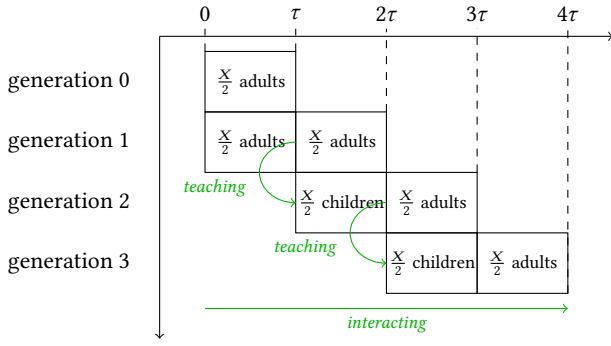


Figure 3: Evolution of a society through 3 agent generations.




Figure 4: Average of probability distributions of mating after each period (logarithmic scale). At each period, agents are numbered from 1 to 20 in the descending order of their income. Data collected with $\tau = 10001$, $X = 40$, $n = 100000$.

To keep the number of agents in the population constant, the initial population is composed of X agents. Half of this population, considered generation 0, dies at the end of the first period and does not reproduce. The other half, generation 1, behaves like the rest of the generations starting from the second period. Figure 3 shows the evolution of a society through 3 generations.

4.3 Agent Mating

In order to reproduce, (adult) agents behave along the following rule: $v = 2$ parents are selected randomly following a distribution s to have $c = 1$ child. As a consequence, individuals may have between 0 and $\frac{X}{2}$ children, with between 1 and $\frac{X}{2} - 1$ partners.

The probability to mate follows a distribution s that may be:

- Maximal (100%) for the v agents having gathered the most income and minimal (0%) for the other agents,
- Proportional to the income gathered from doing their tasks,
- Proportional to their success rate in interactions,
- Inversely proportional to the ontology distance with other agents,
- Equiprobable.

Here, we will only test the maximal (*best*), the income-based (*income*) and the equiprobable (*random*) distributions. The maximal strategy is introduced as a strong selection baseline, instead of [2]’s selection. It corresponds to 10% selection. Figure 4 shows these probability distributions. In the case of the performed experiment, income-based is very different from maximal and closer to equiprobable.

In [2], it is reported that only oblique and not vertical transmission is tested. However, the setting may be interpreted as vertical

transmission of the $v = \frac{X}{2} \times 5\%$ agents with largest income ($s = \textit{best}$) having together $c = \frac{X}{2}$ children.

4.4 Knowledge Transmission

The transmission process goes through two steps (Figure 3).

Initial vertical transmission. In the first stage, each agent of the new population acquires knowledge directly from its v parents. Children may have an initial ontology that is empty, random, the result of merging their parents’ ontologies, or the result of being taught by their parents. Although all options are possible, we only experiment with either random ontologies (abusively denoted by $r = 0$) or ontologies learnt from the parents reusing the technique of [10]. More precisely, $r\%$ of all objects with distinct properties (object types) are randomly selected. Each parent labels half of these objects with the decisions it would make with respect to its own ontology (which may be incorrect). This set of labelled objects is presented to the child as a training sample (S) from which it learns its ontology.

Interactions. Once this initial transmission has been performed, agents interact with other agents following the protocol of [10]. The interactions are constrained such that agents are initially biased towards more interactions with their parents: At its i^{th} interaction, each agent has probability P_i to interact with one of its parents selected randomly. The probability of restricting the agent interaction depends on i and the restricted interaction reduction rate $\epsilon < 1$. It is determined by:

$$P_i = \max(0, 1 - i \times \epsilon)$$

Thus, the bias is maximal at the first interaction, decreases as interactions increase and has no effect after interaction $\lceil \frac{1}{\epsilon} \rceil$. This mimics agents progressively broadening their social circles. In the experiments, we use $\epsilon = .01$. Hence, for each child, this bias disappears after 100 interactions.

When agents are adult, they can interact with anyone:

- their children (vertical transmission),
- the children of others (oblique transmission),
- other agents of their generation (horizontal transmission).

Given the non-oriented character of transmission, it is possible that an adult learns knowledge from a child, as this occurs in real life and contrary to genetic transmission. This learnt piece of knowledge may even be transmitted to future generations if it goes again to the next generation through vertical or oblique transmission.

Agents adapt their ontologies with an operator (*split*) which splits a concept in two according to the description of the object for which both agents do not agree [10].

5 EXPERIMENTS

We perform two experiments to test each of the hypotheses introduced in Section 3.4. They use the same modus operandi, only a few independent variables differ. Table 1 summarises the parameter values considered in these experiment plans. Values of τ end by 1, in order to compute measures at the $\tau - 1$ iteration, i.e. before the population is replaced. Each combination of parameter values is run 10 times.

5.1 Experiment 1: Effect of Inter-Generation Transmission on Ontology Accuracy

As mentioned in Section 2.3, intra-generation transmission converges towards a stable state in which agents always agree on the same decision. However, this decision may not be the correct one but, without feedback from the environment, agents have no reason to know it. Hence the accuracy of their ontology is not maximal.

Typically, this situation may evolve through changing the conditions (adding new agents, modifying the environment). The introduction of new agents, which have to learn their ontologies from imperfect ones (either starting with a random ontology or learning from an imperfect and incomplete sample provided by their genitors), introduces variations in the system.

The first experiment, aims at assessing the effects of introducing agent generations. Thus, we focus on the variables affecting vertical knowledge transmission: the proportion of instances covered by the training sample and the length of population life span, because it constrains the amount of interactions with parents.

As a consequence, we vary the transmission percentage r which corresponds to how complete and how imperfect the inter-generation transfer is. When the transmission sample ratio is 0, agents start with random ontologies. We also vary the length of the period τ which corresponds to agents' half-life. When the period length is greater than the number of iterations ($\tau > n$), the experiment happens within one generation (no variation of agent knowledge).

Hypothesis H_1 can thus be rephrased as *Adding inter-generation transmission leads to higher accuracy than intra-generation transmission alone*.

5.2 Experiment 2: Interaction Between Selection and Interaction Length

Selection mechanisms are assumed to provide a reproductive advantage: the most accurate ontologies will provide higher fitness which itself provides to their bearers the possibility to reproduce more. Such ontologies may spread more in further generations. Previous results confirmed this assumption and showed that selection played an important role in such a spreading (Section 2.2).

In this second experiment, we test the less selective policies. More specifically, we investigate whether intra-generation transmission, a typical cultural evolution mechanism, may compensate for the reduction or absence of parent selection.

This experiment focuses on (1) how parents are selected for reproduction, and (2) how long an agent generation lives: because agents need time to agree with each other on which pieces of knowledge to adopt. Thus, we vary the parent selection policy s as *random*, *income* and *best* and the period length τ as in the first experiment. We also perform the experiment for agents which do not adapt their knowledge after interaction ($op = none$), fully discarding horizontal transmission. To that extent, the adaptation operator *none* is introduced, the operator *split* being that of [10].

Hypothesis H_2 can thus be tested as *with sufficient intra-generation transmission, accuracy obtained with or without selection is similar*.

5.3 Ontology Accuracy

For each experiment, we report the ontology accuracy which is used to test the hypotheses. The accuracy of an ontology is computed as




Figure 5: Average accuracy (over r) by period lengths.

the proportion of object types, i.e. feature combinations, for which the correct decision is made by the ontology. It is averaged over the X ontologies of all agents at a particular iteration and averaged over 10 runs. The test of the hypotheses themselves applies to the final state of the experiment (iteration n). We display the evolution of the measure over all iterations.

6 RESULTS

In the following, the results for the two experiments are presented. For each of them, we first show how the results answer the associated hypothesis. Additional details are discussed regarding the interaction of the different parameters with the results.

6.1 Inter-Generation Transmission Improves Ontology Accuracy

To test Hypothesis H_1 , we compare the average final accuracy of experiments with only one generation ($\tau = 200001$) to experiments with multiple generations (of different period lengths). Figure 5 shows this evolution. By performing an ANOVA (Analysis of Variance) test, the accuracy at the end of the experiment when inter-generation transmission occurs at periods of length 10001 and 20001 is different from when it does not occur at all. Post-hoc, tuckey-hsd (honestly significant difference) test shows, that the difference is significant ($p \ll 0.01$) for both 10001 and 20001 periods compared to a single generation (200001). Thus, Hypothesis H_1 is accepted with $\tau \geq 10001$ when the period is long enough.

Inter-generation transmission needs long interaction periods. It can be observed, in the early iterations of Figure 5, that each generation improves its accuracy over the previous one. In particular, the accuracy obtained at 2τ is strictly superior to the accuracy at τ . This confirms that agents with vertical transmission are able to reach a higher accuracy than horizontal transmission alone. However, when the interaction period is not long enough, agents do not have time to spread relevant knowledge widely. Hence, vertical transmission suffers from the low accuracy of the transmitted knowledge and the short period only allows to recover from this. This explains why, when the period length is 5001, the accuracy does not improve.

Short interaction periods can be compensated for by larger transmission percentage. Table 2 shows the final accuracy of agents per period length and transmission percentage. When the

| Meaning | Variable | Experiment 1 | Experiment 2 |
|-----------------------------|------------|---|-----------------------------|
| Number of Iterations | n | 200000 | |
| Size of the Population | X | 40 | |
| Period Length | τ | 5001, 10001, 20001, 200001 | 5001, 10001, 20001 |
| Transmission Percentage | r | 0 = <i>random</i> , 20, 40, 60, 80, 100 | 40 |
| Parent Selection | s | <i>random</i> | <i>income, random, best</i> |
| Number of Parents | v | 2 | |
| Adaptation Operator | op | <i>split</i> | <i>split, none</i> |
| Rest. Inter. Reduction Rate | ϵ | 0.01 | |

Table 1: Independent variables and experiment values.

period length is 5001, the accuracy gets higher as the transmission percentage gets larger. This is because the shorter the period, the smaller the transmission achieved by interaction. If the initial training sample is small, then agents do not receive enough knowledge from their parents to perpetuate what has been gained during the previous period.

| | | Transmission percentage (r) | | | | | |
|--------------------------|--------|---------------------------------|-------------|-------------|-------------|-------------|-------------|
| | | 0 (random) | 0.2 | 0.4 | 0.6 | 0.8 | 1.0 |
| period length (τ) | 5001 | 0.63 | 0.63 | 0.66 | 0.72 | 0.73 | 0.82 |
| | 10001 | 0.91 | 0.91 | 0.91 | 0.86 | 0.90 | 0.82 |
| | 20001 | 0.95 | 0.91 | 0.89 | 0.94 | 0.91 | 0.86 |
| | 200001 | 0.86 | 0.86 | 0.84 | 0.81 | 0.83 | 0.81 |

Table 2: Final average accuracy grouped by experiment parameter values. In bold, the highest values of the column.

Complete transmission percentage harms variation in long interaction periods. In contrast, when the period is long (10001 or 20001), the transmission percentage only affects accuracy when it is complete ($r = 1.0$). Then accuracy is lower. This might be because faithful transmission reduces variation. As long as there is a small variation, agents improve. When the transmission percentage is small, agents compensate for it by transmitting through interaction. **Transmission percentage and period length interact.** Figure 6 compares the accuracy of agents with ($r \neq 0$) and without ($r = 0$) initial vertical transmission, under different period lengths ($\tau = 5001$ and $\tau = 20001$). When the period is short ($\tau = 5001$), a higher transmission percentage ($r = .8$) yields better results than a lower transmission percentage ($r = 0$ and $r = .2$ provide very close final results). On the contrary, with a long period ($\tau = 20001$), the best results are obtained without initial vertical transmission ($r = 0$), those with initial vertical transmission being very close to each other.

This is explained by the capacity of intra-generation transmission to spread accurate knowledge to the whole population. This knowledge has a chance to be transmitted even with low r and even in absence of initial vertical transmission ($r = 0$) because it can be transmitted from parents through interaction. In this case, a low r provides the variation allowing to further increase accuracy. On the contrary, if there is not enough intra-generation transmission (short τ), the perpetuation of knowledge benefits from a more faithful initial vertical transmission. This shows the delicate balance to be found between r and τ to ensure knowledge improvement.





Figure 6: Average accuracy with (red and brown) or without (blue) initial vertical transmission (blue).

6.2 Intra-Generation Transmission can Compensate for the Lack of Selection

To test Hypothesis H_2 , we first show that the selection of parents without the intra-generation transmission does actually improve knowledge accuracy. Then, we show that this effect does not exist when there is intra-generation transmission. Table 3 summarises these results. Results reported below are those with $\tau = 20001$, the same results are obtained with 5001 and 10001 (20001 provides the least favourable figures).

Selection is efficient. Figure 7 shows in dashed lines the evolution of agent accuracy with only the inter-generation transmission comparing maximal (*best*), income-based (*income*) and equiprobable (*random*) selection policies. In the absence of intra-generation

| $op \setminus s$ | <i>random</i> | <i>income</i> | <i>best</i> |
|------------------|---------------|---------------|-------------|
| <i>none</i> | none | medium | medium |
| <i>split</i> | high | high | high |

Table 3: Accuracy improvement in function of selection (s) and horizontal transmission (op). In absence of horizontal transmission ($op = none$), maximal and income-based selection improves final accuracy; with horizontal transmission ($op = split$), all strategies provide a higher improvement.




Figure 7: Average accuracy by parent selection with (plain lines) and without (dashed lines) horizontal transmission.

transmission, having random parents does not improve the accuracy over periods, though parent selection improves it. Maximal selection provides better results than income-based selection. The ANOVA test on the final accuracy of the three selection methods results in a significant difference ($p \ll 0.01$).

Intra-generation transmission compensates for the lack of selection. As it can be observed in Figure 7, the evolution of agent accuracy when there is intra-generation transmission (solid lines), is significantly higher than when it is not present (dashed lines). Furthermore, contrary to having inter-generation transmission only, when the intra-generation transmission is present, the way parents are selected has little impact on the final accuracy. In the presence of intra-generation transmission ($op = split$), ANOVA returns no significant difference between the three parent selection methods ($p = 0.34$). In this case, Table 4 shows that the difference between no selection (*random*) and maximal and income-based selection policies is close to 0, though it is significantly larger without intra-generation transmission. Thus, we accept Hypothesis H_2 : intra-generation transmission compensates for the lack of selection.

| $op \setminus s$ | <i>income</i> | <i>best</i> |
|------------------|--------------------|--------------------|
| <i>none</i> | -0.115 ± 0.035 | -0.165 ± 0.035 |
| <i>split</i> | -0.01 ± 0.04 | 0.005 ± 0.035 |

Table 4: 95% confidence intervals of mean difference between random and the other selection methods with ($op = split$) and without ($op = none$) intra-generation transmission.

7 DISCUSSION

Such results provide a better understanding of transmission parameters which warrant a proper cumulative cultural evolution.

We have shown that the combination of vertical and horizontal transmission of knowledge over generations of agents improves knowledge accuracy. This confirms previous results [2, 14] under broader conditions: initial knowledge is not necessarily correct, no drastic selection of teachers is applied, no artificial noise is introduced to boost variation. We have also shown that a very important factor in the cumulative improvement of knowledge is the population life span.

Thus this paper can be read as a confirmation, using a different framework and relaxed constraints, of [2]’s results about the efficiency of vertical transmission for improving agent knowledge. It can also be seen as a rebuttal that strong parent selection is an important factor for observing such effects.

In fact, the wider a culture is shared in a population, the less important the selection. The obtained results show that spreading quality knowledge requires time. If agents have a short life span and no selection, then knowledge will not improve because the fittest one will have little chance to be passed to the next generation. But if they have enough time to spread accurate knowledge, then it will improve over generations without parent selection.

It can be questioned whether evolution without selection is still evolution. This is the specificity of cultural evolution that, to the selection of individuals by the environment, is added the selection of culture by these individuals, occurring during horizontal transmission. [2] showed that (1) the intra-generation transmission can introduce variation in culture and (2) its selection occurs in the inter-generation transmission. Contrary to that, we showed how (1) inter-generation transmission can be the one introducing variations (which allow agents to improve further their accuracy as shown in Section 6.1) and (2) intra-generation transmission can select the knowledge that spreads in the agent population (Section 6.2). Contrary to genes, even if parents do not provide the best cultural assets, children are able to acquire them from peers or other sources. These results show the robustness of cultural evolution in which the two transmission modes can balance each other.

The social import of such results is that it is not necessary to have a drastic selection of agents for the society’s culture to improve over generations. The good news is that if we are too ignorant, there is hope that our children can catch good knowledge. Of course, there should be a minimal transmission of what is improved for the evolution to be cumulative. But this is also ensured when efficient culture is widely spread, as culture should be.

DATA AVAILABILITY

All experiments were performed in the *Lazy lavender* software environment [1]. Settings, output and data analysis notebooks are made available at [8] and [9].

ACKNOWLEDGMENTS

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

REFERENCES

- [1] 2020. *Lazy lavender*. <https://gitlab.inria.fr/moex/lazylav>.
- [2] Alberto Acerbi and Domenico Parisi. 2006. Cultural Transmission Between and Within Generations. *Journal of Artificial Societies and Social Simulation* 9, 1 (2006), 1–16.
- [3] Michael Anslow and Michael Rovatsos. 2015. Aligning experientially grounded ontologies using language games. In *Graph Structures for Knowledge Representation and Reasoning – 4th International Workshop, GKR (Lecture Notes in Computer Science, Vol. 9501)*. Springer, New York, 15–31.
- [4] David Anzola and Daniel Rodríguez-Cárdenas. 2018. A model of cultural transmission by direct instruction: An exercise on replication and extension. *Cognitive Systems Research* 52 (2018), 450–465.
- [5] Manuel Atencia and Marco Schorlemmer. 2012. An interaction-based approach to semantic alignment. *Journal of Web Semantics* 12 (2012), 131–147.
- [6] Elhanan Borenstein and Eytan Ruppin. 2003. Enhancing Autonomous Agents Evolution with Learning by Imitation. *Artificial Intelligence and Simulation of Behaviour Journal* 1, 4 (2003), 335–347.
- [7] James Martin Borg, Alastair Channon, and Charles Day. 2011. Discovering and maintaining behaviours inaccessible to incremental genetic evolution through transcription errors and cultural transmission. In *European Conference on the Synthesis and Simulation of Living Systems (Advances in Artificial Life, ECAL 2011)*. MIT Press, Cambridge, 101–108.
- [8] Yasser Bourahla. 2021. 20210601-DOTG: Knowledge is transmitted between agents and across generations. <https://doi.org/10.5281/zenodo.5931033> <https://sake.re/20210601-DOTG>.
- [9] Yasser Bourahla. 2021. 20210927-DOTG: Knowledge is transmitted between agents and across generations with different parent selection methods. <https://doi.org/10.5281/zenodo.5929553> <https://sake.re/20210927-DOTG>.
- [10] Yasser Bourahla, Manuel Atencia, and Jérôme Euzenat. 2021. Knowledge Improvement and Diversity under Interaction-Driven Adaptation of Learned Ontologies. In *Proceedings of the 20th International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS '21)*. International Foundation for Autonomous Agents and Multiagent Systems, Richland, SC, 242–250.
- [11] Robert Boyd and Peter Richerson. 1985. *Culture and the evolutionary process* (1 ed.). The University of Chicago Press, Chicago.
- [12] Aleksander Byrski, Rafal Drezewski, Leszek Siwik, and Marek Kisiel-Dorohinicki. 2015. Evolutionary multi-agent systems. *The Knowledge Engineering Review* 30, 2 (2015), 171–186.
- [13] Luigi Luca Cavalli-Sforza and Marcus William Feldman. 1981. *Cultural transmission and evolution: a quantitative approach* (1 ed.). Princeton University Press, New Jersey.
- [14] Daniele Denaro and Domenico Parisi. 1996. Cultural evolution in a population of neural networks. In *Proceedings of the 8th Italian Workshop on Neural Nets*. Springer London, London, 100–111.
- [15] Jérôme Euzenat. 2017. Interaction-based ontology alignment repair with expansion and relaxation. In *Proceedings of the Twenty-Sixth International Joint Conference on Artificial Intelligence (IJCAI'17)*. AAAI Press, California, 185–191.
- [16] Scott Gerard and Munindar Singh. 2013. Evolving Protocols and Agents in Multiagent Systems. In *Proceedings of the 2013 International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS '13)*. International Foundation for Autonomous Agents and Multiagent Systems, Richland, SC, 997–1004.
- [17] Ben Jolley, James Martin Borg, and Alastair Channon. 2016. Analysis of social learning strategies when discovering and maintaining behaviours inaccessible to incremental genetic evolution. In *International Conference on Simulation of Adaptive Behavior (Lecture Notes in Computer Science, Vol. 9825)*. Springer, New York, 293–304.
- [18] Rachel Kendal, Neeltje Boogert, Luke Rendell, Kevin Laland, Mike Webster, and Patricia Jones. 2018. Social Learning Strategies: Bridge-Building between Fields. *Trends in Cognitive Sciences* 22 (2018), 651–665.
- [19] Stefano Nolfi. 2005. Emergence of communication in embodied agents: Co-adapting communicative and non-communicative behaviours. *Connection Science* 17, 3–4 (2005), 231–248.
- [20] Luke Rendell, Robert Boyd, Daniel Cownden, Magnus Enquist, Kimmo Eriksson, Marcus Feldman, Laurel Fogarty, Stefano Ghirlanda, Timothy Lillicrap, and Kevin Laland. 2010. Why Copy Others? Insights from the Social Learning Strategies Tournament. *Science* 328 (2010), 208–213.
- [21] Robert Reynolds. 1994. An Introduction to Cultural Algorithms. In *Proceedings of the Third Annual Conference on Evolutionary Programming*. World Scientific Publishing, New Jersey, 131–139.
- [22] Peter Richerson and Robert Boyd. 2005. *Not By Genes Alone: How Culture Transformed Human Evolution* (1 ed.). The University of Chicago Press, Chicago.
- [23] Sandip Sen and Stéphane Airiau. 2007. Emergence of Norms through Social Learning. In *Proceedings of the Twentieth International Joint Conference on Artificial Intelligence (IJCAI'07)*. Morgan Kaufmann Publishers Inc., San Francisco, 1507–1512.
- [24] Luc Steels. 2012. *Experiments in cultural language evolution* (1 ed.). John Benjamins Publishing Company, Amsterdam.
- [25] Jurriaan van Diggelen, Robbert-Jan Beun, Frank Dignum, Rogier van Eijk, and John-Jules Meyer. 2006. ANEMONE: An effective minimal ontology negotiation environment. In *5th International Joint Conference on Autonomous Agents and Multi-Agent Systems (AAMAS '06)*. Association for Computing Machinery, New York, 899–906.