

Correspondence Mapping Induced State and Action Metrics for Robotic Imitation

Aris Alissandrakis, Chrystopher L. Nehaniv, and Kerstin Dautenhahn

Abstract—This paper addresses the problem of body mapping in robotic imitation where the demonstrator and imitator may not share the same embodiment [degrees of freedom (DOFs), body morphology, constraints, affordances, and so on]. Body mappings are formalized using a unified (linear) approach via correspondence matrices, which allow one to capture partial, mirror symmetric, one-to-one, one-to-many, many-to-one, and many-to-many associations between various DOFs across dissimilar embodiments. We show how metrics for matching state and action aspects of behavior can be mathematically determined by such correspondence mappings, which may serve to guide a robotic imitator. The approach is illustrated and validated in a number of simulated 3-D robotic examples, using agents described by simple kinematic models and different types of correspondence mappings.

Index Terms—Correspondence problem, imitation and social learning, programming by demonstration, state and action metrics.

I. INTRODUCTION

IMITATION is a powerful learning tool when a number of agents interact in a social context. The demonstrator and imitator agents may or may not belong to the same species (a parent teaching a child, a human training an animal) or even be biological and artificial entities (e.g., in human–robot interaction). The latter is a very interesting paradigm explored in computer science and robotics, with researchers influenced by work on biology, ethology, and psychology working to design controllers that would allow their robots to be programmed and learn more easily and efficiently [10], [11], [15], [16], [18], [22], [25].

A fundamental problem when learning *how* to imitate is to create an appropriate (partial) mapping between the actions afforded by particular embodiments to achieve corresponding states and effects by the model and imitator agents (solving a *correspondence problem*) [21]. The solutions to the correspondence problem will depend on the subgoal granularity and the metrics used to evaluate the similarity between actions, states, and/or effects, resulting in qualitatively different imitative behaviors [1], [2]. The related problem of *what* to imitate addresses the choice of metrics and subgoal granularity

Manuscript received December 12, 2005; revised March 20, 2006. The work described here was conducted within the EU Integrated Project COGNIRON and was supported by the European Commission Division FP6-IST Future and Emerging Technologies under Contract FP6-002020. This paper was recommended by Guest Editor Y. Demiris.

The authors are with the School of Computer Science, Adaptive Systems Research Group, University of Hertfordshire, AL10 9AB Hatfield, U.K. (e-mail: a.alissandrakis@herts.ac.uk).

Digital Object Identifier 10.1109/TSMCB.2006.886947

that should be used for imitating, depending on the context. See [9] and [24] for robotic examples and [8], [12], and [14] for ethological and psychological aspects.

Related to the work on solving the correspondence problem for imitation learning in robotics is the ALICE generic imitation framework [1], [2] and the JABBERWOCKY system [3]–[5]. Both are generic approaches addressing multiple demonstrator and target imitator embodiments using different metrics to achieve different types of social learning, matching different behavioral aspects.

In this paper, we introduce a novel generic approach to the correspondence problem, via body mapping for the cases of state and/or action matching. In particular, partial, relative, and mirror matching all arise as special cases of such correspondence mappings. Moreover, an infinite set of metrics (parameterized by correspondence matrices) for imitation performance are induced via such body correspondences.¹ This contributes toward a characterization of types of matching in social learning. Previously, we studied the space of effect metrics [5], while in this paper, the focus turns to *state* and *action* metrics. The approach is illustrated and validated via a number of simulated 3-D and robot examples mapping across dissimilar embodiments and is applicable to robot programming by human demonstration.

II. DIFFERENT BODIES

Different agent bodies can be described as simplified kinematic models, comprising of a rooted acyclic connected graph of segments (see Fig. 1).² Each segment has a base and a tip end, and is described by the quintuple $(i, \ell_i, p_i, \theta_i, \phi_i)$ as follows:

i	index number of the segment;
ℓ_i	segment length;
p_i	index number of the parent segment;
θ_i and ϕ_i	azimuth and polar angles, respectively, for the spherical coordinates $(\ell_i, \theta_i, \phi_i)$ that indicate how the segment is positioned in 3-D space (relative to the end of its parent segment).

NB: In general the range of the angles θ_i and ϕ_i may be constrained within given respective ranges.

The values of θ_i and ϕ_i are *relative* for each segment, but *absolute* angles for segment i , Θ_i , and Φ_i can be obtained

¹That is, to say that a correspondence mapping “induces” a metric means exactly that it mathematically determines the metric.

²Cf. modeling agents as simple open kinematic chains in [6].

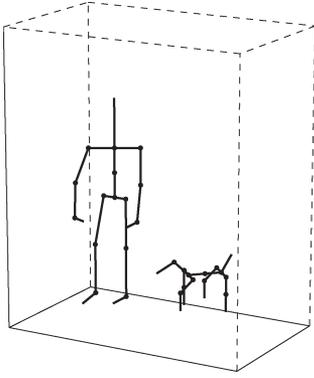


Fig. 1. Some embodiment examples using simple kinematic models. (Left) A human and (right) an AIBO robot.

inductively starting from the segment after the root³

$$\begin{aligned}\Theta_i &= \theta_i + \Theta_{p_i} \\ \Phi_i &= \phi_i + \Phi_{p_i}.\end{aligned}$$

See the Appendix for detailed kinematic equations.

The **state** of such a kinematic model can be defined as the vector containing the values of the degrees of freedom (DOF), i.e., the values of the azimuth and polar angles for each segment in the acyclic graph.

Depending on whether the relative or the absolute angle values are used, for an embodiment with n segments, two different state vectors can be considered

$$\begin{aligned}S_{\text{relative}} &= [\theta_1 \phi_1 \theta_2 \phi_2 \cdots \theta_n \phi_n] \\ S_{\text{absolute}} &= [\Theta_1 \Phi_1 \Theta_2 \Phi_2 \cdots \Theta_n \Phi_n].\end{aligned}$$

Here, complying with the particular kinematic models that will be used, the state vector is composed of the alternating azimuth and polar angles for each of the body segments.⁴ For the rest of this paper, the notation S_j will be used to refer to the state value of the j th DOF of an agent.⁵

An **action** can be defined as the difference between two consecutive state vectors S and S'

$$A = S' - S.$$

Using either the relative or absolute representation of the state vectors for calculating an action vector produces mathematically equivalent results. Note that depending on the embodiment, a change in the relative values of the j th DOF can influence the absolute values of subsequent DOFs (see the

³For mathematical convenience, the root node is treated as a segment of length $\ell_0 = 0$, but θ_0 and ϕ_0 can have nonzero values to orient the entire model. For expository purposes, without loss of generality, in this paper, we ignore the latter possibilities ($\Theta_0 = \theta_0 = 0$, $\Phi_0 = \phi_0 = 0$).

⁴Depending on the particular body representations used, the contents of the vectors (and the ordering of the elements) can of course vary. For example, if Euler angles (ϕ, θ, ψ) were used instead of spherical coordinates, the state vector for an embodiment with n segments could be defined as $S = [\phi_1 \theta_1 \psi_1 \phi_2 \theta_2 \psi_2 \cdots \phi_n \theta_n \psi_n]$.

⁵Here, the number of DOFs for each agent will in general be twice the number of segments, depending on embodiment restrictions, e.g., for $S = [\theta_1 \phi_1 \theta_2 \phi_2]$, $S_3 = \theta_2$.

Appendix for the kinematic equations). For this paper, the actions will be defined using the relative-state vectors.

Effects can be defined as changes to the body-world relationship (e.g., location) of the agent and/or to positions, orientations, and states of external objects. In this paper, we will only consider state and action behavior aspects and metrics. For characterization of the *space of effect metrics* and some mathematical definitions of effect metrics, see [5].

III. SOME FIRST STATE AND ACTION METRICS

To evaluate the similarity of behavior, with respect to states and actions, between an agent β imitating another agent α , we define and use appropriate metrics. For the moment, let us assume that both agent embodiments have the same number of DOFs n . A first global **state metric** can be defined as

$$\mu_{\text{state}}(S^\alpha, S^\beta) = \sum_{j=1}^n |S_j^\alpha - S_j^\beta| \quad (1)$$

where S_j^α and S_j^β are the values of the state vectors for the two agents. Depending on whether the relative or absolute-state vectors are used, we call the state metric *relative* or *absolute*, respectively.

A first global **action metric** can be defined as

$$\mu_{\text{action}}(A^\alpha, A^\beta) = \sum_{j=1}^n |A_j^\alpha - A_j^\beta| \quad (2)$$

where A_j^α and A_j^β are the values of the action vectors for the two agents. Note that instead of absolute value, one could alternatively use any L^p -norm, the choice of which might have consequences for optimization in different applications.

An agent performing actions so as to minimize one (or a weighted combination) of these two metrics would successfully imitate a demonstrator in respect to states and/or actions. Note that it is not necessary and, in general, will not be possible to bring the value of the metric to zero with a matching behavior, especially in the case of dissimilar embodiments. Instead, finding minima is the goal.⁶

For effect metrics, see [5]. In the following sections, some more complex state and action metrics are defined.

IV. CORRESPONDENCE MAPPING

For two agents, demonstrator α and imitator β with n and m DOFs, respectively, an $n \times m$ **correspondence matrix** can be defined as

$$\mathcal{C} = \begin{bmatrix} w_{1,1} & w_{1,2} & \cdots & w_{1,m} \\ w_{2,1} & w_{2,2} & \cdots & w_{2,m} \\ \vdots & \vdots & \ddots & \vdots \\ w_{n,1} & w_{n,2} & \cdots & w_{n,m} \end{bmatrix}$$

⁶Of course, one can replace μ by $\mu' = \mu - m$, where $m = \inf \mu$ and then seek to solve $\mu' = 0$.

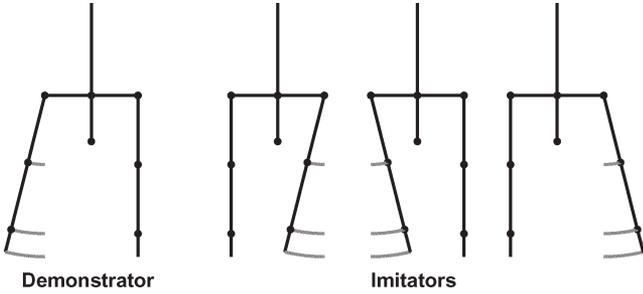


Fig. 2. Examples of symmetry via correspondence mapping. The figure shows (left) a demonstrator and (right) three imitators, facing the reader, each with an upper human torso embodiment. The demonstrator is moving its right arm to its left. Each of the three imitators is using different correspondence mappings: mapping the demonstrator’s right arm to the left arm of the imitator (second from the left), using a weight of minus one, but maintaining the same arm mapping (second from the right), and, finally, both mapping the demonstrator’s right arm to the left arm of the imitator and using a weight of minus one. The gray lines trace the movement of the arms.

where the $w_{i,j}$ values are real-valued weights, determining how the j th DOF of the imitator β depends on the i th DOF of the demonstrator α . The j th column of the matrix can be thought as a vector indicating how the DOFs of the demonstrator influence the j th DOF of the imitator. Depending on how many of the weights have a nonzero value, this correspondence mapping can be one-to-one, one-to-many (or many-to-one), or many-to-many. If *partial* body imitation is desired, some DOF of the imitator (and/or the demonstrator) can be omitted by setting an entire column (respectively, row) to zero in the correspondence matrix.

The choice of the correspondence mapping will in general depend on the particular task. Assuming both agents share the same embodiment (and, as a result, have the same number of DOFs), a simple example of a one-to-one correspondence mapping would be using the identity matrix as a correspondence matrix. Alternatively, if some mirror symmetry is wanted, then the DOFs for the right arm and leg of the demonstrator (see the example in Fig. 2, left) could be mapped to the DOFs for the left arm and leg of the imitator, and vice versa (Fig. 2, second from the left). Another possible form of symmetry results from mapping some of the demonstrator’s DOF using a weight of minus one [e.g., if the demonstrator raises its hand, the imitator should lower its hand, or if the demonstrator turns its head to the left, the imitator should turn to the right (see the example in Fig. 2, second from right)].

If the agents do not have the same number of DOFs (or depending on their particular morphology), it may be useful to map a single DOF to many DOFs. For example, consider correspondences between a human body as model to a dolphin-like imitator⁷: a dolphin, using its mouth corresponding to either human hand (grasping an object) or using its tail to toss a ball back to a human that used both arms, comprise real-world examples of many-to-one mappings. These two examples also illustrate that the correspondence needs not be static—the human hands are mapped to different dolphin body parts in

⁷Different mappings do appear to be employed by real-life dolphins in imitating humans [17].

each case—but can be adapted depending on the context and the tasks involved.

V. INDUCED STATE AND ACTION METRICS

The metric definitions in Section III are appropriate for the most simple one-to-one mapping (the identity mapping), with both agents sharing the same number of DOFs (and probably a very similar morphology). But, in general, using a correspondence matrix, other metric definitions can be induced.

First, the state vector S^α and the action vector A^α of the demonstrator can be multiplied with the correspondence matrix

$$\mathcal{S} = S^\alpha \times \mathcal{C} \quad (3)$$

$$\mathcal{A} = A^\alpha \times \mathcal{C} \quad (4)$$

producing two new vectors in imitator coordinates. Combining (1) and (3) for the state metric gives

$$\mu_{\text{state}}^{\mathcal{C}}(\mathcal{S}, S^\beta) = \sum_{j=1}^m |S_j - S_j^\beta| \epsilon_j \quad (5)$$

where S^β is the imitator’s attempted matching state, and the corrective term

$$\epsilon_j = \begin{cases} 0, & \text{if } \sum_{i=1}^n w_{i,j}^2 = 0 \\ 1, & \text{otherwise} \end{cases} \quad (6)$$

takes the value zero if the j th column of the correspondence matrix contains only zeros (effectively omitting the imitator’s j th DOF). Intuitively, the components of \mathcal{S} and \mathcal{A} (for such $\epsilon_j \neq 0$) can be thought as current subgoal state and action target values. The imitator can match the state \mathcal{S} by assuming state S^β so as to minimize the metric $\mu_{\text{state}}^{\mathcal{C}}$. As in the previous definition, this state metric is called *relative* or *absolute* depending on whether the relative or absolute-state vectors are used, respectively.

Finally, combining (2), (4), and (6) for the action metric gives

$$\mu_{\text{action}}^{\mathcal{C}}(\mathcal{A}, A^\beta) = \sum_{j=1}^m |A_j - A_j^\beta| \epsilon_j \quad (7)$$

where A^β is the imitator’s attempted matching action.

These $\mu_{\text{state}}^{\mathcal{C}}$ and $\mu_{\text{action}}^{\mathcal{C}}$ metrics are called the *induced state* and *action metrics for the linear correspondence C*.

Depending on the correspondence mapping used, a plethora of new complex metrics (also allowing for dissimilar embodiments) can be induced considering state or action aspects. The next section will illustrate a variety of examples.

VI. MAPPING ACROSS DISSIMILAR BODIES

Using a system implemented in MATLAB, we are able to describe a variety of agent embodiments as simple kinematic models (defined in Section II with kinematic equations as in the Appendix). These embodiments include models of robotic platforms, as shown in Fig. 3.

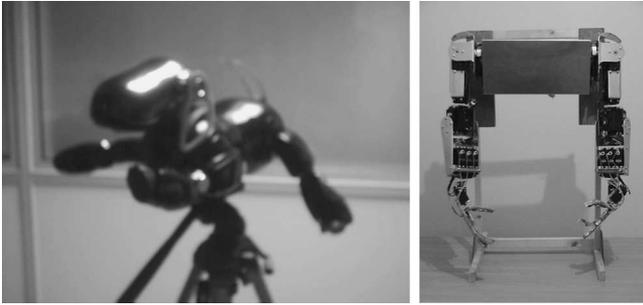


Fig. 3. Example robotic platforms. (Left) A doglike Sony AIBO robot is positioned suspended on a tripod. (Right) A pair of human-size robotic arms is shown attached to a wooden rack (the robot arms were designed and made by Michael Walters of the University of Hertfordshire).

For a given demonstrator and imitator embodiment pair, the imitator attempts to match the behavior of the demonstrator by minimizing a given metric (or a combination of metrics). This can be done continuously (*immediate imitation*) or after the completion of the demonstration (*deferred imitation*) [19]. Moreover, the *granularity* or “fineness” of the matching of actions, states, and/or effects determines a sequence of subgoals for the imitator to achieve, and the appropriate level of granularity may be different depending on the context and task. Different correspondence mappings can be defined between the two agents, yielding qualitatively different types of matching behaviors.

A. Examples

Using the system, we conducted a series of simulation runs, using a variety of agent embodiments and correspondence mappings. The demonstrator performs a series of actions, and the imitator tries to minimize the correspondence induced relative-state metric. Continuously using the components of \mathcal{S} for which $\epsilon_j \neq 0$ as the current subgoals for each DOF j , the imitator performs actions that attempt to reduce the contribution of error in each such component. Here, the rate of change was restricted to half the componentwise error. Of course, many other selection mechanisms are possible for both immediate or deferred imitation.

1) *Identity and Mirror Symmetry Mappings*: Two examples of imitation across similar embodiments are shown in Fig. 4. Both demonstrator and imitator are humanoid. In the first example, the identity correspondence mapping is used. In the second example, using the same demonstration, symmetry is achieved by mapping the left body parts of the demonstrator to the right body parts of the imitator and vice versa (see also examples in Fig. 2).

2) *Multiple Mappings Between Dissimilar Bodies*: The model of an AIBO robot is used as an imitator in the examples shown in Fig. 5. In the first example, the right arm of the demonstrator is mapped on the right front leg of the robot, the left arm to the left front leg, the right leg to the back right leg, and the left leg to the back left leg. As each of the arms and legs of the demonstrator consists of three segments, and the imitator’s legs consist of two segments, only the first two segments are mapped. In the second example, the imitator’s

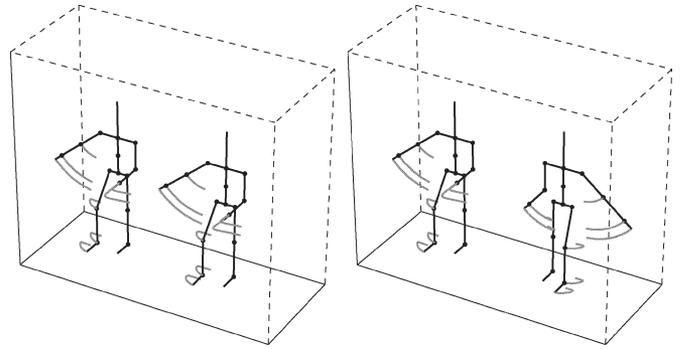


Fig. 4. Two examples of imitation across similar embodiments (humanoid). Both demonstrator (left in both examples) and imitator (right in both examples) share the same humanoid embodiment. In the example on the left, the identity mapping is used as the correspondence mapping. In the example on the right, the left arm and leg of the demonstrator are mapped on the right arm and leg of the imitator (and vice versa) with a weight of minus one, resulting in a mirror symmetry. The gray traces visualize the body-part trajectories.

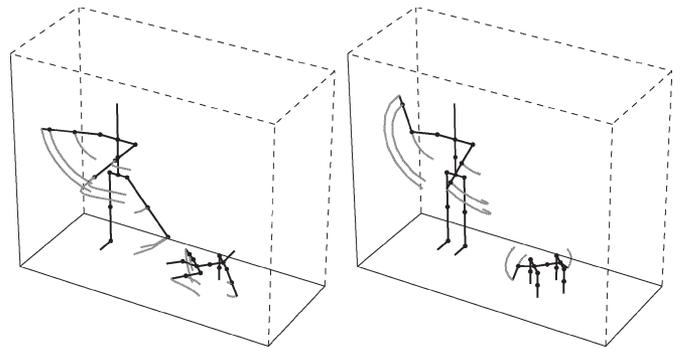


Fig. 5. Two examples of imitation across dissimilar embodiments (humanoid and doglike). The demonstrator (left in both examples) is embodied as a humanoid, while the imitator (right in both examples) is embodied as a doglike AIBO robot (see Fig. 3, left). In the example on the left, a simple one-to-one correspondence mapping is used, mapping the first two (out of three) segments of the demonstrator’s arms and legs to the two segments of the four imitator’s legs. In the second example on the right, the demonstrator’s first segment of the left arm is mapped on the imitator’s tail and the demonstrator’s first two right arm segments to the neck and head of the imitator’s head. This results in the demonstrator controlling “puppeteer-like” the head and tail of the robot. The gray traces visualize the body-part trajectories.

head and tail are controlled “puppeteer-like” by mapping the first two segments of the right arm and the first segment of the left arm of the demonstrator to them, respectively. The latter can also be thought as an example of using a body part of the demonstrator (the left arm) to refer to a body part of the imitator (the tail) that does not have a direct equivalent on the demonstrator’s (human) body. Also, although here both bodies have “heads,” it might be the case that it is more “expressive” (i.e., the motions/posture more easily to be perceived by the imitator and/or performed by the demonstrator) to use the right arm to indicate the head movements.

3) *Partial Mappings*: An example of partial mapping is shown in Fig. 6 (left). As the imitator is an upper torso humanoid, the DOFs in the lower body parts of the (whole humanoid) demonstrator are ignored (via zero rows in the matrix), with a unity one-to-one mapping used for the upper body.

4) *One-to-Many Mappings*: When the perception of the demonstrator by the imitator is limited, a complicated

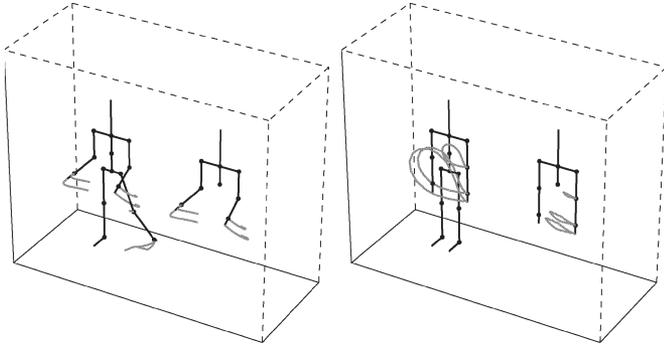


Fig. 6. Two examples of imitation across dissimilar embodiments (whole and upper torso only humanoids). The demonstrator (left in both examples) is embodied as a humanoid, while the imitator (right in both examples) is embodied as an upper torso humanoid robot (similar to the one shown in Fig. 3, right). In the first example on the left, the arms of the demonstrator are mapped using a weight of one to the arms of the imitator. Note that the movement of the demonstrator’s left leg is ignored as these demonstrator’s DOFs are omitted (via a zero row in the correspondence matrix). In the example on the right, the same mapping is used, but the rate of movement of the imitator is severely limited, resulting in *impersistence* (see further discussion in Section VII). The gray traces visualize the body-part trajectories.

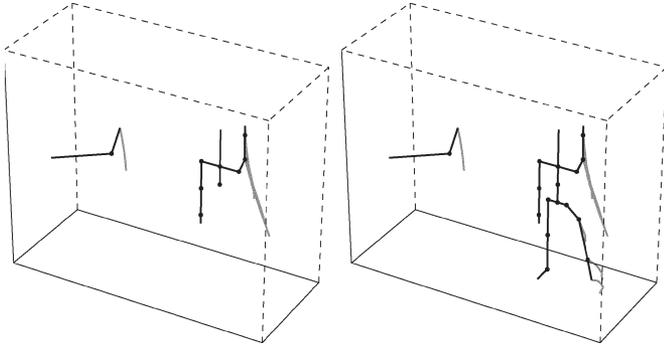


Fig. 7. Two examples of imitation across dissimilar embodiments using one-to-many correspondence mappings. The demonstrator (left in both examples) is embodied as an abstract “letter V” shape, visualizing the three motion sensors attached to a human (one to waist and one on each hand), while the imitator (right in both examples) is embodied as a humanoid. In the example to the left, the left segment of the demonstrator is mapped with different weight values to the imitator’s left-arm segments. In the second example to the right, this mapping is extended by also mapping the left segment of the demonstrator with different weight values to the imitator’s left leg segments. The gray traces visualize the body-part trajectories.

one-to-many correspondence mapping could be used. Assuming a human acting as a demonstrator, but providing the system and the imitator with only the coordinates of three motion sensors, one attached to her/his waist and one in each hand, and filtering perception through this sensory apparatus yields a reduced representation of the demonstrator embodiment that can be modeled as a “V”-letter-shaped kinematic model.⁸ The θ and ϕ of each arm segment of the “V” embodiment can be mapped on each of the segments of the corresponding arms of

⁸Note that as the human moves her/his arms around, the lengths of the two segments of the “V” will change accordingly and not remain constant. But, this can be ignored since, for the correspondence mapping, the important parameters are the azimuth and polar angles. These can be found from the (relative to the waist sensor) Cartesian coordinates of each arm sensor.

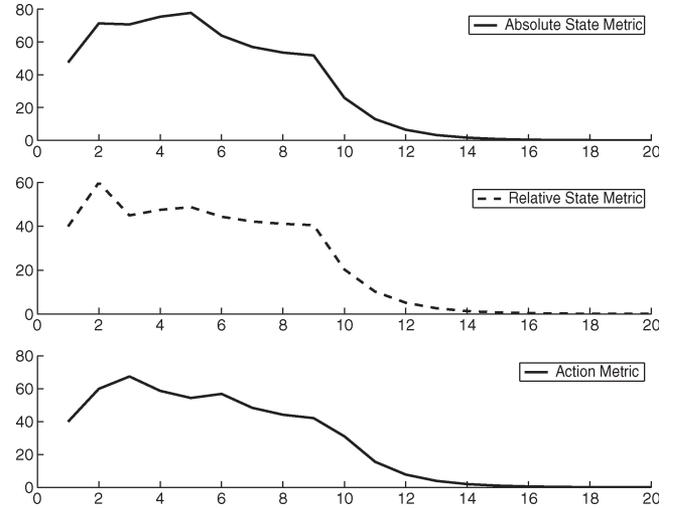


Fig. 8. Induced absolute-state, relative-state, and action metric values for the imitation example shown in Fig. 4 (left), plotted during the simulation. The abscissa and ordinate units are time step and value of the metric, respectively. Note that, here, the system is using the relative-state metric only to produce the imitative behavior.

a humanoid imitator, with different weights. For example, in Fig. 7 (left), the following mapping:

$$\begin{aligned} w_{\theta_{LA}, \theta_{LS}} &= w_{\phi_{LA}, \phi_{LS}} = 1 \\ w_{\theta_{LA}, \theta_{LE}} &= w_{\phi_{LA}, \phi_{LE}} = 0.5 \\ w_{\theta_{LA}, \theta_{LW}} &= w_{\phi_{LA}, \phi_{LW}} = 0.25 \end{aligned}$$

is used, where LA is the demonstrator’s left “V” arm segment, and LS, LE, and LW are the imitator’s left shoulder, elbow, and wrist segments, respectively. As a result, as the human demonstrator (not shown) lifts her/his left arm, the left segment of the “V” model also rises, and the left arm of the humanoid imitator rises as well. In Fig. 7 (right), an extension of this mapping is used with

$$\begin{aligned} w_{\theta_{LA}, \theta_{LH}} &= 0.2 \\ w_{\phi_{LA}, \phi_{LH}} &= 1 \\ w_{\theta_{LA}, \theta_{LK}} &= -1 \\ w_{\theta_{LA}, \theta_{LAN}} &= 0.05 \\ w_{\phi_{LA}, \phi_{LAN}} &= 0.05 \end{aligned}$$

where LH, LK, and LAN are the imitator’s left hip, knee, and ankle segments. As a result, both the left arm and leg of imitator rise when the demonstrator’s left “V” segment moves. These mappings are presented here only as indicative examples of complex one-to-many mappings and were hand-coded.

B. Evaluation Using State and Action Metrics

The values of the induced absolute-state, relative-state, and action metrics during the simulations shown in Figs. 4–7 are plotted in Figs. 8–15.

The rate of movement of the imitator during the simulations was limited to reflect real-world constraints. The current

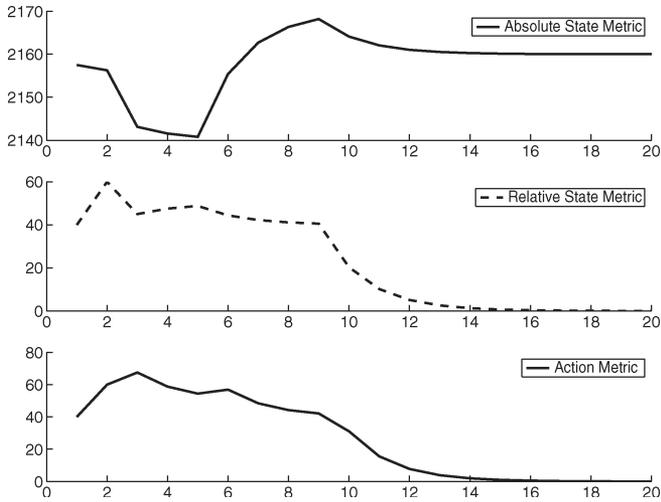


Fig. 9. Induced absolute-state, relative-state, and action metric values for the imitation example shown in Fig. 4 (right), plotted during the simulation. The abscissa and ordinate units are time step and value of the metric, respectively. Note that, here, the system is using the relative-state metric only to produce the imitative behavior.

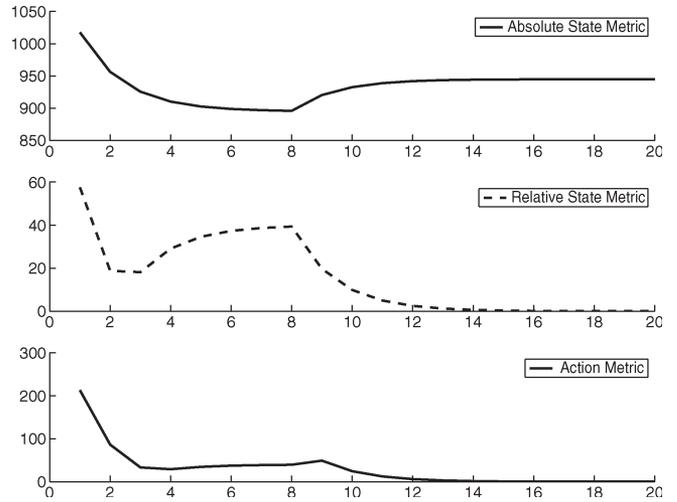


Fig. 11. Induced absolute-state, relative-state, and action metric values for the imitation example shown in Fig. 5 (right), plotted during the simulation. The abscissa and ordinate units are time step and value of the metric, respectively. Note that, here, the system is using the relative-state metric only to produce the imitative behavior.

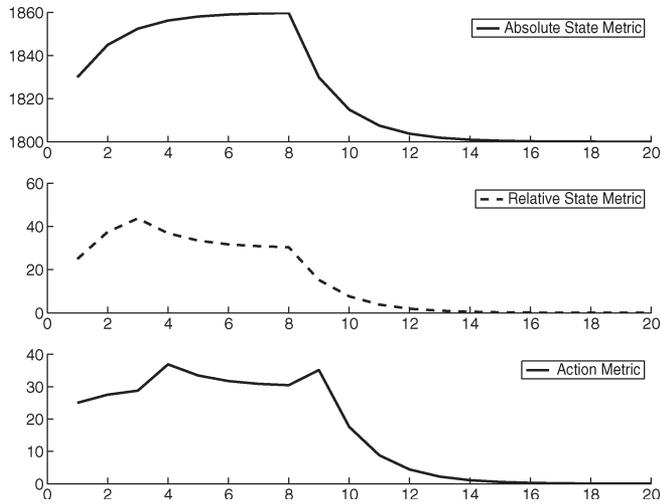


Fig. 10. Induced absolute-state, relative-state, and action metric values for the imitation example shown in Fig. 5 (left), plotted during the simulation. The abscissa and ordinate units are time step and value of the metric, respectively. Note that, here, the system is using the relative-state metric only to produce the imitative behavior.

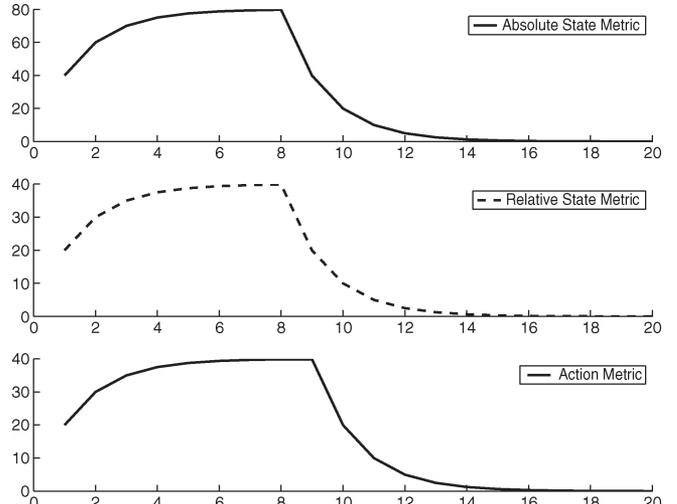


Fig. 12. Induced absolute-state, relative-state, and action metric values for the imitation example shown in Fig. 6 (left), plotted during the simulation. The abscissa and ordinate units are time step and value of the metric, respectively. Note that, here, the system is using the relative-state metric only to produce the imitative behavior.

subgoal is constantly updated using the current state of the demonstrator during the run. When each demonstration finishes, the simulation does not terminate until the imitator manages to reach the final subgoal of the demonstrator.

As the imitator tries to minimize the relative-state errors, the value of the metric used by the system (here, in all cases, the relative-state metric, middle plot) in all figures is shown to converge to zero. Depending on the embodiment and the correspondence mapping, the value of the absolute-state metric (top plot) also converges to a value that may (e.g., see Fig. 8) or may not (e.g., see Fig. 9) be zero. The value of the action metric (bottom plot) trivially converges to zero (since toward the end of each simulation, both the demonstrator and the imitator stop performing any actions).

C. Mapping to Robotic Target Platforms

Capturing human demonstration via the “V”-perception mechanism (see Section VI-A above), implemented using the Polhemus Liberty motion capture system, and mapping to robotic target platforms (AIBO and robot arms—see Fig. 3) according to different correspondence matrices, as described above, allow the achievement of qualitatively dissimilar types of robotic imitation based on human demonstration.

VII. DISCUSSION AND FUTURE WORK

We have shown how partial, mirror symmetric, one-to-one, one-to-many, many-to-one, and many-to-many body mappings can be characterized by (linear) *correspondence matrices*.

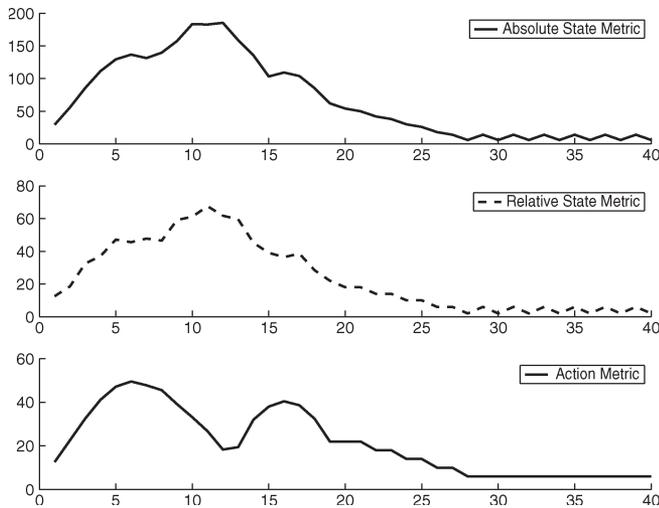


Fig. 13. Induced absolute-state, relative-state, and action metric values for the imitation example shown in Fig. 6 (right), plotted during the simulation. The abscissa and ordinate units are time step and value of the metric, respectively. Note that, here, the system is using the relative-state metric only to produce the imitative behavior.

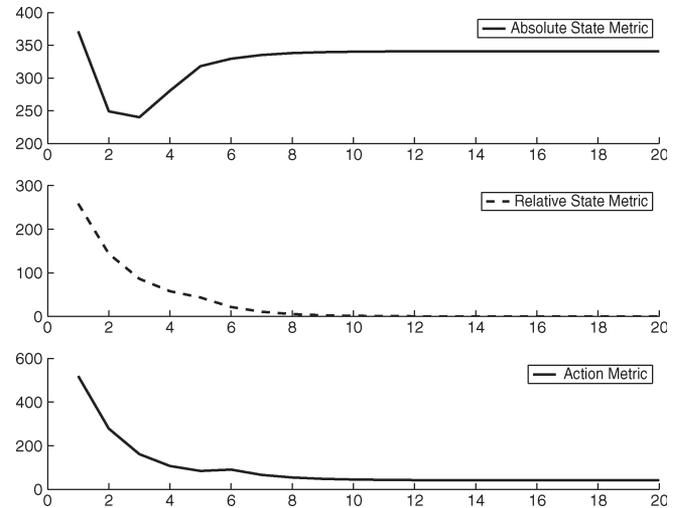


Fig. 15. Induced absolute-state, relative-state, and action metric values for the imitation example shown in Fig. 7 (right), plotted during the simulation. The abscissa and ordinate units are time step and value of the metric, respectively. Note that, here, the system is using the relative-state metric only to produce the imitative behavior.

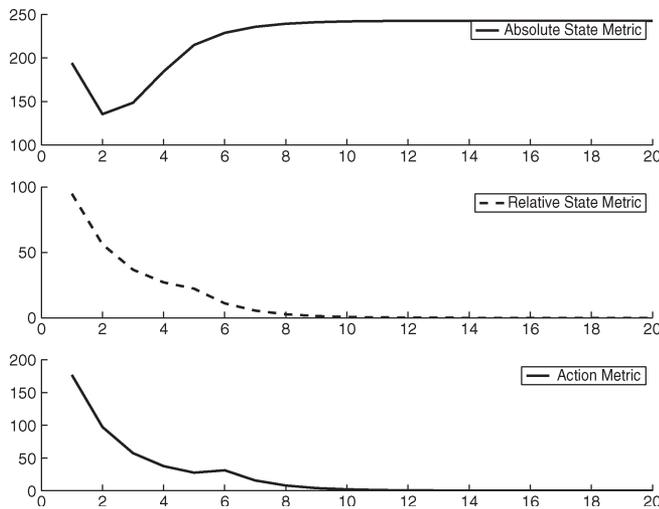


Fig. 14. Induced absolute-state, relative-state, and action metric values for the imitation example shown in Fig. 7 (left), plotted during the simulation. The abscissa and ordinate units are time step and value of the metric, respectively. Note that, here, the system is using the relative-state metric only to produce the imitative behavior.

These correspondences induce an infinite variety of absolute and relative-state and action metrics that can be used to guide robotic imitation across dissimilar embodiments—even radically different ones in which neither the size of body parts, nor their type, nor number of DOFs needs be preserved—enhancing existing approaches to imitation learning. The study of nonlinear correspondences for achieving matching behavior in states, actions, and/or effects would extend this set of metrics.

Note that as discussed in Section VI-B and shown in Figs. 8–15, in some situations, the values of some state or action metrics (that were not used by the system) do not decay to zero, while the imitator still produces a meaningful imitation (either by mirroring or by using different body parts). Depending on the metric used, qualitatively different imitation behaviors will result.

Currently, the correspondence mapping in this system is given for each experimental run, but finding the correspondence can be approached using reinforcement learning and an experiential history (adding memory), as in our previous work with the ALICE generic imitation framework [1], [2]. How such correspondences can be built is one of the hardest problems of imitation learning and comprises an important aspect of the *what* to imitate problem. For related work on visuo-somatic mapping, see [7] and [13].

The scope of this paper is not to compare between different correspondence mappings (as they are currently given); but for future work, it would be an interesting issue to examine how different mappings (derived, e.g., by a genetic algorithm system or via a suitably adapted version of ALICE [1] using correspondence matrices) can be compared and evaluated for a particular task.

The solutions generated by the present system do not take into account joint limits, body mass distribution, or detailed dynamics, which may lead to unfeasible and/or unstable solutions (cf. [20] for an example where the trajectory had to be corrected to ensure proper balance, after mapping). By relying on simulation, we avoided many of the typical problems associated with the interpretation of the demonstrator’s actions and with the computation of the demonstrator’s state. It will, in general, be necessary to adapt a generated action sequence for a particular configuration of the particular physical system.

The current system is purely reactive with no memory (and, as a result, has no learning). This results in certain limitations. For example, in Fig. 6 (right), the demonstrator is moving its left arm in a loop trajectory, but since the imitator is continuously reacting but with a limited rate of movement, the imitator is unable to reach the current subgoal before it changes (until finally the demonstrator completes the entire demonstration). This inability to sustain appropriate actions is called *impersistence* [23] and can be solved by adding memory to the system, containing the sequence of subgoals.

Future work would naturally also address the derivation of, and switching between, appropriate correspondence mappings depending on the needs of the imitator agent in the social and task context. The developed system could eventually serve as a *correspondence engine for imitation learning*, incorporating aspects of discovering *what* to imitate, depending on the context and interaction history.

APPENDIX KINEMATIC EQUATIONS FOR THE MODELS

A kinematic model (as defined in Section II) with n segments can be positioned in the 3-D workspace by assigning the values (x_0, y_0, z_0) as the coordinates of both the base and the tip of the root segment. Inductively, the Cartesian coordinates (x_i, y_i, z_i) of the base of each subsequent segment i are equal to the coordinates of the tip of its parent segment p_i

$$\begin{aligned}x_i &= x'_{p_i} \\y_i &= y'_{p_i} \\z_i &= z'_{p_i}.\end{aligned}$$

The Cartesian coordinates (x'_i, y'_i, z'_i) of the tip of each segment i can be found by using the spherical coordinates $(\ell_i, \theta_i, \phi_i)$ describing the current (relative) position of the segment

$$\begin{aligned}x'_i &= x_i + \ell_i \cos(\phi_i + \Phi_{p_i}) \cos(\theta_i + \Theta_{p_i}) \\y'_i &= y_i + \ell_i \cos(\phi_i + \Phi_{p_i}) \sin(\theta_i + \Theta_{p_i}) \\z'_i &= z_i + \ell_i \sin(\phi_i + \Phi_{p_i}).\end{aligned}$$

ACKNOWLEDGMENT

The authors would like to thank M. Walters and A. Mirza (both with the Adaptive Systems Research Group, University of Hertfordshire) for their help with the robots for this work. The authors would also like to thank the three anonymous reviewers for their constructive suggestions for improving this paper.

REFERENCES

- [1] A. Alissandrakis, C. L. Nehaniv, and K. Dautenhahn, "Imitation with ALICE: Learning to imitate corresponding actions across dissimilar embodiments," *IEEE Trans. Syst., Man, Cybern. A, Syst. Humans*, vol. 32, no. 4, pp. 482–496, Jul. 2002.
- [2] —, "Towards robot cultures?—Learning to imitate in a robotic arm test-bed with dissimilar embodied agents," *Interact. Stud.: Soc. Behav. Commun. Biol. Artif. Syst.*, vol. 5, no. 1, pp. 3–44, 2004.
- [3] A. Alissandrakis, C. L. Nehaniv, K. Dautenhahn, and J. Saunders, "Imitating using JABBERWOCKY to achieve corresponding effects in context," Univ. Hertfordshire, Hatfield, U.K., Computer Science Tech. Rep. 441, Nov. 2005. Journal version forthcoming.
- [4] —, "Achieving corresponding effects on multiple robotic platforms: Imitating using different effect metrics," in *Proc. 3rd Int. Symp. Imitation Animals and Artifacts*, Hatfield, U.K., Apr. 12–14, 2005, pp. 10–19.
- [5] —, "An approach for programming robots by demonstration to manipulate objects: Considerations on metrics to achieve corresponding effects," in *Proc. 6th IEEE Int. Symp. CIRA*, 2005, pp. 61–66.
- [6] R. Amit and M. J. Mataric, "A correspondence metric for imitation," in *Proc. 19th Nat. Conf. Artif. Intell., 16th Conf. Innovative Appl. Artif. Intell.*, D. L. McGuinness and G. Ferguson, Eds., San Jose, CA, Jul. 25–29, 2004, pp. 944–945.

- [7] M. Asada, M. Ogino, S. Matsuyama, and J. Ooga, "Imitation learning based on visuo-somatic mapping," in *Proc. 9th Int. Symp. Exp. Robot.*, O. K. Marcelo and H. Ang, Eds. Berlin, Germany: Springer-Verlag, 2006, vol. 21, pp. 269–278.
- [8] H. Bekkering and W. Prinz, "Goal representations in imitative actions," in *Imitation in Animals and Artifacts*, K. Dautenhahn and C. L. Nehaniv, Eds. Cambridge, MA: MIT Press, 2002, pp. 555–572.
- [9] A. Billard, Y. Epars, S. Calinon, G. Cheng, and S. Schaal, "Discovering optimal imitation strategies," *Robot. Auton. Syst.*, vol. 47, no. 2/3, pp. 68–77, 2004.
- [10] A. Billard, "Learning motor skills by imitation: A biologically inspired robotic model," *Cybern. Syst.*, vol. 32, no. 1/2, pp. 155–193, Jan. 2001.
- [11] C. Breazeal and B. Scassellati, "Robots that imitate humans," *Trends Cogn. Sci.*, vol. 6, no. 11, pp. 481–487, Nov. 2002.
- [12] G. Butterworth, "Pointing is the royal road to language for babies," in *Pointing: Where Language, Culture, and Cognition Meet*, S. Kita, Ed. Mahwah, NJ: Lawrence Erlbaum, 2003, pp. 9–26.
- [13] M. Cabido-Lopes and J. Santos-Victor, "Visual transformations in gesture imitation: What you see is what you do," in *Proc. IEEE ICRA*, Taipei, Taiwan, R.O.C., Sep. 14–19, 2003, pp. 2375–2381.
- [14] J. Call and M. Carpenter, "Three sources of information in social learning," in *Imitation in Animals and Artifacts*, K. Dautenhahn and C. L. Nehaniv, Eds. Cambridge, MA: MIT Press, 2002.
- [15] J. Demiris and G. Hayes, "Imitation as a dual-route process featuring predictive and learning components: A biologically-plausible computational model," in *Imitation in Animals and Artifacts*, K. Dautenhahn and C. L. Nehaniv, Eds. Cambridge, MA: MIT Press, 2002, pp. 327–361.
- [16] W. Erlhagen, A. Mukovskiy, E. Bicho, G. Panin, C. Kiss, A. Knoll, H. van Schie, and H. Bekkering, "Action understanding and imitation learning in a robot-human task," in *Proc. ICANN*, W. Duch, J. Kacprzyk, E. Oja, and S. Zadrozny, Eds. New York: Springer-Verlag, Sep. 11–15, 2005, vol. 3696, pt. 1, pp. 261–268.
- [17] L. M. Herman, "Vocal, social, and self imitation by bottlenosed dolphins," in *Imitation in Animals and Artifacts*, K. Dautenhahn and C. L. Nehaniv, Eds. Cambridge, MA: MIT Press, 2002, pp. 63–108.
- [18] Y. Kuniyoshi, M. Inaba, and H. Inoue, "Learning by watching: Extracting reusable task knowledge from visual observations of human performance," *IEEE Trans. Robot. Autom.*, vol. 10, no. 6, pp. 799–822, Dec. 1994.
- [19] J. Nadel, C. Guerini, A. Peze, and C. Rivet, "The evolving nature of imitation as a format of communication," in *Imitation in Infancy*, J. Nadel and G. Butterworth, Eds. Cambridge, MA: Cambridge Univ. Press, 1999, pp. 209–234.
- [20] S. Nakaoka, A. Nakazawa, K. Yokoi, H. Hirukawa, and K. Ikeuchi, "Generating whole body motions for a biped humanoid robot from captured human dances," in *Proc. ICRA*, Taipei, Taiwan, R.O.C., Sep. 14–19, 2003, pp. 3905–3910.
- [21] C. L. Nehaniv and K. Dautenhahn, "Mapping between dissimilar bodies: Affordances and the algebraic foundations of imitation," in *Proc. EWLR-7*, J. Demiris and A. Birk, Eds., Edinburgh, U.K., Jul. 20, 1998, pp. 64–72.
- [22] M. N. Nicolescu and M. M. Mataric, "Learning and interacting in human-robot domains," *IEEE Trans. Syst., Man, Cybern. A, Syst. Humans*, vol. 31, no. 5, pp. 419–430, Sep. 2001.
- [23] J. Saunders, C. L. Nehaniv, and K. Dautenhahn, "An experimental comparison of imitation paradigms used in social robotics," in *Proc. 13th IEEE RO-MAN*, Okayama, Japan, Sep. 20–22, 2004, pp. 691–696.
- [24] B. Scassellati, "Imitation and mechanisms of joint attention: A developmental structure for building social skills," in *Computation for Metaphors, Analogy and Agents*, vol. 1562, C. L. Nehaniv, Ed. New York: Springer-Verlag, 1999, pp. 176–195.
- [25] S. Schaal, "Is imitation learning the route to humanoid robots?" *Trends Cogn. Sci.*, vol. 3, no. 6, pp. 233–242, 1999.



Aris Alissandrakis received the Ph.D. degree in computer science from the University of Hertfordshire, Hatfield, U.K., in 2003 and the M.Eng. degree in cybernetics from the University of Reading, Reading, U.K., in 1999.

He is currently a Postdoctoral Research Fellow with the Adaptive Systems Research Group, School of Computer Science, at the University of Hertfordshire, where he studies imitation in the context of human-robot interaction.



Christopher L. Nehaniv received the Ph.D. degree from the University of California, Berkeley, in 1992.

He is currently a Professor of Mathematical and Evolutionary Computer Sciences with the School of Computer Science, University of Hertfordshire, Hatfield, U.K. He is also the Director of the Engineering and Physical Science Research Council Network on Evolvability in Biological and Software Systems.



Kerstin Dautenhahn received the Ph.D. degree from the University of Bielefeld, Bielefeld, Germany, in 1993.

She is currently with the University of Hertfordshire, Hatfield, U.K., where she is a Professor of Artificial Intelligence with the School of Computer Science and the Coordinator of the Adaptive Systems Research Group. She has edited several books and frequently organizes international research workshops and conferences. She has published more than 100 research articles on social robotics, robot learning, human–robot interaction, and assistive technology, and is involved in several European robotics projects.