Mephu-Nguifo, E. Baker, M.J. & Dillenbourg, P. (1999). Knowledge Transformations in Agents and Interactions: A comparison of Machine Learning and Dialogue Operators. In P. Dillenbourg (Ed.), *Collaborative Learning : Cognitive and Computational Approaches*, pp. 122-146. Amsterdam : Pergamon / Elsevier Science.

A comparison of learning and dialogue operators for computational models.

E. Mephu Nguifo, M.J. Baker & P. Dillenbourg

Abstract

This paper compares dialogue operators and machine learning operators from the point of view of understanding the mechanisms by which learning takes place as a result of collaboration between agents. Machine Learning operators are operators that make knowledge changes in the knowledge space. Dialogue operators are operators used in collaborative learning dialogues as transformation functions that allow knowledge to be co-constructed in dialogue. We describe the degree of overlap between both sets of operators, by applying learning operators to an example of dialogue. We review several differences between these two set of operators: the number of agents, the coverage of strategical aspects and the distance between what one says or hears and what one knows. We discuss the interest of fusing dialogue and learning operators in the case of person-machine cooperative learning and multi-agent learning systems.

1. Introduction.

Is there something below a machine learning algorithm? Thousands of machine learning systems have been developed, clustered in categories such as "explanation based learning", "similarity based learning", "reinforcement learning", ... Can we describe this variety of algorithms with a restricted set of elementary learning operators ? Can we integrate them into multi-strategy learning systems? Michalski (1993) addresses this question. He proposes a set of operators in knowledge space, termed *knowledge transmutations*. Transmutations are generic patterns of knowledge change. A transmutation may change knowledge, derive new knowledge or perform manipulations on knowledge that do not change its content.

This search for "atoms" is common to many scientific fields. We do observe a similar effort in dialogue studies concerning collaborative learning. Scholars attempt to understand the mechanisms by which dialogue leads to learning. Again, these mechanisms have been clustered into categories such as negotiation, argumentation, explanation, mutual regulation, ... Beyond these labels, can we describe a dialogue with more atomic operators? There exist many classifications of dialogue units, reflecting different theoretical approaches. We use here the classification proposed by Baker (1994). It includes a set of operators, termed *transformation functions*. Transformation functions describe - at the knowledge level - the relation between the contents of two utterances.

A core issue in collaborative learning research is to understand how dialogue and learning are related. Hence, naturally, we look for convergence between the "atoms" isolated in machine learning and in dialogue studies. Therefore, this chapter aims to compare the two above mentioned set of operators, transmutations and transformations. Michalski's and Baker's operator sets are not unanimously recognised in their respective scientific community as the common reference. They are treated here as examples, selected because they are familiar to the authors.

Our goal is mainly scientific: to understand the relationship between dialogue and learning, and express this understanding through a computational model. This work has also practical implications: to integrate dialogue operators into machine learning algorithms to adapt these algorithms, on one hand, to the interactions with a user, and on the other hand, to interactions with other artificial agents (learning in multi-agent systems).

Our comparison method is simple. In section 2 , we analyse collaborative dialogues with the learning operators proposed by Michalski. Section 3 applies Baker's operators for analysing collaborative problem-solving. In section 4, we draw conclusions with respect to the relationship between these sets of operators at the theoretical level. Then, we draw more practical conclusions on the interoperability of dialogue and learning operators with respect to the following goals: modelling collaborative learning (section 5), and implementing human-machine collaborative learning systems (section 6).

2. A taxonomy of learning operators

Michalski (1993) defines learning as follows: *Given* an input knowledge (I), a goal (G), background knowledge (BK) and a set of transmutations (T), *determine* output knowledge (O) that satisfies the goal, by applying transmutations from the set T to input I and/or background knowledge BK. Transmutations perform change of the knowledge space, i.e. the space where can be represented all possible inputs, all of the learner's background knowledge and all knowledge that the learner can generate. A transmutation may change existing knowledge, derive new knowledge or perform certain manipulations on knowledge that do not change its content.

To define these operators, Michalski introduces two concepts: a reference set and a descriptor. A *reference set* of statements is an entity or a set of entities that these statements describe or refer to. A *descriptor* is an attribute, a relation, or a transformation whose instantiation (value) is used to characterise the reference set or the individual entities in it. For example, consider a statement: "Paul is small, has a PhD in Computer Science from Montpellier university, and likes skiing". The reference set here is the singleton "Paul". The sentence uses three descriptors: a one-place attribute "height(person)", a binary relation "likes(person, activity)" and a four-place relation "degree-received(person, degree, topic, university)". The reference set and the descriptors are often fixed once in a machine learning system.

Generalization extends the reference sets of input, i.e it generates a description that characterizes a larger reference set than the input. It is based on inductive, deductive or analogical inference.	*Specialization* narrows the reference set of objects. It is based on deductive, inductive or analogical inference.
Abstraction reduces the amount of detail in a description of the given reference set. It is based on deduction.	*Concretion* generates additional details about the reference set.
Similization derives new knowledge about a reference set on the basis of the similarity between this set and another reference set about which the learner has more knowledge. It is based on analogical inference.	*Dissimilization* derives new knowledge on the basis of the lack of similarity between the compared reference sets. It is also based on analogical inference.
Association determines a dependency between given entities or descriptions based on the observed facts and/or background knowledge. Dependency may be logical, causal, statistical, temporal, etc...	*Disassociation* asserts a lack of dependency. For example, determining that a given instance is not an example of some concept, is a disassociation transmutation
Selection is a transmutation that selects a subset from a set of entities (a set of knowledge components) that satisfies some criteria. For example, choosing a subset of relevant attributes from a set of candidates, or determining the most plausible hypothesis among a set of candidate hypotheses.	*Generation* generates entities of a given type. For example, generating an attribute to characterize a given entity, or creating an alternative hypothesis to the one already generated.
Agglomeration groups entities into larger units according to some goal criterion. If it also hypotheses that the larger units represent general patterns in data, then it is called clustering.	*Decomposition* splits a group (or a structure) of entities into subgroups according to some goal criterion.
Characterization determines a characteristic description of a given set of entities, which differentiates these entities from any other entities. For example, a simple form of such description is a list (or a conjunction) of all properties shared by the entities of the given set.	*Discrimination* determines a description that discriminates the given set of entities from another set of entities.

Table 1: Pairs of opposite knowledge generation transmutations (proposed by Michalski, 1993)

Transmutations are bi-directional operations: they are grouped into pairs of opposite operators, except for derivation that span a range of transmutations. Two categories of transmutations are defined:

- Knowledge generation transmutations change informational content of the input knowledge. They are performed on statements that have a truth status. These transmutations are generally based on deductive, inductive, and/or analogical inference.
- Knowledge manipulation transmutations are operators that view input knowledge as data or objects to be manipulated. There is no change of the informational content of the knowledge. Examples are insertion, deletion, sorting or unsorting operators.

In the following, we restrict ourselves to the first category (table 1): changes at the knowledge level, which can later be compared to Baker's knowledge level operators. It is more difficult to relate operators which concern the form since the form of an utterance is very different from the AI knowledge representation scheme.

Derivations are knowledge generation transmutations that derive one piece of knowledge from another piece of knowledge (based on some dependency between them), but do not fall into the special categories described above. Because the dependency between knowledge components can range from logical equivalence to random relationship, derivations can be classified on the basis of the strength of dependency into a wide range of forms.

- *Reformulation* transforms a segment of knowledge into a logically equivalent segment of knowledge.
- Deductive derivation, Abductive Explanation and Prediction can be viewed as *intermediate derivations*. A weak intermediate derivation is the cross-over operator in genetic algorithm (Goldberg, 1989). Mathematical or logical transformations of knowledge also represents forms of derivations.
- *Randomization* transforms one knowledge segment to another one by making random changes. For example, the mutation operation in a genetic algorithm (Goldberg, 1989).

	(…)
[1] A1:	So what can we say if there's an inelastic impact ?
[2] A2:	Well, that the energy .... all the energy ....
[3] A1:	Well, that the kinetic energy is theoretically nil !
[4] A2 :	It's nil on arrival, in fact …
[5] A1:	Since ... since the object stops, in fact, oh yes, especially since there it doesn't move, uh...
[6] A2:	it's nil at the start and nil on arrival … about energy … yes, but at the moment of an inelastic impact, what is it that …
[7] A1:	we've been doing that for a while now ! <sighs>
[8] A2:	but we've also …
[9] A1:	wait, … inelastic impact, right, you've conservation of momentum, but … the kinetic energy isn't conserved ! I think that's what we've seen … with elastic impact, by contrast, both are conserved
[10] A2:	Yes, elastic impact, there's the total energy which is conserved …
[11] A1:	Yes
	(…)

Figure 1. Example of mutual refinement strategy in physics collaborative problem-solving dialogue.

3. A taxonomy of dialogue operators

A key issue in the study of collaborative problem-solving is to understand how jointly agreed solutions are generated in dialogue. The solutions that are jointly produced can rarely be reduced to simple 'accumulations' of individual proposed solution elements. Rather, solutions *emerge* by an interactive process in which each agent (student) *transforms* the contributions of the other, in order to arrive at a mutually satisfactory solution element. This process may be described as one by which *knowledge* is *co-constructed* by a process of *negotiation* (where the term 'knowledge' is relativised to the agents concerned, in the absence of a higher authority or arbitrator).

A model for collaborative problem-solving in dialogue based on the notion of negotiation has been described by Baker (1994). We have chosen to discuss this model because it is of course more familiar to us, but also because it was developed on the basis of detailed analysis of collaborative problem-solving dialogues (it has been validated with respect to dialogue corpora for several different tasks in physics problem-solving). Most other taxonomies of dialogue operators have been based on extensions of Mann and Thompson's (1988) model for "rhetorical relations" in texts to interactions[1].

Although we can not discuss this model in detail here, the basic idea is that collaborative problem-solving proceeds by a negotiation process, defined as a type of interaction where the agents have the mutual goal of achieving agreement with respect to an as yet unspecified set of *negotia*, under certain constraints (relating to the

[1] But see also the earlier work of Hobbs (1982) on "coherence relations" in discourse, and Sanders, Spooren & Noordman (1992) for a synthesis of different approaches on textual and dialogual relations.

problem, the social situation, the knowledge states of each agent, …). Such a final state may be achieved by three possible strategies : *mutual refinement* (each agent makes proposals, each of which are transformed by the other), *stand pat* (one agent only makes proposals, with different forms of feedback, encouragement, discouragement, …, from the other) and *argumentation* (conflict in proposals is made explicit and mutually recognised, each tries to persuade the other to accept their proposals). Although knowledge may in fact by more or less indirectly co-constructed during each strategy (e.g. during 'constructive argumentation'), here we shall concentrate on the most frequent and typical strategy that is used : *mutual refinement*..

Each strategy is defined in terms of a set of *communicative acts* and sets of *relations* (created by dialogue operators) that are established between the propositions that they express. The basic communicative acts for the mutual refinement strategy are OFFER and ACCEPTANCE or REJECTION. These are defined using Bunt's (1989) model for dialogue. OFFER's have the following most important pertinence condition (when uttered by agent A1) : "accept(A2,p) -> accept(A1,p)". In other words, they are *conditional* : A1 will accept the proposition p (a problem solution, an action …) iff A2 will do so ("I will if you will"). Acceptances and rejections have the function of allowing the agent that made the original offer to accept its own offer or not (on the basis that the other does so).

However, OFFERs and ACCEPTANCE/REJECTIONs rarely occur in isolation, but rather in sequences, and the sequential position of communicative acts produce additional secondary effects on the contexts of agents. For example, if A1 offers "We are in France", then A2 offers "we are in Lyon", then the second offer indirectly communicates acceptance of the first, in virtue of the informational (logico-semantic) *relations* between the contents of the two offers ("Lyon is in France" & "in Lyon" -> "in France"). Similarly, "We are in Lyon", followed by "We are in France" could, in certain contexts, communicate *rejection* (i.e. we are in France, but I don't agree that we are in Lyon). This is why it is also important to study the relations between communicative acts in this strategy, that - at least on the *knowledge* level - may be defined in terms of dialogue operators, or *transformation functions*.

Transformation functions (TFs) are described in terms of the *logico-semantic relations* that are established between the propositions expressed in pairs of communicative acts, either of the same speaker or between speakers. The two communicative acts do not have to directly follow each other in the dialogue. The claim that relations exist between propositions expressed in communicative acts is of course a simplification - but one that most often works - since a given proposed proposition relates in fact to *the previous context*, from the agents' own point of view. This point will be taken up in discussion.

3.1 ***Expansion TFs***

These transform the initially offered proposition by extending it in some way, either its 'scope' along a generality/specificity dimension, or in informational terms (adding a proposition, inferring a new one on the basis of the initial offer). The following are typical example TFs in this class.

Expansion TFs		
TF name	Transformation	Examples from the corpus
Generalisation	p =TF=> generalisation of p	"well, … the energy ..." =TF=> "... all the energy ..."
Conjunction	p =TF=> $p \wedge q$	"it's nil on arrival ..." =TF=> "it's nil at the start, and it's nil on arrival .."
Disjunction	p =TF=> $p \vee q$	"the volume acts" =TF=> "the volume or the mass"
Specific-value	A(?_x) =TF=> ?_x = φ (A - predicate; ?_x - variable)	"...kinetic energy " =TF=> "is theoretically nil!"
Inference	p =TF=> $p \rightarrow q$	"potential energy will increase" =TF=> "therefore the rebound force increases"

Table 2: Set of expansion transformation functions (Baker, 1994)

It should be noted that quite often, the transformation takes place in a way which is left partially *implicit* . For example, if one student offers p1, then the second may apply the conjunction TF simply by offering p2 ; if it is not mutually believed that p2 is contradictory with p1, then the second offer may be interpreted as "OFFER($p2 \wedge p1$)".

3.2 Contraction TFs

These functions are usually the inverse of expansion functions : they restrict the previously offered proposition, or render it more specific (less general). However, this is not always the case. For example, the inverse of inferring a new proposition q from a previous proposition p (p =TF=> $p \rightarrow q$) is to propose that q implies p (p =TF=> $q \rightarrow p$), in other words, to give a reason for p. This case therefore comes into the *foundational* class (see below).

Contraction TFs		
TF name	Transformation	Example
Exclusive-disjunction-choice	$p \vee q$ =TF=> p	"the mass or the density" =TF=> "rather the density"
Contra-inference	$p \rightarrow q$ =TF=> $p \wedge \neg q$	"since it rebounds lower that shows it's the friction" =TF=> "it's lower, but it's not the friction that's involved !"
Subtype	p =TF=> sub-type p	"... all the energy, ..." =TF=>"... the kinetic energy ..."
Class-restriction	p =TF=> domain of validity of p restricted	" … do you really think kilogramme ball made of rubber would rebound a lot ?" =TF=> "yes, but only in the case of a perfectly elastic rebound"

Table 3: Set of contraction transformation functions (Baker, 1994)

This contraction strategy corresponds to knowledge *deconstruction*. Although theoretically, inverses of all expansion functions could exist, in reality examples are hard to find. One possible explanation is in terms of cognitive economy : one the students have co-constructed a possible solution by expansion, if they recognise "that's not it, it can't be that !", it is easier to simply let the solution drop and start again (perhaps taking part of the previous solution) rather than to deconstruct (contract) it piece-by-piece.

3.3 Foundational TFs

These provide foundations (reasons for/against, explanations) for offered propositions. Often, this occurs at the end of an expansion phase, when the students 'step back' and attempt to verify or check the current joint proposal. For example, in Figure 1:

"it's nil on arrival, in fact..." =TF-> "… since the object stops"

Usually, *counter-reasons* indicate a shift to the argumentation strategy, although more isolated occurrences may occur within the mutual refinement strategy when weighing the 'pros and cons' of a proposal that is in fact mutually agreed.

3.4 Neutral TFs

These leave the content of the initially offered proposition either completely unchanged or else transform its meaning, expression in language or conceptualisation. They often operate at the level of negotiating *understanding* and *agreement*. This is why exact *repetitions* (nil transformation at knowledge level) usually function to confirm understanding and agreement. For example,

"… the object stops, ..." =TF=> "… it doesn't move, …".

Often this occurs as a 'termination' to an expansion phase, when the students try to summaries the current joint solution. Some transformations on the language (or terminology) are very important from a learning point of view. For example, in

"with the weight" =TF=> "… the mass"

the students pass from everyday language, and conceptualisations ("weight") to the target (scientific) language/conceptions ("mass").

4. Applying Michalski's operators for analyzing a dialogue

To evaluate the interoperability of learning and dialogue operators, we attempted to apply Michalski's operators to a short dialogue. It illustrates the way in which knowledge is transformed (negotiated) in dialogue using the mutual refinement strategy. Necessarily, it only illustrates some of the transformation functions that occur. The task involves students conducting experiments (figure 1) where they try to discover the properties of balls of different substances that could explain their rebound behaviour - in fact, the co-efficient of restitution (A1,A2 = agents/students ; [1], etc., = line numbers).

Figure 2 shows a graphical analysis of the extract in Figure 1, using Michalski's operators. The different propositions expressed by each agent (student) are analysed in separate columns; relations between them, according to the Michalski operators, are shown as labels on arrows.

General remarks. The extract begins by A1 determining the focus of the discussion : "inelastic impact" ([1]). This is accepted by A2, and remains the common focus until the second part of [9], when it shifts to the case of "elastic impact". In this first discussion on the case of inelastic impacts, the joint solution is successively transformed basically by successions of *generalization* and *specialization* of what is to be discussed within the "inelastic impact" focus - energy, kinetic energy - and by adding more details to the reference set (class of inelastic impacts) by *concretion* - kinetic energy is nil on arrival at the ground, energy but not momentum is conserved. Once this solution element as been jointly transformed as far as the agents deem necessary, A1 then moves on to consider the inverse case - elastic impact.

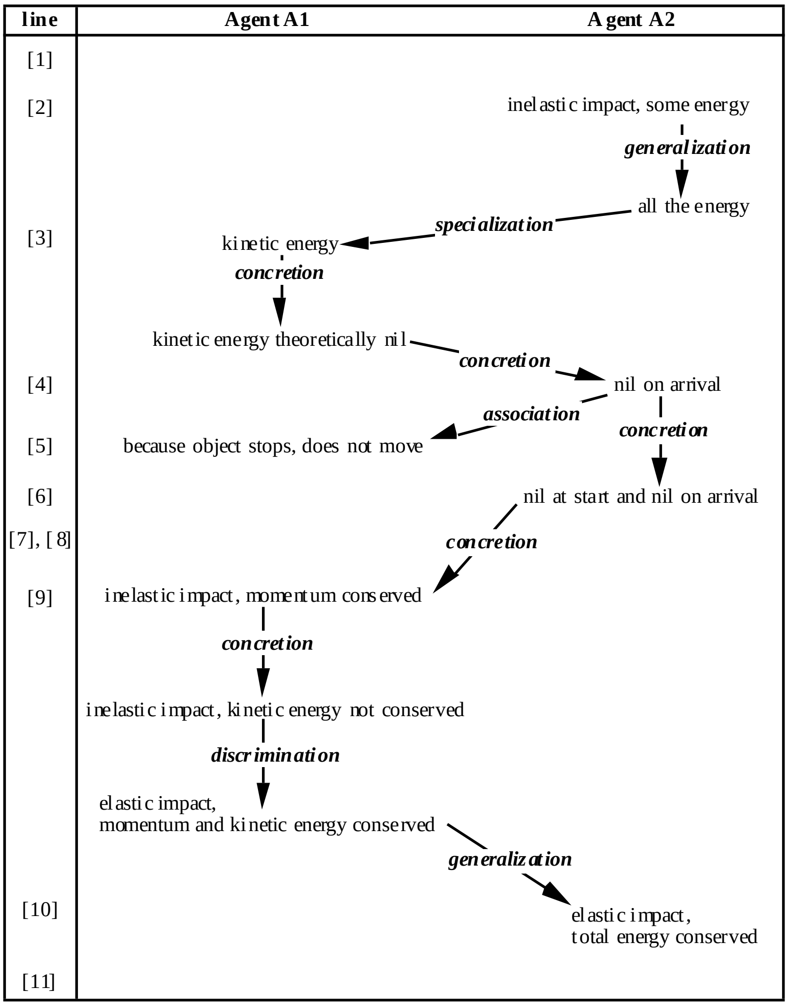

Figure 2. Analysis of example physics collaborative problem-solving dialogue using some Michalski's machine-learning operators.

Transformations and collaboration. The expressed propositions and transformations are distributed across both agents. An important aspect is 'who transforms what, expressed by whom?'. Thus, agents may transform propositions that they themselves have expressed, either within their turn or across turns. More importantly - for the study of collaborative learning - they transform propositions expressed by others. Such a graphical analysis thus provides one way of mapping out the extent to which agents are really 'working together', i.e. collaborating (Baker, 1995). Roschelle and Teasley (1995) have described some similar phenomena as "collaborative produced production rules". When agents are not collaborating (resolving in parallel), transformations are usually performed by agents on their own contributions across turns.

What is the degree of goodness of fit ? In each transformation analysed, the most appropriate operator has been chosen. Nevertheless, certain aspects are left out of the analysis, even when restricting consideration purely to the knowledge level, and more fine-grained distinctions can be made. Thus, within line [5], although agent A1 appears, on the surface, to give a reason for why the kinetic energy is nil on arrival at the ground, this does not really transform knowledge itself into new propositions ("because the object stops, it does not move"). Rather, the sequence "nil on arrival"-"object stops"-"does not move" is a sequence of *reformulations* of the meaning of "kinetic energy nil", in more everyday language. A second difficulty with applying the operators can be seen from line [9], where A1 concludes discussion of "inelastic impact" and begins to consider the case of "elastic impact". Is this transformation really discrimination (apparently the best fit), i.e. " determines a description that discriminates the given set of entities from another set of entities" ? It is analysed as discrimination here since, by stating that energy is conserved with an elastic impact, this discriminates this case from that of an inelastic impact where energy is *not* conserved. However, in the case considered here, what is important is that the agent moves from considering a class of entities "impacts that are inelastic" to considering the negation of it, i.e. "impacts that are not inelastic (= elastic)". This does not actually *transform* the knowledge base, but rather shifts attention so that all of the set of "impacts" will be described.

What is left out of the analysis ? In addition to there being finer knowledge-level distinctions that can be made with the propositions that are analysed, some utterances are of course left out entirely, since they do not express or transform new knowledge. It is interesting at this point to briefly mention what they are, for this specific example, and to describe their functions:

Elements left out of analysis :

[1] A1: So what can we say if there’s an inelastic impact ?

[7] A1: we've been doing that for a while now ! <sighs>
[8] A2: but we've also …

[10] A2: Yes<…>
[11] A1: Yes

Line [1] is left out of the analysis since it does not express a new proposition, nor transform knowledge. Rather, in proposes/introduces a new sub-problem to be *focussed* on. Lines [7]-[8] express frustration with lack of progress, from the point of view of A1, of joint problem-solving. In these terms it can be viewed as part of metacognitive control of 'knowledge transformation'. Finally, lines [10] and [11] act as a 'punctuation' to the knowledge transformation sequence. They do this given that the agents must reach agreement, that a sequence of transformations is to be terminated, and that it is so due to joint agreement. All of these aspects missed out of analysis relate to a single fact about dialogue and collaborative problem solving : it needs to be *controlled* or *managed* (Bunt, 1995).

In summary, analysis of collaborative problem solving dialogues with a specific set of machine learning operators can in fact give us a picture of how collaboration has taken place, on a certain level of generality. Aspects that are specific to (human) interactions such as negotiation (reformulation) of meaning and interaction management are of course not included in such an analysis. This fact may have important implications given that the effort required to ground interaction, to create joint meaning may be important for collaborative learning (Schwartz, 1995; Baker et al, this volume) and that interaction management problems may hinder it.

5. Theoretical comparison

The analysis presented in the previous section shows that the set of learning operators can be applied, at least at the superficial level, for describing knowledge refinement along collaborative dialogues. However, it does also reveal three differences that we analyse below: the number of agents, the multidimensionality of communicative acts and the relation between knowledge and utterances.

5.1 The number of agents

An obvious difference between dialogue operators and machine learning operators is the former relate contents of communicative acts uttered by different agents, whilst the latter relate knowledge states of the same agent.. Nevertheless - at the knowledge level at least - this difference is shallow. First, dialogue operators do function very well as monologue operators as in lines [2] of figure 2 (A2 generalizes his own statement) or in lines [3] (A1 concretizes his own statement). Conversely, single-agent learning operators can be viewed as dialogue operators in multi-agent learning: An agent A, using operator X to solve a conflict between two divergent goals, can be re-implemented into a multi-agent systems, in which two (sub-)agents A1 and A2 have respectively each of these goals and negotiate with the same operator X.

Dialogue and learning operators can be adapted for the case of, respectively two or one agent, because they are intrinsically binary operators, i.e. operators which describe the differences between two knowledge states (as stored in a system or expressed by agents), without indication whether these knowledge entities belong to one or more agents. The notion of agent is anyway completely arbitrary in DAI. An agent can be any functional unit inside the system: an 'edge detector' agent in a vision system, a grammatical parser in a language processing system,... Sometimes, a single rule is labelled as an agent, sometimes it is a whole rulebase. In psychology, there is a 'natural' notion of agent, the human individual. Nevertheless, the notion of functional agent (versus physical) emerged as well, for instance in Minsky's (1987) metaphor of 'the society of mind'. Similarly, in *distributed cognition* theories (Salomon, 1990), some agents which are physically absent from the interaction, are considered as functionally present in individual or group activities because their intelligence in embodied in the tools used by this individual or this group. Within these theories, a group of individuals itself is viewed as a single cognitive system. In other words, psychologist can also arbitrarily determine what is the "unit". Viewing a group as one or more agents, viewing an individual as one or more agents, ... are methodological choices. Different levels of granularity reveal different aspects of cognitive processes, the observer has to choose a level of analysis where can be observed the phenomena (s)he is looking for.

In summary, the fact that operators relate the knowledge states of one or two agents does constitute an intractable problem, since what is counted as one agent appears more and more to be an arbitrary decision, both in DAI and in psychology.

5.2 The multidimensionality of communicative acts

Does an utterance conveys a unambiguous knowledge change? No. Given the conversational context, the same utterance may imply different knowledge transformations. Let us consider the example below. If one simply considers the utterances, A2 seems simply to give an example of one class. But *why* did A2 provide this example ? Assuming that A1 knows that whales do not live on the ground, A2s utterance repairs A1's over-generalisation. Hence, A1's knowledge state should now refine the "mammal" concept. The *learning* transformation in A1 is therefore *generalisation* (suppressing a non-relevant feature of the 'mammal' concept), whereas on the dialogue level, the topic of discussion is made more *specific.*

[1] A1 : All mammals live on the ground.
[2] A2 : Whales are mammals.

This example illustrates the fundamental issue of the multidimensionality of speech acts and reminds us that, in Baker's framework, the content transformation is only one dimension for describing a negotiation process. In addition to domain-level relations, established by transformations, agents also transform the meaning of utterances, in order to check what was meant, to establish mutual understanding.

The gap which appears here between dialogue and learning is not intrinsic to these processes, but rather related to the taxonomies at hand. In dialogue studies, a great attention has been paid to the management of dialogue, and this attention is reflected by the definition of specific operators. At the opposite, Michalski's taxonomy does not include 'management' operators. Control is however critical in many learning processes. For instance, in incremental learning from examples, the selection of relevant examples and counter-examples directly influences the learning outcomes. The issue is to know if the operators used in dialogue management and in learning management are similar. The interoperability of learning and dialogue operators should hence distinctively discussed at two levels, the knowledge level and the control level.

5.3 The relation between knowledge states and utterances.

The previous point reveals the difficulty to match the knowledge expressed in utterances with internal knowledge states. We face 3 fundamental issues.

- Does an agent believes what he says? In general, yes. However, there is no infallible algorithm for inferring mental states from utterances (although there may be non-monotonic 'default' methods - e.g. Perrault, 1990). For example, if one student says *" … it's the different densities that explain the rebounds*", what does this mean ? That the student *believes* that proposition ? Not necessarily. In fact, we often observed that students make "proposals" - "might it be x ?" - without necessary commitment to believing them (Baker, 1994).

- Does an agent answer to a particular utterance ? Not necessarily. The agents knowledge spaces are almost never rendered entirely explicit in communication, simply because this would make communication highly uneconomical. Hence, a given utterance often does not relate *directly* to a previous one, but rather to the agents perception of its underlying context. In the example below, A2 responds in fact to the following proposition attributed to A1 (what A1 is perceived to be implying/implicating): "the steel ball rebounded higher than the rubber one, so the differences in mass explain the different rebound behaviour".

[1] A1 : look, the steel ball rebounded higher than the rubber one
[2] A2 : it's not a question of mass !

- Does an agent believes what he hears? Of course not. There is a probability that the utterance receiver (1) does not hear correctly, (2) does not understand, (3) understands but disagrees, (4) understands and agrees but cannot integrate in his knowledge base, ... There is no guarantee that agents directly and completely *internalise* (in a Vygotskian sense) the knowledge publicly constructed in dialogue.

All of these points raise deep linguistic and philosophical questions on the nature of interaction, and the relationship between communication and mental states. Hence, what we do in the remaining of this chapter is to see whether we can bypass these theoretical dead ends if we consider how these differences can be overcomed for practical goals.

6. Practical implications

6.1 Multi-agent learning

Despite these fundamental issues between beliefs and utterances, can the (partial) interoperability of dialogue and learning operators contribute to implemented learning in multi-agent systems (Weiss et al., this volume)? We examine this issue, with our own perspective where modelling human collaboration prevails on engineering issue.

What is the specificity of collaborative learning with respect to learning alone? Some well-known mechanisms such as sharing the cognitive load can be easily be translated into a multi-agent architecture. But, the more fundamental point is that, during collaborative learning, different types of *interactions* occur between learners, which seem to have cognitive effects. The real challenge is therefore to model such verbal and non-verbal interactions and their cognitive effects. This cannot be done with simple 'information flow' models, where Agent-A learns X simply because Agent-B communicates X to Agent-A. Such a model would for instance contradict the facts that the explainer often gets more benefit than the explainee (Ploetzner & al, *this volume).* Hence, it is interesting to look for models which deeply integrate dialogue and learning (Dillenbourg, 1996).

Unfortunately, dialogue and learning have traditionally been studied in different branches of artificial intelligence. Traditionally, AI has been clustered into disciplines such as problem solving, machine learning, computational linguistic, robotics, vision, tutoring systems, ... Agent architectures reflect this history: they generally dissociate some core reasoning layer (knowledge representation, problem solving, learning,...) from the interface layer (incoming and outgoing messages, vision, ...). There has been a few attempts to merge different techniques and efforts to unify AI sectors e.g. around SOAR (Rosenbloom & al, 1993). Nevertheless, the traditional AI clustering is not a good starting point when one looks for what is common between learning and dialogue.

The first step to develop an algorithm which unifies learning and dialogue is to improve the interoperability of learning/dialogue operators. People Power (Dillenbourg & Self, 1992) illustrates this principle. The artificial agent uses the same operators for reasoning and for dialoguing: agree or refute. Dialogue was based on a binary tree of argument where any argument could be agreed or refuted, where refutations could then on their turn be agreed or refuted, and so on. When the artificial agent reasoned alone, it used the same operators with itself, i.e. it was able to refute its own arguments, its own refutations, and so forth. In other words, reasoning was implemented as an inner dialogue. Learning resulted from the fact that the artificial agent replayed - *modus modendi* - some parts of previous dialogue during its inner monologues. The principle of similarity between dialogue and reasoning was applied here in its simplest way, the set of operators being extremely simple (agree/disagree). Real dialogues are of course more complex. An avenue for research is to design similar dialogue/learning mechanisms but with a richer set of dialogue/learning operators?

The "reasoning as a dialogue with oneself" illustrates the applicability of dialogue operators as learning operators. The reverse applicability, using learning operators to describe group interactions, can be implemented if one sees collaborative dialogues as the process of building a shared knowledge set. Individual utterances can hence be seen as learning operators with transform the shared knowledge set.

However, the shared knowledge space is not a publicly accessible entity, as in the case of knowledge states in a machine learning system - representations of it exist in each agent. Some discrepancies are actually important for learning, provided that they are detected and resolved. This brings us back to a point mentioned earlier concerning the multidimensionality of negotiation. In addition to domain-level relations, established by transformations, agents also transform the meaning of utterances, in order to check what was meant, to establish mutual understanding of the joint space - a process that has been described as "social grounding" (Clark & Schaefer, 1989). Such grounding or negotiation of meaning is required even in the case where agents attempted to make all of their relevant knowledge states publicly accessible.

In other words, there is something "stereoscopic" in collaboration: Agent A exploits differences between his own knowledge and his representation of Agent-B's partner knowledge. This inter-partner comparison is different from what we saw so far. Learning operators describe differences between successive knowledge states, dialogue operators between more or less successive utterances. In both cases, this sequentiality limits somewhat the scope of potential differences. At the opposite, the differential reasoning on two independent knowledge states must cover any possible difference, until an empty intersection. The set of operators used both in learning and dialogue should hence be extended to cover the mutual modelling process which supports the construction of shared knowledge.

6.2 Integrating user interaction in machine learning

There is a growing interest for integrating interaction with an expert user in a machine learning system. The goals is not properly to implement collaborative learning systems. The main motivation is to gain robustness by adding user-system interactions. Valid algorithms may produce incorrect results from correct data, simply because there is some mismatch between the way in which data is provided and the way it is processed. This is especially important as we deal with imperfect domain theories, where initial data may be incomplete or may contain many biases. Knowing this, the user can modify data that they provide to the system in order to allow it to refine the knowledge that it has already built. Of course, most machine learning systems allow very restricted types of interactions with the user.

LEGAL (Mephu Nguifo, 1994) is one of such machine learning systems. It has its foundations from a learning model reported by (Mephu Nguifo, 1995). In such model, the learning system interacts with an expert-user in order to: (1) learn knowledge from initial data provided by the expert, (2) derive new results from learned knowledge, and (3) also help the expert-user during the interpretation of learning results. LEGAL receives initial input (binary data and learning parameters) from the user. The user is considered to be an expert of the domain. The learned knowledge arises from the successive application of different transmutation functions such as specialization, selection, characterization, generalization and discrimination over the lattice structure of the binary data.

The derivation of new results is a way to validate the learned knowledge. Unseen data are provided by the user to test the efficiency of the learned knowledge. LEGAL uses different kinds of derivation transmutations for this purpose. It can either use deductive derivation or analogy inference by combining reformulation and similization transmutation functions.

As the user works on an incomplete model, the last step becomes very important since both expert and system can change their knowledge state depend on the results interpretation. To allow this, the learning model includes the notion of proofs and refutations (Lakatos, 1984; Hayes-Roth, 1986) through the mechanisms of objections. Objections are built by the system as a way of proving its results. An objection is true until the user refutes it as an explanation of the system decision. Whilst the acceptance of objection could change the user knowledge, the refutation of an objection should allow to modify the knowledge learned by the system.

This system is based on a static and cyclic process for controlling learning. It has been extended to a dynamic model of control (Mephu Nguifo, 1997). This extended model adds an indirect dialogue between the user and the system, by integrating various tools with different purposes, and which are linked together in order to dynamically control the learning process.

Nevertheless, these systems basically remain cooperative since there is a fixed division of labour between the system and the user: the former has to learn, explain, or recognize, and the latter must choose the initial data and analyse the system's results. We are seeking more collaborative systems where the user and the system can *negotiate* the data, the amount of noise, the heuristics, ... In collaboration, the system and the human share roughly the same set of actions (Dillenbourg & Baker, 1996). Hence, any reasoning step performed by the machine learning system has to be available also to the user through some dialogue function. The interoperability of dialogue and learning operators may improve the collaboration between a human agent and the system agent in two ways: (1) by increasing the modularity of functions, each agent being able to contribute of a smaller step of the learning process, (2) by increasing the symmetry of the distribution of functions, most functions being allocable either to the human or to the machine agent.

7. Conclusions

This chapter compared two sets of operators which come from different research communities. Learning operators have been proposed in machine learning where knowledge states are directly inspectable. Dialogue operators come from psychology and linguistics, they describe the content conveyed by the subjects actions, namely by their utterances. The gap between dialogue and learning operators reflects the gap between knowledge and action. For instance, we pointed out that dialogue operators cover the 'strategy' aspects of dialogue (e.g. saying something one does not believe in order to check one's partner agreement), while Michalski's learning operators do not cover these 'strategy' aspects (e.g. selecting which example to consider next). While the interoperability of learning and dialogue operators seems feasible at the knowledge level, it seems more difficult to achieve it at the strategical level.

We mentionned that learning operators could be applied to the analysis collaborative dialogues, not by modelling individual knowledge states, but by tracing the emergence of a body of shared knowledge. Our recent experiments on computer-supported collaborative work (Dillenbourg & Traum, 1997) show that this shared knowledge is - to a large extent - reified on the whiteboard, probably because the information displayed on the whiteboard is persistent and hence the best place to co-construct knowledge. Less persistent knowledge, such as decision with regard to the strategy, are not displayed on the whiteboard. Hence, this shared knowledge set becomes - to some extent - observable, which brings the dialogue setting closer to the machine learning situation.

8. References

Baker, M. (1994). A Model of Negotiation in Teaching-Learning Dialogues. In *The Journal of Artificial Intelligence in Education,* **5**(2), 199-254.

Baker, M. (1995).

Baker, M., Hansen, T., Joiner, R. & Traum, D. (1997). Learning to understand each other, the tool and the case at hand: the role of grounding in computer-mediated collaborative learning. (In this volume).

Bunt, H.C. (1989). Information dialogues as communicative action in relation to partner modelling and information processing. In *The structure of multimodal dialogue,* 47-74, Eds. Taylor M.M, Néel F, & Bouwhuis D.G., North-Holland: Elsevier Sciences Publishers.

Bunt, H.C. (1995). Dialogue Control Functions and Interaction Design. In R.J. Beun, M. Baker, M. Reiner (Eds.), Dialogue and Instruction, Modeling Interaction in Intelligent Tutoring Systems. Proceedings of the NATO Advanced Research Workshop on Natural Dialogue and Interactive Student Modeling, (pp. 197-214). Berlin, Germany: Springer.

Clark, H.H. & Schaefer, E.F. (1989), Contributing to discourse. *Cognitive Science*, 13:259-294.

Dillenbourg, P. & Self, J.A. (1992). A computational approach to socially distributed cognition. *European Journal of Psychology of Education*, 3(4):353-372.

Dillenbourg, P. (1996) Some technical implications of distributed cognition on the design of interactive learning environments. *Journal of Artificial Intelligence in Education*, 7 (2), pp. 161-179.

Dillenbourg, P., Baker, M., Blaye, A. & O'Malley, C. (1995) The evolution of research on collaborative learning. In E. Spada & P. Reiman (Eds) *Learning in*

Humans and Machine: Towards an interdisciplinary learning science. (Pp. 189-211) Oxford: Elsevier.

Dillenbourg, P. & Baker, M (1996) Negotiation spaces in Human-Computer Collaborative Learning. Proceedings of the *International Conference on Cooperative Systems* (COOP')6), juan-Les-Pins (France), June 12-14 1996

Goldberg, D.E. (1989). *Genetic algorithms in search, optimization, and machine learning*, Reading, MA, Addison-Wesley.

Hayes-Roth, F. (1986). Using Proofs and Refutations to Learn from Experience. *Machine Learning: An A.I. Approach*, T.M.Mitchell & al. Eds, San Mateo, CA, p221-240.

Hobbs, J. R. (1982). Towards and Understanding of Coherence in Discourse. In *Strategies for Natural Language Processing* (eds.) Lehnert, W.G. & Ringle, M.H. , pp. 223-244. Hillsdale, New Jersey : Lawrence Erlbaum Associates.

Lakatos, I. (1984). *Preuves et Refutations*. Hermann Ed.

Mann, W.C., & Thompson, S.A. (1988). Rhetorical Structure Theory: Toward a functional theory of text organisation. *Text* 8(3), 243-281.

Mephu Nguifo, E. (1994). Galois Lattice: a Framework for Concept Learning. Design, Evaluation and Refinement. In Proceedings of 6th Intl. Conf. on Tools with Artificial Intelligence, New Orleans (LA), November 6-9 th, IEEE Press.

Mephu Nguifo, E. (1995) Multi-Level Concept Modelling for SBL. In Proceedings of MLNet (ECML) workshop. on Knowledge Level and Machine Learning, Heraklion (Crete - Greece), April 28-29th, Eds. Dieter Fensel & al.

Mephu Nguifo, E. (1997) An interactive environment for dynamic control of machine learning systems. In Proceedings of 7th Intl. Conf. on Human Computer Interaction, San Franscisco (CA), August 24-29th. (To appear)

Michalski, R.S. (1993). Inferential Theory of Learning as a Conceptual Basis for Multistrategy Learning. *Machine Learning* 11(2-3): 111-151.

Ploezner, R., Dillenbourg, P., Praier, M., and Traum, D. (1997) Learning by explaining to oneselve and to others. (this volume)

Rosenbloom., Laird, & Newell, A. (1993) *The Soar Papers: Readings on Integrated Intelligence*, MIT Press.

Salomon, G. (1990) Cognitive effects with and of computer technology. *Communication research*, 17(1):26-44.

Schwartz, D.L. (1995). The emergence of abstract dyad representations in dyad problem solving. *The Journal of the Learning Sciences*, 4 (3), pp. 321-354.