

Synchrony and Perception in Robotic Imitation across Embodiments

Aris Alissandrakis, Chrystopher L. Nehaniv, and Kerstin Dautenhahn

Adaptive Systems Research Group
Faculty of Engineering & Information Sciences
University of Hertfordshire
Hatfield Herts AL10 9AB
United Kingdom

<http://homepages.feis.herts.ac.uk/~nehaniv/ASRG/>

email: {A.Alissandrakis,C.L.Nehaniv,K.Dautenhahn@herts.ac.uk}@herts.ac.uk

Keywords: social learning, imitation, synchronization, correspondence problem, loose perceptual matching, cultural transmission

Abstract Social robotics opens up the possibility of **individualized social intelligence** in member robots of a community, and allows us to harness not only individual learning by the individual robot, but also the acquisition of new skills by observing other members of the community (robot, human, or virtual).

We describe **ALICE (Action Learning for Imitation via Correspondences between Embodiments)**, an implemented generic mechanism for solving the **correspondence problem** between differently embodied robots. ALICE enables a robotic agent to learn a behavioral repertoire suitable to performing a task by observing a model agent, possibly having a different type of body, joints, different number of degrees of freedom, etc. Previously we demonstrated that the character of imitation achieved will depend on the granularity of subgoal matching, and on the metrics used to evaluate success.

In this work, we implement ALICE for simple robotic arm agents in simulation using various metrics for evaluating success according to actions, states, or effects or weighted combinations. We examine the roles of synchronization, looseness of perceptual match, and of proprioceptive matching by a series of experiments. As a complement to the social developmental aspects suggested by developmental psychology, our results show that **synchronization** and **loose perceptual matching** also allow for faster acquisition of behavioral competencies at low error rates.

We also discuss the use of social learning mechanisms like ALICE for transmission of skills between robots, and give the first example of transmission of a skill through a chain of robots, despite differences in embodiment of agents involved. This simple example demonstrates that by using social learning and imitation, **cultural transmission** is possible among robots, even heterogeneous groups of robots.

1 Introduction: Acquiring Skills via Social Learning by Imitation

Imitation is an important means to acquire new competencies in a social context. Human cultures use imitation in a variety of ways. This may include learning of new motor skills, such as learning how to tie shoe laces, how to play tennis, or learning a dance. Imitation may also be involved when we learn a language, learn tool use, or learn how to behave in particular social contexts, such as learning dress codes or greetings. In animal sciences the question of which animal species imitate is under hot debate. Imitation is different from other mechanisms of social learning, e.g. when animals learn mainly due to the *presence* of conspecifics, or when conspecifics draw attention to certain *features* of the environment that are involved in the behavior. In the latter case biologists would call this form of social learning *local* or *stimulus enhancement*. Other types of social learning (also called observational learning) are variously taxonomized and include *goal emulation*, *social learning of affordances*, etc. Some authors require that ‘do as I do’ (the common sense interpretation of imitation or ‘apeing’) should only be called imitation if it involves learning a *novel* behavior [27]. See [30, 8] for an in depth discussion of different types of social learning, definitions of imitation, and the main research questions involved.

1.1 The Correspondence Problem

Roboticians have been interested in imitation since the early 1990’s [14, 7]. In robotics, learning by imitation involving a following strategy (the imitator following the model around in an environment) has been used widely [11, 4]. The acquisition of motor skills is another very active area of research, involving learning control policies to match via-points in the motion trajectory of a model [25] or, e.g. in work where a virtual humanoid agent learns to dance the Macarena [15]. While most previous work has engineered ad hoc mechanisms to achieve imitation, general mechanisms for solving the correspondence problem, i.e. *how an imitating agent can imitate a model with possibly dissimilar embodiment*, are our focus in this paper. An informal definition of the **correspondence problem** [19, 20, 21, 22] is as follows:

Given an observed behavior of the model, which, from a given starting state, leads the model through a sequence (or hierarchy) of

subgoals (in states, action, and/or effects, while possibly responding to sensory stimuli and external events), find and execute a sequence of actions using one's own (possibly dissimilar) embodiment, which, from a corresponding starting state, leads through corresponding subgoals (in corresponding states, actions, and/or effects, while possibly responding to corresponding events). [22, p. 49]¹

The notion of ‘corresponding’ states, actions and effects here is formalized by a choice of metrics, and the choice of subgoals to be matched defines the granularity and program-structure of the candidate matching behavior [19]. Our previous work studying the correspondence problem included a **chessworld** scenario where the results showed that the newly developed mechanism called **Action Learning for Imitation via Correspondences between Embodiments (ALICE)** could solve the correspondence problem for agents with dissimilar embodiment [1, 3]. This work also demonstrated that the metric and the level of subgoal granularity can each dramatically affect the character of imitative behavior that is generated, and that one metric is not in general universally “better” than another, but various choices of metrics contribute to which aspects of a behavior are to be imitated [3]. The ALICE mechanism is related to statistical string-parsing models of social learning from ethology [6] and associative sequence learning theory from psychology [13]. In our current work we use a different test-bed, namely a scenario where robotic arms imitate each other.

1.2 Immediate Imitation: Synchrony and Social Intelligence

Imitation can have another important role in robotics, besides skill acquisition. Developmental psychologists have revealed the crucial role of imitation in how humans become social beings, e.g. how they identify others as persons, and how they recognize individuals (cf. work on neonatal imitation, e.g. [16] and others).

Synchronization of behavior plays a fundamental role in child-caretaker interactions, as becomes evident in developmental studies with babies and infants [26, 17]. The contingencies and dynamic aspects of interaction and communication are stepping stones in the social development of infants, and they are prerequisites of immediate imitation. It has been argued by Nadel [17] that immediate imitation creates intersubjectivity and is the first step by which infants make ‘contact’ to other human beings. Individualized social intelligence in humans and social animals may rely on a common core of these and related mechanisms (cf. [28, 23]).

This and other evidence from the study of animal social complexity (e.g. [12]) suggest that synchronization and immediate imitation might also be key ingredients for the development of individualized social intelligence in robots [9].

We show below that use of synchronization of behavior in a robotic test-bed can also dramatically speed up solution of the correspondence problem. Our results also show that the use of loose perceptual matching speeds up solution of the correspondence problem.

¹See that reference and [19, 20, 21] for the formal statement of the correspondence problem relating to the use of different error metrics, and for other applications.

2 The ALICE Mechanism in a Robotic Arm Test-bed

In order to study the correspondence problem we developed the ALICE (Action Learning via Imitation between Corresponding Embodiments) generic imitation mechanism. This mechanism is intended as a controller for the actions of an imitating agent, making use of a **correspondence library**. The **keys** to the entries of this library consist of some combination of actions/states/effects of the model agent, and proprioceptive information concerning the imitator's own state. Perceptions of the model and possibly proprioceptions are converted to the form of a key for the imitator to look-up corresponding actions for that key in the imitator's correspondence library. These actions are the actions that the imitating agent should perform in order to achieve a matching behavior, according to an evaluation metric. As new actions corresponding to the perceptual keys are learned they are added to the imitator's library, which is initially empty.

These possible actions can be generated using any kind of generating algorithm to propose actions (e.g. inverse kinematics). In our work we simply use a random generating algorithm, since we are not concerned about the precise nature of the generating mechanism here. Proposed actions are then evaluated according to a metric and will either update an existing entry with more fitting solutions, or create a new entry of their own if the current state/action/effect or proprioceptive aspects comprising the key have not been observed previously up to that point.

The type of the resulting imitating behavior will depend on the metric used, whether the imitator will try to match the perceived model actions/states/effects or some combination of them. For more details on ALICE see [3] and below.

2.1 The Robotic Arm Test-Bed

The current test-bed was created as a simple, yet ‘rich enough’ environment that would allow for several interacting models and imitator agents, having dissimilar embodiments [2]. Each agent (Fig. 1) occupies a two-dimensional workspace and is embodied as a robotic arm that can have any number of rotary joints, each of varying length. The agent's embodiment can thus be described by a vector $L = [\ell_1, \ell_2, \ell_3 \dots, \ell_n]$, where ℓ_i is the length of the i^{th} segment of the arm.

There are no complex physics in the workspace and the movement of the arms is simulated using simple forward kinematics but without collision detection or any static constraints (in other words, the arms can bend into each other). Our intention is to demonstrate the features of the imitative mechanism and not to build a faithful simulator.

An **action** of a given agent is defined as a vector describing the change of angle for each of the joints, $A = [\alpha_1, \alpha_2, \alpha_3, \dots, \alpha_n]$, where n is the number of its joints. These angles are relative to the previous state of the arm and can only have three possible values, $+10^\circ$ (anti-clockwise), 0° or -10° (clockwise).

A **state** of an agent is defined as the absolute angle for each of the joints, $S = [\sigma_1, \sigma_2, \sigma_3, \dots, \sigma_n]$, where n is the number of its joints. We can distinguish between the **previous state** and the **current state** (the state of the arm after the current action was executed). As a result of the possible actions, the absolute angle at




Figure 1: **Example Embodiment.** A two-joint robotic arm, with segments of length ℓ_1 and ℓ_2 , moving from state S_0 (arm completely outstretched along the horizontal axis) to state S to state S' to state S'' , as it sequentially performs actions A , A' , and A'' . Note that the effects are not shown in this figure.

each joint can be anywhere in the range of 0° to 360° (modulo 360°) but only in multiples of 10° .

The end tip of the arm can leave a trail of paint as it moves along the workspace. The **effect** is defined as a directed straight line segment connecting the end tip of the previous and the current states of the arm (approximating a paint trail). The effect is internally implemented as a vector of displacement $E = (x_c - x_p, y_c - y_p)$, where (x_p, y_p) and (x_c, y_c) are the end tip coordinates of the arm for the previous and current state respectively.

2.2 Metrics

The imitating agents can perceive the actions, states and effects of the model agents, and also their own actions, states and effects, and therefore we define several metrics to evaluate the similarity between them. Metrics are scaled to take values from 0 to 100%. Ideally the metric value should be zero, indicating a perfect match.

State metric

The state metric calculates the averaged distance between the various joints of an agent (posed in a particular state) and the corresponding joints of another agent² (posed in a different state) as if they were occupying the same workspace. Ideally this distance should be zero when the arms take corresponding poses, but this may not be possible due to embodiment differences. Using forward kinematics, the coordinates of the ends for each

²The state metric can be used not only between different agents, but also to evaluate the similarity between two states of the same agent. This is true for the action and the effect metric as well.

joint are found.

$$x_i = \sum_{j=1}^{i-1} x_j + \ell_i \cos\left(\sum_{j=1}^i \sigma_j\right)$$

$$y_i = \sum_{j=1}^{i-1} y_j + \ell_i \sin\left(\sum_{j=1}^i \sigma_j\right)$$

If both agents have the same number of joints the correspondence between them is straightforward; the Euclidean distance for each pair is calculated, the distances are then all summed and divided by the number of joints to give the metric value.

$$d_i = \sqrt{(x_i^{\text{model}} - x_i^{\text{imitator}})^2 + (y_i^{\text{model}} - y_i^{\text{imitator}})^2}$$

$$\mu = \frac{1}{n} \sum_{i=1}^n d_i$$

If the agents have a different number of joints, then some of the joints of the agent with more joints are ignored. To find which joint corresponds with which, the *ratio* of the larger number of joints over the smaller number of joints is calculated, and if not integer, is rounded to the nearest one. In computing the metric, the i^{th} joint of the agent with the smaller number of joints, will correspond to the $(\text{ratio} \times i)^{\text{th}}$ joint of the agent with the larger number of joints. For example if one of the agents has twice the number of joints, only every second joint will be considered.

Action metric

For the action metric, the same algorithm as the one described above for the state metric is used, but considering the action vectors instead of the state vectors.

The value in the case of the state metric represents an absolute positional error; for the action metric, it represents the relative error between the change of the state angles caused by the compared actions.

Effect metric

The **effect metric** is defined as Euclidean length $\mu = \sqrt{(x_1 - x_2)^2 + (y_1 - y_2)^2}$ of the vector difference between two effects (x_1, y_1) and (x_2, y_2) .

2.3 How versus What to Imitate

The choice of the metric determines, in part, *what* will be imitated, whereas solving the correspondence problem concerns *how* to imitate [8]. In general, aspects of state, action and effect as well as level of granularity of matching could all play roles in the choice of metric for solving the problem of how to imitate [21, 3, 18], and the metrics described above were chosen to measure these aspects (at fine granularity) in our test-bed. Ongoing research is addressing the complementary problem of how to extract agent subgoals and generate suitable metrics automatically [21, 18, 5].

2.4 Growth of the Correspondence Library

The model's behavioral pattern may naturally be broken down as a sequence of actions that move the robotic arm of the agent from the previous state to the current state, while leaving behind a trail of paint as the effect.

The nature of the experimental test-bed with the fixed-base rotary robotic arms favours cyclical looping effects and the model patterns used in the experiments were designed as such (Fig. 2). Each complete behavioral pattern that returns the arm to its initial state observed by the imitator is called an **exposure**, and the imitator is exposed to repeated instances of the same behavioral pattern. At the beginning of each new exposure it is possible to reset the imitating agent to the initial state. This resetting is called **synchronization** in our experiments.

The correspondence library is initially empty. At each time step, i.e. for each action of the model, the imitator agent may be able to perceive the model's action, previous and current state and also the effect. The agent might perceive any of those aspects or a combination depending on the metric it is using.

The first time a percept occurs, a new entry is created in the correspondence library with that percept as its indexing key.

W. S. ti o m o i s, l s d l k b - A E s c o d i m e u s e. E s o - s r q u i r i s e. N E. b




Figure 3: **An example of social transmission.** The original model `model0` ($L = [20, 20, 20]$) is shown to the left. In the middle, a two-joint imitator `imitator0` ($L = [30, 30]$) acts also as a model for `imitator1` on the right ($L = [20, 20, 20]$). Due to the different embodiment of the agent `imitator0`, the replication of the model pattern is similar, but not exact. `imitator1` has the same embodiment as the original model `model0` and, although indirectly transmitted, the resulting pattern is closer to that of the original model than is the behavior of the intermediate agent `imitator0` used as a model by this second imitator. Both imitators use the action metric.

The metric used by both imitators was the action metric. In order to speed up the learning and the build-up of the correspondence library, the generating mechanism was creating five random candidate actions to choose from using the action metric.

After transmission through the intermediary, the behavioral pattern that has been acquired by the second imitator in (Fig. 3, right) is quite similar to the original despite differences in embodiment in the chain of transmission.

This first example shows transmission of a behavioral pattern through a chain of robotic agents, despite differences in embodiment of agents involved. This simple example serves as proof of the concept that by using social learning and imitation, rudimentary cultural transmission with variability is possible among robots, even heterogeneous ones.

4 Experiments and Results

4.1 Loose Perceptual Matching

When the ALICE mechanism looks up a perceptual key in the correspondence library to find the relevant entry to the currently perceived model actions, states and effects, it is possible not to require an exact match of the entry keys, but one that is close enough, depending on a threshold. We conducted ten experimental runs under the same conditions. Each run consisted of ten exposures to the model behavior for two imitating agents trying to imitate a model agent, one of them requiring an exact match for the trigger keys and the other one accepting a 10% margin of looseness. The metric used by both agents was a weighted half-half combination of the action and the state metrics. Both agents synchronized after each exposure to the model. Each of the ten runs lasted eleven exposures and the average metric value (that can be seen as the error) for each exposure was logged. The value of error metric for the agent using loose perceptual matching is plotted in Fig. 4 (top panel), and that for the agent using exact matching in Fig. 4 (middle panel).

The ratio of the average error of the imitating agent that uses loose matching over the average error of the imitating agent that requires an exact match can be seen

in Fig. 4 (bottom panel), constantly decreasing and below 1. This indicates that the numerator (error with loose matching) is minimized faster than the denominator (error with exact matching) and is explained by the fact that there are fewer and more generic entries in the correspondence library of the imitator with the loose matching, resulting in a faster improvement of performance.

When exact matching is used, a significantly longer learning period is required, and therefore loose matching to within 10% was used in the the rest of the particular experiments reported here.

4.2 Synchronization

Inspired by the biological and psychological importance of synchronization (sec. 1.2), we implemented synchronization in our test-bed as follows and did a series of experiments to assess its efficacy. At the end of each exposure of the imitating agent to the model, it is possible to reset the imitator arm to the same initial position, thus synchronizing the imitation attempt to the model's behavior. We conducted ten experimental runs, each with two imitating agents trying to imitate a model agent, one of them synchronizing to an outstretched initial state after each exposure to match the initial state of the model, and the other not. The metric used by both agents was a weighted half-half combination of the action and the state metrics. Both agents used a loose entry key matching of 10%. Each run lasted eleven exposures and the average metric value for each exposure was logged.

The ratio of the average error of the imitating agent that uses synchronization over the average error of the imitating agent that does not synchronize back to the start position at the end of trying to imitate an observed behavioral pattern can be seen in Fig. 5 (bottom panel) constantly decreasing and below 1. This indicates that the numerator (error when using synchronization) is minimized faster than the denominator (error when not using synchronization) and indicates that it is difficult for an imitating agent that does not synchronize to reach again states relevant to the model pattern if the initial imitation attempts are not successful. Soon after the start of the first exposure to the behavior pattern, the imitator not using synchronization becomes 'lost' due to cumulative errors that are not corrected by resetting (Fig. 5 (top panel)), while the imitator using synchronization shows steady improvement (Fig. 5 (middle panel)). As a result the non-synchronizing agent might require a far greater number of exposures to return (via the random walk of the generating mechanism) back 'on track' and successfully imitate.

4.3 Proprioceptive Matching

Proprioception is always used by the ALICE mechanism whenever perceptual keys include a state or effect component, since the imitator's own state is taken into account when calculating the metric values for the different possible actions; but is not used if the perceptual key consists of only the action component.

The correspondence library entry keys may contain both perceptual (the model's action, state and effect) and proprioceptive (the imitator's own state at the time of the observation) data. It is possible to exclude this proprioceptive component from the keys and to trigger the keys based only on the perception. We conducted ten experimental runs, each with two imitating agents

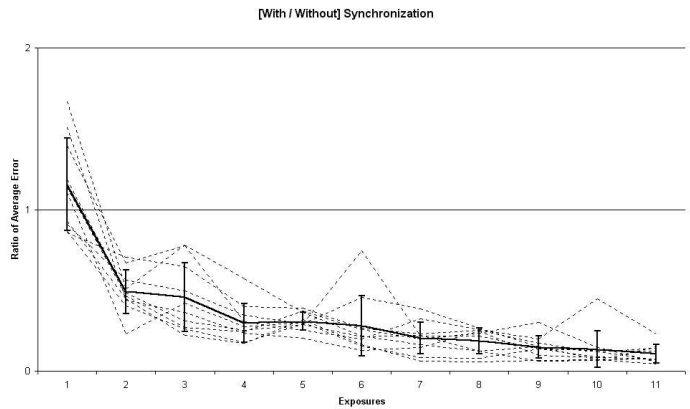
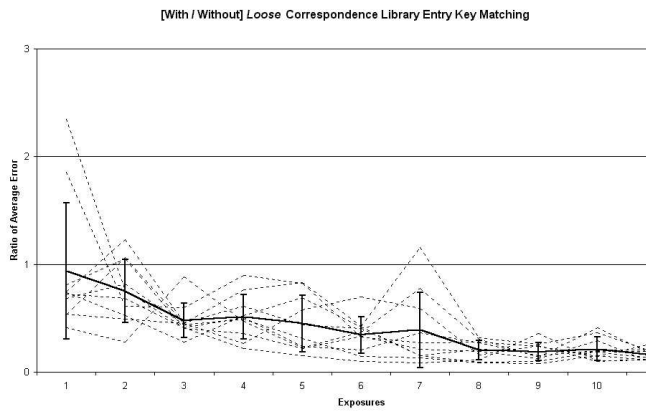
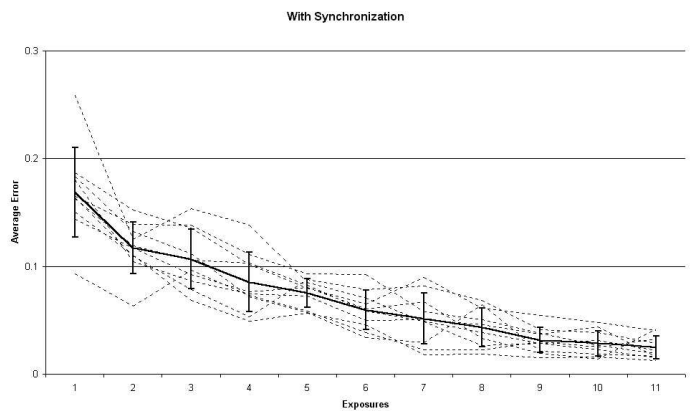
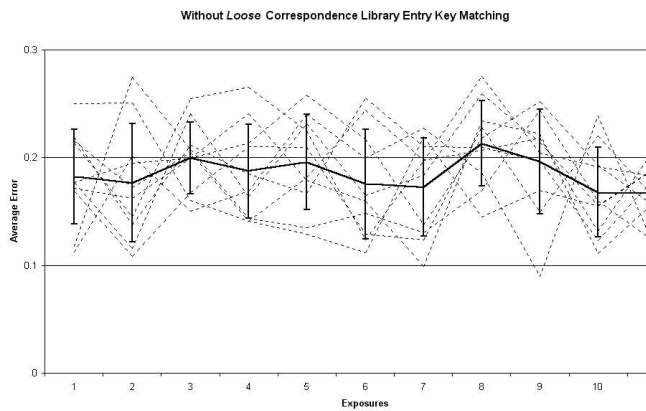
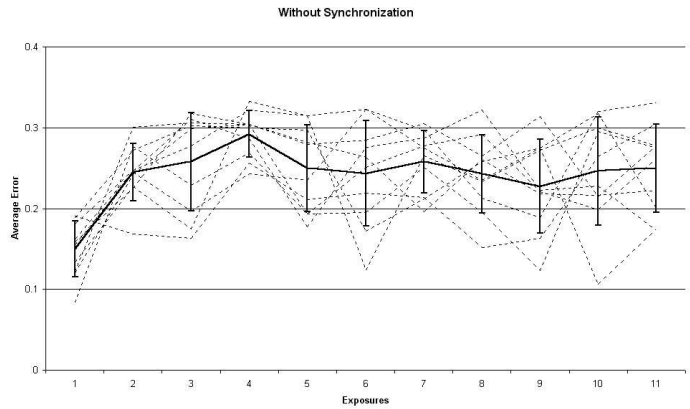
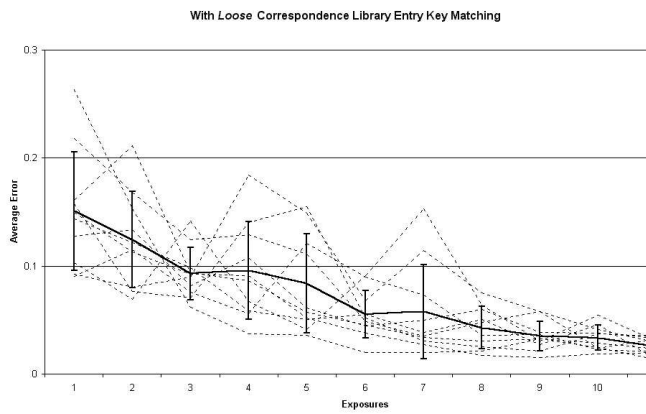


Figure 4: Loose Perceptual Matching Experiments.

The error metric value of robotic agents over 11 exposures when using loose perceptual matching (top panel) and when using exact matching (middle panel). The ratio of the average error per exposure of the imitating agent using loose perceptual matching over the average error of the imitating agent that does not (bottom panel) indicates a comparative many-fold reduction of error with use of loose perceptual matching. In each panel, the thicker line shows average values of all the ten experiments, with the bars indicating the standard deviation. Dotted lines indicate values during individual runs.

Figure 5: Synchronization Experiments.

The error metric value of robotic agents over 11 exposures when not using synchronization (top panel) and when using synchronization (middle panel). The ratio of the average error per exposure of the imitating agent using synchronization over the average error of the imitating agent that does not (bottom panel) indicates a dramatic reduction of error with synchronization. In each panel, the thicker line shows average values of all the ten experiments, with the bars indicating the standard deviation. Dotted lines indicate values during individual runs.

trying to imitate a model agent, one of them using proprioceptive matching, the other not. The metric used by both agents was a weighted half-half combination of the action and the state metrics. Both agents used a loose entry key matching of 10% (for the perception component) and the generating mechanism was creating five random actions to choose from. Each run lasted eleven exposures and the average metric value for each exposure was logged.

The ratio of the average error per exposure of the imitating agent that does not use proprioceptive matching over the average error of the imitating agent that does can be seen in Fig. 6, constantly decreasing and below 1. This indicates that the numerator (error when not using proprioceptive matching) is minimized faster than the denominator (error when using proprioceptive matching). Similar to the experimental results for loose matching described above, ignoring the proprioceptive matching for keys reduces the number of entries in the library, thus allowing them to update more frequently, resulting in a faster improvement of performance.⁶

5 Conclusions and Outlook

The results of our experiments using ALICE in differently embodied robotic arm agents show that (1) cultural transmission of behavioral patterns is possible in a heterogeneous community of robots, (2) loose perceptual matching increases the rate of solving the correspondence problem significantly, (3) synchronization dramatically increases the rate of solving the correspondence problem, (4) utilizing proprioceptive matching for keys does not, at least for early stages of learning, aid in rate of the solution of this problem within our experiments (although it certainly did not prevent its solution).

The potential for cultural transmission of skills through a heterogeneous population of robots using our methods might be applied to the acquisition and transmission of skills in more complex populations of robots, involved in carrying out useful tasks, e.g. on the shop-floor of a factory, with new robots coming and going, acquiring behaviors by observation without having to be explicitly programmed and without humans having to develop different control programs for different types of robots that need to perform the same task. Instead, the robots would autonomously create their how programs using social learning and a correspondence library.

This together with previous work [3] using a chess-world test-bed serves to establish the generalizability of the ALICE framework. Scalability in different settings depends on particularities of the embodiments, the sophistication of the generating mechanism used (here, only random actions were needed) to propose candidate matching actions or action sequences, processing speed, and optimization problems for the specific platforms. Future work in solving the correspondence problem will also involve applications to fault-tolerance and self-repair by imitating agents, as well as new methods for subgoal extraction and the automatic generation of metrics.

⁶Utilizing a proprioceptive matching component for keys in the correspondence library increases the number of keys K , on which the search space depends exponentially, by a factor equal to the number of all possible states of the imitator (cf. footnote 3 above).

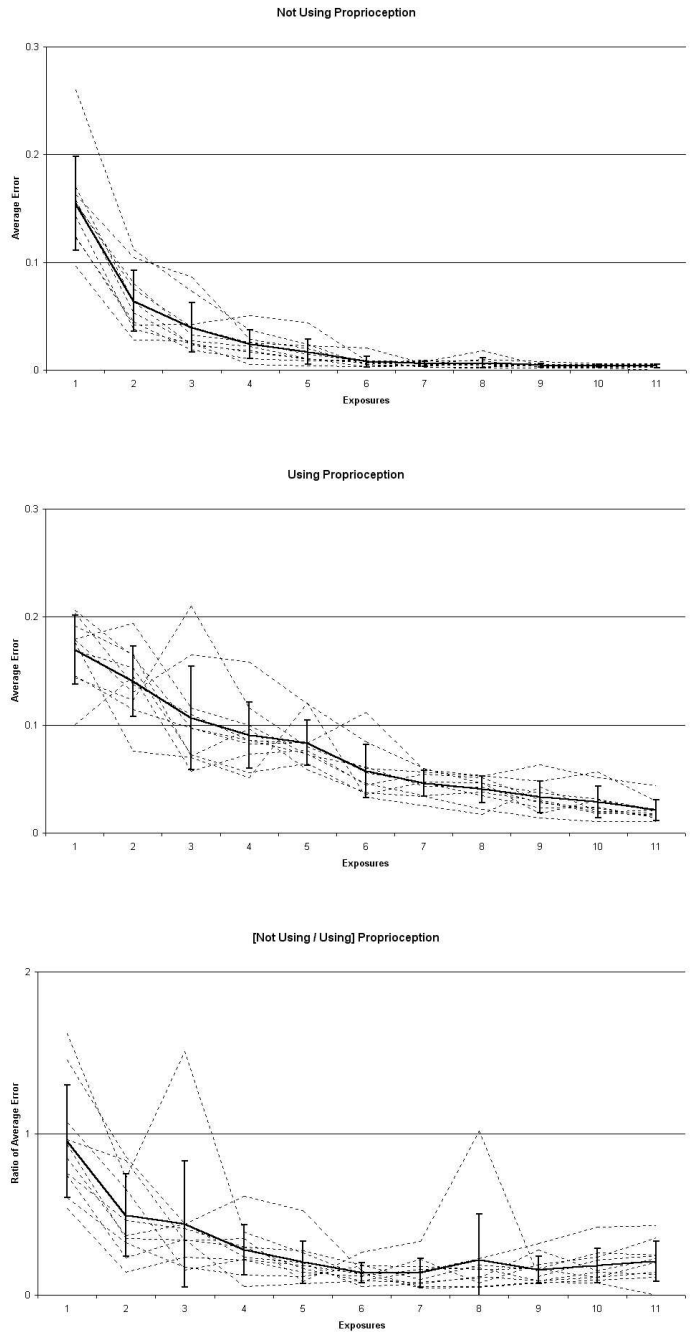


Figure 6: **Proprioceptive Matching Experiments.**

The error metric value of robotic agents over 11 exposures when using proprioceptive matching for keys in the correspondence library (top panel) vs. not using proprioception in this way (middle panel) for 10 runs. The ratio of the average error per exposure of the imitating agent not employing proprioceptive matching over the average error of the imitating agent that does (bottom panel) indicates some comparative reduction of error when not using proprioceptive matching. In each panel, the thicker line shows the average values of all the ten experiments, with the bars indicating the standard deviation. Dotted lines indicate values during individual runs.

References

- [1] A. Alissandrakis, C. L. Nehaniv, and K. Dautenhahn. Through the looking-glass with ALICE - trying to imitate using correspondences. In *Proc. First International Workshop on Epigenetic Robotics: Modeling Cognitive Development in Robotic Systems, September 17-18, 2001, Lund, Sweden*, volume 85, pages 115–122. Lund University Cognitive Studies, 2001.
- [2] A. Alissandrakis, C. L. Nehaniv, and K. Dautenhahn. Do as I do: Correspondences across different robotic embodiments. In *Proc. Fifth German Workshop on Artificial Life*, pages 143–152. IOS Press, 2002.
- [3] A. Alissandrakis, C. L. Nehaniv, and K. Dautenhahn. Imitation with ALICE: Learning to imitate corresponding actions across dissimilar embodiments. *IEEE Trans. Systems, Man & Cybernetics: Part A*, 32(4):482–496, 2002.
- [4] A. Billard and K. Dautenhahn. Grounding communication in autonomous robots: an experimental study. *Robotics and Autonomous Systems*, 24(1-2):71–81, 1998.
- [5] A. Billard and S. Schaal. Computational elements of robot learning by imitation. In *American Mathematical Society Central Meeting, Madison, October 12-13, 2002*, 2002.
- [6] R. W. Byrne. Imitation without intentionality: Using string parsing to copy the organization of behaviour. *Animal Cognition*, 2:63–72, 1999.
- [7] K. Dautenhahn. Trying to imitate – a step towards releasing robots from social isolation. In P. Gaussier and J.-D. Nicoud, editors, *Proc. From Perception to Action Conference, Lausanne, Switzerland*, pages 290–301. IEEE Computer Society Press, 1994.
- [8] K. Dautenhahn and C. L. Nehaniv. An agent-based perspective on imitation. In K. Dautenhahn and C. L. Nehaniv, editors, *Imitation in Animals and Artifacts*. MIT Press, 2002.
- [9] Kerstin Dautenhahn. Getting to know each other – artificial social intelligence for autonomous robots. *Robotics and Autonomous Systems*, 16:333–356, 1995.
- [10] R. Dawkins. *The Selfish Gene*. Oxford, 1976.
- [11] J. Demiris and G. Hayes. Imitative learning mechanisms in robots and humans. In *Proc. of the 5th European Workshop on Learning Robots, Bari, Italy, July 1996*, pages 9–16, 1996.
- [12] L. M. Herman. Vocal, social, and self imitation by bottlenosed dolphins. In K. Dautenhahn and C. L. Nehaniv, editors, *Imitation in Animals and Artifacts*. MIT Press, 2002.
- [13] C. M. Heyes and E. D. Ray. What is the significance of imitation in animals? *Advances in the Study of Behavior*, 29:215–245, 2000.
- [14] Y. Kuniyoshi, H. Inoue, and M. Inaba. Design and implementation of a system that generates assembly programs from visual recognition of human action sequences. In *Proc. IEEE International Workshop on Intelligent Robots and Systems IROS '90*, pages 567–574, 1990.
- [15] M. J. Matarić, V. Zordan, and Z. Mason. Movement control methods for complex, dynamically simulated agents: Adonis dances the macarena. In *Proceedings, Autonomous Agents '98, Minneapolis/St. Paul, May 10-13*, pages 317–324. ACM Press, 1998.
- [16] A. Meltzoff. The human infant as imitative generalist: a 20-year progress report on infant imitation with implications for comparative psychology. In B. G. Galef and C. M. Heyes, editors, *Social Learning in Animals: the Roots of Culture*, pages 347–370. Academic Press, New York, 1996.
- [17] J. Nadel, C. Guerini, A. Peze, and C. Rivet. The evolving nature of imitation as a format of communication. In J. Nadel and G. Butterworth, editors, *Imitation in Infancy*, pages 209–234. Cambridge, 1999.
- [18] C. L. Nehaniv. Nine billion correspondence problems and some methods for solving them. In *Proceedings of the 2nd International Symposium on Imitation in Animals and Artifacts – Aberystwyth, Wales, 7-11 April 2003*. Society for the Study of Artificial Intelligence and Simulation of Behaviour, 2003.
- [19] C. L. Nehaniv and K. Dautenhahn. Mapping between dissimilar bodies: Affordances and the algebraic foundations of imitation. In John Demiris and Andreas Birk, editors, *Proceedings European Workshop on Learning Robots 1998 (EWLR-7), Edinburgh, 20 July 1998*, pages 64–72, 1998.
- [20] C. L. Nehaniv and K. Dautenhahn. Of hummingbirds and helicopters: An algebraic framework for interdisciplinary studies of imitation and its applications. In J. Demiris and A. Birk, editors, *Interdisciplinary Approaches to Robot Learning*, pages 136–161. World Scientific Series in Robotics and Intelligent Systems, 2000.
- [21] C. L. Nehaniv and K. Dautenhahn. Like me? - measures of correspondence and imitation. *Cybernetics and Systems*, 32(1-2):11–51, 2001.
- [22] C. L. Nehaniv and K. Dautenhahn. The correspondence problem. In K. Dautenhahn and C. L. Nehaniv, editors, *Imitation in Animals and Artifacts*. MIT Press, 2002.
- [23] I. M. Pepperberg. Allospecific referential speech acquisition in grey parrots (*Psittacus erithacus*): Evidence for multiple levels of avian vocal imitation. In K. Dautenhahn and C. L. Nehaniv, editors, *Imitation in Animals and Artifacts*. MIT Press, 2002.
- [24] L. Rendell and H. Whitehead. Culture in whales and dolphins. *Behavioral and Brain Sciences*, 24(2):., 2001.
- [25] S. Schaal. Is imitation learning the route to humanoid robots? *Trends in Cognitive Sciences*, 3(6):233–242, 1999.
- [26] C. Trevarthen, T. Kokkinaki, and G. A. Friamenghi, Jr. What infants' imitations communicate: With mothers, with fathers, and with peers. In J. Nadel and G. Butterworth, editors, *Imitation in Infancy*, pages 128–185. Cambridge, 1999.
- [27] E. Visalberghi and D. Frigaszy. Do monkeys ape? - ten years after. In K. Dautenhahn and C. L. Nehaniv, editors, *Imitation in Animals and Artifacts*. MIT Press, 2002.
- [28] A. Whiten and R. W. Byrne. The Machiavellian intelligence hypotheses: editorial. In R. W. Byrne and A. Whiten, editors, *Machiavellian intelligence*, chapter 1. Clarendon Press, 1988.
- [29] A. Whiten, J. Goodall, W.C. McGrew, T. Nishida, V. Reynolds, Y. Sugiyama, C.E.G. Tutin, R.W. Wrangham, and C. Boesch. Cultures in chimpanzees. *Nature*, 399:682–685, 1999.
- [30] T. R. Zentall. Imitation in animals: Evidence, function and mechanisms. *Cybernetics and Systems*, 32(1-2):53–96, 2001.