

Human-Agent Interaction and Machine Learning

Michael Kaiser^{1*}, Volker Klingspor,² and Holger Friedrich¹

¹ University of Karlsruhe, Institute for Real-Time Computer Systems & Robotics,
D-76128 Karlsruhe, Germany

² University of Dortmund, Lehrstuhl Informatik VIII, D-44221 Dortmund, Germany

Abstract. Human-Agent Interaction as a specific area of Human-Computer Interaction is of primary importance for the development of systems that should cooperate with humans. The ability to learn, i.e., to adapt to preferences, abilities and behaviour of a user and to peculiarities of the task at hand, should provide for both a wider range of application and a higher degree of acceptance of agent technology. In this paper, we discuss the role of Machine Learning as a basic technology for human-agent interaction and motivate the need for interdisciplinary approaches to solve problems related to communication with artificial agents for task specification, teaching, or information retrieval purposes.

1 Introduction

The concept of "intelligent agents" has recently found growing interest both in theory-oriented research as well as in applications such as robotics, manufacturing, information retrieval, and human-computer interaction. Especially, the idea of agents as "intelligent systems" that act in cooperation with or on behalf of a human user becomes increasingly important in today's information-oriented society. Research issues in this context comprise, among others:

1. Which tasks are really suited for agents (robots and softbots)?
2. How can artificial agents be designed and developed in a systematic manner?
3. How can humans communicate with artificial agents (and vice versa)?
4. How can artificial agents adapt to changing user preferences, a changing environment, etc.?

For Machine Learning, the field of intelligent agents offers a wide area of both research and applications. Specifically, the following questions ought to be answered if intelligent agents are to become useful for everyday users:

1. What and how can artificial agents learn from humans?
2. What and how can artificial agents learn from each other?
3. What and how can artificial agents learn without external guidance or instruction, i.e., on their own?

* Now with ABB Corporate Research Ltd., CHCRC.C, CH-5405 Baden-Dättwil, Switzerland

Throughout this paper, we will try to set the stage for an interdisciplinary treatment of these questions. To this aim, we will shortly discuss the basics of human-agent interaction, and identify learning tasks that are to be solved. Afterwards, we will give a brief overview on the role current Machine Learning techniques can play to solve these tasks. Finally, some examples will be presented that illustrate the application of Machine Learning in Interaction tasks. The examples are related to adaptive user interfaces, user-adaptive information retrieval, and human-robot interaction.

2 The Psychology of Human-Agent Interaction

Agents should act in cooperation with or on behalf of a user. They should aid her to perform a difficult task, or take over a task completely which she doesn't want to do or isn't capable of. The more knowing and autonomous an agent gets, the closer the relationship between user and agent approaches that of two equal partners. However, in all cases the underlying relationship can be characterized as the human user **making use** of the artificial agent. For the agent, the user is the reference w.r.t. task specification and evaluation of performance.

Especially w.r.t. the actual communication (for task specification, monitoring, and performance evaluation), agent designers (and the agents themselves) should therefore take a quite objectivistic point of view, assuming that the reference for the meaning of symbols used for communication will be the user's understanding of these symbols. The agent should not – at least for communication purposes – construct its own symbols, but should ground user-defined symbols onto its own perceptions and actions.

Within this setting, the need for an agent to be adaptive w.r.t. the input it receives from its human user is obvious. The specific learning tasks, which are related to both communication and task execution, are discussed in the next section.

3 Learning Tasks in Human-Agent Interaction

Principally, an agent may have to adapt to two different things: the preferences of the user and variations in its tasks, including variations in its operating environment ³ In the first case, we can distinguish learning **for** communication (i.e., to learn how to communicate with the human user) and learning **from** communication (which includes, for example, learning by observation).

To learn for communication requires the agent to build relations between user-supplied symbols to its own knowledge representation and, finally, to ground these symbols on its own perceptions and actions. This is the classical problem of symbol grounding [Harnad, 1990]. It has found explicit treatment especially in human-robot interaction [Klingspor et al., 1996] and distributed AI

³ One might even say that an "agent-worthy task" should **require** an agent to adapt itself to it. [Beale and Wood, 1994].

[Kaiser et al., 1996b], but it is inherent in any learning task that requires the agent to relate its own behaviour to a user-defined concept, just as *usefulness* or *interestingness*. To solve this task, supervised inductive techniques are mostly being used, ranging from ILP [Klingspor et al., 1996] to neural networks.




Fig. 1. Duality of abstraction and explanation.

The basic idea is always the same (Fig. 1): The meaning of language symbols is **explained** to the agents without using the language. The task of the agent is to abstract from the observations made during the explanation to the general concept represented by the user-given language symbol.

When learning from communication, the agent receives instructions from the user that are related to its behaviour. To understand these instructions, a common language between the agent and the user is mandatory, learning for communication may be required to provide it.

In the most simple case, instructions given to the agent are complete in that sense that they relate each possible state the agent may encounter to an appropriate action. Then, no induction is necessary. However, in the more general case only for some states information about how to behave are given to the agent, such that **incremental supervised learning** has to take place. The user may also be able to supply additional information to guide the agent's learning process, such that learning needs not to be purely inductive. This is the typical setting of Programming by Demonstration (see section 4.1).

Another prototypical setting is that of **reinforcement learning**, in which the agent receives a possibly delayed qualitative feedback from the user. Hence, the agent doesn't get explicit instructions. Therefore, the agent must explore the space of possible actions, in order to adapt to the user's performance criteria.

In any case, if learning takes place in a supervised manner, such that the given feedback consists of an optimal action or a quantitative indication of the error made by the agent, either

- the teacher must know the action space of the instructed agent and formulate

the advice appropriately,

or

- the instructed agent must be able to map the teacher’s advice onto its own action space.

Both requirements are not trivial, especially if agents should learn from other agents that are not structurally identical. For agents that should learn from human users, appropriate interfaces are therefore mandatory.

For learning on the basis of a scalar reward, the situation is very similar. Either

- the teacher must know the range of possible rewards used by the instructed agent (i.e., what is the ”good” and ”bad” in terms of the pupil),

or

- the instructed agent must know the mapping between the teacher’s reward and its own range of rewards (what does the teacher mean by ”good” and ”bad”).

In all cases, teacher or pupil must initiate the learning process. An important requirement is also that the teacher knows the limits of the instructed agent, since it makes no sense to try to teach an agent to go beyond its maximum capabilities (e.g., to try to position a robot with a higher precision than it is capable of). To enable the teacher to take care of this aspect requires the teacher to query these limits from an agent and to correctly interpret the agent’s answer. Similarly, the instructed agent must understand the teacher’s request and relate its capabilities properly to the task specified by the teacher.

4 Applications of Machine Learning in Human-Agent Interaction

4.1 Interface Agents

By interface agents, we mean agents that assist users with computer-based tasks. Simple examples are guides that have been given a fixed behaviour (such as Microsoft’s ”Wizards”). More sophisticated agents that learn by watching a human user are being developed under the heading of Programming by Demonstration (PbD)[Cypher, 1993].

PbD has been applied successfully several domains such as graphic editors like the MONDRIAN [Lieberman, 1993] system that learns new graphical editing procedures using a form of explanation based generalization. Instructible software agents [Maulsby, 1994] learn tasks like autonomous text formatting or graph editing by observing the user doing the task. The knowledge is adapted and corrected by processing corrections issued by the user during automatic task execution. Intelligent specification interfaces such as the ViCCS interface

[Minton et al., 1995] and the interface agent assisting with completion of repetitive forms of Hermens and Schlimmer [Hermens and Schlimmer, 1994] are other examples of PbD applications.

Finally, an interface agent may also directly control the feedback given to a human user, even in an application-independent manner. An example of this is the system described in [Münch and Stangenberg, 1996], which provides selective haptic feedback to users of an X-Windows application.

4.2 Agents for Information Retrieval

The quickly growing World Wide Web demands for agents collecting and presenting interesting information to a user. NewsWeeder [Lang, 1995] is an agent that learns which Netnews are interesting for the reader. It learns from explicit specifications of the user preferences (the reader rates each read article), and NewsWeeder clusters highly rated articles by calculating the frequency of the words occurring in the article.

PZ additionally uses the linguistic structure of the text to find interesting articles of several daily newspapers in the Web. It generates a personalized newspaper, which consists only of those articles that are of interest for a specific user.

WebWatcher [Armstrong et al., 1995] learns, which pages in the WWW are related to each other by observing the user when surfing in the Web. It generalizes the structural information contained in the sequence of pages being visited. Similarly, Lieberman's Letizia [Lieberman, 1995] aims at hypothesizing on the user's interests by observing the user's actions, in order to find possibly interesting pages, starting from the current page.

Knowledge Discovery in Databases (KDD) is another information retrieval task in which the human user can greatly benefit from support by an intelligent agent. Apart from performance issues [Kargupta et al., 1997], we also find that the main task for such agents is to perform information filtering, such that only those patterns are being presented that are really of interest for the human user [Davies and Edwards, 1995b, Davies and Edwards, 1995a]. In addition, the user may also communicate with the actual mining agent, in order to guide its search [Wrobel et al., 1996].

4.3 Human-Robot Interaction

In robotics, the idea of the "personal robot" or "personal robotic assistant" (e.g., for aiding the elderly or disabled) is lately receiving a lot of attention. Robot manufacturers all over the world are expecting to open a new, consumer-oriented market that will allow for sales figures far beyond those that can be obtained from today's industrial robots. However, to enable new robot applications with emphasis on service tasks, it is necessary to develop techniques which allow untrained users to make efficient and safe use of a robot.

In brief, what is required is

1. an interface that allows a user to intuitively instruct the robot, and

2. informative feedback such that the user can immediately understand what's happening on the robot's side.

Especially for users who are not experts in robot programming, **Programming by Demonstration (PbD)** [Cypher, 1993] has a considerable potential to become a suitable programming technique and to replace conventional robot programming languages. As in the case of interface agents, **Robot Programming by Demonstration (RPD)**, [Heise, 1989]) relies on demonstrations of the task. Those demonstrations are used as the primary input and can be used to transfer different kinds of knowledge to a robot.

- Demonstrations were proven to be suitable for the acquisition of *new program schemata on task level* [Segre, 1989]. In [Kuniyoshi et al., 1994], sequences of video images were analyzed in order to generate assembly plans. [Andreae, 1984] presented NODDY, a system which generates generalized programs by fusing several demonstrations. Single demonstrations and acquired user intentions are the basis for the robot programs generated by the system described in [Friedrich et al., 1996].
- On the *control level*, demonstrations can be used as the basis for learning both, *open-loop and closed-loop elementary skills*. The acquisition of open-loop skills is mostly focused on the reconstruction of trajectories from a sequence of demonstrated states (positions) [Delson and West, 1994], [Ude, 1993]. Systems supporting the acquisition of closed-loop elementary skills comprise acquisition techniques for manipulation tasks such as deburring operations [Asada and Liu, 1991] and assembly [Kaiser and Dillmann, 1996] as well as for vehicle control [Pomerleau, 1991] and autonomous robot navigation [Reignier et al., 1995], [Kaiser et al., 1996a].
- Learning new perceptive skills for object and landmark recognition can also take place on several system control levels. [Accame and Natale, 1995] present an approach to learn sensor parameterizations from demonstrations. Learning active perception skills, i.e., the combination of actions and sensing for the purposes of object recognition and communication, is the topic of work presented in [Klingspor et al., 1996].

In Machine Learning, **behavioural cloning** has become popular as a synonym for "skill acquisition via human demonstration." Typical applications are the cart-pole balancing tasks [Guez and Selinsky, 1988], [Dzeroski et al., 1995], as well as the work on "Learning to fly" [Sammut et al., 1992] and recent work on crane control [Urbancic and Bratko, 1994], [Bratko et al., 1995]. In contrast to work in robotics, these approaches focus on the evaluation of a specific learning technique for cloning. Also **imitation learning** [Hayes and Demiris, 1994], [Demiris and Hayes, 1996] must be considered in the context of Robot Programming by Demonstration. It is concerned with learning by imitating another agent (another robot), and as thus can be considered a special case of PbD. An important difference is that the actions that the robot performs due to the imitation mechanism form the basis for learning. While this technique puts less demands

on the actual interaction interfaces, the learning agent must be much more skilled in interpreting its observations.

5 Summary and Conclusion

Throughout this paper, we discussed the importance of Machine Learning for human-agent interaction. We identified the basic learning tasks that need to be solved and indicated possible solutions to these tasks. Several application examples have shown that learning capabilities are really required, and that agents can exhibit such capabilities by using state of the art ML techniques.

However, we have also seen that to design learning agents is a task that should be solved in a strongly user-oriented manner. Therefore, the discussion of technical issues must be complemented by treatment of the underlying interaction mechanisms and actual usability concerns. To provide a forum for this kind of interdisciplinary approach to human-agent interaction and especially the learning issues in this area, has been the major motivation for proposing the workshop.

References

- [Accame and Natale, 1995] Accame, M. and Natale, F. D. (1995). Neural tuned edge extraction in visual sensing. In *Third European Workshop on Learning Robots*, Heraklion, Kreta.
- [Andreae, 1984] Andreae, P. M. (1984). Constraint limited generalization: Acquiring procedures from examples. In *Proceedings of the National Conference on Artificial Intelligence (AAAI)*, pages 6 – 10.
- [Armstrong et al., 1995] Armstrong, R., Freitag, D., Joachims, T., and Mitchell, T. (1995). Webwatcher: A learning apprentice for the world wide web. In *1995 AAAI Spring Symposium on Information Gathering from Heterogeneous, Distributed Environments*.
- [Asada and Liu, 1991] Asada, H. and Liu, S. (1991). Transfer of human skills to neural net robot controllers. In *IEEE International Conference on Robotics and Automation*, pages 2442 – 2448, Sacramento.
- [Beale and Wood, 1994] Beale, R. and Wood, A. (1994). Agent-based interaction. In *Proceedings of the HCI 94*, Glasgow, UK.
- [Bratko et al., 1995] Bratko, I., Urbancic, T., and Sammut, C. (1995). Behavioural cloning: Phenomena, results, and problems. In *5th IFAC Symposium on Automated Systems based on Human Skill*, Berlin.
- [Cypher, 1993] Cypher, A. I. (1993). *Watch what I do – Programming by Demonstration*. MIT Press, Cambridge.
- [Davies and Edwards, 1995a] Davies, W. H. E. and Edwards, P. (1995a). Agent-based knowledge discovery. In *1995 AAAI Spring Symposium on Information Gathering from Heterogeneous, Distributed Environments*.
- [Davies and Edwards, 1995b] Davies, W. H. E. and Edwards, P. (1995b). Distributed learning: An agent-based approach to data-mining. In *11th International Conference on Machine Learning, Workshop on Agents that learn from other agents*, Tahoe City, California.

- [Delson and West, 1994] Delson, N. and West, H. (1994). The use of human inconsistency in improving 3D robot trajectories. In *IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 1248 – 1255, München.
- [Demiris and Hayes, 1996] Demiris, J. and Hayes, G. (1996). Imitative learning mechanisms in robots and humans. In *Fifth European Workshop on Learning Robots*, pages 9 – 16, Bari.
- [Dzeroski et al., 1995] Dzeroski, S., Todorovski, L., and Urbancic, T. (1995). Handling real numbers in ILP: A step towards better behavioural clones. In *European Conference on Machine Learning*, pages 283 – 286, Heraklion, Kreta.
- [Friedrich et al., 1996] Friedrich, H., Münch, S., Dillmann, R., Bocionek, S., and Sassin, M. (1996). Robot programming by demonstration: Supporting the induction by human interaction. *Machine Learning*, 23:163 – 189.
- [Guez and Selinsky, 1988] Guez, A. and Selinsky, J. (1988). A neuromorphic controller with a human teacher. In *IEEE International Conference on Neural Networks*, pages 595 – 602, San Diego.
- [Harnad, 1990] Harnad, S. (1990). The symbol grounding problem. *Physica D*, 42:335–346.
- [Hayes and Demiris, 1994] Hayes, G. and Demiris, J. (1994). A robot controller using learning by imitation. In *International Symposium on Intelligent Robotic Systems*, pages 198 – 204, Grenoble.
- [Heise, 1989] Heise, R. (1989). Demonstration instead of programming: Focussing attention in robot task acquisition. Technical Report 89/360/22, Department of Computer Science, University of Calgary.
- [Hermens and Schlimmer, 1994] Hermens, L. A. and Schlimmer, J. C. (1994). A machine-learning apprentice for the completion of repetitive forms. *IEEE Expert*, pages 28 – 33.
- [Kaiser et al., 1996a] Kaiser, M., Deck, M., and Dillmann, R. (1996a). Acquisition of basic mobility skills from human demonstrations. *Studies in Informatics and Control*, 5(3):223 – 234.
- [Kaiser and Dillmann, 1996] Kaiser, M. and Dillmann, R. (1996). Building elementary robot skills from human demonstration. In *IEEE International Conference on Robotics and Automation*, Minneapolis, Minnesota, USA.
- [Kaiser et al., 1996b] Kaiser, M., Rogalla, O., and Dillmann, R. (1996b). Communication as the basis for learning in multi-agent systems. In *ECAI '96 Workshop on Learning in Distributed AI Systems*, pages 50 – 59, Budapest. Wird in erweiterter Form erscheinen in: G. Weiss, Hrsg., *Learning in Distributed AI Systems*, Springer Lecture Notes in AI, 1997.
- [Kargupta et al., 1997] Kargupta, H., Hamzaoglu, I., Stafford, B., Hanagandi, V., and Buescher, K. (1997). PADMA: An architecture for scalable text classification. In *High Performance Computing*. to appear.
- [Klingspor et al., 1996] Klingspor, V., Morik, K., and Rieger, A. (1996). Learning concepts from sensor data of a mobile robot. *Machine Learning*, 23(2/3):305 – 332.
- [Kuniyoshi et al., 1994] Kuniyoshi, Y., Inaba, M., and Inoue, H. (1994). Learning by watching: Extracting reusable task knowledge from visual observation of human performance. *IEEE Transactions on Robotics and Automation*, 10(6):799 – 822.
- [Lang, 1995] Lang, K. (1995). Newsweeder: Learning to filter net news. In *Proceedings of the 12th International Conference on Machine Learning*. Morgan Kaufmann.
- [Lieberman, 1993] Lieberman, H. (1993). *Watch what I do*, chapter MONDRIAN: A teachable graphical editor.

- [Lieberman, 1995] Lieberman, H. (1995). Letizia: An agent that assists web browsing. In *Proceedings of the International Joint Conference on Artificial Intelligence*, Montreal, Canada.
- [Maulsby, 1994] Maulsby, D. (1994). *Instructible agents*. PhD thesis, University of Calgary, Canada.
- [Minton et al., 1995] Minton, S., Philpot, A., and Wolfe, S. (1995). Specification by demonstration: The ViCCS interface. In *International Conference on Machine Learning, Workshop on Programming by Demonstration*, Tahoe City.
- [Münch and Stangenberg, 1996] Münch, S. and Stangenberg, M. (1996). Intelligent control for haptic displays. In *Eurographics '96*, Poitiers, France.
- [Pomerleau, 1991] Pomerleau, D. A. (1991). Efficient training of artificial neural networks for autonomous navigation. *Neural Computation*, 15:88 – 97.
- [Reignier et al., 1995] Reignier, P., Hansen, V., and Crowley, J. (1995). Incremental supervised learning for mobile robot reactive control. In *Intelligent Autonomous Systems 4*, pages 287 – 294, Karlsruhe.
- [Sammut et al., 1992] Sammut, C., Hurst, S., Kedzier, D., and Michie, D. (1992). Learning to fly. In *Ninth International Conference on Machine Learning*, pages 385 – 393, Aberdeen.
- [Segre, 1989] Segre, A. M. (1989). *Machine Learning of Robot Assembly Plans*. Kluwer Academic Publishers.
- [Ude, 1993] Ude, A. (1993). Trajectory generation from noisy positions of object features for teaching robot paths. *Robotics and Autonomous Systems*, 11:113 – 127.
- [Urbancic and Bratko, 1994] Urbancic, T. and Bratko, I. (1994). Reconstructing human skill with machine learning. In *European Conference on Artificial Intelligence*, pages 498 – 502, Amsterdam.
- [Wrobel et al., 1996] Wrobel, S., Wettschereck, D., Verkamo, I., Siebes, A., Mannila, H., Kwakkel, F., and Klösken, W. (1996). User interactivity in very large scale data mining. In *9. Fachgruppentreffen der GI Fachgruppe 1.1.3 Maschinelles Lernen*, pages 125 – 130, Chemnitz.