Social Learning for Social Robots

Aris Alissandrakis
a.alissandrakis@herts.ac.uk
http://homepages.feis.herts.ac.uk/~comqaa1/

Adaptive Systems Research Group
http://adapsys.feis.herts.ac.uk
University of Hertfordshire

HCT Seminar
University of Sussex, 24 Nov 2006

The Adaptive Systems Research Group is a multidisciplinary group of faculty, students, and friends of the University of Hertfordshire who have a connection to research in adaptive systems and related areas: computer science, robotics, biology, mathematics, from psychology to engineering.

Areas of research:

- Artificial Life
- Socially Intelligent Agents/Robots
- Artificial Intelligence (Embodied)

Key Academics:

- Prof. Kerstin Dautenhahn (coordinator)
- Prof. Chrystopher Nehaniv
- Dr. René te Boekhorst
- Dr. Lola Cañamero
- Dr. Daniel Polani

$\sim$ 14 Postdocs/Research Assistants

$\sim$ 15 PhD students

One of the great challenges in bringing robots out of laboratories into real-world domestic and public environments is to address the fact that humans and robots have different modalities, and exploit this knowledge for successfully operate in the same space.

This talk will present current work by the Adaptive Systems research group at the University of Hertfordshire on social interaction and social learning between humans and robots, in shared environments.

Imitation is a powerful learning tool when a number of agents interact in a social context. Robots capable of imitation (or other, simpler, forms of social learning) would allow humans to interact with them and transfer knowledge in a more adaptive and ”natural” (for the humans) way.

Learning how to do an act from seeing it done. [Thorndike 1898]

Copying of a novel or otherwise improbable act or utterance, or some act for which there is clearly no instinctive tendency. [Thorpe 1963]

Something **C**(opy) is produced by an organism or machine, where **C** is similar to something else **M**(odel), registration of **M** is necessary for the production of **C**, and **C** is designed to be similar to **M**. [Mitchell 1987]

Learning how to do an act from seeing it done. [Thorndike 1898]

Copying of a novel or otherwise improbable act or utterance, or some act for which there is clearly no instinctive tendency. [Thorpe 1963]

Something **C**(opy) is produced by an organism or machine, where **C** is similar to something else **M**(odel), registration of **M** is necessary for the production of **C**, and **C** is designed to be similar to **M**. [Mitchell 1987]

Learning how to do an act from seeing it done. [Thorndike 1898]

Copying of a novel or otherwise improbable act or utterance, or some act for which there is clearly no instinctive tendency. [Thorpe 1963]

Something **C**(opy) is produced by an organism or machine, where **C** is similar to something else **M**(odel), registration of **M** is necessary for the production of **C**, and **C** is designed to be similar to **M**. [Mitchell 1987]

First-level: Similar to *mimicry*, the copy is produced unintentionally without the registration of the model, achieved through morphogenesis and evolutionary selection processes.

Second-level: At this level the organism perceives the model and produces a copy, but is not concerned about the accuracy and the correspondence of this copy.

Third-level: Comparisons between the produced copy and the model can be made by the organism. Learning is involved and if the replication is repeated, the organism will try to better match the model.

Fourth-level: Variation into the copy is introduced by the self-aware organism, changing some aspects to conform to the model and others to extend it in some way.

Fifth-level: The organism has knowledge of another's perspective and is able to plan the imitation, employing it as either parody or deceit (drawing or avoiding attention to it respectively). [Mitchel 1987]

First-level: Similar to *mimicry*, the copy is produced unintentionally without the registration of the model, achieved through morphogenesis and evolutionary selection processes.

Second-level: At this level the organism perceives the model and produces a copy, but is not concerned about the accuracy and the correspondence of this copy.

Third-level: Comparisons between the produced copy and the model can be made by the organism. Learning is involved and if the replication is repeated, the organism will try to better match the model.

Fourth-level: Variation into the copy is introduced by the self-aware organism, changing some aspects to conform to the model and others to extend it in some way.

Fifth-level: The organism has knowledge of another's perspective and is able to plan the imitation, employing it as either parody or deceit (drawing or avoiding attention to it respectively). [Mitchel 1987]

First-level: Similar to *mimicry*, the copy is produced unintentionally without the registration of the model, achieved through morphogenesis and evolutionary selection processes.
Second-level: At this level the organism perceives the model and produces a copy, but is not concerned about the accuracy and the correspondence of this copy.
Third-level: Comparisons between the produced copy and the model can be made by the organism. Learning is involved and if the replication is repeated, the organism will try to better match the model.
Fourth-level: Variation into the copy is introduced by the self-aware organism, changing some aspects to conform to the model and others to extend it in some way.
Fifth-level: The organism has knowledge of another's perspective and is able to plan the imitation, employing it as either parody or deceit (drawing or avoiding attention to it respectively). [Mitchel 1987]

First-level: Similar to *mimicry*, the copy is produced unintentionally without the registration of the model, achieved through morphogenesis and evolutionary selection processes.

Second-level: At this level the organism perceives the model and produces a copy, but is not concerned about the accuracy and the correspondence of this copy.

Third-level: Comparisons between the produced copy and the model can be made by the organism. Learning is involved and if the replication is repeated, the organism will try to better match the model.

Fourth-level: Variation into the copy is introduced by the self-aware organism, changing some aspects to conform to the model and others to extend it in some way.

Fifth-level: The organism has knowledge of another's perspective and is able to plan the imitation, employing it as either parody or deceit (drawing or avoiding attention to it respectively). [Mitchel 1987]

First-level: Similar to *mimicry*, the copy is produced unintentionally without the registration of the model, achieved through morphogenesis and evolutionary selection processes.

Second-level: At this level the organism perceives the model and produces a copy, but is not concerned about the accuracy and the correspondence of this copy.

Third-level: Comparisons between the produced copy and the model can be made by the organism. Learning is involved and if the replication is repeated, the organism will try to better match the model.

Fourth-level: Variation into the copy is introduced by the self-aware organism, changing some aspects to conform to the model and others to extend it in some way.

Fifth-level: The organism has knowledge of another's perspective and is able to plan the imitation, employing it as either parody or deceit (drawing or avoiding attention to it respectively). [Mitchel 1987]

Imitation is best considered as the behaviour of an autonomous agent in relation to its environment, including other autonomous agents.
The mechanisms underlying imitation are not separated from the behaviour-in-context, including the social and non-social environments, motivations, relationships among the agents, the agent's individual and learning history etc. [Dautenhahn and Nehaniv 2002]

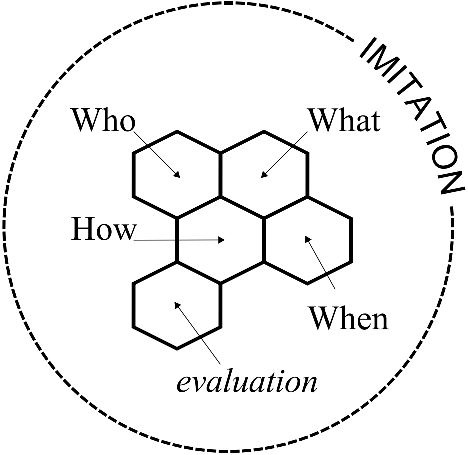

- **Who to imitate?**
- *When* to imitate?
- *What* to imitate?
- *How* to imitate?
- *Evaluation* . . .

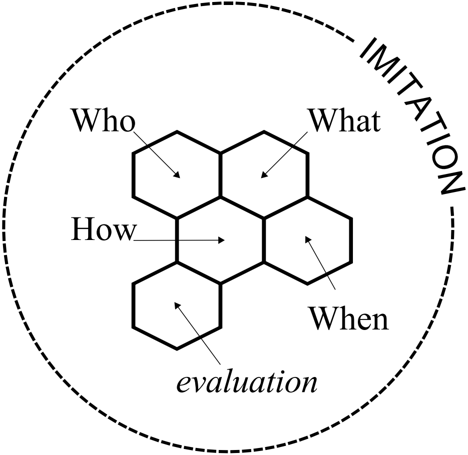

- *Who* to imitate?
- **When to imitate?**
- *What* to imitate?
- *How* to imitate?
- *Evaluation* . . .

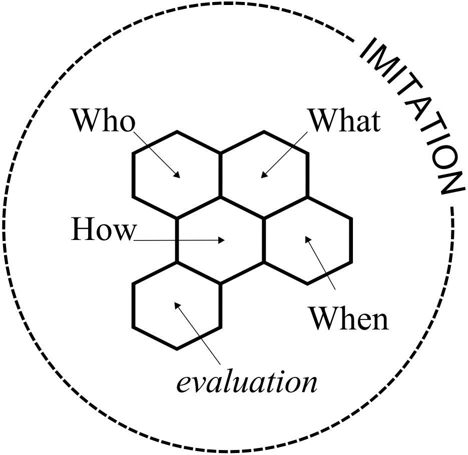

- *Who* to imitate?
- *When* to imitate?
- **What to imitate?**
- *How* to imitate?
- *Evaluation* . . .

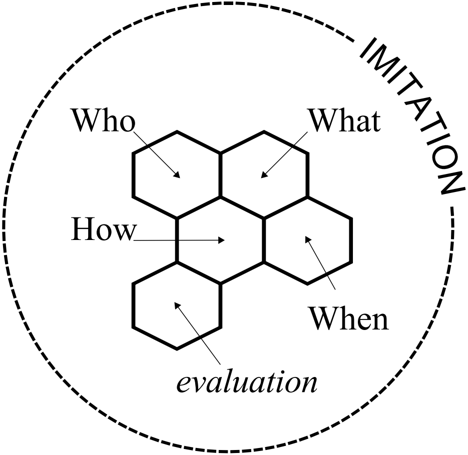

- *Who* to imitate?
- *When* to imitate?
- *What* to imitate?
- **How to imitate?**
  - *Evaluation . . .*

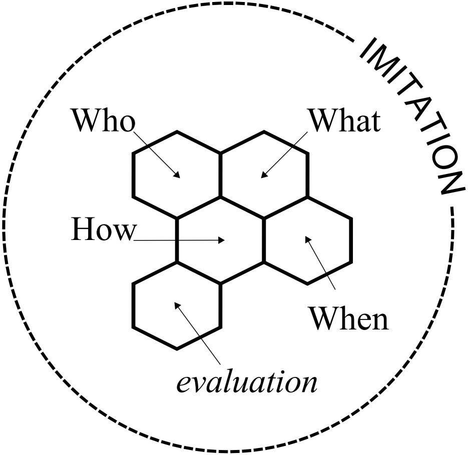

- *Who* to imitate?
- *When* to imitate?
- *What* to imitate?
- *How* to imitate?
- **Evaluation . . .**

Given an observed behaviour of the model, which from a given starting state leads the model through a sequence (or hierarchy [or program]) of *sub-goals* in *states*, *actions* and/or *effects*, one must find and execute a sequence of actions using one's own (**possibly dissimilar**) embodiment, which from a corresponding starting state, leads through corresponding sub-goals - in corresponding states, actions, and/or effects, while possibly responding to corresponding events.
[Nehaniv and Dautenhahn 2000, 2001, 2002]

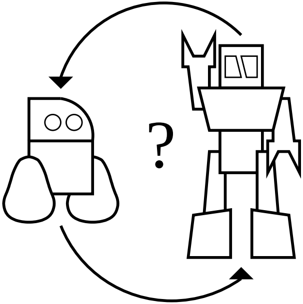

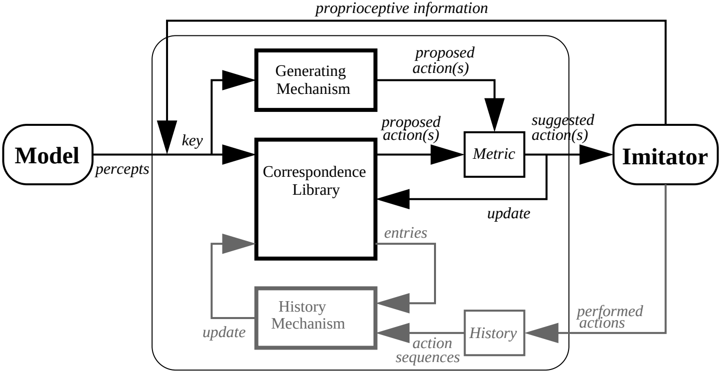

We have developed ALICE (**A**ction **L**earning via **I**mitating **C**orresponding **E**mbodiments), a generic framework for solving the correspondence problem. The ALICE framework builds up a library of actions from the repertoire of an imitator agent that can be executed to achieve corresponding actions, states and effects to those of a model agent (according to given metrics and granularity); it provides a functional architecture that informs the design of robotic systems that can learn socially from a human demonstrator. [Alissandrakis et al. 2002, 2003, 2004]

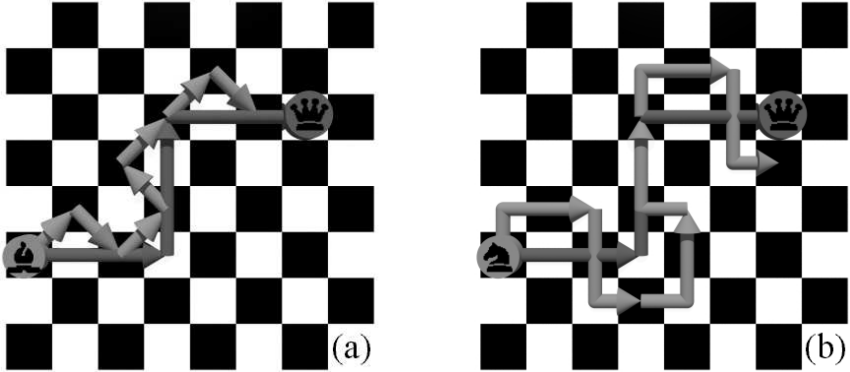
(a)
(b)

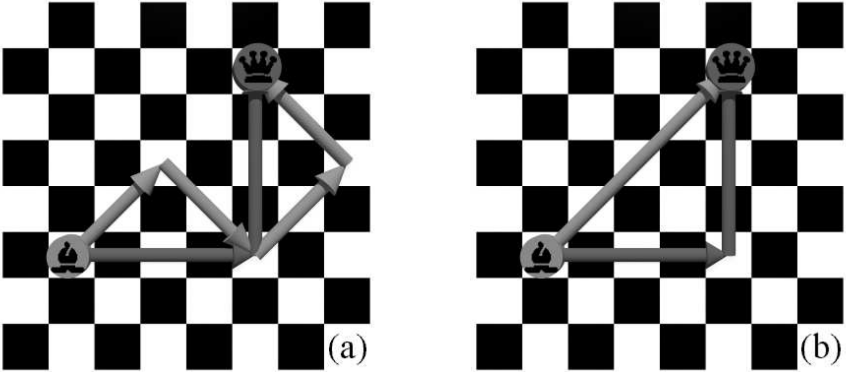
(a)
(b)

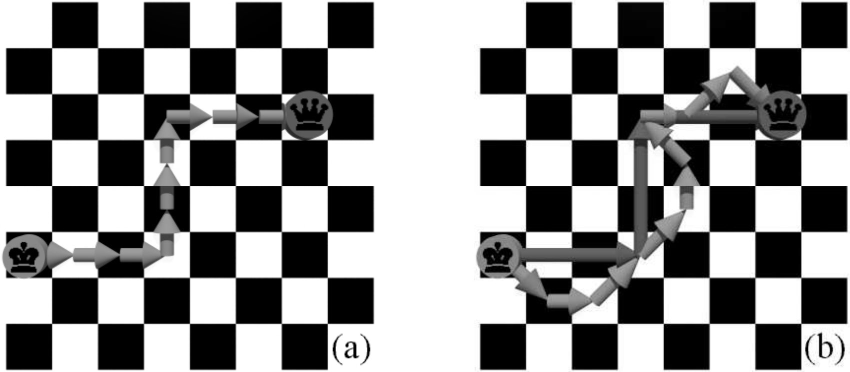
(a)
(b)

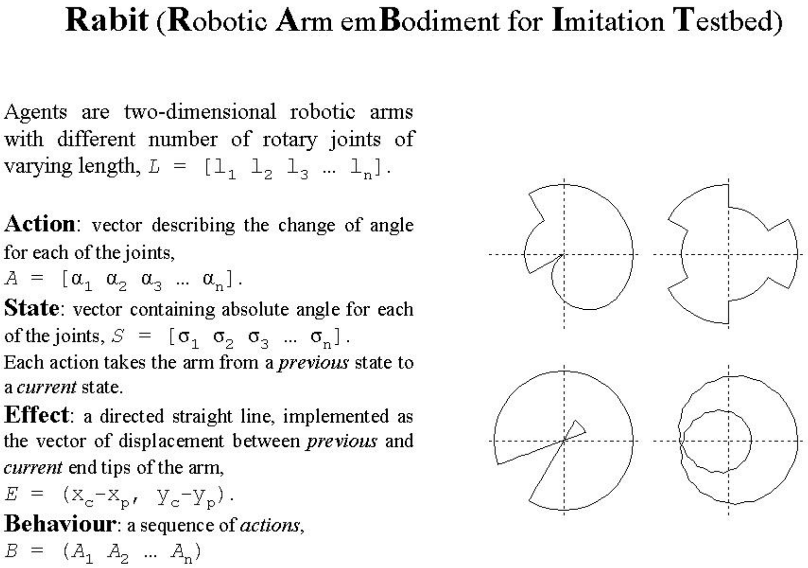

Rabit (**R**obotic **A**rm em**B**odiment for **I**mitation **T**estbed)

Agents are two-dimensional robotic arms with different number of rotary joints of varying length, $L = [l_1\ l_2\ l_3\ \dots\ l_n]$.

Action: vector describing the change of angle for each of the joints,
$A = [\alpha_1\ \alpha_2\ \alpha_3\ \dots\ \alpha_n]$.

State: vector containing absolute angle for each of the joints, $S = [\sigma_1\ \sigma_2\ \sigma_3\ \dots\ \sigma_n]$.
Each action takes the arm from a *previous* state to a *current* state.

Effect: a directed straight line, implemented as the vector of displacement between *previous* and *current* end tips of the arm,
$E = (x_c - x_p,\ y_c - y_p)$.

Behaviour: a sequence of *actions*,
$B = (A_1\ A_2\ \dots\ A_n)$

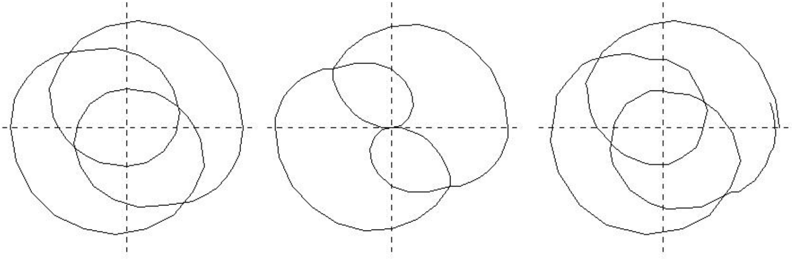

(original) Model
L= [20 20 20]

(intermediate) Model
L= [30 30]

Imitator
L=[20 20 20]

An example of cultural transmission among heterogeneous agents.

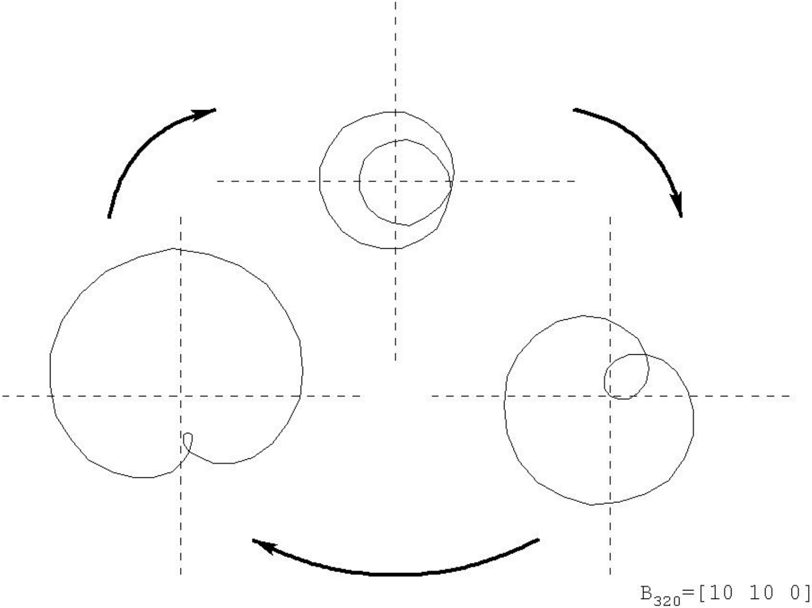
B_{320}=[10 10 0]

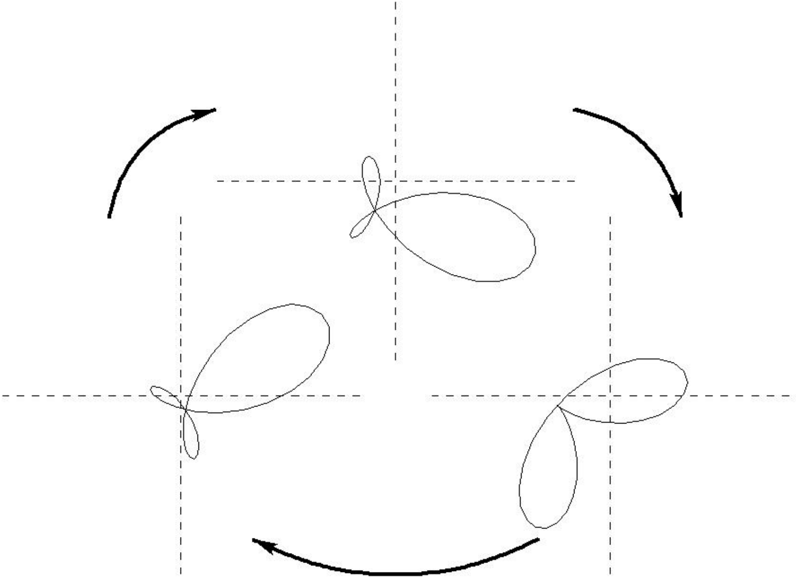

B_{415}=[10 -10 -10 -10]

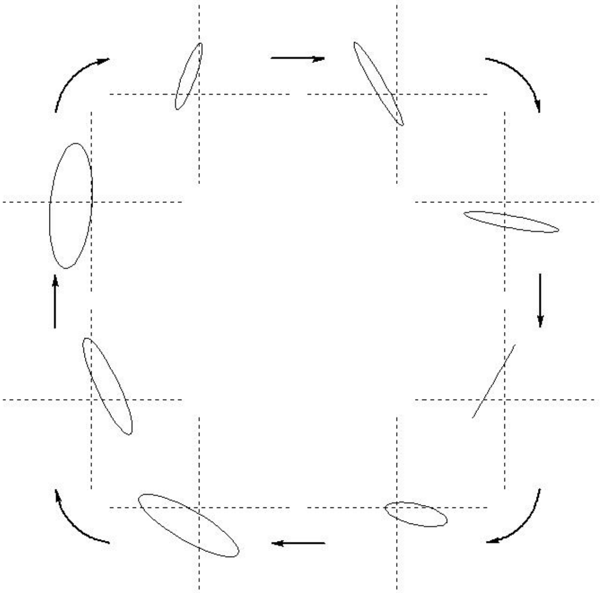

B_{415}=[10 -10 -10 0]

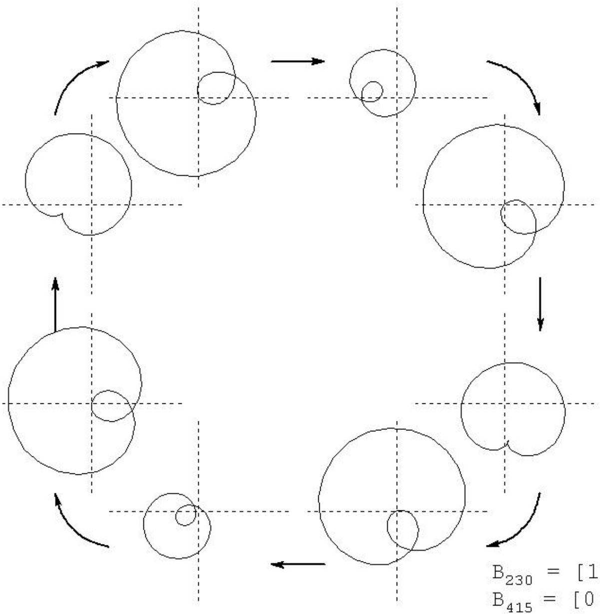

$B_{230} = [10\ 10]$
$B_{415} = [0\ 10\ 0\ 10]$

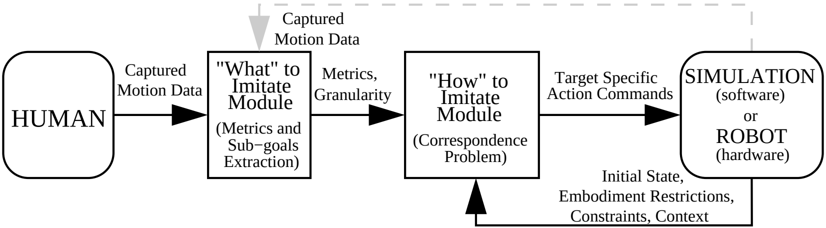

Using data captured from a human and given appropriate metrics and sub-goal granularity, the multi-target system can produce action command sequences that when executed by a software or hardware agent can achieve corresponding actions, states and/or effects. The corresponding actions, states and effects as demonstrated by the imitator can also be captured and used as a demonstration for another imitating agent. Differently embodied and constrained target systems in various contexts need to be supported. [Alissandrakis et al. 2005]

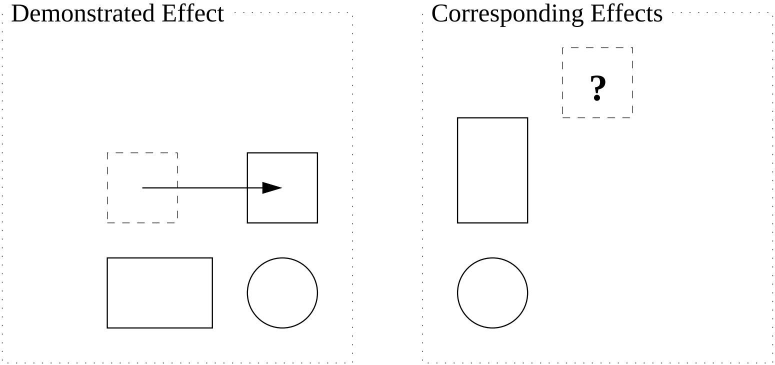

Given the demonstration, in a different context (here initial configuration), where on the workspace should the object be moved to?

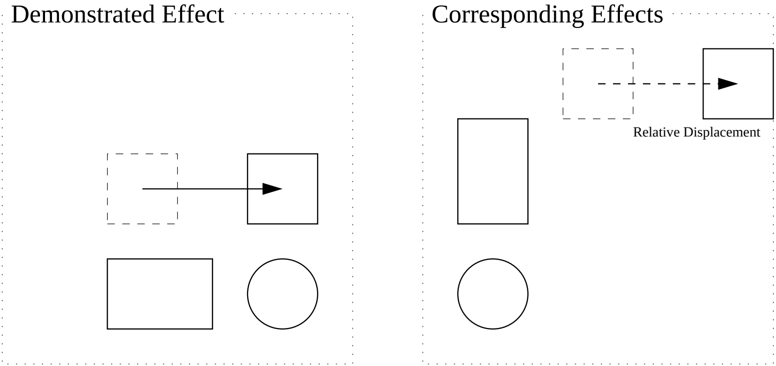

Should the same *relative displacement* be achieved?

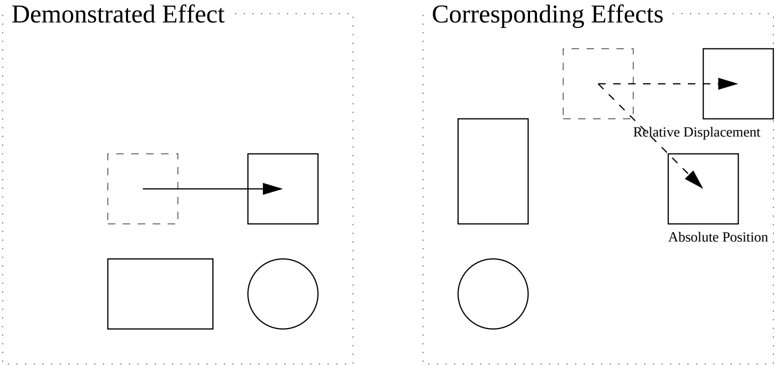

Should the same *absolute position* on the workspace be reached?

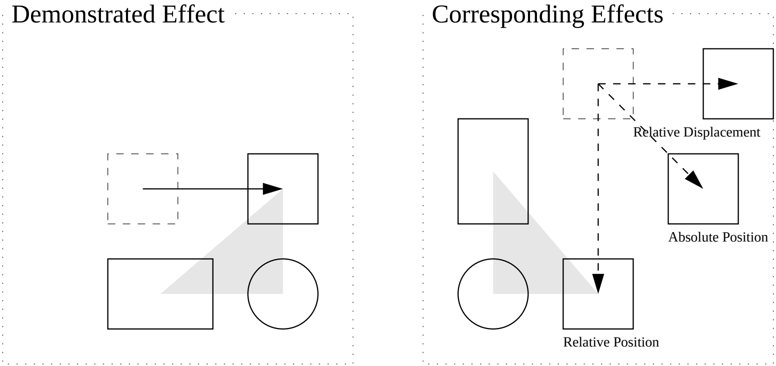

Or should the same *relative* (to the other objects) *position* be reached?

. . . complement the JABBERWOCKY work (towards a characterization of the space of effect metrics) with work on the formalization of body mappings via correspondence matrices [Alissandrakis et al. 2006].

Different agent bodies can be described as simplified kinematic models, comprising of a rooted acyclic connected graph of *segments*.

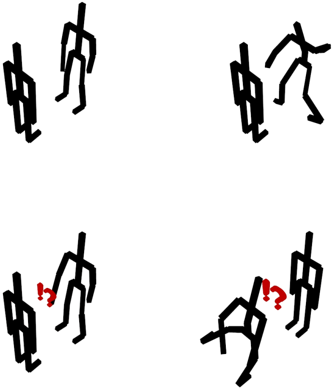

State

(or posture) of such a kinematic model can be defined the vector S containing the values of the degrees of freedom (DOF).

Action

can be thought as motion, or the amount of change in the DOFs required so that a posture transforms to another, and can be defined as the difference between two consecutive state vectors S and S': $A = S' - S$.

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.

State

(or posture) of such a kinematic model can be defined the vector S containing the values of the degrees of freedom (DOF).

Action

can be thought as motion, or the amount of change in the DOFs required so that a posture transforms to another, and can be defined as the difference between two consecutive state vectors S and S': $A = S' - S$.

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.

State

(or posture) of such a kinematic model can be defined the vector S containing the values of the degrees of freedom (DOF).

Action

can be thought as motion, or the amount of change in the DOFs required so that a posture transforms to another, and can be defined as the difference between two consecutive state vectors S and S': $A = S' - S$.

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.

A first global[1] **state metric** can be defined as

$$\mu_{state} = \sum_{j=1}^{n} |S_j^{\alpha} - S_j^{\beta}|,$$

where S_j^{α} and S_j^{β} are the values of the state vectors for the two agents.
Similarly, a first global **action metric** can be defined as

$$\mu_{action} = \sum_{j=1}^{n} |A_j^{\alpha} - A_j^{\beta}|,$$

where A_j^{α} and A_j^{β} are the values of the action vectors for the two agents.

---

[1] Here, 'global' implies that both agents have the *same embodiment*, i.e. the same morphology and number of corresponding DOFs.

A first global[1] **state metric** can be defined as

$$\mu_{state} = \sum_{j=1}^{n} |S_j^{\alpha} - S_j^{\beta}|,$$

where S_j^{α} and S_j^{β} are the values of the state vectors for the two agents. Similarly, a first global **action metric** can be defined as

$$\mu_{action} = \sum_{j=1}^{n} |A_j^{\alpha} - A_j^{\beta}|,$$

where A_j^{α} and A_j^{β} are the values of the action vectors for the two agents.

---

[1] Here, 'global' implies that both agents have the *same embodiment*, i.e. the same morphology and number of corresponding DOFs.

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

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

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.

By multiplying an appropriate correspondence matrix $\mathcal{C}$ with the state and action vectors S^α and A^α of the demonstrator respectively, two new vectors in imitator coordinates can be produced:

$$\mathcal{S} = S^\alpha \times \mathcal{C}$$
$$\mathcal{A} = A^\alpha \times \mathcal{C}$$

Combining the global metric definitions with the induced vectors gives

$$\mu^{\mathcal{C}}_{state} = \sum_{j=1}^{m} |\mathcal{S}_j - S^{\beta}_j| \epsilon_j$$

$$\mu^{\mathcal{C}}_{action} = \sum_{j=1}^{m} |\mathcal{A}_j - A^{\beta}_j| \epsilon_j$$

where the corrective term

$$\epsilon_j = \begin{cases} 0 \text{ if } \sum_{i=1}^{n} w^2_{i,j} = 0 \\ 1 \text{ otherwise} \end{cases}$$

takes the value zero if the j^{th} column of the correspondence matrix contains only zeros (effectively omitting the imitator's j^{th} DOF).

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

Intuitively, the components of $\mathcal{S}$ and $\mathcal{A}$ (for such $\epsilon_j \neq 0$) can be thought as the current subgoal state and action target values.

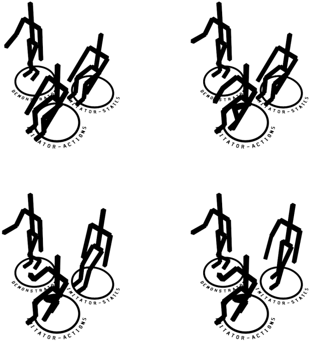

Identity mapping.

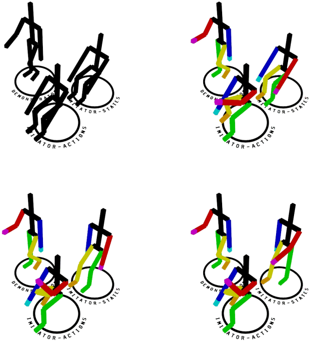

Mirror (right → left, controlateral) mapping.

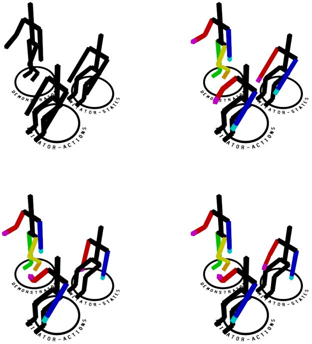

Partial human → human mapping.

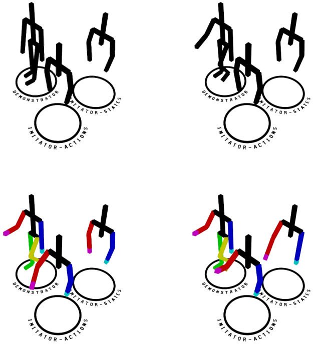

Human → upper torso mapping.

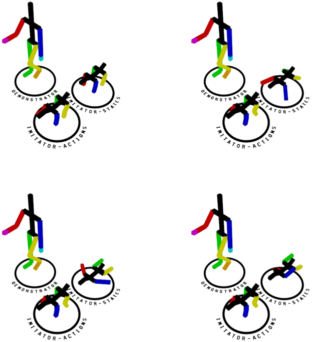

Human → dog mapping.

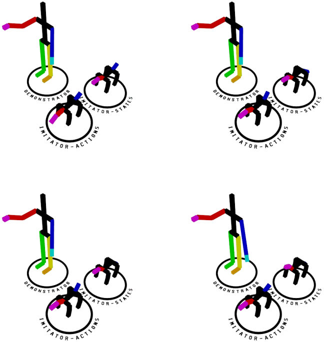

"Puppeteer" human → dog mapping.

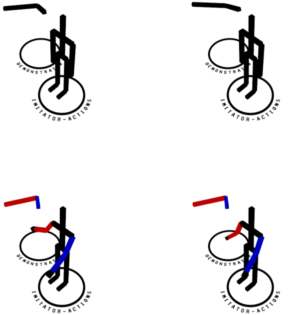

“Abstract human” → human mapping.

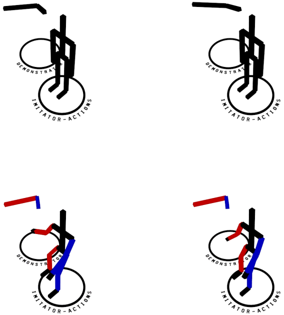

A different “abstract human” → human mapping.

ROSSUM - **RO**bot **S**elf-imitation and **S**caffolding **U**tility **M**echanism

[Saunders et al. 2006a, 2006b]

A robot social learning architecture that uses self-imitation with task and environmental scaffolding; the architecture is based on ideas from cognitive development with examples from social animals. *Zone of Proximal Development* (ZPD) - the gap between what the child can learn unaided and what can be learnt with the help a more capable peer - this guiding and support is called **scaffolding**. Vygotsky argued[2] that the learner learns based on their own sensorimotor experiences, *their own activity is at the centre of the learning process* (**self-imitation**). The approach is closest to that of *Extended Ideomotor Theory*[3]: *Similarity* between an event *perceived* and an event *learned* from the imitators own actions - *will induce that event*.

---

[2] J.N.Wertsch (1985) - Vygotsky and the social formation of mind.

[3] Wolfgang Prinz (1995-2005), William James (1890)

ROSSUM - **RO**bot **S**elf-imitation and **S**caffolding **U**tility **M**echanism
[Saunders et al. 2006a, 2006b]

A robot social learning architecture that uses self-imitation with task and environmental scaffolding; the architecture is based on ideas from cognitive development with examples from social animals. *Zone of Proximal Development* (ZPD) - the gap between what the child can learn unaided and what can be learnt with the help a more capable peer - this guiding and support is called **scaffolding**. Vygotsky argued[2] that the learner learns based on their own sensorimotor experiences, *their own activity is at the centre of the learning process* (**self-imitation**). The approach is closest to that of *Extended Ideomotor Theory*[3]: *Similarity* between an event *perceived* and an event *learned* from the imitators own actions - *will induce that event*.

---

[2] J.N.Wertsch (1985) - Vygotsky and the social formation of mind.

[3] Wolfgang Prinz (1995-2005), William James (1890)

ROSSUM - **RO**bot **S**elf-imitation and **S**caffolding **U**tility **M**echanism
[Saunders et al. 2006a, 2006b]

A robot social learning architecture that uses self-imitation with task and environmental scaffolding; the architecture is based on ideas from cognitive development with examples from social animals. *Zone of Proximal Development* (ZPD) - the gap between what the child can learn unaided and what can be learnt with the help a more capable peer - this guiding and support is called **scaffolding**. Vygotsky argued[2] that the learner learns based on their own sensorimotor experiences, *their own activity is at the centre of the learning process* (**self-imitation**). The approach is closest to that of *Extended Ideomotor Theory*[3]: *Similarity* between an event *perceived* and an event *learned* from the imitators own actions - *will induce that event*.

---

[2] J.N.Wertsch (1985) - Vygotsky and the social formation of mind.

[3] Wolfgang Prinz (1995-2005), William James (1890)

ROSSUM - **RO**bot **S**elf-imitation and **S**caffolding **U**tility **M**echanism
[Saunders et al. 2006a, 2006b]

A robot social learning architecture that uses self-imitation with task and environmental scaffolding; the architecture is based on ideas from cognitive development with examples from social animals. *Zone of Proximal Development* (ZPD) - the gap between what the child can learn unaided and what can be learnt with the help a more capable peer - this guiding and support is called **scaffolding**. Vygotsky argued[2] that the learner learns based on their own sensorimotor experiences, *their own activity is at the centre of the learning process* (**self-imitation**). The approach is closest to that of *Extended Ideomotor Theory*[3]: *Similarity* between an event *perceived* and an event *learned* from the imitators own actions - *will induce that event*.

---

[2] J.N.Wertsch (1985) - Vygotsky and the social formation of mind.

[3] Wolfgang Prinz (1995-2005), William James (1890)

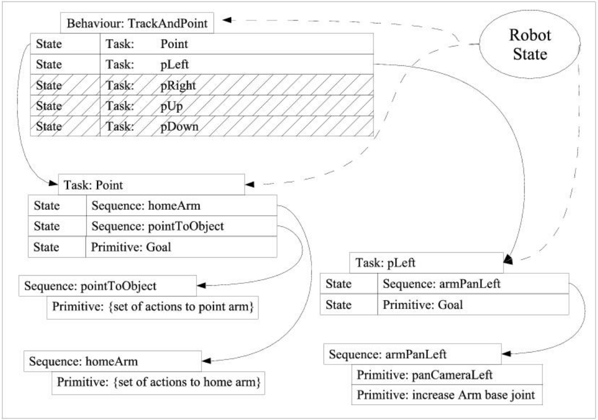

The architecture *environmental scaffolding* to select relevant attributes and *task scaffolding* to build competencies.

The robot learning process is platform independent, constructed during learning and extensible.

A *prediction mechanism* is used to inform trainer of existing competencies.

Thank You!

Please check `http://homepages.feis.herts.ac.uk/~comqaa1/` and feel free to send an email to `a.alissandrakis@herts.ac.uk` for further information.

- Alissandrakis, A., Nehaniv, C. L. and Dautenhahn, K.: 2002, Imitation with ALICE: Learning to imitate corresponding actions across dissimilar embodiments, IEEE Trans. Systems, Man & Cybernetics: Part A 32(4), 482–496.
- Alissandrakis, A., Nehaniv, C. L. and Dautenhahn, K.: 2003, Synchrony and perception in robotic imitation across embodiments, Proc. IEEE International Symposium on Computational Intelligence in Robotics and Automation (CIRA'03), pp. 923–930.
- Alissandrakis, A., Nehaniv, C. L. and Dautenhahn, K.: 2004, Towards Robot Cultures? - Learning to Imitate in a Robotic Arm Test-bed with Dissimilar Embodied Agents, Interaction Studies: Social Behaviour and Communication in Biological and Artificial Systems 5(1), 3–44.
- Alissandrakis, A., Nehaniv, C. L. and Dautenhahn, K.: 2006, Action, State and Effect Metrics for Robot Imitation, 15th IEEE International Symposium on Robot and Human Interactive Communication (RO-MAN 06), University of Hertfordshire, Hatfield, United Kingdom, 6-8 September, 2006, pp. 232–237.
- Alissandrakis, A., Nehaniv, C. L., Dautenhahn, K. and Saunders, J.: 2005, Achieving Corresponding Effects on Multiple Robotic Platforms: Imitating in Context Using Different Effect Metrics, Proceedings of the Third International Symposium on Imitation in Animals & Artifacts, The Society for the Study of Artificial Intelligence and Simulation of Behaviour, pp. 10–19.

- Dautenhahn, K. and Nehaniv, C. L.: 2002, An agent-based perspective on imitation, in K. Dautenhahn and C. L. Nehaniv (eds), Imitation in Animals and Artifacts, MIT Press, pp. 1–40.
- Mitchell, R. W.: 1987, A comparative-developmental approach to understanding imitation, in P. P. G. Bateson and P. H. Klopfer (eds), Perspectives in Ethology 7: Alternatives, Plenum Press, pp. 183–215.
- Nehaniv, C. L. and Dautenhahn, K.: 2000, Of hummingbirds and helicopters: An algebraic framework for interdisciplinary studies of imitation and its applications, in J. Demiris and A. Birk (eds), Interdisciplinary Approaches to Robot Learning,World Scientific Series in Robotics and Intelligent Systems, pp. 136–161.
- Nehaniv, C. L. and Dautenhahn, K.: 2001, Like me? - measures of correspondence and imitation, Cybernetics and Systems 32(1-2), 11–51.
- Nehaniv, C. L. and Dautenhahn, K.: 2002, The correspondence problem, in K. Dautenhahn and C. L. Nehaniv (eds), Imitation in Animals and Artifacts, MIT Press, pp. 41–61.
- Saunders, J., Nehaniv, C. L. and Dautenhahn, K.: 2006a, Teaching Robots by Moulding Behavior and Scaffolding the Environment, 1st Annual Conference on Human-Robot Interaction (HRI2006), Salt Lake City, Utah, USA, March 2-4, 2006, pp. 142–150.

- Saunders, J., Nehaniv, C. L. and Dautenhahn, K.: 2006b, Using Self-Imitation to Direct Learning Special Session on Learning Through Imitation, 15th IEEE International Symposium on Robot and Human Interactive Communication (RO-MAN 06), University of Hertfordshire, Hatfield, United Kingdom, 6-8 September, 2006, pp. 244–250.
- Thorndike, E. L.: 1898, Animal intelligence: An experimental study of the associative process in animals, Psychol. Rev. Monogr. 2, 551–553.
- Thorpe, W. H.: 1963, Learning and Instinct in Animals, London: Methuen.