3rd International Conference on the Learning Sciences, Atlanta, GA, December 16-19, 1998

Using Agents to Investigate Strategies for Human Collaborative Learning

Dan L. Grecu, Lee A. Becker
Department of Computer Science
Worcester Polytechnic Institute
100 Institute Road
Worcester, MA 01609, U.S.A.
E-mail: (dgrecu, lab)@cs.wpi.edu

Abstract: Given the cost and effort involved in investigating different methods or protocols for human collaborative learning, this paper proposes the use of software learning agents to simulate and test various protocols. The best performing protocols would then be tested on human subjects. The paper presents a new form of cooperative learning, called coactive learning. After arguing that humans could be taught to learn coactively, experiments are conducted with various coactive learning schemes. The experiments demonstrate situations in which machine learning agents using coactive learning can perform better than individual machine learning agents. Collaborative learning through coaction is shown to generate properties that individual learners can achieve only through computationally intensive strategies. It is suggested that human learners be trained to learn coactively using the most successful agent coacting protocols, and compared, on the same type of learning task, to individual learners and groups of learners not trained in these protocols.

1. Introduction

The benefits of peer interaction for learning tasks have been theorized since the time of Piaget (1928) and Vygotsky (1978, originally 1930, 1933, 1935). Recent years have seen a great increase in interest and research on collaborative learning, group learning, and cooperative induction ([Hoyles & Forman 1995], [Bruffee 1993], [Laughlin et al. 1991], [Laughlin 1996]). Collaborative learning applies very well to situations where the knowledge emerges from the interaction of knowledgeable peers, like in communities of scientists. This paradigm is known as non-foundational social construction of knowledge [Knorr-Cetina 1981]. In contrast, foundational knowledge construction takes place in domains for which there exists a basic theory that can be applied to reason and to solve problems. Such theories can be communicated or shared in a traditional way by an authority, for example a professor or a teacher. Non-foundational domains can be regarded as equivalent to the Weak Theory Domains, a term used by the machine learning community [Porter, Bareiss & Holte 1990]. Weak theories and non-foundational knowledge, such as experiential knowledge, are modeled in artificial intelligence using heuristic rules or case-based reasoning [Kolodner 1993]. Decision tasks with reduced information availability, or method selection for problem-solving are typical application domains for weak theories. In general, weak theories are useful for tasks that require to construct mappings between heuristics and classes of problems.

In the following we describe coactive learning as a collaboration method that provides support for learning in non-foundational or weak theory domains. We illustrate how this collaborative learning technique can be tested and evaluated in an agent environment prior to attempting its transfer in the human domain. For the experiments to be relevant for human collaborative learners several conditions must be met: First, the task domain must be suitable for collaborative learning. Second, any assumptions about the nature of the individual learners, in terms of algorithms and knowledge representations, must not be contradicted by what we know about human learners. Finally, the collaboration protocols must be practical and specific enough for human learners to follow.

The range of protocols for interaction among collaborative learners is very wide. The testing performed with software learning agents can be of benefit in pruning ineffective methods. It should be noted that different methods or protocols could be effective by different measures. As illustrated below, agents allow us to look into some of the trade-offs that need to be made.

The remainder of this paper is organized as follows: We start with a description of coactive learning and the assumptions it imposes on individual learners. We then argue that these assumptions are reasonable for humans, and thus humans potentially could be taught to learn coactively. We next describe the experimental task and knowledge representation, and then show how coactive learning can be done with this knowledge representation. We investigate different protocols and present results of experiments with artificial and real world domains. Finally, we discuss these results and our conclusions.

2. Coactive Learning

The term coacting is used in the psychological literature to refer to an individual's performance in the presence of other individuals performing the same task [Hill 1982]. We use the term to indicate a more active interaction between the learners, but still one in which each learner produces its own result of learning, rather than a group result. Figure 1 depicts the general scheme for coactive inductive learning. The learners may be trained on different examples [Figure 1a], perhaps at a distance, or on the same training set [Figure 1b].

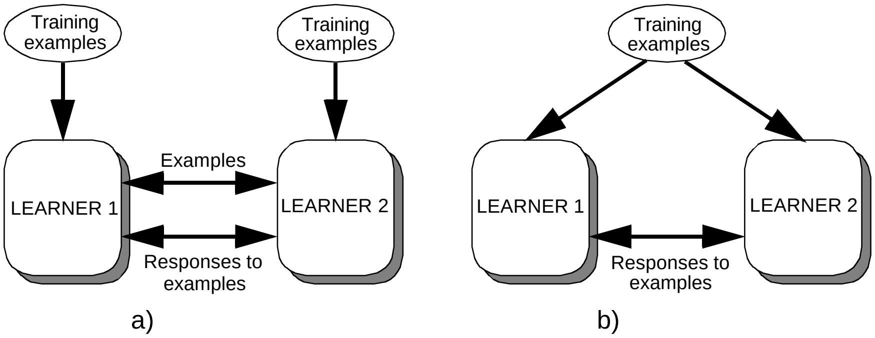

Figure 1: General coactive learning schemes

The communication between learners is limited to the exchange of relevant examples and of ‘responses’ to examples, as discussed below. A learner may communicate an example to a coactor to ask for an opinion or exclusively for the benefit of the coactor. Obviously, critical training examples may be of significant use to a learner, affording some of the advantages of supervised learning with examples being provided by a knowledgeable or experienced teacher or peer tutor.

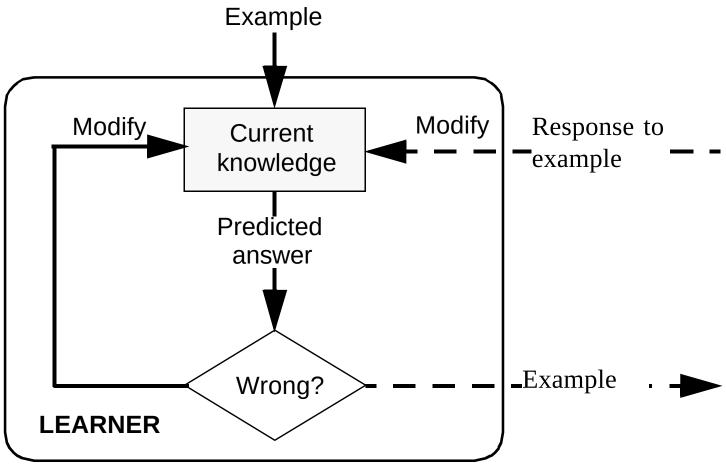

Figure 2: Generic coactive learner model

Our model of coactive learning is based on the assumption that the individual learners are performing incremental learning through error recovery [Figure 2]. Papert (1980), following Piaget (1928), stresses the importance of learning through errors. This model does not commit learners to any particular representation of knowledge or learning algorithm.

3. The Experimental Task and Knowledge Representation

The task which the learning agents in our experiments were learning to perform was classification [Clancey 1985]. This is a very widely applicable task, and as suggested above is suitable for collaborative learning. The knowledge was represented as a collection of past instances or cases. A large number of psychologists and cognitive scientists have proposed models or conducted experiments showing that humans use stored instances or cases in problem solving ([Medin 1978], [Brooks 1978], [Medin, Dewey & Murphy 1983], [Ross 1984], [Holyoak & Kob 1987], [Medin & Ross 1989], [Brooks et al. 1989], [Ross 1989]).

Based on their relevance for human learning, we have chosen to implement our experiments with instance-based learning (IBL) algorithms. IBL is an inductive learning model that generates concept descriptions by storing specific training instances and is used for classification tasks. Aha, Kibler & Albert (1991) present an compare three IBL algorithms. Each generates a prediction about a new instance based on the classifications given by the most similar *k* stored instances.

The algorithms differ in their storage strategies. IB1 stores all observed instances. IB2 stores only the misclassified instances, which reduces storage requirements, but makes the algorithm sensitive to noise. IB3 stores all misclassified training instances, but maintains a classification record of correct vs. incorrect predictions, discarding poor predictors as potentially stored noise. Due to its computational requirements, IB3 is less likely to be accessible to human learners.

Agent experiments were run on two types of domains: The artificial domain consisted of instances representing points with integer coordinates lying inside a 100x100 square, subdivided into four equal-sized subsquares. The prediction task required to determine the subsquares to which the test instances belong. This corresponds to a concept description where the examples have two attributes and are partitioned into four classes. The artificial domain was designed for an extensive computational testing of the agent environment. We have also run some of the experiments on four of the real-world databases available from the University of California at Irvine Machine Learning Repository. The Hungarian and the Cleveland databases represent heart disease diagnosis information based on 14 classification features. The tumor database represents tumor location information (22 possible location classes) based on 17 symptom features. Finally, the voting database includes the votes of all the members of the U.S. Congress on 16 key issues (boolean values), and the predicted class attribute is the party affiliation of the congressmen. We used the following training/testing ratios: Cleveland database – 250/53, Hungarian database – 250/44, Tumor database – 275/64, Voting database – 350/85. The real-world databases are representative for data sets used in typical human tasks, such as diagnosis, and are characterized by large numbers of attributes, by sparsely distributed training instances, by the presence of unknown attribute values, and by numeric, nominal, and boolean attribute values.

4. Coactive Instance-Based Learning

In the following we describe experiments that explore three coactive learning settings: learning on the same training set using different learning strategies, learning with similar strategies from different data, and ‘parallelized’ learning in groups. All experiments are carried out with IB1 and IB2 learners, or versions of these strategies. IB3 results are presented for comparison purposes.

4.1 Learning on the same training set using different individual strategies

We consider first a single IB1 learner coacting with a peer IB2 learner, learning on the same set of training instances [Figure 1.a]. The artificial domain experiments have been run on no-noise data, as well as on data sets with 10% noise (i.e., each instance attribute had a chance of 10% of being corrupted). Noisy data corresponds to situations where human learners are provided with flawed examples. For reference, Table 1 shows the performance of each learner, including IB3. The artificial domain included 400 instances, out of which 80% were randomly chosen for training.

Learner	Artficial domain (instances stored)		Real-world domains			
	No noise	10% noise	Cleveland	Hungarian	Tumor	Voting
IB1	97.4% (320)	86.9% (320)	54.3 (250.0)	77.9 (250.0)	37.7 (275.0)	92.5 (350.0)
IB2	95.2% (28.3)	74.1% (101.5)	50.4 (132.6)	71.5 (63.5)	34.1 (190.2)	89.9 (36.5)
IB3	94.5% (37.1)	92.5% (51.3)	52.0 (86.1)	79.1 (29.3)	37.8 (145.8)	92.8 (24.3)

Table 1: Performance of IB1, IB2 and IB3 learners on artificial domain and real-world domains

Table 2 presents schemes that illustrate properties of coacting improving on the basic performance of each learner. Each scheme adds new modifications to the individual strategies. The common feature for this setting is that each learner asks its peer for its classification result for the current training instance. The experiments test modifications in the learning strategies that enable learners to take advantage of coacting. Italics indicate modifications from the standard IBL algorithms.

Scheme	Learner	Coacting protocol	Artificial domain		Real-world domains			
			No noise	10% noise	Cleveland	Hungarian	Tumor	Voting
1	IB1	–S –S +S +S	97.4 (320)	86.9 (320)	54.3 (250.0)	77.9 (250.0)	37.7 (275.0)	92.5 (350.0)
	IB2	–*NS* +NS –S +NS	94.6 (24.3)	91.2 (25.3)	51.7 (31.3)	78.7 (24.7)	38.3 (30.8)	92.0 (16.5)
2	IB1	–S –S +S ***P*** +S	97.4 (320)	87 (320)	54.3 (250.0)	77.9 (250.0)	37.7 (275.0)	92.5 (350.0)
	IB2	–*NS* +NS –S +NS	94.5 (45)	91.9 (44.9)	54.3 (53.4)	77.0 (45.6)	39.4 (56.8)	93.5 (24.3)
3	IB1	–S –S ***D*** +S ***P*** +S ***P***	97.4 (315.5)	90.5 (301.6)	54.8 (231)	78.4 (237.1)	38.2 (250.7)	92.7 (345.4)
	IB2	–*NS* +NS –S ***D*** +NS	96.3 (160.3)	95.5 (140.7)	57.2 (76.4)	80.6 (100.7)	38.7 (63.0)	93.0 (137.3)
4	IB1	–S –S ***D*** +S ***P*** +*NS* ***P***	95.4 (39)	85.1 (72)	50.3 (116.9)	75.7 (58.6)	33.9 (198.0)	89.9 (36.4)
	IB2	–*NS* +NS –S ***D*** +NS	94 (34.3)	91.9 (48.3)	55.0 (39.0)	77.7 (33.5)	37.3 (53.0)	91.2 (29.0)

Table 2: IB1-IB2 coacting results (*D*/*P* means 'drop/pass classifying instance')

In **scheme 1** the IB2 learner has been modified not to store a misclassified instance that is also misclassified by IB1, assuming that it may be noise. This improves IB2's prediction performance on noisy data by more than 15%, and brings it close to the IB3 performance on all data sets. In addition, IB2's storage is, in some cases, less than 25% that of IB3. In **scheme 2** when the IB2 learner misclassifies a training instance that is correctly classified by the coacting IB1 learner, the IB2 learner stores the instance which IB1 used to make the correct prediction. This modification further improves IB2's performance on noisy data, while on 3 of the 4 real-world domains its accuracy surpasses that of IB3. **Scheme 3** introduces a modification in IB1, in that it drops its classifying instance if it incorrectly predicts the current instance, where IB2 correctly predicts. IB1 assumes it has stored a noisy instance that it is currently using for classification. When both IB1 and IB2 correctly classify a training instance, IB1 passes the most similar stored instance, which was used to correctly classify the new instance, and IB2 stores this passed instance. IB2 also eliminates its classifying instance, when it misclassifies the training instance and IB1 classifies it correctly. Scheme 4 exhibits the best prediction accuracy on noisy examples, surpassing that of IB3 by 3%. Its storage is increased very significantly, however, by this modification. It should be noted that the modified IB2 in this scheme also performs very well in the no noise condition. It provides the averaged highest performance over the noise and no noise conditions, as well as on the real-world data. In **scheme 4** IB1 differs from the previous scheme in that it does not store the new training instance whenever both it and IB2 correctly classify the new instance. Although IB1's prediction accuracy, and consequently IB2's, is reduced, IB1 stores significantly less cases. The IB2 produced by the scheme greatly surpasses an independent IB2 in the noise condition and comes within 0.6% of an independent IB3.

In summary, the different schemes for coacting between an IB1 and an IB2 agent allow one to choose or make trade-offs among different performance measures. They prove how simple collaboration schemes can enhance the performance of individual learners to a level that is comparative to IB3, a strategy that is less plausible for human learners, due to its significant computational demands. They also show how a noise reduction property can result from the interaction, whereas none of the individual learners possessed a noise filtering ability.

4.2 Learning with similar strategies and different data

In this section we look at learning situations where the subjects use similar algorithms (IB2), but train on different examples. Given two learners, A and B, only A takes advantage of coacting. A's IB2 is modified not to store cases misclassified by B as a precaution against noise. The learners use *different* training sets, the 320 training instances being equally divided between them. The cells with an enhanced border in Table 3 represent situations where subject A consults its coactor B on how it would classify A's current instance. Depending on the response on the forwarded training instance, A may request B to provide the classifying instance it has used for prediction. In **scheme 1** subject A consults B only when it misclassifies the training instance. If B correctly classifies A's training instance, A will request B's classifying instance and store it. On noisy data sets, subject A outperforms B in terms of accuracy by 9.3%, while storing only half as many instances as agent B (28 vs. 56 instances). In **scheme 2** A consults B on all its training instances and stores B's classifying instance whenever B performs correctly. Scheme 2 increases the noisy-data performance of agent A by another 4.2%, due to the reinforcement A receives on the instances that it classifies correctly.

Scheme	**Learner**	**Coacting protocol**				**Artificial domain**		**Real-world domains**			
						No noise	**10% noise**	**Cleveland**	**Hungarian**	**Tumor**	**Voting**
1	A	–NS	–S ***G***	+NS	+NS	92.7 (23.6)	84 (28.1)	51.3 (32.5)	77.4 (24.0)	33.1 (32.8)	91.6 (18.0)
	B	–S	+NS	–S	+NS	94.5 (20.6)	74.7 (56.2)	50.0 (66.3)	72.3 (34.4)	31.7 (99.0)	90.8 (21.5)
2	A	–NS	–S ***G***	+NS	+NS ***G***	94.1 (26.5)	88.2 (39.2)	53.1 (37.5)	78.8 (28.3)	32.3 (40.0)	91.8 (23.3)
	B	–S	+NS	–S	+NS	94.5 (20.6)	74.7 (56.2)	50.0 (66.3)	72.3 (34.4)	31.7 (99.0)	90.8 (21.5)

Table 3: IB2-IB2 coacting results (*G* means 'get classifying instance from coactor')

In these schemes A eliminates noise from its concept description, even though it uses a noise sensitive peer as a 'consultant'. A uses only half of the training instances for learning, while B uses the other half to construct a description that provides feedback to A. The cumulative effect of this collaboration yields better performance than a single IB2 learner that is provided with all the training data. This proves to be true both on the artificial domain, as on the real-world data sets.

4.3 Dividing up the learning task in groups

Finally, we summarize experiments on groups of 2, 4, and 8 identical learners with distributed data [Grecu & Becker 1998]. The experiments that we have run allow a group of IB2 learners to split large data sets into equal subsets and to achieve the same performance as an individual learner that has seen all the data. The basic protocol requires each learner to pass the instance it has misclassified to its peers, which will then store it. The experiments on non noisy data show that even when a data set comprised of 4000 cases is distributed among 8 learners, the performance of each learner is not less than that of an individual IB2 that sees all the instances. Collaboration can thus bring a task which is less affordable for an individual, in terms of the processing requirements, into a feasible range, with virtually no loss of performance. Furthermore, we show that even if each learner is trained on skewed data, with up to 60% of its training data coming from the same class, and with each learning set being biased towards a different class, the resulting concept description (the same for all learners) provides a performance that is very close to the one of learners that receive uniformly distributed data.

5. Discussion

There has been considerable debate about the relative effectiveness and importance of learning together with peers versus learning under interaction with teachers or adults. Hughes and Greenough (1995), Crook (1995) and Piaget (1928; 1965) stress the importance of peer or child-child learning, while Vygotsky (1978; 1991) and Wood (1976; 1978) emphasize the importance of adult-child learning. Peer learning is more similar to the coactive learning schemes where there is a relatively balanced example and response exchange between the two agents. Adult-child teaching features are prevalent in the schemes where an agent requests and receives considerably more examples and information about examples from its coactor than the coactor itself. It is also worth noting that the success of some of the schemes relies in the complementary features that are contributed by the two agents: A conservative, noise-sensitive IB1 learner and a strongly instance averse IB2 strategy, such as those in scheme 3 in the IB1-IB2 setting, can still provide an excellent accuracy due to the compensatory effect of IB1 over IB2.

While the interaction strategies between learners described in this paper should pose no difficulty in their application within human groups, this might be facilitated by computer-mediated interaction support [Healy et al. 1995]. Schemes such as those described in the IB1-IB2 or in the IB2-IB2 coaction can be implemented, stored and retrieved in networked environments that connect the workstations of remotely located human learners. The interaction being provided in an automated way, the human learners can each concentrate on the learning task by combining their own data and hypothesis with the support provided through the computer-supported cooperative work environment [Greif 1997].

We see coactive learning as being relevant in human problem-solving tasks where the individuals have to provide solutions without being previously trained or with little training on the situations they might face. Examples of tasks in this category are diagnosis and design. Under such circumstances problem-solving is relatively slow and expensive as the individuals have no (or few) previous cases on which to rely. The incremental and early acquisition of skills is essential, as every new learned piece of knowledge can be put to use as soon as it has been learned. Coaction can amplify the learning speed and accuracy based on the support received from other individuals that have the same learning target, even though they might have been exposed to different experiences (of the same type), and even though they might rely on different individual learning techniques and representations.

6. References

[Aha, Kibler & Albert, 1991] Aha, D. W., Kibler, D., & Albert, M. K. (1991). Instance based learning algorithms. *Machine Learning, 6*, 37-66.

[Brooks 1978] Brooks, L. (1978). Non-analytic concept formation and the memory of instances. In E. Rosch & B. B. Lloyd (Eds.), *Cognition and Categorization* (pp. 169-211). Hillsdale, NJ: Lawrence Erlbaum.

[Brooks, Allen & Norman 1989] Brooks, L. R., Allen, S. W., & Norman, G. (1989). Multiple and variable availability of familiar cases. *Case-Based Reasoning Workshop*, Pensacola, FL.

[Bruffee 1993] Bruffee, K. A. (1993). *Collaborative Learning: Higher Education, Interdependence and the Authority of Knowledge*. Baltimore and London: The John Hopkins University Press.

[Clancey 1985] Clancey, W. J. (1985). Heuristic classification. *Artificial Intelligence*, *27*, 289-350.

[Cover & Hart 1967] Cover, T. M., & Hart, P. E. (1967). Nearest neighbor pattern classification. *IEEE Transactions on Information Theory, 13*(1), 21-27.

[Crook 1995] Crook, C. (1995). On resourcing a concern for collaboration within peer interactions. *Cognition and Instruction*, *13*(4), 541-547.

[Davies & Edwards 1996] Davies, W., & Edwards, P. (1996). The communication of inductive inferences. In G. Weiss (Ed.), *Distributed Artificial Intelligence Meets Machine Learning* (pp. 223-241). Berlin: Springer Verlag.

[Grecu & Becker 1998] Grecu, D.L., & Becker, L.A. (1998). Coactive Learning for Distributed Data Mining, *4th International Conference on Knowledge Discovery and Data Mining*, Menlo Park, CA: AAAI Press.

[Greif 1997] Greif, I. (Ed.). (1997). *Computer-Supported Cooperative Work: A Book of Readings*. San Francisco, CA: Morgan Kaufmann.

[Healy, Pozzi & Hoyles 1995] Healy, L., Pozzi, S., & Hoyles, C. (1995). Making sense of groups, computers and mathematics. *Cognition and Instruction*, *13*(4), 505-523.

[Hill 1982] Hill, G. W. (1982). Group versus individual performance: Are N+1 heads better than one? *Psychological Bulletin*, *91*(3), 517-539.

[Holyoak & Koh 1987] Holyoak, K. J., & Koh, K. (1987). Surface and similarity in analogical transfer. *Memory and Cognition*, *15*, 332-440.

[Hoyles & Forman 1995] Hoyles, C., & Forman, E.A. (Eds.) (1995). *Cognition and Instruction*, Special issue: Processes and Products of Collaborative Problem Solving: Some Interdisciplinary Perspectives, *13*(4).

[Hughes & Greenough 1995] Hughes, M., & Greenough, P. (1995). Feedback, adult intervention, and peer collaboration in initial logo learning. *Cognition and Instruction*, *13*(4), 525-539.

[Kolodner 1993] Kolodner, J. (1993). *Case-Based Reasoning*. San Mateo, CA: Morgan Kaufmann.

[Knorr-Cetina 1981] Knorr-Cetina, K. D. (1981). *The Manufacture of Knowledge: An Essay on the Constructivist and Contextual Nature of Science*. Oxford: Pergamon Press.

[Laughlin, VanderStoep & Hollingshead 1991] Laughlin, P. R., VanderStoep, S. W., & Hollingshead, A. B. (1991). Collective versus individual induction: Recognition of truth, rejection of error, and collective information processing. *Journal of Personality and Social Psychology*, *61*(1), 50-67.

[Laughlin 1996] Laughlin, P. R. (1996). Group Decision Making and Collective Induction. In E. H. Witte & J. E. Davis (Eds.), *Understanding Group Behavior* (Vol. 1, pp. 61-80). Mahwah, NJ: Lawrence Erlbaum.

[Medin 1978] Medin, D. L., & Schaffer, M. M. (1978). Context theory of classification learning. *Psychological Review*, *85*, 207-238.

[Medin, Dewey & Murphy 1983] Medin, D. L., Dewey, G. I., & Murphy, T. D. (1983). Relationships between item and category learning: Evidence that abstraction is not automatic. *Journal of Experimental Psychology: Learning, Memory and Cognition*, *9*(4), 607-625.

[Medin & Ross 1989] Medin, D. L., & Ross, B. H. (1989). The specific character of abstract thought: Categorization, problem solving, and induction. In R. J. Sternberg (Ed.), *Advances in the psychology of human intelligence*. Hillsdale, NJ: Lawrence Erlbaum.

[Papert 1980] Papert, S. (1980). *Mindstorms: Children, computers and powerful ideas*. Brighton, England: Harvester Press.

[Piaget 1928] Piaget, J. (1928). *Judgement and reasoning in the child*. London: Routledge and Kegan Paul.

[Piaget 1965] Piaget, J. (1965). Les operation logiques et la vie sociale, *Etudes sociologiques* (pp. 143-171). Geneva: Librairie Droz.

[Porter, Bareiss & Holte 1990] Porter, B. W., Bareiss, R., & Holte, R. C. (1990). Concept learning and heuristic classification in weak-theory domains. *Artificial Intelligence*, *45*(1-2), 229-263.

[Rao & Georgeff 1991] Rao, A., & Georgeff, M. P. (1991). Modeling rational agents within a BDI-architecture. *International Conference on Principles of Knowledge Representation and Reasoning*.

[Ross 1984] Ross, B. H. (1984). Remindings and their effects in learning a cognitive skill. *Cognitive Psychology*, *16*, 371-416.

[Ross 1989] Ross, B. H. (1989). Some psychological results on case-based reasoning. *Proc. of the Case-Based Reasoning Workshop*, Pensacola Beach, FL.

[Vygotsky 1978] Vygotsky, L. S. (1978). *Mind in society: The development of higher psychological processes.* Cambridge, MA: Harvard University Press.

[Vygotsky 1991] Vygotsky, L. S. (1991). Genesis of higher mental functions. In P. Light, S. Sheldon, & M. Woodhead (Eds.), *Learning to think* (pp. 32-41). London: Routledge.

[Wood, Bruner & Ross 1976] Wood, D., Bruner, J. S., & Ross, G. (1976). The role of tutoring in problem-solving. *Journal of Child Psychology and Psychiatry*, *17*, 89-100.

[Wood, Wood & Middleton 1978] Wood, D., Wood, H., & Middleton, D. (1978). An experimental evaluation of our face-to-face teaching strategies. *International Journal of Behavioral Development*, *1*, 131-147.