

Influencing Robot Learning Through Design and Social Interactions: A Balancing Framework

Yuval Marom

School of Computer Science
and Software Engineering
Monash University
Victoria 3800, AUSTRALIA
yuvalm@csse.monash.edu.au

Gillian Hayes

School of Informatics
University of Edinburgh
Edinburgh EH9 3JZ, Scotland, UK
gmh@inf.ed.ac.uk

Abstract

We present a framework for addressing a challenging trade-off between influencing the learning of a robot through design and through social interactions. We identify different kinds of influences that a designer can introduce at design time, and that an expert can introduce using social interactions, and we use these to characterise a two-dimensional design space. As well as discussing how the two sources of influence affect each other, we propose how learning performance typically varies as a result, and present some empirical findings.

1. Introduction

One of the challenges in epigenetic robotics is finding a balance between designing innate knowledge and enabling robots to learn and acquire knowledge through development and interaction. On the one hand programming a robot through design is desirable because (if done correctly) it makes for a robust and reliable control architecture. On the other hand programming a robot through situated interaction makes for a more general and adaptive control architecture. Addressing this trade-off is difficult because robots require sufficient reliable control, but also sufficient flexibility to adapt to noisy and changing environments.

A robot can acquire knowledge through *individual* physical interactions with the environment, but this would require it to have some internal rewarding mechanism that favours certain sensorimotor experiences over others. Programming such a mechanism can be as difficult, inaccurate, and time-consuming as programming the robot to perform the task in the first place. Further, it might take the robot a long time to go through all the possible experiences and discover those relevant to the task. Alternatively (or additionally) the robot can acquire knowledge through *social* interactions with another,

more experienced agent situated in the environment, who takes the role of exposing the robot to the relevant experiences. Our work involves this latter approach, which has been widely recognised in the literature as providing task-relevance to a learner robot and speeding-up learning (Demiris and Hayes, 1996; Schaal, 1999; Mataric, 2000; Gaussier et al., 1998).

Thus our work addresses the trade-off mentioned above in situations where a robot learns a task from an experienced agent — an expert, robotic or human. We address the trade-off explicitly by identifying increasing amounts of influence on the robot’s learning that can be applied by both the designer and the (social) expert, and then characterising these sources of influence as a two-dimensional design space, shown in Figure 1. Our aim in this paper is to propose the design space, and how learning performance typically varies within it.

More specifically, the kind of influence we are referring to is related to biasing the robot’s notion of *saliency*. By ‘saliency’ we mean the level of granularity at which significant differences in the sensorimotor data are assessed, and we argue that this notion is dependent on the particular task to be learned and the environment it is learned in. For example, if the task involves manipulating objects, the robot needs to detect finer differences between objects than if the task involves pushing these objects. Of course, the notion of saliency also depends on other factors such as the robot’s morphology (*e.g.* sensor capabilities) and its learning architecture. It is important to stress that we consider such other factors as making up the robot’s *existing* capabilities, and we are interested in how the robot’s learning can be influenced *given* its existing capabilities.

2. Influencing Robot Learning Through Design

Finding a desirable amount of influence from the designer corresponds to the well known and difficult problem in Artificial Intelligence of choosing a level of abstraction (see, for example, Marr, 1982). A learn-

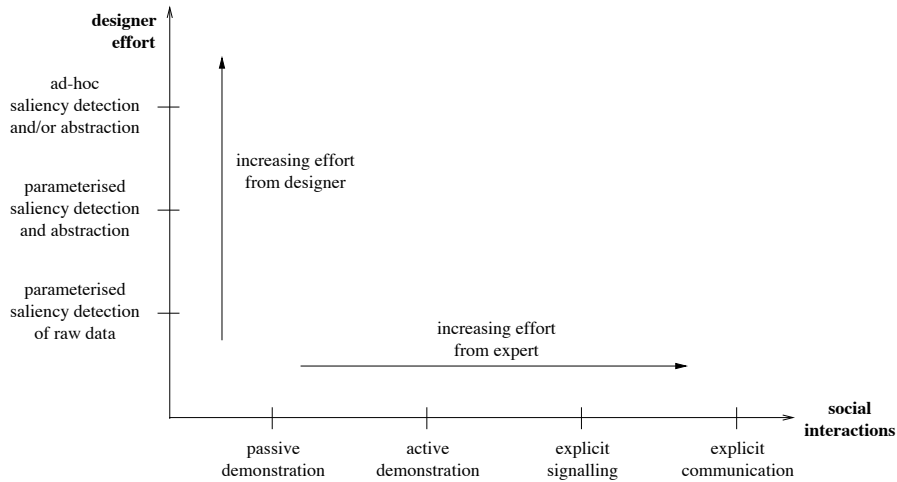


Figure 1: Levels of social interactions and designer effort. The designer effort axis corresponds to different levels of effort from the designer, required for learning at different levels of abstraction; the social interactions axis specifies interactions of different complexities, corresponding to different levels of effort by a social expert. We argue that the expert’s effort can be used to balance the effort required in programming the robot at design time.

ing system might be set up where the robot learns from its low-level ‘raw’ sensorimotor data (again, the designer could also design what constitutes ‘raw’ data by manipulating the robot’s morphology, but such design is not considered here), or from some higher-level abstraction of the data. For example, a robot could learn a phototaxis task by learning to react to the *continuous* data from its light-detecting sensors, or it could learn to react to pre-defined *ranges* of these sensors (*e.g.* low, mid, high), specifically chosen by the designer to work well for this task in a particular environment.

When real robots are involved, learning from low-level raw data is difficult due to their unstructured and noisy nature, so either or both of some kind of saliency detection and abstraction is necessary. The vertical axis in Figure 1 refers to the amount of influence from the designer in abstracting the sensorimotor data with which the robot learns a task. This could correspond, for example, to setting saliency parameter values, setting parameters for a self-organising abstraction mechanism, or manually pre-structuring the data as in the example above. The latter is what we mean by ‘ad-hoc’ abstraction in Figure 1. Ad-hoc saliency detection corresponds to a detection component (for example, the sensor-value that signals a change of activation) being embedded rigidly in the control architecture.

Increasing points on this axis correspond to more ‘effort’ given by the designer to make the learning work well for the *particular* purpose in mind — that is, more biased and less flexible to learn other tasks in other situations. We are using the word ‘effort’ loosely to refer to the kind of design activities involved in fine-tuning a control architecture to suit very specific needs. Thus, for example, an increas-

ing effort can correspond to spending more time on trial-and-error programming.

3. Influencing Robot Learning Through Social Interactions

In a recent survey, Fong et al. (2003) demonstrate a wide range of robotic and software systems that utilise social interactions of various types, and where different design issues are addressed. For example, robots that are capable of expressing emotions must be able to communicate through speech and facial expressions; in robots where embodiment is important the morphology of the robot must be carefully designed; and in robots where human-centered perception is important the robot must be able to track people, and perform speech, gesture, and face recognition. These distinctions are made with reference to the capabilities of the robot (or software agent).

In contrast, the different types of social interactions identified here refer to different capabilities *of the expert*. More precisely, they refer to different ways in which the expert can interact with the robot to *purposely* influence the robot’s learning, and we claim these different interactions to have increasing complexities — they require different levels of ‘effort’ by the expert. Four types of interactions are identified, and will be discussed below: passive demonstrations, active demonstrations, explicit signalling, and explicit communication. In the first two the teaching is implicit — there is no explicit transfer of information between the teacher and the learner; the teacher merely demonstrates a task, and the learner learns in terms of its own experiences. In the latter two types of social interactions the demonstrations are enhanced with an explicit one-directional influence from the teacher on the learning, or with

two-directional exchange of information between the teacher and learner, respectively.

With regards to design issues, the relevant ones identified in the survey by Fong et al. (2003) are those concerning human-centered perception. The different types of social interactions mentioned above rely on the fact that the robot has the appropriate mechanisms for these interactions. That is, it must be able to track and copy the actions of the teacher who is demonstrating a task, and it must be able to receive and send explicit information, as required. Further, these capabilities should be task-independent, that is, they should be useful for interacting with the expert regardless of what task is involved. As long as this is the case, the learning setup can be applied to different situations. Then, the effort from the expert can be used to balance the effort from the designer.

As with designer effort, the word ‘effort’ attributed to the social expert quantifies different activities used by the expert to bias the robot’s learning to particular needs, either directly or indirectly. As this influence is transferred from the designer to the social expert who is situated in the environment together with the learner, the resulting learning setup is more general and adaptive to different tasks.

3.1 Increasing Complexities of Social Interactions

Passive Demonstrations

The minimal amount of effort required by an expert wishing to teach some task to a learner robot is to execute the task as if there is no learner. In other words, the expert ‘demonstrates’ the task *passively* and independently of the states or actions of the learner. In the early work on learning by imitation (Hayes and Demiris, 1994), this kind of minimalistic effort from the expert is actually promoted as one of the advantages of programming robots through demonstration, especially when the demonstrator is another robot, and one wishes to capitalise on its existing knowledge in training another robot with the least effort. This is indeed an advantage if learning from such passive demonstrations is possible. However, a passive demonstration is not always a sensible demonstration strategy because, for example, the ability of the learner robot to learn could be hampered if it loses sight of the teacher or if it struggles to copy the actions of the teacher. Indeed, when the expert is a human, he/she inevitably takes more care in demonstrating the task, whereas examples of passive demonstrations are generally attributed to robotic demonstrators.

This problem of the learner having the ability to match the actions of the teacher is very challenging, especially when we consider that the learner and teacher can have different morphologies (see, for example, the ‘correspondence problem’, formulated by

Nehaniv and Dautenhahn, 2000). However in the majority of related work, the researchers design their learners and teachers in such a way that they can bypass this problem, and therefore ensure that the learner is able to copy the actions appropriately for the particular task. In contrast, Alissandrakis et al. (2000) address this issue with simulated agents of different morphologies that imitate each other (this work is revisited later in the paper).

Active Demonstrations

There are at least three ways in which the expert might tailor the demonstration according to the states or actions of the learner, and therefore demonstrate the task more actively. Firstly, the teacher could adapt the demonstrations in order to make it easier for the learner to match the teacher’s actions, for example, by slowing down the demonstration. In the mobile robot experiments by Billard and Hayes (1999), a learner robot follows a teacher around an environment, but the teacher can also detect the learner and align itself in front of it, thus reducing the possibility that the learner loses the teacher.

Secondly, the teacher might perform the demonstration in such a way as to ensure that not only does the learner not get lost, but that it is actually exposed to ‘clean’, consistent, and distinct experiences. By observing and inferring the action-copying behaviour of the learner, the teacher can manipulate the learner’s experiences. For example, Gaussier et al. (1998) report that in physical experiments involving a human teaching a mobile robot various ‘dances’, the demonstrations are inevitably more adaptive than in similar simulated experiments involving a simulated teacher; the human teacher ensures the learner passes exactly through correct edges in the trajectories of the dance, and that the timings of the learner’s actions are precise, by adapting his (the teacher’s) own trajectory and speed.

The third way in which the expert can influence the demonstrations is by deviating from the ‘natural’ demonstration in order to ‘exaggerate’, or accentuate, the important differences between the components of the task. These kinds of interactions are not reported in related work, although a human demonstrator could be performing such demonstrations without realising it. We give an example in Section 5.

Kaplan et al. (2001) suggest other active demonstration methods, which are inspired from techniques used by humans to train animals, especially dogs. For example, they suggest more physical interactions (termed ‘modelling’, or ‘moulding’), where the trainer physically manipulates the animal into the desired positions. In robotics, this could correspond, for example, to manipulating the robot with a joystick (Kaiser and Dillmann, 1996; Hugues and Drogoul, 2001).

These various kinds of active demonstrations are difficult to program for a robotic teacher, whereas for a human teacher such demonstrations are not only easier, but also more intuitive and adaptive. The reason for this is that a human teacher is situated in the environment *together* with the learner, *while* the learner is learning, and can therefore tailor the interactions ‘on-line’ in response to what the robot is doing. In contrast, in order to program an equivalent robotic teacher, the designer would have to guess what would be a good active strategy *before* the interactions begin. Also, it can be argued that human demonstrations are inevitably adaptive to the robot’s interests.

Explicit Signalling

In the examples of social interactions presented so far, there has been no explicit transfer of information between the expert and the learner. There is, however, an *implicit* transfer of information, because the learner learns to perform a novel task through the social interactions. That is, because the learner’s actions are influenced by the demonstrator’s, and the learner learns from these situated actions, information can be thought of as being transmitted *indirectly* through the environment. If the expert has the ability to send direct explicit signals to the learner as well as demonstrate the task, and the learner has the ability to detect and interpret these signals, this more complex type of interaction can have various uses.

The signals from the expert can form part of the stimulus that the learner learns from, with the aim of learning a symbolic representation of the sensorimotor data, such as a language for communication (Billard and Hayes, 1999; Kaplan et al., 2001). In other approaches where symbolic learning is not required, and thus the signals do not form part of the learning, the signals can still be used to directly *affect* the learning. One purpose of such signalling is to explicitly draw the learner’s attention to salient experiences (Moukas and Hayes, 1996; Nicolescu and Matarić, 2003), and another purpose is to provide the learner with feedback about its actions. This latter kind of signalling is generally used when it is the sole source of social interactions, that is, when demonstrations are not available; instead, the learner already has basic sensorimotor skills, and the expert teaches a task utilising these skills by rewarding the relevant ones (Nehmzow and McGonigle, 1994; Kaplan et al., 2001).

Explicit Communication

With explicit signalling (mentioned above), the expert sends the learner signals that influence the learning directly and explicitly. When sending these kinds of signals, it might be useful for the expert to

know how the learning is being influenced, for example how the learning is progressing. If the learner could send signals back to the expert about its internal states, for example how familiar experiences are, then the expert could use such information to determine how to proceed with the demonstration and signalling. For example, if a particular experience is not familiar enough then the expert might demonstrate and signal it more frequently. Explicit signals from the expert might also be used by the robot to tune saliency parameters that determine how it perceives and learns from experiences, in which case it might be beneficial for the expert to have an idea of the effect that the signals have on such parameters.

Kaplan et al. (2001) utilise a combination of explicit (verbal) and implicit (non-verbal) communication between an expert and a robot to refine the robot’s learned sequence of behaviours. The robot demonstrates the sequence of behaviours that it has just learned (non-verbal communication), which might include some irrelevant behaviours, and the expert only rewards (verbal communication) the relevant behaviours; the robot then updates its internal measure of similarity between behaviours which influences their sequencing. Klingspor et al. (1997) identify different types of verbal and non-verbal communication strategies, with which a robot gives feedback to its user about its perceptions and actions when it executes a learned task. However, they do not discuss how this feedback is used by the user to influence further demonstrations. These two approaches use communication to refine a behaviour which is already learned, whereas communication is proposed here as an on-line approach for influencing the learning *while the robot is learning*.

In our work we have implemented all but this type of social interactions. However, we believe that the most influence an expert can have on a robot’s learning is through explicit communication.

3.2 Using Social Interactions to Balance Designer Effort

We argue that in order for the stronger types of social interactions to have the kind of influence on the learning that we are addressing here, that is, an influence on the robot’s notion of saliency, saliency must be treated explicitly and in a parameterised fashion. If the learning setup is made flexible enough at design time, then it can be the responsibility of the social expert to bias the robot’s learning to the particular task. This is particularly important for the latter two types of social interactions, which influence saliency more explicitly.

In the following examples from the literature, the notion of saliency is set flexibly through the use of different levels of granularity to control the learning. We have already mentioned the work of Alissandrakis et al. (2000), involving agents of different morpholo-

gies, who therefore cannot copy each other’s actions exactly. Using different imitation strategies they attempt to either copy exact movements, selected landmarks in the movements, or simply the end position, each involving imitating at different levels of granularity.

Similarly, in the work by Billard et al. (2003), the robot is equipped with a number of imitation strategies that are applicable for recognising demonstrations at different levels of granularity. For example, the task might involve only moving a specific type of object in any direction and using any hand, moving any box in a specific direction, moving the boxes in a particular sequence, or moving the boxes always with the same hand-box relationship (*e.g.* always the left hand, or always the hand closest to the object). By demonstrating the particular task a number of times, the demonstrator is able to highlight those features that are salient (or ‘time-invariant’, as the authors refer to them).

Another example is the work by Gaussier et al. (1998) mentioned earlier, where the architecture is set up to detect saliency at different levels of granularity, as controlled through a number of time constants, or parameters. However, it is unclear if and how the demonstrations influence these parameters.

4. A Balancing Framework

The identification of different complexities of social interactions and designer effort is crucial for investigating the interactions between them. Such an investigation is missing in current research, but is necessary in order to address the trade-off mentioned at the start of this paper: balancing between the influences from the designer on the robot’s learning at design time, and the influences from a social expert during situated interactions. The designer must guess the learning dynamics at design time, while the social expert is situated in the same environment and performs the same task while the robot is learning the task. Thus a desirable balance is one where the learning set up is sufficiently robust, but also sufficiently general, adaptive, and faithful to the robot’s experiences.

The two-dimensional design space shown in Figure 1 is proposed as a framework for characterising this balance. In (Marom, 2003), we show the usefulness of this characterisation, firstly by demonstrating how it applies to related work, secondly by presenting an empirical investigation of how performance varies as a function of the two dimensions, and lastly by considering the implications of this investigation to the related work. As mentioned in Section 1, our aim in this paper is to propose the design space, and how learning performance varies within it. For more details, discussion, and examples of related work, the reader is referred to (Marom, 2003).

Figure 2 shows our proposal for how performance

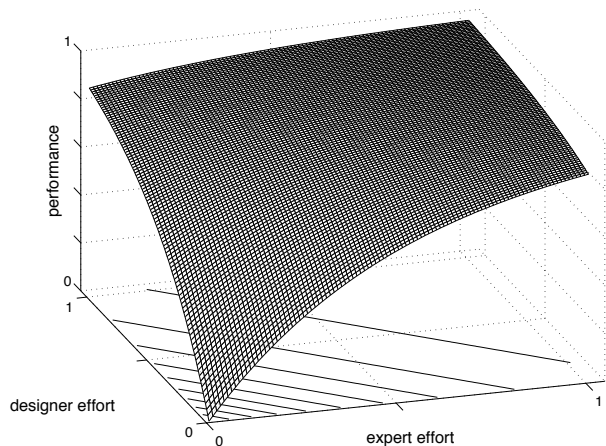


Figure 2: This performance surface characterises learning performance within the design space identified in Figure 1.

typically varies in the design space. This performance surface is a generalisation from our empirical results, and is intended purely as a qualitative demonstration — the actual shape chosen in the figure is arbitrary. The general features that should be noted from the figure are that the performance as a function of each dimension increases and levels out, and that one can maintain a particular level of performance (see projections on the horizontal plane) by compensating for some decrease in one dimension by some increase in the other.

We highlight the following issues with regards to the performance surface:

- Designer effort can be balanced with more effort from an expert during social interactions.
- Performance can be improved by increasing either of the two types of effort.
- Increasing the amount of effort of one type reduces the potential for improving the performance with the other type. For example, if the designer has spent a lot of effort in abstracting the sensorimotor data of the robot, there is very little improvement an expert can provide through active or explicit demonstrations.
- The points where performance converges depend on other design issues not considered by the characterisation of designer effort in this paper, such as the robot’s morphology and the learning architecture. For example, a particular choice of sensors, or a particularly good learning architecture that can learn from ‘raw’ data, can mean that there is very little abstraction the designer can introduce, and thus the performance surface levels out faster in the ‘designer effort’ axis.

This performance surface depicts our view that given a robot’s existing set of learning capabilities, it is always possible to improve the robot’s learning performance by increasing the influence or bias on the robot’s abstraction of its experiences; and further,



Figure 3: The robot Gillespie learns a wall-following task from a human demonstrator who actively demonstrates the task and explicitly influences the learning using hand signals.

that this influence can be transferred from design time to situated social interactions, depending on the interaction capabilities of the robot and the expert.

5. Empirical Findings

As mentioned previously, our aim in this paper is to propose the balancing framework, and so we only provide a brief account of a subset of our empirical findings. The most comprehensive and detailed account of this work is given in (Marom, 2003), and individual papers are also referenced below for some specific experiments.

The examples we present here involve our Real World Interface (RWI) B21 robot, Gillespie, learning a wall-following task from a human demonstrator (Figure 3). Gillespie is programmed to detect and follow the human using its on-board video camera through a simple colour-tracking algorithm — the demonstrator wears a green shirt which is easily detectable.

The input that the robot uses to learn the task with comes from its sonar sensors (*not* from the camera). Prior to learning, this input goes through an *attention* system. This attention system provides the explicit formalisation of saliency that was argued for in Section 3.2. It contains a saliency parameter for novelty detection, and it is also responsible for abstracting the robot’s experiences through self-organisation. This is a Kohonen network with habituating nodes, based on the Grow When Required algorithm (Marsland et al., 2002) — see (Marom et al., 2001) for more details.

The demonstrator demonstrates the task by moving around the arena following walls, and turning into the middle of the arena occasionally (4-5 times during a run) and making random turns. The demonstration strategy taken here is an ‘active demonstration’ one: the demonstrator faces the robot and adapts his movements to ensure the

robot’s tracking system is able to keep up (first aspect of active demonstrations discussed in Section 3.1); and further, he ‘exaggerates’ the differences between the components of the task (third aspect of active demonstrations discussed in Section 3.1), making large random turns while in the middle of the arena to distinguish them from making straight forward moves while parallel to a wall.

The demonstrator can also signal to the robot at very specific times during the task using a red glove, which is easily contrasted with the green shirt used in the teacher-following behaviour. The role of this signalling is to force novelty detection, thereby triggering the attention system (causing an update of the self-organising map), and the learning system. In effect, the signalling takes the role of the saliency parameter.

5.1 Learning at Different Levels of Abstraction

We present two learning setups, involving different amounts of designer effort, and show the implications for increasing the effort from the demonstrator. In both learning setups, the active demonstration strategy is utilised, so the increase in effort from the demonstrator corresponds to the signalling. It is evaluated by comparing performance when the signals are used as described above to trigger attention and learning, with performance when the signals are ignored.

In the first learning setup, the attention system is used only to detect novelty in the robot’s sonar data, and the robot subsequently learns through a separate learning system whenever novelty is detected. Learning involves associating the robot’s ‘raw’ sonar input with its motor values, using a Multi-Layer Perceptron (MLP) with back-propagation weight updates. For more details on these experiments see (Marom and Hayes, 2001). Here the designer effort corresponds to setting the value of the detection parameter (the lowest category on the vertical axis in Figure 1), and our experiments show the importance of this task by comparing the learning performance with different parameter values. The complete range of parameter values are tested, including a very insensitive value where the attention system never detects novelty, and a very sensitive value where the attention system always detects novelty.

In the second learning setup the self-organisation capability of the attention system is used directly in the learning, as follows. The output of the attention system is a discrete set of nodes (nodes in a Kohonen map), each node representing a region in the robot’s sensory space, discovered through self-organisation while the robot follows behind the demonstrator. Thus this set of nodes corresponds to an abstraction of the robot’s perceptions. The learning consists of associating these set of nodes with a set of pre-

defined basic motor skills. For more details on these experiments see (Marom et al., 2001, 2002). This setup involves learning at a higher level of abstraction, where there is an information loss from what is, in essence, a compression of the raw data. Therefore, relative to the first learning setup, where learning occurs at a lower level of abstraction, here more care must be taken by the designer in setting the saliency parameters (middle category on the vertical axis in Figure 1) in order to achieve a representation that is useful for the robot to learn the task with.

Both learning setups are evaluated in a separate recall phase, where Gillespie is placed in the arena on its own, and its ability to execute the learned task is measured numerically. The way the robot executes the learned task differs for the two setups. In the first, the sonar input is propagated through the learned MLP, and the output is converted to motor commands. In the second, the input activates one of the nodes in the self-organising map, and the associated motor action for that node is executed. The task evaluation measure is calculated from the robot’s sonar sensors, favouring particular sensed configurations (*e.g.* highest reward for sensing the wall parallel on either side). Upper and lower baseline scores are also calculated with the robot executing a hand-crafted wall-following behaviour, and a random wandering behaviour, respectively.

5.2 Comparative Performances

With the first learning setup, the best performance is achieved when the signals from the demonstrator are used. Further, the significance of the signalling increases as the signals start to dominate over the attention system in triggering the learning system. This is achieved when the detection parameter is set very insensitively, and consequently saliency detection is mainly triggered by the demonstrator. In contrast, in the second learning setup the robot learns equally well with and without the signals from the demonstrator, regardless of the value of the detection parameter.

The message to take from the comparison of these two learning setups is as follows:

- When much effort is given by the designer in abstracting the perception of the robot (second learning setup), there is little potential for improvement in performance through stronger social interactions. With a careful setting of the saliency detection parameter, the designer is able to achieve a desirable representation in the self-organising map to influence the learning.
- When the learning occurs at a lower level of abstraction (first learning setup), not forced by the designer, there is potential for stronger influences from social interactions. The responsibility of usefully influencing the robot’s learning can be transferred from the designer to the social expert.

In terms of the performance surface (Figure 2), these two experiments demonstrate that performance can be improved by increasing the influence from either the designer or the social expert, and that if the influence from the designer is high enough, the effect of the influence from the expert is diminished.

6. Discussion

This paper proposes a framework for addressing a balance between influencing the learning of a robot through design and through social interactions. The kind of influences it addresses are related to saliency, or level of granularity at which a robot learns a task. Other researchers have recognised that the issue of saliency and granularity should be made flexible at design, and instead determined during social interactions. However, the balance that we propose here is not formalised in related work.

We have identified different ways in which either a designer or a social expert can influence the learning of a robot, and then characterised a two-dimensional design space. We have also suggested how learning performance typically varies within the space, in other words, how the two sources of influence interact to affect performance. In particular, our characterisation shows that a particular level of performance can be maintained by using less designer effort and stronger social interactions, and thus a learning setup results that is less biased to a particular task, and more general and adaptive to different tasks.

The performance surface we have suggested is a qualitative generalisation of our various empirical findings, some of which we presented in this paper. Making such a generalisation is difficult, because different experiments involve different robots, tasks, learning architectures, and performance evaluation measures. Comparing experiments between different researchers is even more difficult, and sometimes the only possible comparison is a qualitative one. Such a qualitative framework cannot give an absolute indication of performance, but it can help make design decisions. For example, the ability to move along the horizontal axis of the design space requires that the social expert has a reasonably good knowledge of how the robot learns, which might not always be the case. Thus the choice of what kind of social interactions to use is a design decision that is based, among other things, on what kind of social expert is available. This can lead, in turn, to a decision on the level of abstraction to enforce through design.

Therefore, we believe the characterisation we have presented here can provide a useful framework for organising related work on robotic systems that learn from social interactions. Also, the characterisation can itself benefit from future work that will help fine-tune the different features it proposes, and expose more testing conditions (*e.g.* more categories in either of the two dimensions of the design space).

References

- Alissandrakis, A., Nehaniv, C. L., and Dautenhahn, K. (2000). Learning how to do things with imitation. In *AAAI Fall Symposium on Learning How to Do Things*.
- Billard, A., Epars, Y., Schaal, S., and Cheng, G. (2003). Discovering imitation strategies through categorization of multi-dimensional data. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS'2003)*.
- Billard, A. and Hayes, G. (1999). DRAMA, a connectionist architecture for control and learning in autonomous robots. *Adaptive Behaviour Journal*, 7(1):35–64.
- Demiris, J. and Hayes, G. (1996). Imitative learning mechanisms in robots and humans. In *Proceedings of the 5th European Workshop on Learning Robots*, pages 9–16.
- Fong, T., Nourbakhsh, I., and Dautenhahn, K. (2003). A survey of socially interactive robots. *Robotics and Autonomous Systems, Special issue on Socially Interactive Robots*, 42(3–4):143–166.
- Gaussier, P., Moga, S., Banquet, J., and Quoy, M. (1998). From perception-action loops to imitation processes: A bottom-up approach of learning by imitation. *Applied Artificial Intelligence*, 7–8(12):701–727.
- Hayes, G. M. and Demiris, J. (1994). A robot controller using learning by imitation. In Borkowski, A. and Crowley, J. L., editors, *Proceedings of the 2nd International Symposium on Intelligent Robotic Systems*, pages 198–204. LIFIA-IMAG, Grenoble.
- Hugues, L. and Drogoul, A. (2001). Shaping of robot behaviors by demonstration. In *Proceedings of the First International Workshop on Epigenetic Robotics: Modeling Cognitive Development in Robotic Systems*, pages 23–29. Lund University Cognitive Studies.
- Kaiser, M. and Dillmann, R. (1996). Building elementary robot skills from human demonstration. In *Proceedings of the IEEE International Conference on Robotics and Automation*, volume 3, pages 2700–2705.
- Kaplan, F., Oudeyer, P.-Y., Kubinyi, E., and Miklosi, A. (2001). Taming robots with clicker training : a solution for teaching complex behaviors. In Quoy, M., Gaussier, P., and Wyatt, J. L., editors, *Proceedings of the 9th European workshop on learning robots*, LNAI. Springer.
- Klingspor, V., Demiris, J., and Kaiser, M. (1997). Human-robot-communication and machine learning. *Applied Artificial Intelligence Journal*, 11:719–746.
- Marom, Y. (2003). *Influencing Robot Learning Through Design and Social Interactions — A Framework for Balancing Designer Effort with Active and Explicit Interactions*. PhD thesis, University of Edinburgh.
- Marom, Y. and Hayes, G. M. (2001). Attention and social situatedness for skill acquisition. In *Proceedings of the First International Workshop on Epigenetic Robotics: Modeling Cognitive Development in Robotic Systems*, pages 105–114.
- Marom, Y., Maistros, G., and Hayes, G. (2001). Experiments with a social learning model. *Adaptive Behavior*, 9(3–4):209–240.
- Marom, Y., Maistros, G., and Hayes, G. M. (2002). Towards a mirror system for the development of socially-mediated skills. In *Proceedings of the Second International Workshop on Epigenetic Robotics: Modeling Cognitive Development in Robotic Systems*.
- Marr, D. (1982). *Vision: A Computational Investigation into the Human Representation and Processing of Visual Information*. W H Freeman & Co.
- Marsland, S., Shapiro, J., and Nehmzow, U. (2002). A self-organising network that grows when required,. *Neural Networks*, 15(8–9):1041–1058.
- Matarić, M. J. (2000). Getting humanoids to move and imitate. *IEEE Intelligent Systems*, 14(4):18–24.
- Moukas, A. and Hayes, G. (1996). Synthetic robotic language acquisition by observation. In *Proceedings of the Fourth Conference on Simulation of Adaptive Behavior (SAB'96)*, Cape Cod.
- Nehaniv, C. and Dautenhahn, K. (2000). Of hummingbirds and helicopters: An algebraic framework for interdisciplinary studies of imitation and its applications. In Demiris, J. and Birk, A., editors, *Learning Robots: An Interdisciplinary Approach*. World Scientific. Robotics and Intelligent Systems Series, Vol. 24.
- Nehmzow, U. and McGonigle, B. (1994). Achieving rapid adaptations in robots by means of external tuition. *From Animals to Animats*, 3.
- Nicolescu, M. and Matarić, M. J. (2003). Natural methods for learning and generalization in human-robot domains. In *Proceedings of the Second International Joint Conference on Autonomous Agents and Multi-Agent Systems*.
- Schaal, S. (1999). Is imitation learning the route to humanoid robots? *Trends in Cognitive Sciences*, 3:233–242.