



HAL
open science

Mentoring, Metacognition and Music: Interaction Analyses and Implications for Intelligent Learning Environments

John Cook

► **To cite this version:**

John Cook. Mentoring, Metacognition and Music: Interaction Analyses and Implications for Intelligent Learning Environments. *International Journal of Artificial Intelligence in Education*, Springer, 1998, 9, pp.45-87. hal-00197350

HAL Id: hal-00197350

<https://telearn.archives-ouvertes.fr/hal-00197350>

Submitted on 14 Dec 2007

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Mentoring, Metacognition and Music: Interaction Analyses and Implications for Intelligent Learning Environments

John Cook, *University Centre for Complementary Learning, Thames Valley University, Ealing, London W5 5RF, UK and Computing Department, The Open University, Milton Keynes MK7 6BJ, UK*

Abstract. This paper reports on an empirical study of the interactive means by which a human teacher supported higher-order, musical thinking in learners. The aim of the work presented here was to create guidelines on how to make practical use of empirical research in teaching agent development. The research overview given in the introduction is followed by an illustrative example of mentoring interactions in musical learning. An elaboration of a musical education problem is then given and a description of the mentoring approach that has been proposed to address this problem is presented. This work uses a theoretical approach to modelling mentoring interactions called the Knowledge Mentoring framework (KMf). An empirical study designed to investigate the musical education problem is reported. In the study, a human teacher attempted to support, in the learners, creative, metacognitive and critical thinking by making use of a specially designed computer-based learning environment called Coleridge. The KMf was used to guide the analysis and modelling of the study interactions. The four main results were a pause taxonomy, reciprocal modelling, seven distinct mentoring stages (with associated state transition networks), and a script of the most frequent mentoring interactions. The last section of the paper presents a discussion of the main implications of these results for Intelligent Learning Environment (ILE) design.

INTRODUCTION

This paper describes an empirical study that investigated the way in which a human teacher supported higher-order, musical thinking in learners. The study involved four sessions. In each session a teacher and one student interacted with each other and a computer-based learning environment. By using a framework derived from the literature, the interaction analysis involved a categorisation of the study data (transcriptions of the sessions) into goals and communicative acts. Frequency counts were generated for each category. Further analyses of this categorised data generated various results.

The interaction analysis framework and the study described below are part of a teaching agent design approach, described in detail elsewhere (Cook, in preparation), that aims to make practical use of empirical research in teaching agent development. Because no empirical work has yet been reported in the area of mentoring, metacognition and musical composition in higher education, an empirical study to inform the design of an artificial agent was conducted. The detailed results of such empirical work can, for example, be used to determine the weightings in preference mechanisms for making decisions about what intervention a computer-based teaching agent should make at a particular point in time. However, details of teaching agent implementation are not the focus of this paper (see Cook, in preparation). This paper describes theoretical and related empirical work, and examines the implications of this work for Intelligent Learning Environment (ILE) design. Specifically, the last section of the paper explores the implications of our empirical results for user modelling and diagnosis, instructional planning, and descriptive-prescriptive models of decision making.

The term interaction analysis is used, rather than the more common dialogue analysis, because it includes linguistic communication (i.e. speech) and non-linguistic forms of communication (e.g. musical actions that have communicative intent or pointing at the computer screen). This research takes as a starting point the following assumption: that two key

requirements for any dialogue or interaction analysis framework in AI-ED research are to (i) provide categories of goals and communicative acts that are sufficiently expressive to describe the necessary educational interactions, while at the same time (ii) being sufficiently constrained to allow computational advantage from their use. Research that aims to meet these key requirements is presented. By analysing a small corpus of educational interactions (two hours of interaction data that incorporates aspects of one and a half hours of interview data), the work detailed below aims to begin to uncover the goal-based, educational structure, for the corpus, of a small aspect of interactive music composition teaching (transforming a musical phrase) and to explore how such findings can be used in the design of teaching agents. The term Intelligent Learning Environment or ILE will be used to refer to a particular type of agent-based learning environment. ILEs emphasise the role of higher-order thinking and have an objective of engaging the student in some goal-directed, problem solving activity that the ILE 'knows' something about (known in the sense that it is believed to be correct). ILEs place a stress on learning by reflection, i.e. metacognition. In the context of ILEs, Self (1993) has proposed a conceptual architecture that can be used as a common language when investigating interaction and metacognition. The architecture, called DORMORBILE (DOMain, REAsoning, MONitoring and REflexion Basis for Intelligent Learning Environments), distinguishes four levels of agent knowledge for student modelling purposes. DORMORBILE provides a useful framework for conceptualising how interaction can take place with an ILE and is used below (see also Pilkington and Parker-Jones, 1996, for another example of its use).

Metacognition can be defined as the understanding of knowledge, an understanding that can be reflected in either effective use or overt description of the knowledge in question (Brown, 1987). By drawing on the four levels of DORMORBILE and a related paper (Self, 1992), it is possible to refine Brown's definition of metacognition and say that it has four components: the domain and reasoning levels are the object level and equate to cognition. The monitoring and reflection levels are the meta-level, and are what we mean in this research by metacognition. Before a composer can engage in metacognitive (i.e. meta-level) activities they usually have to undertake extensive training in object-level cognition in areas as diverse as musical genres, notation, harmony, melody, rhythm, counterpoint and aural skills. However, the object-level is not the focus of this research (although acquiring such object-level knowledge will have probably entailed some meta-level activity at some stage).

An agent is considered to have each of four levels. The domain level contains domain-specific knowledge (as in learner modelling). It contains facts (e.g. "register is the range of a human voice or of a musical instrument") and rules (e.g. "a distinctive transition between two motives can be achieved by a change in some musical dimension like register"). The reasoning level represents processes that operate on the domain level to produce new elements in the domain level. One reasoning level process in music is 'local grouping boundaries' (Lerdahl, 1988). A grouping can be a musical motive, several motives combine to make a phrase or a section. The process of establishing local grouping boundaries (e.g. establishing a boundary between one phrase and the next) requires the presence of salient distinctive transitions at the musical surface (i.e. the hearer must be able to perceive a change in the music). Applying the reasoning-level process 'local grouping boundaries' (to achieve distinctive transitions at the musical surface) might involve the domain-level knowledge related to 'register', specifically that of getting the violin to take over the playing of a melody from the double bass in order to achieve a shift in register, and thus indicate the transition from one phrase to the next. The outcome of this reasoning-level activity is the rule 'changing instrument from bass to violin can achieve a distinctive transition', which could then itself be pushed down to the domain level. The monitoring level acts on lower levels. Thus, a monitor may ask of a reasoning-level process 'does repetition give rise to local grouping boundaries?'. At the reflection level we try, for example, to make generalisations, e.g. 'where else in my compositions can I use repetition to achieve a distinctive transition?' The outcomes of monitoring and reflection can be pushed down to lower levels. Thus, meta-level knowledge can get pushed down the levels to become object-level knowledge.

Creative reflection is a constrained type of metacognition, and is defined in this paper as the ability of a learner to imagine musical opportunities in novel situations and to then make accurate predictions (verbally) about these opportunities (see motivating example below for an

elaboration). An example of 'imagining a musical opportunity' would be an utterance like: 'I don't want to repeat the same note twice in my phrase'. In the empirical study described below, a human teacher had an overall goal of supporting learner creative reflection. The teacher adopted a specific educational approach, termed "mentoring". Mentoring pedagogical goals were used by the teacher that intended to promote creative, metacognitive and critical thinking.

A rich source of data for analysis can be obtained by studying human teachers and learners. Face-to-face dialogues provide an unconstrained source of phenomena in pedagogical interactions. Studies of such learning interactions can potentially expand our knowledge of the detail of how a human teacher interacts to support learning. Because fine-grained data is required from interaction analysis to inform ILE design, the findings can help uncover some of the processes and mechanisms involved in learning. Commonly used goals, and actions to achieve those goals, can then be modelled in a pedagogical agent on the basis of empirical data. Alternatively, such an understanding can be used to structure the design of a computer-based learning environment so that it supports the desired learning goal. However, a very basic question raises its head: how, or to what extent, can studies of dialogue and interaction be exploited in a concrete way by designers of ILEs? It is difficult to extrapolate from human-human interaction to what can be supported in, or simulated for human-computer interaction in ILEs. Very often, type-written dialogues are not as rich as face-to-face dialogues. A good example of studies that use the former approach is an explanation from an expert to a learner via screen and keyboard in the Wizard of Oz Technique; see for example Winkels (1992, pp. 28). The study dialogue is more likely to be of a broader "bandwidth" (Fox, 1993, Chapter 8), i.e. contain more communicative information, if it is face-to-face (as in the case of the study described below). In both study approaches (type-written or face-to-face) the human expert is asked to perform the functions of a future interactive system. However, humans are able to perform teaching functions that a future system is very unlikely to be able to mimic. Thus, going from a determinate corpus (which represents a set of behaviours of what is done by one teacher) to a future system (what should be done) is problematic: what findings should be used in the future system? The work described in this paper addresses this problem by presenting a theoretically motivated framework for interaction analysis and modelling. The framework is then linked to the concrete (framework category fit to empirical data generated by the study) in order to generate results (i.e. models of interactions) that can be used to motivate the design of a future system. This work takes the view that it is not simply a question of direct transference of the corpus to a system, rather, it is one of incorporating some functionalities from the corpus, and then, extending them to what an artificial agent can do in addition.

Very little work has been done on how computers can be used to support creative, metacognitive and critical thinking in the subject area of musical composition. Although a similar point was originally made in the late 1980s by Baker (1989) and Holland (1989), it nevertheless remains true today. Furthermore, the widespread use of computer-based sequencers in music composition education has brought with it some problems. Sequencers are software based recorders of musical performance data that allow musical material to be assembled from the bottom-up, layer by layer. When using sequencers, a student composer may tend to engage in a cycle of playing, listening and editing. This process has a tendency to relieve the student composer of the need to memorise and internalise successively imagined musical material (Morgan, 1992). Traditionally, however, music composition has always required the development of memory, reflection, critical judgement and analysis, all skills that require higher-order thinking about musical materials (Sloboda, 1985, p. 118). Thus, the use of sequencers in the current training of composers in higher-education (the focus of this research) appears to limit higher-order, metacognitive thinking in some learners.

The main body of this paper is organised into seven sections. The first section gives an example of mentoring interactions and creative reflection in order to clarify their meaning and to motivate further discussion. The second section provides an elaboration of the above musical education problem and a description of the mentoring approach that has been proposed to address this problem. Part three describes the theoretical approach to analysing and modelling mentoring interactions used in this research, which is called the Knowledge Mentoring framework (KMf). The KMf draws on Speech Act theory, (Austin, 1962; Searle, 1969) which is usually taken to mean the illocutionary act: the performance of some speech action indicating

an act that the speaker makes in relation to another (e.g. an assertion or question). Part four reports on an empirical study that aimed to answer the following research question: what are the interactive means by which a music composition teacher stimulates creative reflection? In part five, seven detailed interaction analyses approaches are described. Part six presents the results and findings of applying the seven analyses techniques to the corpus. The interaction analyses generated some general findings and four main results. The first result was a taxonomy of pauses. As we describe in the later sections of the paper, this analysis of the different functions of pauses can be used as one source of evidence for the extent to which the students are actually engaging in reflective activity (creative reflection), this being the teacher's (mentor's) main objective in the study described here. Second, post-experimental cue data was used to show that learner and teacher may carry out reciprocal modelling, which involves each agent building up beliefs about a partner's beliefs. Thirdly, evidence of seven distinct mentoring stages was found. Because a detailed elaboration of these seven mentoring stages could prove an invaluable resource to an ILE, state transition networks were constructed that represent, for each stage, the sequence in which goals are pursued. Finally, a fourth result was a script of the most frequent mentoring interactions (which were abstracted from the results of other analyses). A script is defined as an interaction sequence that is likely to happen and is presented as a series of goals or communicative acts, with each agent taking turns. The last section of the paper, part seven, explores the implications of the empirical results for user modelling, instructional planning and descriptive-prescriptive models of decision making.

MOTIVATING EXAMPLE OF MENTORING INTERACTIONS AND CREATIVE REFLECTION

A small interaction is now described in an attempt to illustrate how mentoring interventions can be made to verify that a learner has engaged in creative reflection (the extract is taken from session 4 of the study described below). The following points should be noted about this example interaction. A few minutes earlier (to the interaction extract shown), the teacher had asked the learner to create a phrase that was "radically different" to the learner's first attempt (i.e. to imagine an opportunity). The learner did this and entered (into the learning environment being used) a list of transposition values that would produce a musical phrase. The learner's list was: 0 8 1 7 -12 -5 0 -4 3 4 1 8 -10 -2 -1 (the sub-section below on Coleridge explains this in more detail). Briefly, the first four numbers in the list, which are the main subject of the dialogue below, would be played back by the computer, using a regular rhythm, as follows: 0 plays C C# F# G, 8 plays G# A D D#, 1 plays C# D G G#, and 7 plays G G# C# D. The teacher then asked the learner to describe the phrase to him (i.e. the learner was asked to make a prediction), an extract from the learner's prediction is given below (the teacher's comments have been taken out):

Learner: Err, well instead of being in groups of 4 [i.e. four notes to a motive], I've made it sort of groups of 8, in places. [PLAYS AN EIGHT NOTE MOTIF ON THE KEYBOARD: C C# F# G G# A D D#.] Hopefully. Not sure if it's going to come out like that. Intervals of 8 and have two carrying on from each other, hopefully. And then there's just a couple of red herrings as well ...

Following a brief discussion about the learner's prediction, the teacher then played the musical phrase on the computer-based learning environment. The interaction shown below then took place (numbers in brackets are pauses in seconds):

Teacher: What do you think?

Learner: (1.0) It wasn't actually quite as (2.5) as I expected. But doesn't mean that

Teacher: What did you expect?

Learner: Err, I got mi, I got mi countin' wrong should be an 0 1 because I wanted it to have a sort of [PLAYS SCALE ON KEYBOARD: C C# F# G G# A D D#]

Teacher: Got your counting wrong. Yeah, so that's

Learner: I didn't want it to play the same note twice.

The first question by the teacher ("What do you think?") is a very open-ended question. The teacher is checking that the learner is really thinking about what he is doing (i.e. that he had an intention when he composed the phrase, that he is not just leaving everything to chance, that he had in fact imagined some opportunity). The learner pauses for one second before answering with an evaluative comment: "It wasn't actually quite as (2.5) as I expected". This comment itself has a longer pause of two and a half seconds embedded within it. The pauses may indicate that the learner really is thinking in some way about the question and his own creative intentions. The learner attempts to continue talking but the teacher interrupts with the question "What did you expect?". The teacher wants the learner to say why the phrase they have just listened to did not meet his (the learner's) expectations. Thus, the teacher appears to have established that the learner is thinking in an intentional manner, and now wants the learner to give an explanation about why the imagined opportunity, as specified earlier in the prediction, did not match the musical outcome (the phrase when played back on the computer). The learner responds to this continued line of questioning by the teacher with an explanation (a diagnosis) of why his phrase did not meet his expectations ("I got mi, I got mi countin' wrong ..."). The learner then plays something on the piano, the point he is trying to make is not yet clear. The teacher therefore attempts to encourage the learner to keep explaining by giving a repetition of the explanation previously made by the learner ("Got your counting wrong"). The teacher then starts to ask another question, however, the learner cuts the teacher short with an explanation of what he had intended ("I didn't want it to play the same note twice"). "It" in the learner's explanation refers to the learner's phrase.

Following questioning by the teacher, the learner has now given, verbally, a clear problem definition (of his creative intention). Although not shown above, discussion then took place between the teacher and learner in which it was clarified that the fourth note in the motive provided by the computer-based learning environment is 7 semitones above the first note (and not an interval of 8 as the learner states in his prediction above) and that it is not the 0 8, in the learner's list, that is wrong; the phrase should have a 9 (which plays A A# C# D) instead of the 7 (the first four numbers in the list would then meet the learner's creative intention of not playing the same notes twice in an 8 note motive). The learner went on to make the alteration to his list (changing the 7 to a 9). The teacher then commented favourably on the novelty of the learner's idea. The exchange shown above was, therefore, successful in helping the learner to compose in an intentional and reflective manner. That is to say, the mentoring interactions assisted the learner's attempts at creative reflection. By getting the student to verbalise how he predicted the phrase would sound, and by the use of focused questioning to find out if the student was composing in an intentional manner, the teacher used mentoring interactions to support the student's attempts at building a stronger image of his imagined opportunity.

MUSICAL EDUCATION PROBLEM AND MENTORING APPROACH

Musical education problem.

Earlier studies (Morgan, 1992; Cook, 1994) have found that some learners tend to have poorly developed recall abilities of the structure of a musical piece that they have created, or of a piece they have just heard. Sloboda (1985, p. 190) suggests that, if listeners do not develop the appropriate internal 'mnemonic' to link elements of the structure of a composition together, the limits on memory will not be overcome (i.e. normally we can memorise no more than about ten unrelated musical items). On the basis of these restricted findings it is possible to claim that, for some learners, sitting at the computer and using computer-based sequencer programs (described

in the introduction) to compose with may be the cause of a *problem*. It appears that some students may tend to let the sequencer do much of the work and do not develop creative reflection abilities and memory recall of composition structure abilities. In one of the earlier studies, Morgan (1992), the action research method was used to observe twelve successive students working with a sequencer to compose. The subjects were either 2nd or 3rd year BA (Hons) Music students who were all studying composition. The teacher-observer (Morgan) only answered technical questions from the students. In post-observational questioning, it was found that the students had poor memory recall of the structure and the detail of a piece they had just been working on. The finding that learners have poor recall of a piece they have just heard is supported by the interview extract below, which is taken from a post-experimental interview with a teacher in a pilot study (Cook, 1994) that was conducted for the research reported in this paper:

Teacher: The way she heard it. Err, and it was, I was a little taken aback by her first response. That's why I asked her to do the drawing. And I was interested to see, when she was doing the drawing, that she was thinking very much in phrases rather than in section structures. Her memory, I was very conscious, that her memory was not very well developed in that respect.

The initial argument for the use of computers in musical education was that the computer would do the mechanical work, freeing up cognitive capacity for the student to reflect (i.e. for metacognition). What we are claiming here is that this reflection does not automatically then take place; the student has to be encouraged in some way. This is similar to the findings of Foss (1987). Foss (described in Self, 1992) outlines experiments with Algebraland, a system that provides a small set of operators ('add', 'combine-term', etc.) for problem solving. Algebraland provides a trace of the student's solution path which students are supposed to reflect upon. However, in practice they rarely did this unless some further activity was devised to encourage reflection. Algebraland does not make direct interventions to encourage reflection.

Mentoring.

Mentoring is proposed here as a way of addressing the above musical education problem (i.e. as an approach to encouraging creative reflection in learners). In the literature, the term mentoring is used to mean anything from a classical conception of a 1-to-1, highly idealised relationship to a conception where the mentor aims to provide either a role model or in-depth guidance (Freedman, 1993). The word mentor first appeared as a character's name in Homer's *The Odyssey*, where 'Mentor' represented the embodiment of wisdom and acted as a guardian of the young Telemachus. In a formal learning situation, "mentoring functions can be understood as variously providing *support*, *challenge*, and *vision*." (Daloz, 1990, p. 223, original italics). The notion of 1-to-1 tutoring by a mentor or musical master has been common practice in music composition teaching since medieval times. In this work, the meaning of the term 'mentoring' has been expanded. Mentoring goals (termed pedagogical goals) are achieved by intermediate subgoals that will plan different types of interaction depending on the goal type. The six mentoring pedagogical goals (shown in italics) intend to (i) promote in the learner vision and *creative thinking*, (ii and iii) use *metacognitive interventions* to promote learner *metacognitive thinking* (monitoring and reflection), (iv) challenge the learner to think *critically*, (v) encourage in the learner *motivation* to learn, and (vi) support the *task*. Each of these mentoring pedagogical goals are now discussed in detail below.

The *creative thinking* goal draws on the metacognitive concept of going-beyond and going-above; we will now expand upon these terms. Brown (1987) has made the following crucial distinction about metacognition, which other researchers still draw upon: metacognitive knowledge-of cognition (i.e. what one knows about cognition), and metacognitive regulation-of processes (i.e. how one uses that knowledge to regulate cognition). Using an extension of this distinction, provided by Adey and Shayer (1994, p. 70-71), we can say that knowledge-of cognition refers to what individuals know about their own cognition, but that it can also involve a subprocess of *going-beyond* the present learning behaviour (e.g. mental experimentation with

ones own thoughts, the ability to imagine possible worlds or opportunities). Going-beyond refers to the what-if aspect of knowledge-of cognition. Creative thinking has aspects of going-beyond when a learner imagines a novel opportunity. A learner needs to go-beyond, as it were, before they can then go-above and make a prediction (the latter involves the regulation-of knowledge). An example of going-beyond and satisfying the mentoring subgoal 'creative imagine opportunity' (the next section gives a definition of this subgoal) would be a learner utterance like "I think I'd like to make my phrase a bit more chromatic actually".

Although a number of regulatory skills (i.e. regulation-of knowledge) have been described in the literature, two essential skills are included in all accounts: planning and monitoring (Schraw and Moshman, 1995, p. 354). Planning as part of the regulation-of cognition involves the selection of appropriate strategies and the allocation of cognitive resources that affect performance. Examples include making predictions before reading, making a prediction before hearing a musical phrase just composed played back, strategy sequencing, and allocating time or attention selectively before beginning a task. Regulation-of cognition is like *going-above*, so to speak, and looking down on one's own thinking, an internal, conscious reflective awareness about strategies and thinking about-thinking and action. An example of going-above would be what mentoring calls 'creative make prediction'; an example was given above in the motivating example.

The creative thinking goal also has some similarities with the metacognitive process of 'reflective access' (Brown, 1987; Self, 1993). Self (1993, p. 2) describes 'reflective access' as follows: "Let us say that an agent has reflective access to a piece of knowledge if it can access, describe and discuss that knowledge in a way that maps onto its actual use by that agent." If a learner is able to describe and then correctly play a series of transpositions to a motif, they would be said to have reflective access to that knowledge. That is to say, the ability to make correct predictions is equivalent to an ability to discuss knowledge in a way that maps onto its actual use. However, although the learner may be said to have reflective access to the knowledge, this access may not yet be creative reflection, which has a further requirement for imagined opportunities and predictions to occur in novel situations. Novelty as part of creativity is a contentious issue (Boden, 1990, p. 30). However, creative reflection is more than 'mere novelty'. As part of the learning process creative reflection is viewed here as a prerequisite ability that allows for the potential of 'genuine creativity' in musical composition.

Metacognitive intervention and *metacognitive thinking*, as mentoring pedagogical goals, extend Self's (1993) DORMORBILE architecture for reflection and monitoring; there is some overlap between the previously mentioned two pedagogical goals and the creative thinking goal described above. The only subgoal for *metacognitive intervention* is 'target M or Ref', which is described in detail in the next section. There are two types of *metacognitive thinking*: monitoring and reflection. Monitoring has two subgoals: 'monitoring evaluate' and 'monitoring diagnose' (both monitoring subgoals require the learner to go-above and regulating their own knowledge). Monitoring refers to one's on-line awareness of comprehension and task performance. The ability to engage in self-explanation when learning is a good example. Several recent studies have found a link between metacognitive knowledge and monitoring activity. For example, Schraw (1994) found that adults' ability to estimate how well they would understand a passage prior to reading was related to monitoring accuracy on a post-reading comprehension test. Studies also suggest that monitoring ability improves with training and practice. Monitoring goals (M) are usually indicated by communicative acts that are either evaluative (e.g. 'that doesn't sound right') or a diagnostic utterance (e.g. 'I got my counting wrong'). Reflection (Ref) by a learner may not always be translated into a linguistic act and hence poses a serious problem for interaction analysis. A silence or some vocalisation like 'umm' may be indicative of reflection depending on the context. Reflective thinking may involve 'going-beyond' type thinking and hence tends to overlap with the creative thinking goal.

Burrows (1989) has identified three approaches to *critical thinking*: the philosophical, psychological and educational. Critical thinking has traditionally been viewed from the philosophical perspective of logic and argumentation (see Walton, 1989, for example), where arguments are viewed as the representatives of internal belief-formation processes. However, there are some concerns regarding the existence of a correlation between the public aspects of arguments, and internal reasoning processes. The second view of critical thinking is held by the

psychologists, who combine problem solving with argument analysis. For example, Sternberg's (1985) definition of critical thinking requires the breaking down of intelligence into components (metacomponents, performance and knowledge-acquisition components). Once they are understood, the strategic use of such components in a particular situation is critical thinking. The third approach to critical thinking is where most educators fit. There is a serious lack of agreement about critical thinking amongst them (Burrows, 1989). Dewey's work on 'reflective inquiry' (Dewey, 1933) is often cited as the last word in critical thinking by many educators. The critical thinking goal in mentoring is based around the work of Lipman (1991). Sternberg (1985) reviews three approaches that he sees as being able to teach his three components of critical thinking; one of these is Lipman's Philosophy for Children (PC): "... no program I am aware of is more likely to teach durable, transferable skills than PC." (Sternberg, 1985, p. 23). Critical thinking is defined here as thinking that facilitates judgement because it relies on criteria, is self-correcting, and is sensitive to context (Lipman, 1991, p. 116). (See Table A, in the appendix for an elaboration of this definition.) The critical thinking goal, taken with the creative and metacognitive thinking goals described above, are compatible with the three components of critical thinking identified above by Sternberg.

Ashman and Conway (1997, p. 53) have proposed that cognitive, metacognitive, and motivational factors predict 'good' and 'poor' performance in learning and problem solving. Hence, *motivation* is included as a mentoring pedagogical goal. Motivation is the sense of a willingness to pursue activities and in mentoring has three subgoals: 'motivation extrinsic' (e.g. qualifications dependent on passing exams), and 'motivation intrinsic' (e.g. a prior interest in the subject) (Draper, 1994, p. 11), plus 'motivation encouragement', which is meant to keep interactions flowing smoothly.

A *task goal* is a description of what the teaching and learning interactions will aim to achieve, it provides a statement against which learning outcomes can be measured. Explaining the nature of a task goal would be achieved directly by communicative acts or action. Actually attaining task goals would be achieved by other mentoring pedagogical goals. That is to say, the task goal may have other pedagogical goals nested within it, and that the converse may also be true. In interaction, there will be an underlying task structure from which, at certain points, other mentoring goals branch off or become achievable. For example, at one point a critical thinking goal may be dominant (e.g. a subgoal 'probing'), however, the goal of that segment of interaction may in fact be to get the student to think about or adopt the desired task goal, e.g. use interval leaps to transpose the motive. Once the student is pursuing a task goal, other mentoring goals may then become applicable.

THEORETICAL APPROACH: KNOWLEDGE MENTORING FRAMEWORK (KMF)

The Knowledge Mentoring framework (KMF) described in this section is a theoretical approach designed by the author to address the music education problem described above. This problem can be restated as being that, when interacting with a computer-based music system, reflection does not automatically take place; the student has to be encouraged in some way. In the KMF, interactions in musical composition learning are viewed as having two related aspects: internal 'dialogue' (i.e. agent monitoring and reflective pauses) plus external interaction with other agents. The KMF thus has a very Vygotskyian (1978) conception of learning, where the teacher mediates knowledge about the society and culture so that it can be internalised by the learner. This author has proposed (Cook, 1996) that some teaching interventions have 'implicit intentions' that will vary in their purpose, depending on the metacognitive level in the learner being targeted: (i) some of which are designed to promote Reflection (Ref) in a learner, e.g. What else could you do with that scale? Can you generalise this to another area of your compositional work?, and (ii) some of which are designed to set up Monitoring goals (M) in a learner, i.e. to help them assess the progress of their own learning, e.g. Is that what you intended? Thus, for example, a teacher may reflect (internal dialogue) on which intervention (external dialogue) will best promote monitoring in a learner (implicit intention of adding a goal in the learner's head). In the motivating example given above, the teacher made various question interventions that seemed designed to verify if the learner was composing in an

intentional manner and that were *targeted* at the learner's monitoring level (an attempt was made to get the learner to accept the goal of explaining why the musical outcome did not match the expectations as described in an earlier prediction by the learner).

Pursuing the pedagogical goal of creative reflection will entail the teacher encouraging the student to verbalise their compositional ideas, i.e. to give a self-explanation (Chi, Bassok et al., 1989; Chi, de Leeuw et al., 1994) in the form of a prediction about how a phrase just created will sound when played. An example (from session 3 of the study) of a learner making a prediction about their phrase would be:

Learner: Emm, the first one [phrase] is just coming down in semi-tones, so it's going to be the motif coming down a semi-tone each time. The next one [phrase] I've done it so it leaps down back to the original, the middle C. Then up.

Becoming competent at creative reflection involves increasing success by the learner at the elaboration of detailed mental structures of musical phrases, which implies motivation to practice the building of these structures. Mozart, apparently, could be simultaneously aware both of a composition's articulated inner structure and of its overall form (Boden, 1990, p. 251). Although we can not expect every student to have the abilities of Mozart, it is reasonable to assume that their abilities at creative reflection can be improved with training. Mozart was put through a strict musical training programme by his father (and mentor) from a very early age, which may in part account for his genius: "In short, a person needs time, and enormous effort, to amass mental structures and to explore their potential." (Boden, 1990, p. 254)

The teaching agent can be viewed in this research as valuing creative reflection. Values are affective agent attitudes. Affective attitudes express values with respect to other attitudes or agent states in terms of liking or disliking (Kiss, Clark et al., 1988). The teaching agent valuing creative reflection means that the agent 'likes' to encourage the learner to 'verbalise their compositional plans and to describe their mental structures'. An agent is understood to be an integrated natural or AI system where, in "order to satisfy their values, agents derive goals from them and then form intentions to take action to reach these goals" (Kiss, Domingue et al., 1991).

In interactive dialogues like those proposed in the KMf, the goals and plans that a teaching agent may generate to promote creative reflection may get disrupted or may require adaptation to meet the needs of the learner. Co-operative interaction between a teaching and learning agent may, therefore, involve the negotiation of plans and the sharing of responsibility when executing plans. Interactions in learning can, therefore, be seen as social actions in the sense that the action is intended to have an effect on the other agent (i.e. there may be an implicit intention behind an intervention, as the motivating example above illustrated). Power (1979) was amongst the first to discuss the relationship between goals, planning and communicative actions in purposeful dialogues. Power introduced the idea of a shared goal tree, where the responsibility for various nodes is either shared or distributed between two agents engaging in dialogue relating to the joint solution of a problem. Power also emphasised the goal-directed nature of speech acts:

The general form of my solution is this: the utterance X is regarded as a goal in the speaker's mind, and this goal is related to the higher goal G by a planning tree which specifies the intermediate goals ... My account therefore has the advantage that it relates speech acts to those non-speech acts with which they have an affinity. (Power, 1979, p. 140-144)

The KMf draws on Power's suggestion of shared and distributed hierarchies of goals and speech acts. *There are three levels in the KMf*. First, we have the Pedagogical Goal level (an agent's internal goals, what Power calls "the higher goal G"), which has six (mentoring) goals that relate to: creative thinking, metacognitive interventions, metacognitive thinking, critical thinking, motivation and the task. Second is the Intermediate Subgoal level (a turn), which can, for example, include a 'probing' subgoal that is related to the critical thinking pedagogical goal, or the 'monitoring evaluate' subgoal that relates to the metacognitive thinking pedagogical goal. The third level, the utterance level, includes a 'Communicative Act' or CA (Bunt, 1989; Baker,

1994). The main CAs in the KMf are: assert, question, offer, request, accept and reject; as well as actions like playing music and pointing at the computer screen.

The general approach taken to analysis in the KMf is to take an interaction, divide it into goals and subgoals (each goal may consist of a number of agent subgoals, i.e. turns), and to then formalise each subgoal (a turn) into an utterance. An utterance is composed of a move function and a CA (Sinclair and Coulthard, 1975). A CA is the smallest monological unit, an illocutionary act realised by verbal or physical action. A move is a CA once put into the context of an interaction itself. A move has a function that specifies its purpose, i.e. the move function specifies the function of a CA in a particular context. An utterance is usually a string or unit of linguistic signs that can be separated from another unit in some way, e.g. by syntax, semantics, pauses, etc. In the KMf, an utterance is a unit that corresponds to a KMf subgoal (turn) higher up the hierarchy (at the intermediate level). A turn may be realised by a set of utterances. Each utterance is composed of a CA, which may in itself realise different move functions. The goal trees for the teacher (which are motivated by empirical work) are given in Figure 1 (parts 1 and 2).

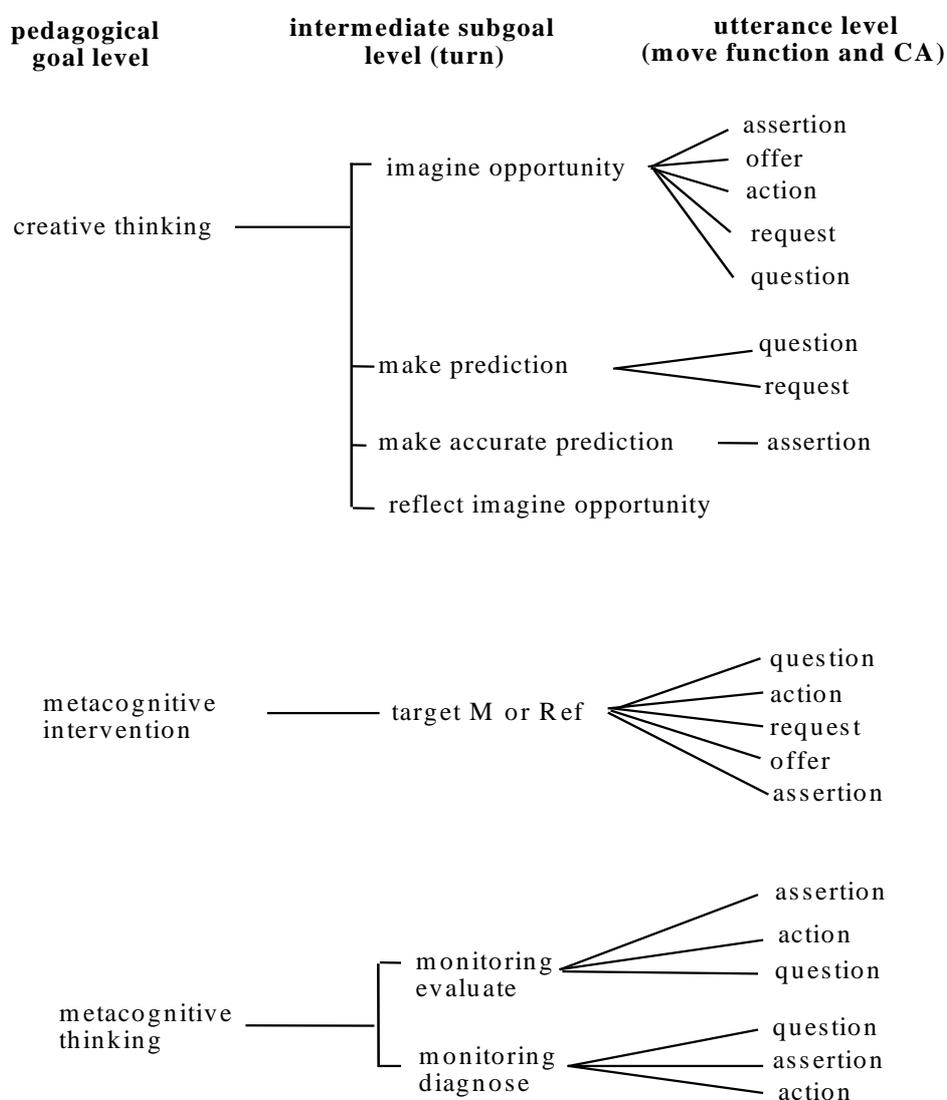


Figure 1 (part 1). Goal trees for the teacher

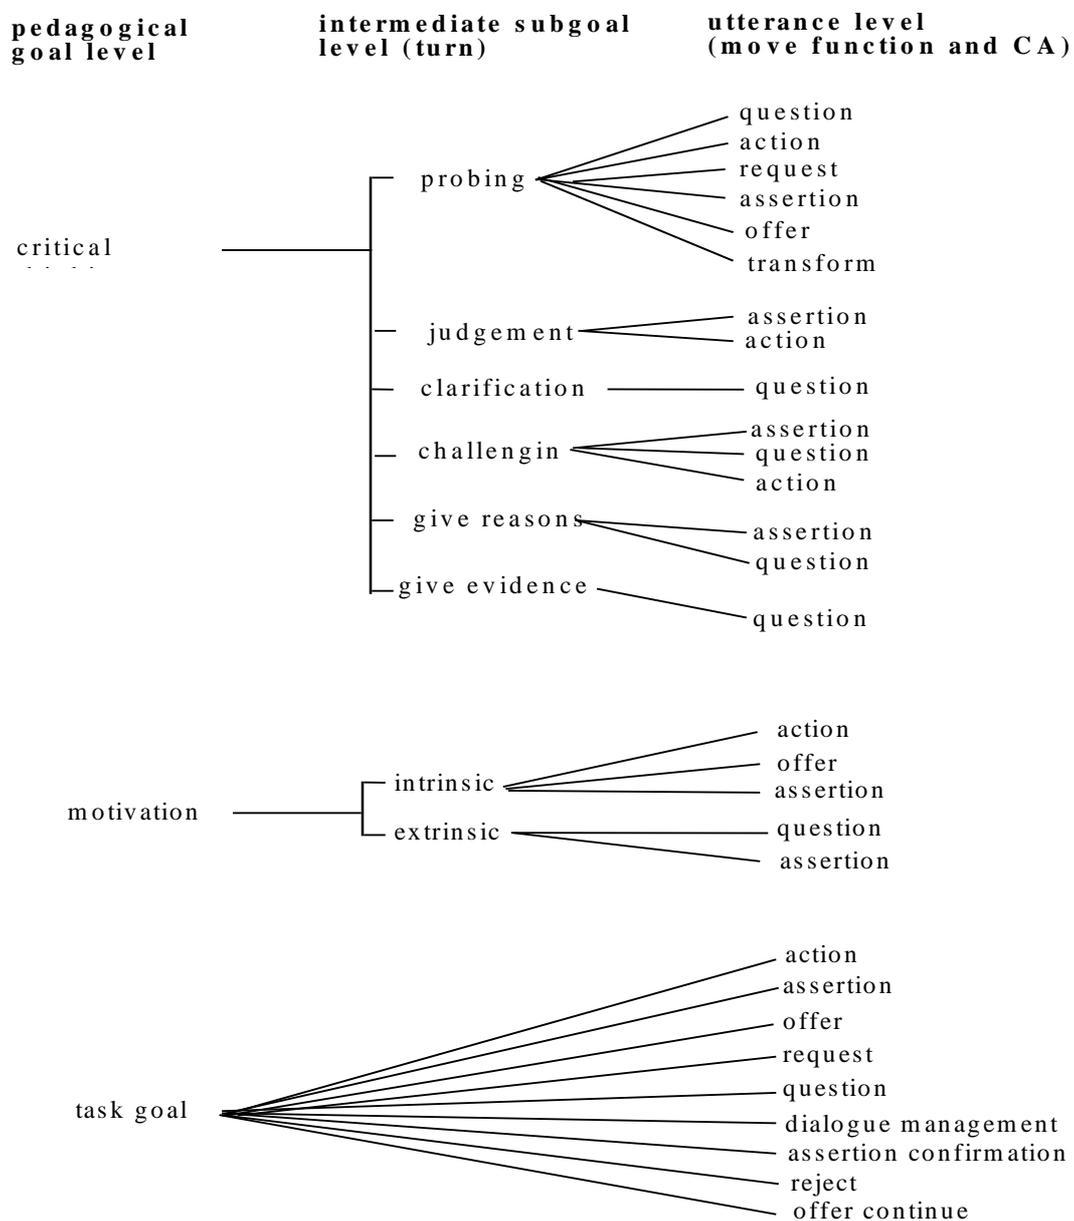


Figure 1 (part 2). Goal trees for the teacher

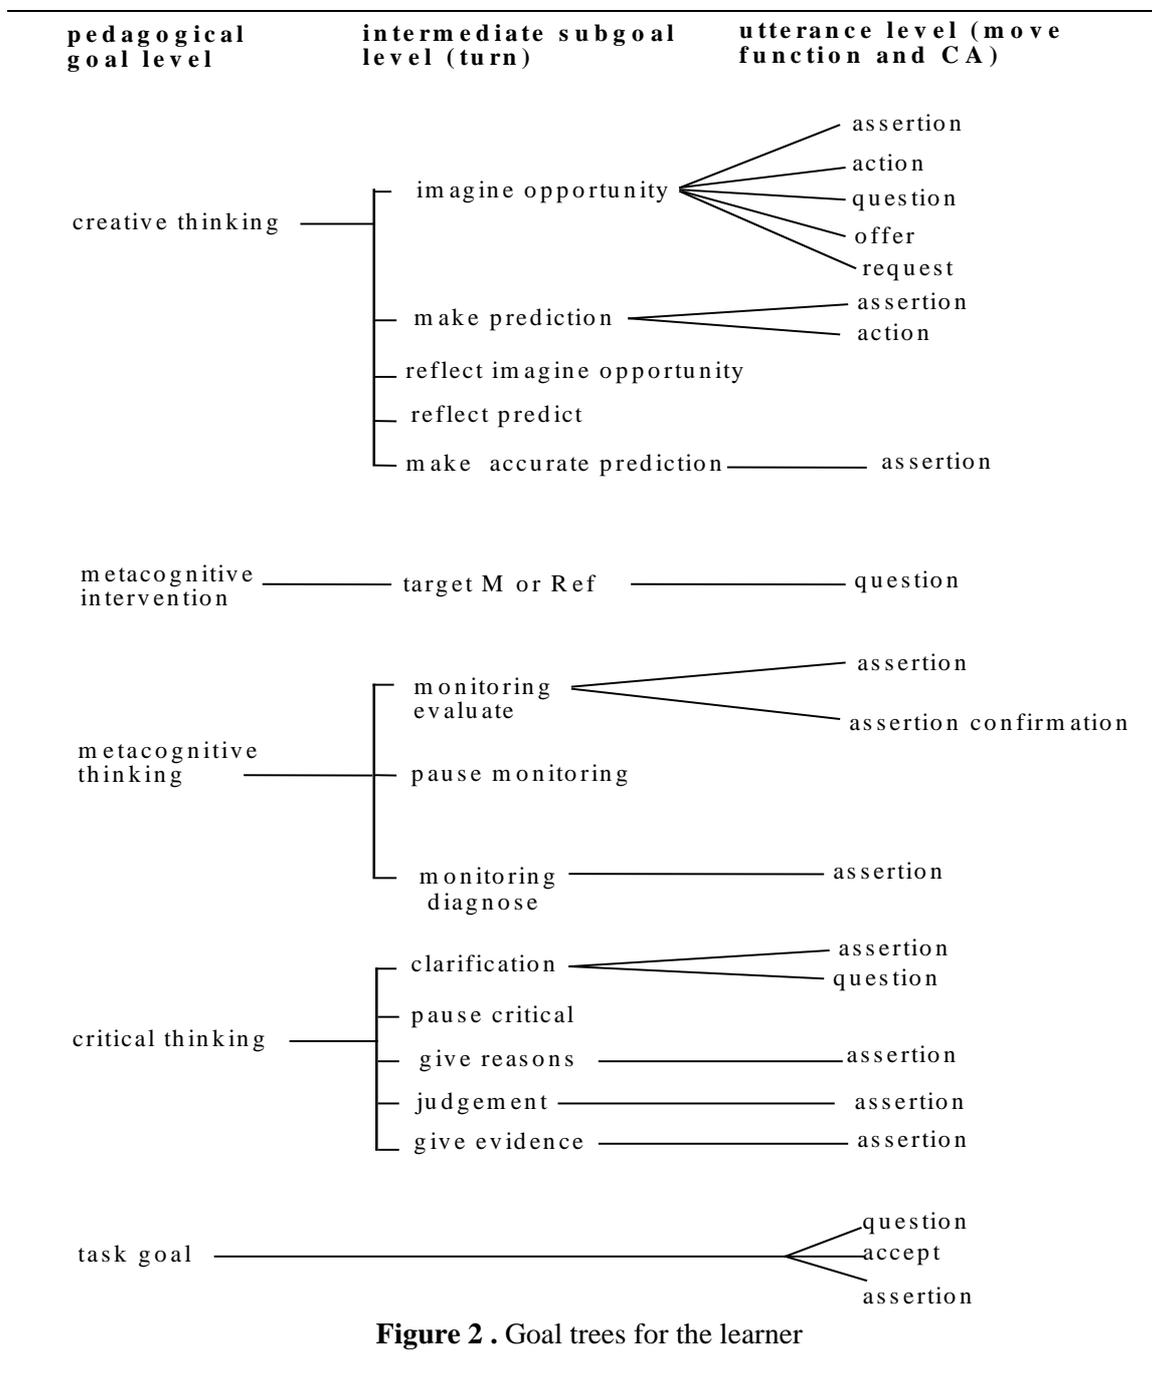


Figure 2 . Goal trees for the learner

What Figure 1 shows is a set of partial hierarchies in that no link is specified for joining the different goal trees shown (this link between goals is addressed in the empirical work below). Note that one or more different utterances may be used in an agent turn to introduce a goal, and that this feature is captured in the goal trees at the utterance level. (See the appendix for a summary of the goal, subgoal and utterance categories used in the goal trees.)

The basic structure of the learner goal trees and the utterances used to achieve them (which are motivated by empirical work) are shown in Figure 2. Again, this is actually a set of partial hierarchies in that no link is specified for joining the different goal trees shown in Figure 2. Figures 1 and 2 show that some goals may be used exclusively by one agent. For example, probing and challenging, plus motivation intrinsic and extrinsic, are used only by the teacher. At the utterance level there are some interesting differences. For example, when first discussing the task goal (Figure 2) the learner would only use three acts (question, accept and assertion). However, the teacher uses nine utterance categories when introducing a task goal (Figure 1, part 2).

In the motivating example given earlier, the teacher asks the question (his second utterance) "What did you expect?". This is a question CA at the utterance level that links to the 'target M or Ref' intermediate subgoal level (shown in Figure 1 part 1, where M = Monitoring and Ref = Reflection). It may also have the move function of "the teacher wants the learner to say why the phrase they have just listened to did not meet his (the learner's) expectations". This move function links it to the teaching agent's turn (intermediate subgoal level) of targeting the monitoring level of the learner ('target M or Ref'), which is itself linked to a higher level pedagogical goal which intends to make metacognitive intervention. In the motivating example, the learner replies "Err, I got mi, I got mi countin' wrong". This may be an assertion CA, but in the context of the teacher's previous CA, it can take on an explanation move function. This learner explanation 'move function' is also linked to the learner's turn (intermediate subgoal level) of 'monitoring diagnose', (which is shown in Figure 2) which is itself linked to a metacognitive thinking goal (pedagogical goal level) even further up the learner's goal hierarchy.

'Offer' as a communicative act in the KMf keeps the essence of Baker's (1994) negotiation in that they can be accepted or rejected. However, in the KMf there is less negotiation about what to do and more emphasis is placed on how to proceed with a task if accepted. In this respect 'offer' in the KMf has some similarities with 'offer' in Fox's (1993) approach.

There are only six primary acts in the KMf (because we are more concerned, in this work, with the goal structure of interactions). The six acts are assertion, question, request, offer, accept and reject. The use of primary acts means that certain aspects of an utterance may not be captured (the move function). For example, utterances that are acts of explaining or answering would be coded as assertions in the KMf. They are assertion CAs that realise the move functions of explanation or answering. However, the use of a limited number of acts in a framework is not unusual. In Cohen and Perrault's (1979) work, agents are represented in terms of beliefs and wants (goals). Kiss (1986, p. 15) points out that as a consequence, the scope of their work is confined to speech acts which can be characterised using beliefs and wants, i.e. requests, informs, and questions. They explicitly exclude promises and warnings.

The approach taken to representing communicative acts in the KMf thus differs from the communicative actions proposed by some other researchers. For example, Bunt (1989, p. 63-64) identifies 22 acts, or communicative functions as he calls them, which have associated with them various appropriateness conditions that specify when an act is relevant for selection. Bunt would probably call primary acts 'general functions' (Bunt, 1989, p. 64). The reason for taking this more constrained approach to the number of acts used in this analysis is, as was pointed out above, that this work is mainly concern with identifying aspects of high-level mentoring interactions. High-level interaction is concerned with the overall goal structure of interactions rather than the detailed structure of a communicative act. Furthermore, in the KMf we took the decision to use a limited number of acts in an attempt to gain future computational advantage (i.e. to reduce complexity and hence increase the potential for the goal trees to be implemented in an ILE). To summarise, one goal can be used to represent the several communicative acts; therefore analysing at the goal level would, it was hoped, enable us to analyse the structure of interactions at a useful level of detail, but for an extended length of time (over two hours of

interaction data was analysed). However, once a goal is at a sufficiently low enough level in Kmf hierarchy, it forms a commitment at the utterance level to communicate, and so a limited number of primary communicative acts are included in the analysis framework.

EMPIRICAL STUDY BASED AROUND COLERIDGE ENVIRONMENT

This section describes an empirical study that aimed to answer the following research question (which relates to the educational problem described above): *what are the interactive means by which a music composition teacher stimulates creative reflection?* Specifically, the discovery of some causal link between a mentoring intervention and learner creative reflection was hoped for.

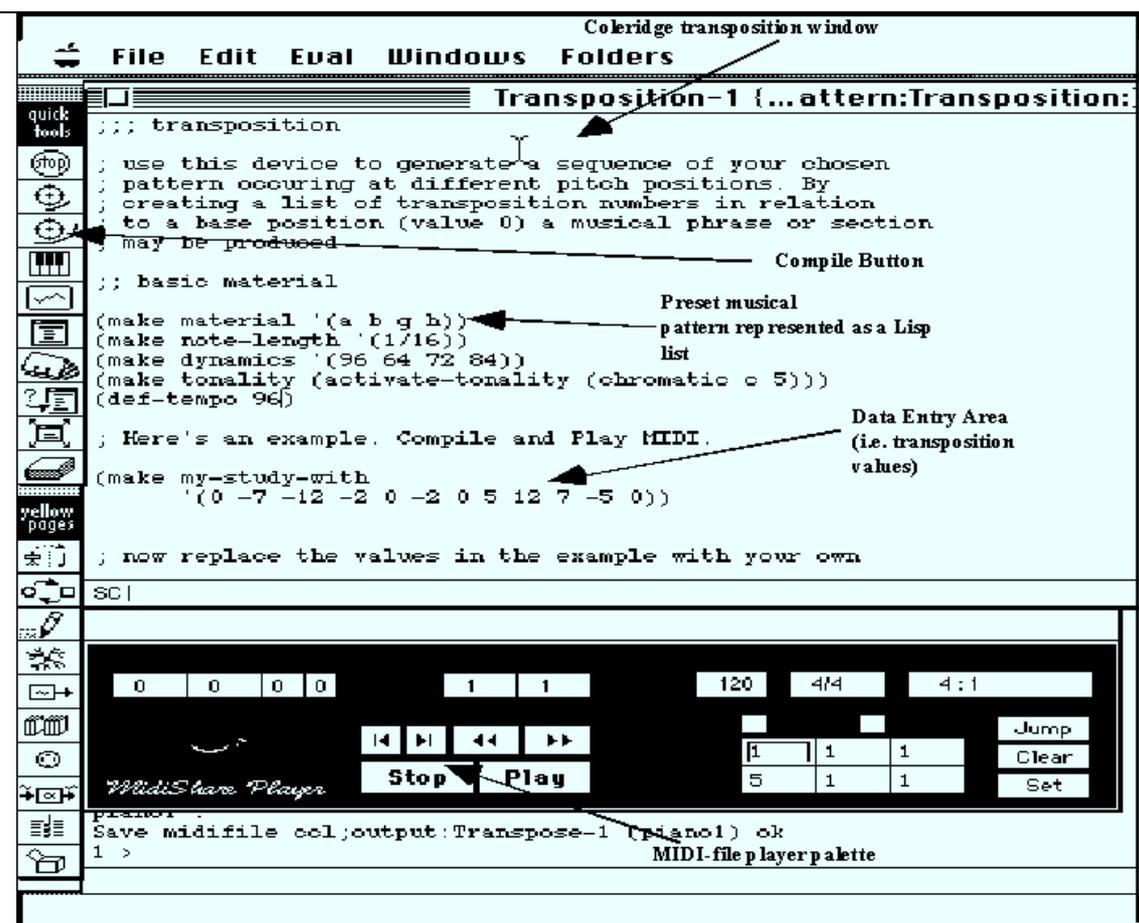


Figure 3. Annotated screen-shot of Coleridge

A learning environment for creative reflection: Coleridge.

On the basis of the considerations described above (see the sub-section 'Musical education problem'), a computer-based learning environment called Coleridge (Cook and Morgan, 1998) was built. Coleridge was developed in collaboration with an experienced composer-teacher and was design to provide fast playback of musical ideas, thus freeing up time for interaction and metacognition (it was used in the study for this reason; non-musical readers may wish to skip to the next sub-section). Coleridge was built in a Common Lisp based music composition language called Symbolic Composer (Morgan and Tolonen, 1995) and is a constrained environment in that it deals with only one small aspect of musical composition (the transformation of a musical pattern into a phrase or section). An annotated screen-shot of Coleridge is shown in Figure 3.

The technique used in Coleridge for transposing a pre-set pattern (C C# F# G is initially given in the form of the Lisp list 'a b g h') is simply to use a transposition number (which represent semi-tone steps, i.e. chromatic pitch transposition). By creating a list of transposition numbers in relation to a base position (value 0) a musical phrase or section may be produced. Zero in the list (i.e. the data entry area shown in Figure 3) will simply give a repetition of the pre-set pattern (i.e. C C# F# G), -7 transposes the pre-set pattern down a fifth (i.e. it produces F F# B C), -12 plays the whole pattern an octave lower, and so on. The 'MIDI-file player palette' shown in Figure 3 can be generated for a list of transposition numbers by clicking on the 'Compile Button'. Compiling a list thus produces a phrase or section ready for playback. Compiling and playing the list '0 -7' would produce the phrase C C# F# G F F# B C.

Research method.

The research method that was used for the study was a combination of complementary methods, involving systematic observation, post-experimental interviews, data transcription and qualitative data analysis to generate quantitative data. These methods were chosen because they provide an approach that is able to investigate in great detail the research question given at the beginning of this section.

Systematic observation is "the process whereby an observer or a group of observers devise a systematic set of rules for recording and classifying classroom events" (Croll, 1986, p. 1). The approach is characterised by an observer who is non-participative. Such a method would, it was hoped, provide fine-grained details of agent interactions appropriate for informing future ILE design.

Interviews were used as part of the research method to provide answers to pre-set questions and to elicit elaborations on 'incidents' in the interactions that the researcher-observer found interesting. In the study, the researcher and interviewer (the author) would take detailed notes, as a session progressed, of interactions that seemed interesting (e.g. interactions that seemed aimed at promoting metacognitive thinking). These notes were used as cues in the form of questions in the post-experimental interviews with the learners and the teacher. Ericsson and Simon (1993, p. xlix) have suggested that following the completion of a task, cueing distinct 'thought episodes' is a useful way to approach the gathering retrospective verbal reports. This involves constraining the retrospective report by the subject to the recall of distinct thought episodes. Post-experimental questioning runs the danger of allowing the subject to report more than their past thoughts, they might resort to speculation and inferences. However, using cues taken from the experiment can diminish this problem. Such a method would, it was hoped, provide aspects of the cognitive and metacognitive dimension of interacting agents, e.g. why did the teacher choose to make one intervention and not another, what knowledge and reasoning (internal dialogue) did the teacher employ when reaching a decision or adapting a plan to meet the user's needs?

Protocol transcription (coding) was based on the approach used by Fox (1993), although in our study an utterance could be verbal or musical. Table 1 below gives an example of the transcription approach (see column 1).

An important point to note, about our approach to coding, is that an agent turn (intermediate subgoal level) could potentially result in more than one category in a goal tree receiving a score. If more than one level in the hierarchies described in the above section on the KMf was involved in an agent's turn, then that turn (or part of that turn, i.e. a communicative act at the utterance level) received a score of 1 for each goal and act category involved. For example, if we were to code turn TA3 in Table 1, then the teaching agent would score as follows: 'target M or Ref' would receive a score of 1 and 'question' would also receive a score of 1.

If the occurrence of a turn (intermediate subgoal level) was identified in the data, it was coded only once and a communicative act or action was also associated with that turn. The subsequent interactions that aimed to satisfy a turn subgoal were coded as communicative acts only, until a new goal was encountered. Utterances LA7 and LA8 in Table 1 provide an example of this approach to coding. At LA7 'creative imagine opportunity' scores 1 and 'question' scores 1. At LA8, the 'creative imagine opportunity' turn subgoal is still active and

does not score, 'request' scores 1. Thus, if an agent's turn makes use of more than one communicative act (i.e. several utterances), then each individual act receives a score of 1.

The motivation goal was an exception to this rule (each occurrence as a turn received a score of 1 because the content of each utterance tended to be similar). This decision was taken based on the authors own intuitions. However, with respect to motivation encouragement (which is meant to keep the conversation moving smoothly, i.e. it serves a dialogue management function), this decision can be supported by drawing on Bunt's (1989, p. 70) observation that

there is some evidence suggesting that dialogue control acts should perhaps not be treated as changing the speaker's and hearers states in exactly the same way as factually-informative acts. In contrast with factually-informative acts, dialogue control acts only have a "local" function in a dialogue, losing their significance almost immediately after they have been performed.

Turns that did not easily fit into one of the predetermined goal tree categories (i.e. the categories shown Tables A and B in the appendix) were coded 'other' (and a category was created that described the apparent function of the turn, e.g. 'incomplete utterance'). These 'other' categories represented adjustments that had to be made, once analysis had commenced, due to a poor fit between interaction data and the predetermined analysis categories (i.e. Tables A and B in the appendix).

HyperRESEARCH™, a qualitative analysis software package, and various spreadsheets were used in the data analyses. The analyses of qualitative data (transcripts of the sessions) using computer-assisted methods, generated quantitative data. Various approaches were used to analyse the data. These different approaches are described in the next section.

Table 1. Example of coding[†]

Agent turn-utterance	Subgoal (turn)	Act (utterance)	Other	Commentary and post-experimental cue data
TA3: Is the tempo, satis//factory?	target M or Ref	question		Comment: The teacher prompts the learner with a rephrasing of the TA2 question (not shown). Transcription comment: // indicate an overlap utterance at LA4
LA4: Yeah, well, well // I'll hear it.		offer		Comment: The learner suggests that the way in which he would like to proceed is to hear the piece first. Transcription comment: // indicate an overlap utterance at TA5.
TA5: Lets hear it. [PLAYS PHRASE]		accept	action	Comment: The teacher accepts the learner's offer of the next step to take and clicks on play.

LA6: (9.07) I think, (1.3) reflect imagine opportunity

Comment:
These pauses are allocated to the category shown in the goal column on the basis of the context, i.e. in the next two utterance at LA7 & LA8, the learner attempts to satisfy the turn subgoal of 'creative imagine opportunity'. The learner is therefore deemed to be reflecting about this opportunity in this turn.

LA6:
(Continued)

Post-experimental cue interview data:
The learner: "Umm, soon as he said 'what about the tempo' I started considering different tempos, how it would sound with different tempos. And although my initial reaction [was] 'what do you mean by a tempo?', what was meant by that was: do you mean in terms of lengthening the phrase in parts. Now I just used the one full phrase and something that I normally do whether experiment, suddenly hearing that phrase but say twice as, lasting twice as long as before and that can have an impact as well. And I think in terms of my phrases, which were just plonked in very high, I think it would add to the impact of the surprise if they were suddenly long and sort of add to the suspense as well. So that's what I meant. And yes I considered, I thought at first when he said tempo I thought 'he's thinking that that I would like it slower', would I like it slower? And thought about it and my answer was that I actually think that it would be better faster."

LA7: depends whether you're talking about tempo across the whole thing? creative imagine opportunity question

Post-experimental cue interview data (this is what the teacher reports he was thinking when the learner made this, LA7, utterance):

The teacher: "Well I evaluated that response as meaning are you talking, are you talking about the general tempo of the piece, or the tempo within particular sections, phrases in the piece. Umm, that's what it meant."

Comment:
The above data gives an indication of the teacher monitoring that occurs whilst the learner is talking. This represents the beliefs that the teacher has about the learner

LA8: If your (LA7 goal request
gonna change still active)
tempo // try it,
try it a bit
quicker.

Transcription comment:
// indicate an overlap by the action at
utterance TA10 (not shown).

Comment:
As the teacher would have had to reset the
tempo in Coleridge himself the 'request'
CA is correct in this context (it has a move
function of a learner directive for teacher
to do something).

† TA = Teaching Agent; LA = Learning Agent

The study.

The study took place at a British University College in November 1996. The aim of the study was to answer the following research question: what are the interactive means by which a music composition teacher stimulates creative reflection? Four teacher-learner sessions were conducted with third year BA Combined Studies students (each of who were taking a music specialism). Out of a total of four students, three were male and one female. Each learner-teacher session lasted about 30 minutes and was recorded on two video cameras and audio tape. Teacher and student sat in front of a composer workstation, which consisted of a Macintosh PowerPC with Coleridge installed, an electronic keyboard and a mixing desk attached to speakers. The learners were interviewed immediately after a session for 10 to 15 minutes. The teacher was then interviewed for 10 to 15 minutes. The following instructions were given to the teacher a few weeks before the session: "Please interact with the students for 30 minutes and try to promote the learning outcome of creative reflection by using a process of mentoring. The interactions should assist the learner in the process of problem seeking. Use the task described below and Coleridge as the basis for interaction with the students." The teacher was familiar with these terms but was nevertheless given definitions. (See Appendix 1 of Cook, 1997, for a detailed description of the materials used in the study, which includes details of the instructions and definitions given to the teacher.)

The task was to ask the learner to generate, by transposition of a 4 note pre-set pattern, a musical phrase (specifically chromatic transformation of the pattern). Slonimsky pattern No 1 (1947), which is C C# F# G, was given by Coleridge at first. No alteration of the rhythm was allowed (although such a possibility could be discussed). The overall tutorial task goal was 'reflecting on the inner structure of a musical pattern'. There were three task subgoals associated with the overall tutorial task goal: for the teacher to elicit an example of structural content from the learner's phrase, second to critically analyse (jointly) the phrase, and finally that of encouraging the learner to place this phrase in the context of a whole musical section.

The following goals were involved in the study. The teacher intended to promote learner creative reflection by using mentoring pedagogical goals. The learner would, it was hoped, accept the task goals and other goals related to mentoring; the learner would then go on to make attempts at creative reflection. The discovery of some causal link between a mentoring intervention and learner creative reflection was hoped for.

In post-experimental interviews the teacher and three of the four learners said that they did not feel that the observation setting had exerted an undue influence on their behaviour. However, learner 2 reported that he did feel that the observation setting had exerted an undue influence on him. Given this suggestive evidence, it is assumed that the corpus collected was a reliable record of teacher-learner interactions. The students normally received composition tuition on a one-to-one basis (often in the room where the study took place) and the sessions observed will not have been too unusual for them. However, this must be weighed against the knowledge that being observed will tend to exert some change of behaviour on the object of observation.

Approximately two hours of teacher-learner interactions were transcribed and analysed using the approach described above. One and a half hours of post-experimental interview data was also transcribed and extracts incorporated into analysis four (which is described below).

ANALYSES

The empirical work involved seven detailed analyses of the corpus collected in the study described above. By using a framework derived from the literature (the KMf), the first interaction analysis entailed a categorisation of the study data (transcriptions of the sessions) into intermediate level subgoals and communicative acts. Frequency counts were generated for each category. Further analyses of this categorised data generated various results. Each of the (interlinked) analysis types are now described below.

First analysis, quantitative results from qualitative data

The aim of this analysis was to identify the extent to which different mentoring goals and subgoals were pursued. In the first analysis, interaction data from the four sessions was analysed using the approach described above in the research methods section. The only categories coded were subgoals at the intermediate level (turns) and communicative acts (at the utterance level, see Figures 1 and 2). The above sub-section called 'Mentoring' described these categories. See the appendix for a summary of the goal, subgoal and communicative act categories used in the interaction analysis. Quantitative counts of the total number of occurrences of each category in the qualitative interaction data were generated by analysis one.

Second analysis, steps leading to learner creative reflection

The question being examined here was: what teacher interventions lead to learner creative reflection? The second analysis involved an examination of two important learning agent subgoals (chosen because they relate to creative reflection): 'creative imagine opportunity' and 'monitoring diagnose'. The categorised data produced by analysis one was re-analysed to locate each occurrence of these two important learning subgoals. For each occurrence of these subgoals, we backtracked through the categorised data to the 6th utterances (this was an arbitrary choice) preceding the occurrences of these important learning subgoals. (Note that we excluded 'actions' from this analysis, except for data entry by the learner.) Teaching interventions in these '6 utterance chunks' were then analysed in an attempt to detect any recurring patterns (i.e. common teacher intervention(s) that may have lead to either learner 'creative imagine opportunity' or learner 'monitoring diagnose', or both).

Third analysis, learner activity leading to all teacher responses

The third analysis of data examined teacher interventions to ascertain which learner intermediate subgoal or communicative act preceded it. Such an analysis would, it was hoped, throw light on the secondary research question: Given a particular student intervention, what are the common forms of teacher responses? This third analysis involved an examination of all teacher turn subgoals and communicative acts (of the categorised data produced by analysis one), and then stepping back through the analysed data to find the first occurrence of a learner turn and/or act. If an act was associated with a turn (e.g. 'question' associated to 'critical probing') only the turn subgoal was recorded (in the previous example 'critical probing'). This relationship was then recorded on a spreadsheet (i.e. learner intervention x leading to teacher response y would increment by 1 the cell xy). Learner utterances and actions that lead to a teacher intervention can be seen as providing detail of the exact way in which a teacher adaptively promotes creative reflection.

Fourth analysis, post-experimental cue data

In the fourth analysis, post-experimental cue data was incorporated into the interaction data (the categorised data produced by analysis one) in an attempt to enhance the analysis (i.e. to include the cognitive and metacognitive dimension of the interacting agents). For a discussion of cue data see the 'Research method' sub-section. Various large extracts were analysed using this analysis approach, a small example of which is provided in Table 1 at LA6 and LA7.

Fifth analysis, mentoring stages

Analyses one to four provided considerable detail of interactions from a micro-level. The fifth analysis attempted to pull out a macro view of interactions (of the categorised data produced by analysis one). The turn subgoals (communicative acts were excluded from this analysis) for each session were analysed to see if any stages within a session could be detected.

Sixth analysis, state transition networks

The sixth analysis involved the mapping-out of various state transition networks to represent the sequence in which goals and subgoals are pursued in interactions. Analysis six took as its starting point the result from analysis five (i.e. the seven mentoring stages, which are discussed in the results section below).

Winograd and Flores (1986, pp. 64) demonstrate an approach to representing networks of speech acts which they call 'conversations for action'. These state transition networks can form the basis for computer tools (Winograd, 1987-1988) and is thus a useful approach, which has been adapted here to meet the analysis needs of this research. The approach provides a formal representation of the interplay between speech-act illocutionary point (Searle, 1976) 'directives' (in KMf questions and requests) and speech-act 'commissives' (in KMf offer, accept and accept-confirm) that are directed towards some explicit co-operative, teaching action. It is important to note that the networks presented below are descriptive and not normative, as are Winograd's (1986).

In the state transition networks (e.g. see the illustrative example in Figure 4) the course of an interaction can be plotted using circles to represent a possible state of the agent interactions and lines to represent goals and communicative acts (Winograd and Flores only represent acts). The diagrams do not represent a model of the mental state of a speaker or hearer, but show the interactions as a 'dance' between agents using acts or subgoals to achieve a goal. The lines indicate goals and communicative acts that can be taken by the teaching agent (TA) and the learning agent (LA). Each act in turn leads to a different state, with its own space of possibilities. So for example, in Figure 4, the teacher may have made an 'offer' of how to proceed or some 'request' regarding a task goal (the arrow from A to B). In the normal course of events, the learner (LA) 'accepts' the teacher's 'offer/request' regarding a task goal (perhaps performing some 'action' and moving to the state labelled C in Figure 4). If the teacher declares satisfaction with the act 'accept confirm', the interaction episode reaches a successful completion and both agents reach state D. If the learner (LA) 'rejects' the teacher's 'offer/request' regarding a task goal, the interaction episode reaches a state of failure (state E).

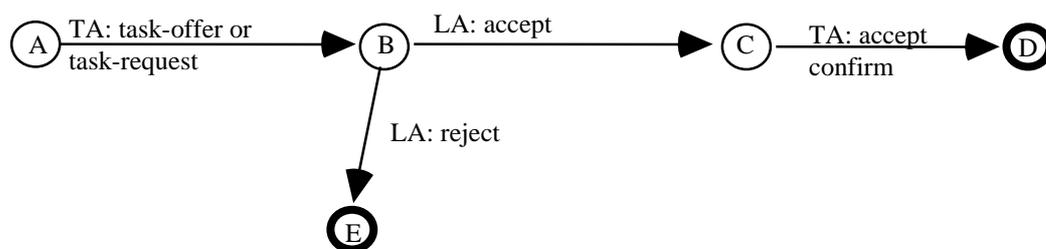


Figure 4. Partial network of acts and relations for achieving task goal

However, at state B the learner may 'reject' the offer or a request made by the teacher. There are a few states of completion from which no further actions can be taken (these are the heavy circles in Figure 4). All other states represent an incomplete interaction. Completion does not guarantee satisfaction, for example if the 'offer' is 'rejected' at state B by the learner in Figure 4.

Seventh analysis, mentoring script.

By drawing on empirical data from some of the earlier analyses, this analysis attempted to generate a script of the most frequent interactions. A script is defined here as an interaction sequence that is likely to happen (the script result is presented below in the next section). A script is presented as a series of intermediate level subgoals (a turn) or communicative acts, with each agent taking turns. The likelihood (a frequency) of a goal occurring is given as a percentage score. The scores were derived from analyses 1 to 3.

RESULTS AND FINDINGS

A brief discussion of some general findings is given below. Four main results are then reported and discussed in detail. The four main results were a pause taxonomy, reciprocal modelling, seven distinct mentoring stages (with associated state transition networks), and a script of the most frequent mentoring interactions.

General findings.

The scores in Tables 2 and 3 below were generated by analysis 1 and show the total occurrences of a category for all four sessions. The score for one category is split between that gained for the learner (shown in the LA column), that gained for the teacher (shown in the TA column) and the total occurrences for that category for all four sessions (the total column).

Some general findings (from *analysis one*) relating to the interactive means by which the teacher attempted to support creative reflection can be reported.

Some attempts were made by learners at 'creative imagine opportunity' (score = 26, Table 2). The mentoring approach taken and the support given by Coleridge appears to have encouraged some learner creativity. Students did not seem able to make accurate predictions, in spite of teacher support and a computer environment design to assist this process. The students did make some attempts at making a prediction; Table 2 shows 9 attempts by the learners at 'creative make prediction'.

However, Table 2 also shows that only 1 out of 4 learners met with success at 'creative accurate prediction'. One possible reason for this may be that, because this form of creative reflection training was new to the students, it was only by the end of the 30 minute sessions that the students had become accustomed to the idea of going-above their own thinking and making predictions (i.e. the regulation-of cognition in the form of planning and predictions).

A popular teacher intervention was 'target M or Ref'; Table 2 shows that it was used on 29 occasions. 'Target M or Ref' is an attempt to get the learner to accept a goal and pursue it, e.g. to get the learner to give an explanation by pursuing the goal 'monitoring evaluate'.

Some internalised self-monitoring took place. The interaction analysis approach recognises two metacognitive thinking subgoals that were related to monitoring. 'Monitoring evaluate' was dialogue that involved some evaluative comment by the learner about the match between a prediction and an outcome (score = 40, Table 2). 'Monitoring diagnose' was an attempt by the learner to diagnose why something did or did not work (score = 12, Table 2). Table 3 shows that there were 29 occasions that were coded as 'pause monitoring' (which is where the context strongly suggests that pauses are indicative of learner metacognitive activity related to monitoring). These findings are encouraging in that the mentoring approach seems to have promoted the monitoring effect, which we would claim is the first step towards creative reflection, and in particular the ability to make accurate predictions.

Table 2[†]

Mentoring subgoals occurrence scores

Goal and subgoal(turn)	LA	TA	total
<u>Creative thinking</u>			
creative imagine opportunity	26	13	39
creative make prediction	9	6	15
creative accurate prediction	1	2	3
<u>Metacognitive intervention</u>			
target M or Ref	1	29	30
<u>Metacognitive thinking</u>			
monitoring evaluate	40	3	43
monitoring diagnose	12	1	13
reflect predict	4	0	4
reflect imagine opportunity	8	1	9
<u>Critical thinking</u>			
critical judgement	2	28	30
critical probing	0	44	44
critical challenging	0	3	3
critical clarification	25	4	29
critical give reasons	9	3	12
critical give evidence	2	2	4
<u>Motivation</u>			
motivation intrinsic	0	11	11
motivation extrinsic	0	4	4
motivation encouragement	7	70	77
<u>Task goals</u>			
task	8	43	51

[†] TA = Teaching Agent and LA = Learning Agent

Table 3[†]

Communicative acts and other occurrence scores

Utterance	LA	TA	total
<u>Communicative Acts</u>			
assertion	243	237	480
assertion confirmation	106	186	292
question	41	150	191
request	11	62	73
offer	6	71	77
offer continue	2	34	36
accept	51	28	79
accept confirm	20	3	23
reject	3	5	8
<u>Other</u>			
pause critical	22	---	22
action	87	344	431
no category obvious	---	---	18
continuation	---	---	17
dialogue management	20	23	43
incomplete utterance	37	35	72
not accept yet LA	13	---	13
pause monitoring	29	---	29
retraction	1	---	1

[†] TA = Teaching Agent and LA = Learning Agent

The results show that the two most frequently used teaching interventions related to critical thinking were 'critical probing' (score = 44 for the teacher, zero for the learner, Table 2) and 'critical judgement' (score = 28 for the teacher, 2 for the learner, Table 2). Learner critical thinking involved 'critical clarification' and 'critical give reasons', which Table 2 shows occurred 25 and 9 times respectively. Interaction relating to critical thinking appeared to be

teacher led. The teacher would initiate such a pedagogical goal with say a 'critical give evidence' turn subgoal and related 'question' communicative act. The learner would typically respond with a 'critical clarification' turn and related 'assertion' communicative act. Thus, some goals belong exclusively to certain agents (e.g. 'critical probe' was only used by the teacher) and others belong almost exclusively to one agent (e.g. 'critical clarification' was used by the learner on 25 occasions but by the teacher only on 4 occasions). 'Critical give evidence' was used by both agents on 2 occasions each.

'Action' (one of the 'other' turn categories in Table 3) is where an agent plays music on a keyboard, uses the mouse to point to something under discussion on the screen or to click on an icon, it excludes all action performed by the computer (e.g. displaying a window). With a score of 344 'actions' (Table 3) the teacher was clearly very active and with 87 'actions' (Table 3) the learner was certainly busy. At face value the teacher may be open to criticism for not letting the learner be more active (there is an imbalance in the scores). However, it has to be said that music, like sports, is a coaching-like subject and the coach-teacher-mentor may tend to automatically reinforce what they are saying with musical actions (e.g. playing a motive whilst discussing it). By re-analysing all 'actions' it became apparent that the teacher was making many musical actions (i.e., either playing back a musical phrase on Coleridge, playing at the keyboard, singing or humming). The teacher made a total of 109 musical 'actions' in all four sessions (i.e. 31.6% of all teacher actions were musical).

The total number of teaching goal interventions by the teacher is the sum of the TA column in Table 2 (excluding motivation encouragement, which is seen as part of dialogue control). The total number of teacher goal interventions was 197. The total number of goal interventions by the learner is the sum of the LA column in Table 2 (again, excluding motivation encouragement, which is seen as part of dialogue control), plus the scores for 'pause critical' and 'pause monitoring' in Table 3 (these are 'other' categories, which means they were added after analysis was started). The total number of learner goal interventions was 198. This shows an equality of interaction that was not anticipated (the scores are similar). However, it should be noted that the learner's total includes a pause score of 63, whereas the teacher scores only one on pauses. This highlights the fact that a teacher's utterance nearly always selects the learner as next speaker, pauses in mentoring tend therefore to belong to the learner.

Table 4 Analysis of six utterance 'chunks'

<u>learner 'creative imagine opportunity' (score = 26) was preceded by:</u>	<u>score</u>
teacher 'target M or Ref'	16
teacher 'target M or Ref' plus learner 'monitoring'	11
teacher 'target M or Ref' THEN learner 'monitoring'	9
<u>learner 'monitoring diagnose' (score = 12) was preceded by:</u>	<u>score</u>
teacher 'target M or Ref'	4
teacher 'question'	4
learner 'monitoring evaluate'	6

Some general findings relating to *analysis two* can be reported, some of which are summarised in Table 4. Table 4 shows that it was found that: (i) 16 out of 26 occurrences of 'creative imagine opportunity' had in the previous six utterances the subgoal 'target M or Ref' by the teacher, and that (ii) 11 out of 26 occurrences of 'creative imagine opportunity' had in the previous six utterances the combination 'target M or Ref' by the teacher and 'monitoring' of either type by the learner, and that (iii) 9 out of 26 occurrences of 'creative imagine opportunity' had in the previous six utterances the subgoal sequence 'target M or Ref' by the teacher followed by 'monitoring' by the learner.

Other less common variations included the use of the critical goals that led to 'creative imagine opportunity'. In particular, 'critical probing' by the teacher, in combination with 'target M or Ref' eventually led to the only occurrence of 'creative accurate prediction' identified in the analysis.

Furthermore, Table 4 also shows that it was found that: (i) 4 out of 12 occurrences of 'monitoring diagnose' had 'target M or Ref' preceding it, and that (ii) another 4 out of 12 occurrences of 'monitoring diagnose' had the 'question' communicative acts on its own preceding it, and that (iii) 6 out of the 12 occurrences of 'monitoring diagnose' were preceded by 'monitoring evaluate' by the student.

Analysis two indicates that students seem to perform monitoring before they imagine a creative opportunity. Some heuristics emerged from analysis two. 'Critical probe' was the teacher's most often used interactive goal (score in Table 2 = 44). 'Critical probing' was not only used as an intervention for eliciting an accurate prediction from one learner, it was also used if the learner was not responding with monitoring and creative like utterances (i.e. as a repair strategy). Striking the right balance between the more open ended 'target M or Ref' questions (which leave space for a learner to integrate new knowledge with existing knowledge for themselves) and the more precise 'critical probing' (which gives direction on how the learner may integrate new knowledge by making reference perhaps to some propositional content) will depend on the student and the task involved. Clearly it will also depend on the teacher's own preferences.

Pause taxonomy result and discussion.

The major result of analysis one was a pause taxonomy, which is shown in Figure 5. A pause is usually a silence (the absence of vocalisation or musical acts); however, there were four occasions that were coded as pauses but where there was no silences (these pauses contained an 'emm' or 'err'). Pauses of 9 tenths of a second upwards were coded. A pause may indicate that the learner is simply having a rest or staring out of the window. However, a pause may indicate something else: pauses may have functions and different pauses may have different functions. This research identified four categories of pauses (a turn was allocated to one of these on the basis of the interaction context, see LA6 in Table 1 for an example).

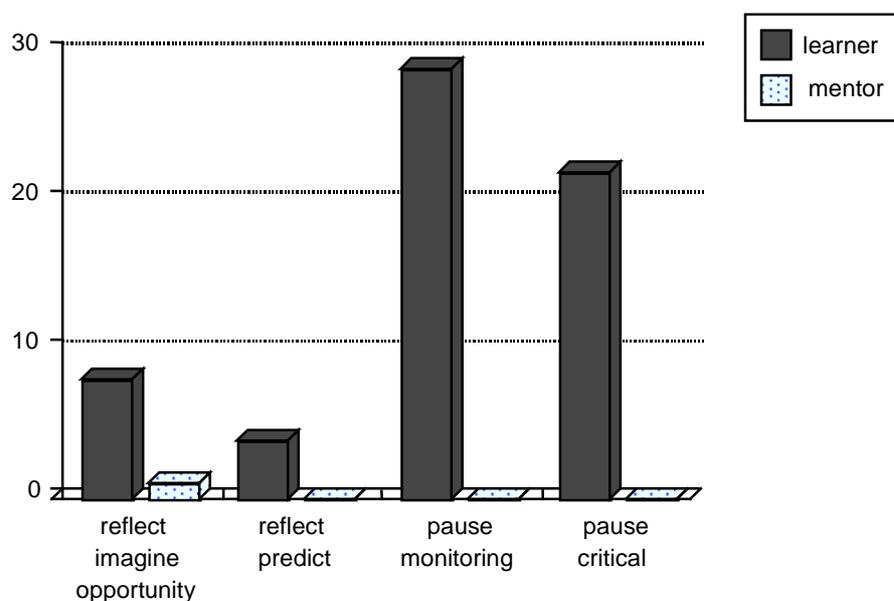


Figure 5. Pause taxonomy scores

The first pause type, *reflect imagine opportunity*, which happen before or during communicative acts related to the 'creative imagine opportunity' intermediate subgoal (score = 8). These are either silences or utterances like 'umm', where the learner reflects about an opportunity before actually using a communicative act to state what that opportunity was. The second pause type was *reflect predict*, which happen before or during communicative acts related to 'creative make prediction' or 'creative accurate prediction' intermediate subgoals (score = 4). Third was *pause monitoring* (score = 29) which was associated with communicative acts indicating monitoring activity by the learner, or a pause left by the teacher (for the learner to reflect in) when making a 'target M or Ref' intervention (the latter happened on six occasions). The fourth pause type was *pause critical*, was related to communicative acts by a learner intending to achieve critical goals (score = 22). If we tabulate the occurrences of pause type by student, we get the results shown in Table 5. Table 5 provides suggestive evidence of a link between pause length and learner ability. At the time of the study the teacher commented that he rated the first student highly and that he thought he would go far professionally (subsequently, in July 1997, student 1 was the only one of the four students in the study to obtain a first class honours degree).

Table 5. Student pause scores by pause type

<u>student</u>	<u>reflect imagine opportunity</u>	<u>reflect predict</u>	<u>pause monitoring</u>	<u>pause critical</u>	<u>total</u>
student 1	5	3	10	10	28
student 2	0	0	3	4	7
student 3	2	1	9	7	19
student 4	1	0	7	1	9
Total	8	4	29	22	63

The teacher rated the second student as one of the weakest. In post experimental interview the teacher made the following comment about student 2:

Teacher: I mean he is very good with purely sonic material, but he is very weak in, if you like, in the pitch domain. And he confirmed that with this exercise. Although what it did show is that in fact [student 2] really needs to do, to work at this level.

The above teacher assessment about students 1 and 2 are matched by the total pause scores in Table 5. Student 1, assessed as the strongest by the teacher, also received the top total pause score (28). Student 1 appears stronger at the creative thinking related pauses ('reflect imagine opportunity' score = 5, 'reflect predict' score = 3). In contrast, student 2 achieves the lowest total pause score (7) and spends no time on pauses related to creative thinking. An interesting line of future work would be to confirm the role of pauses in creative tasks. If this tentative result could be confirmed then this approach would be appropriate for computer-based assessment of learner progress in creative thinking tasks. One caveat is that although student 3 was the only student to make an accurate prediction, she did not achieve the highest pause score. However, her total pause score was the second highest.

The pause taxonomy presented above may be at odds with the conception of a pause used by others. Brown and Yule (1983) use a pause as boundary indicators for a discourse analysis unit. The use of a pause as a boundary unit was used by Ng and Bereiter (1995) in a study of different levels of goal orientation in learning. Levinson (1983, p. 299-300) describes a system where a pause is used as a general cover term for the various periods of non-speech. The term

silence (the absence of vocalisation) is used in a technical sense and assigned to a category, depending on rules of turn taking, of either: a gap, which is not attributable to any party by the turn taking system (e.g. the current speaker has not selected the next speaker, the silence therefore belongs to no one); a lapse (which tend to be longer gaps); or a next speaker's significant (or attributable) silence. The following example has two examples of attributable silence due to the fact that A's utterances select B as next speaker according to the rules of turn taking:

A: Is there something bothering you or not?

=> (1.0)

A. Yes or no

=> (1.5)

A. Eh?

B: No.

from Levinson (1983, p. 299-300)

The pauses identified in our taxonomy (Figure 5) are all attributable types, but of a more fine-grained nature than Levinson's. In human-machine studies a pause may be seen as breaking contact. In work on impasse learning a pause may be seen as indicating that a learner is having a problem. Fox (1993) sees a pause following a question as indicative of a student impasse:

If the student stares blankly for 1 or 2 seconds, the tutor is almost certain to intervene with a hint, or another question; if the student appears to be "working on it," the tutor is likely to give him or her leeway until either the correct answer was satisfied or the student gets stuck. (Fox, 1993, p. 72)

A student may well be in trouble if they stare blankly for 1 or 2 seconds. However, in our study the 'good' student (student 1) paused 5 times when reflecting on a creative opportunity, and achieved an average pause length of 8.37 seconds (average taken from results not presented here, see Cook, in preparation). The outcome of the task for student 1 was described as successful by the teacher (in post-experimental interview comments about the learner's composition). Fox suggests that subtle cues can help interpret the meaning of pauses. However, what the (tentative) empirical finding of this study may imply is that longer or more frequent pauses on certain tasks may be desirable.

Reciprocal modelling result and discussion.

Analysis four found that reciprocal modelling (where participants in learning interactions build up models of the other participants' expectations) may be more important in creative learning activities than has been so far recognised. The interview data (taken from analysis four) at LA6 in Table 1 shows that the learner brings into his considerations 'what he thinks the teacher might be expecting him to do', i.e. the student is building up a model of the teacher's expectations. This was also found in session 4 and is what Dillenbourg (1993, p. 2) called reciprocal modelling and Bunt (1989) calls partner modelling. In the example given in Table 1, once the teacher's expectations have been inferred by the learner (i.e. beliefs about a partner's beliefs), the approach seems to be one of trying confound those expectations, to surprise. This is evidenced at LA8 (Table 1) where the learner suggests that the tempo could be "quicker", the opposite to what the learner believes the teacher wants him to do. The teacher appears to value

such activity. In the interview data associated with TA15 (not shown in Table 1) the teacher comments about the learner's large interval leaps:

Teacher: He did something that I just didn't expect. Which is a very healthy sign in any composer, prepared to take those sort of risks.

This tells us something about what the teacher values from his students, i.e. novelty, surprise, etc., and that his model of the learner indicates that he believes the learner has committed to this value. Of course, not all teachers will value the same thing and surprises have to be the right kind of surprises.

Partner modelling was identified in 2 out of 4 of the sessions using the post-experimental cue technique. Partner modelling may have happened in all four sessions, a question (in post-experimental interview) that would have ascertained this was simply not asked in 2 cases. In one instance the learner had a goal of causing surprise:

(Session 1 example)

Learner: I thought at first, when he said tempo, I thought 'he's thinking that I would like it slower', would I like it slower? And thought about it, and my answer was that I actually think that it would be better faster.

In session four the learner's goal was to do what was expected of him, but with an increased intensity

(Session 4 example)

Learner: So now when [the teacher] says: 'Make it more irregular', I go out of my way to make it more irregular.

The first session was deemed by the teacher as more successful than the fourth, possibly because he valued surprise. One implication for teaching is that learners should not be allowed to become too entrenched in reacting on the basis of one partner model. It is the author's feeling that, although the student in session 4 was very capable and indeed outspoken, his stated goal of going "out of my way to make it more irregular" may have been counter-productive. Session 4 accounted for all 13 occurrences of 'not accept yet' shown in Table 3, which is where the learner does not appear to accept the validity of the tutorial 'task' or 'offers' made by the teacher on how to proceed in the session.

Seven stages of mentoring result and discussion.

Table 6, a result from the fifth analysis, shows that there was evidence for seven mentoring stages (in the order shown) in most of the sessions. An 'x' in the session column means that the stage indicated by the associated row was identified in that session. An '-' means that a stage was not identified for that session.

In this analysis, print-outs of the goals (communicative act were excluded from this analysis) identified by analysis 1, for each session, were analysed to see if any sequence or clusters of goal-types could be detected within a session. All analysis was done by hand because the qualitative analysis software package used did not provide easy list manipulation or text export facilities. Initially 10 stages were identified, however, because 3 of these stages were identified as only occurring in two or less of the four sessions, they were excluded from the stages shown in Table 6. The mentoring stages result represents an empirically based plan for a mentoring session, with the purpose column in Table 6 giving preferences for particular intermediate goals that were identified as being used in that stage. A detailed presentation of the goals that were and were not identified in each stage, for each session, can be found in the appendices of Cook (in preparation).

Because a detailed elaboration of these seven mentoring stages could prove an invaluable resource to an ILE, analysis six spent some time on an in-depth exploration of each stage. State transition networks have been generated for all seven stages of mentoring identified in analysis five (Table 6). Unlike the goals used to achieve transitions in mentoring stages 1 and 7 (i.e. open session and end session), which are relatively simple, the networks for mentoring stages 2 to 6 are more complicated. The seven state transition networks are all empirically derived (i.e. they represent all the goals identified by analysis 1, with a few exceptions that are noted). Thus, the seven networks describe the goals and subgoals identified in the interaction analysis. That is to say, the sequencing of goals identified in analysis 1 (for all four sessions) is directly represented in the networks

Table 6. Seven stages of mentoring

<u>Stage</u>	<u>Purpose</u>	<u>Session</u>			
		1	2	3	4
1. Open session with extrinsic motivation.	Influence learner's affective response to session	x	-	x	x
2. Introduce task and initial probe	Making sure that the task is understood.	x	x	x	x
3. Initial use of Target M or Ref	Promoting monitoring and 'creative imagine opportunity'.	x	x	x	x
4. Teacher led critical thinking using probing and judgement	Giving direction on how to give predictions and imagined opportunities.	x	x	x	x
5. Target M or Ref used (iterate 3-4 times)	Leaving space for the learner to give predictions and imagined opportunities.	x	x	x	x
6. Return to Teacher probing	Getting 'accurate prediction' (if not already achieved) or more 'creative imagine opportunity'.	x	x	-	x
7. End session.		x	x	x	x

To illustrate the approach, the network for mentoring stage 3 is given in Figure 6 (all seven networks and associated discussions are available in the appendices of Cook, in preparation). The network shown in Figure 6 covers all the goals observed in sessions 1 to 4 for the third mentoring stage (i.e. initial use of target M or Ref) with the minor provisos that, although occasionally a goal is repeated in the interaction data, this is not represented in the network. In session 1 the final state was 12, in session 2 this stage ended quickly at state 10, in session 3 the final state was 19, and in session 4 the final state was 16.

A common goal sequence was identified in stage 3, i.e. 'target M or Ref' (state transition 8 -> 9 in Figure 6) leading to 'monitoring evaluate' (9 -> 10) then 'monitoring diagnose' (10 -> 11). This goal sequence was also identified in analysis 2 (see discussion in sub-section 'General findings'), and was also evident in mentoring stage 5 (not shown).

Mentoring script result and discussion.

Note that in analysis 1 the total number of teacher goal interventions was identified as 197. The total number of learner goal interventions was 198. These totals can be used in one approach to calculating the likelihood (i.e. frequency) of a goal occurring in the script discussed below.

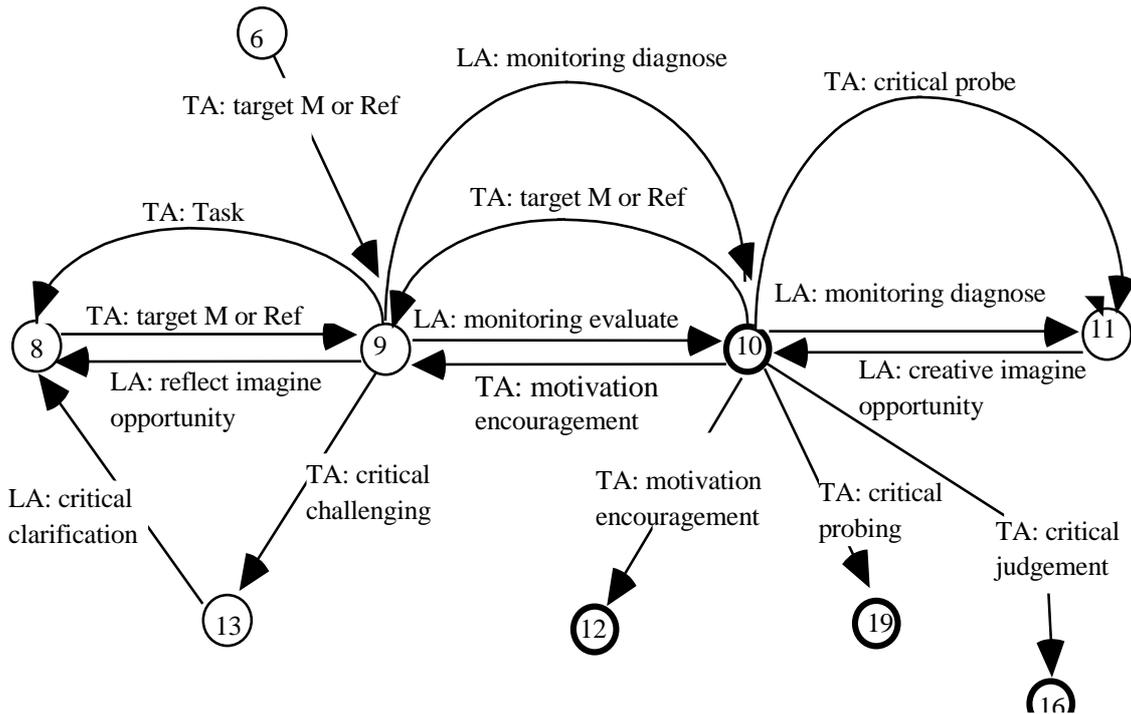


Figure 6. † Mentoring stage 3 network: 'initial use of target M or Ref'

† TA = Teaching Agent and LA = Learning Agent

In the script shown in Figure 7, in the left column we have the teacher intervention, the central column represents a number of possible learner responses, and the right-hand column shows the teacher's counter response. In Figure 7, a box is a goal or subgoal with associated likelihood of that goal occurring in a dialogue. For example, for the subgoal 'target M or Ref', (Table 2 score = 29), the likelihood of this teacher subgoal occurring in a dialogue = $(29 \times 100) / 197 = 14.7\%$. Circles are communicative acts. The arrows show the likelihood of the sequence of goals, linked by the arrow, happening.

The sequences shown in the mentoring script have percentage weightings (on the arrows shown in Figure 7), which are derived from analyses 2 and 3 (an example of this approach to calculating often occurring sequences is given below). We can cross check the mentoring script in Figure 7 against the mentoring stage 3 state transition network (Figure 6), which involved the initial use of 'target M or Ref' by the teacher.

It should be noted that 'assertion confirmation' (in the script) was often associated with the goal 'motivation encouragement' (used in the network). In the network in Figure 6, we can see the sequence identified in the mentoring script (Figure 7). For example, in Figure 6 we have the following two routes through the network that are 'predicted' in the script.

Route one is: 8 to 9(TA: target M or Ref); 9 to 10(LA: monitoring evaluate); 10 to 9(TA: motivation encouragement); 9 to 10(LA: monitoring diagnose); END or 10 to 12(TA: motivation encouragement); END

Route two is as follows: 6 to 9(TA: target M or Ref); 9 to 10(LA: monitoring evaluate); 10 to 11(LA: monitoring diagnose); 11 to 10(LA: Creative Imagine opportunity); END or 10 to 12(TA: motivation encouragement); END. These routes were also in evidence the Stage 5 network (which is not shown).

The above example, of how aspects of the script were derived, is only meant to be illustrative. One approach used to generating weightings on sequences (arrows) in the script is to use the spreadsheet, generated by the analysis 3, of intervention-response counts (the spreadsheet is available in the appendices of Cook, in preparation). For example, the link

shown in Figure 7, from learner 'monitoring evaluate' to mentor 'assertion confirmation' was derived from this analysis 3 spreadsheet. The spreadsheet showed that the most common mentor response to a learner 'monitor evaluate' was 'assertion confirmation' (spreadsheet cell score = 20). Therefore, the likelihood of this link occurring = $(20 \times 100) / 59 = 33.8\%$. Where 59 is the total number of different mentor responses to the learner intervention 'monitoring evaluate' (i.e. the total of the appropriate column on the spreadsheet). This descriptive model of decision making is discussed in the next section and has been used in the implementation of a prototype teaching agent (see Cook, in preparation).

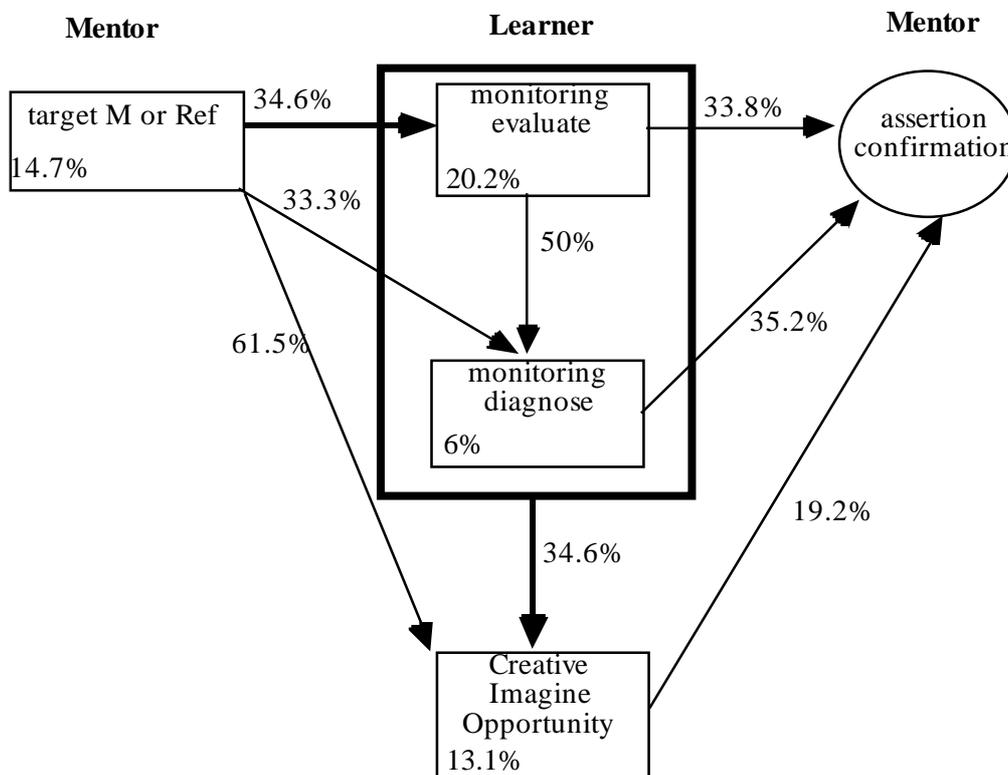


Figure 7. Mentoring script

Conclusion on results

The research question, posed in part 4 of this paper ('Empirical Study Based Around Coleridge Environment'), was used to guide the study of music composition learning: what are the interactive means by which a music composition teacher stimulates creative reflection? If the main findings relating to the pause taxonomy and reciprocal modelling are taken together with the mentoring networks and mentoring script, then it is claimed that the analyses presented above have provided a detailed answer to this question (although only for one small and specific instance of teaching).

IMPLICATIONS OF RESULTS FOR AGENT-BASED ILES

The research question that was used in the introduction to focus this research project was the wider agent design question: how, or to what extent, can studies of dialogue and interaction be exploited in a concrete way by designers of ILES? This section discusses some possible practical uses of the results of the study, and consequently addresses this wider research question. Specifically, this section explores the implications of the empirical results for student

modelling and diagnosis, instructional planning and descriptive-prescriptive models of decision making.

Implications for student modelling and diagnosis

Student modelling is one of the basic mechanisms by which an ILE can individualise interactions to a particular learner. Within Intelligent Tutoring Systems, student models are mainly used to support the learning of content; whereas within ILEs a student model would be used to manage the learning process and to ensure that learners engage in the desired metacognitive processes of monitoring and reflection (Self, 1992).

In the introduction to this paper, we have already described how a DORMORBILE-like student model (Self, 1992; Self, 1993) would be structured for musical composition. Although this explanation was given in order to clarify the difference between musical cognition and musical metacognition, in doing this we also provided an indication of what facts, knowledge and processes would be found at each of the four levels; in addition we described how musical knowledge might be 'compiled down' from a higher level to a lower level as a learner becomes familiar with what is being learnt. Other implications of our results for student modelling are now discussed below.

Although the SE Coach (Conati, Larkin et al., 1997) appears in agreement with our own research position (that time is needed by a learner to explain and reflect), no precise definitions of low viewing time (of a domain related problem) or high viewing time are given by Conati et al. The process of forming and updating the student model by analysing data made available to the system is often called diagnosis. The empirically generated values related to pauses (Table 5) could serve as a diagnosis metric to govern ILE inference about how well a learner is doing. This analysis of the different functions of pauses can be used as one source of evidence, although not infallible, for the extent to which the students are actually engaging in reflective activity (creative reflection).

Bull and Broady (1997) have found that students who are allowed to share each others' student model will often spontaneously collaborate or tutor each other and subsequently may develop a deeper understanding of the target domain. However, Bull and Broady also found that in some cases performance scores decreased. The reasons for this are unclear. Bull and Broady speculate that this may be due to a stronger partner not wishing to appear too dominant. One result presented in this paper has suggested that student and mentor engage in reciprocal modelling, i.e. they build models of each other, and that this can have a varying effect on performance. These models could be made explicit in the manner described by Bull and Broady by using an approach to partner modelling proposed by Bunt (1989). This approach to partner modelling is now described.

Reciprocal modelling was identified in 2 out of 4 of the sessions using the post-experimental cue technique (it may have happened in all four session, a question that would have ascertained this was simply not asked in 2 cases). In session 1 the learner had a goal of causing surprise. In session four the learner's goal was to do what was expected of him, but with increased intensity. By adapting Bunt's (1989, pp. 55) approach, we can specify the following appropriateness condition, i.e. we can say when it is relevant to develop a particular model, for 'partner modelling' of the teacher from the learner's perspective. The general model of appropriateness conditions of partner modelling is shown in Figure 8.

In Figure 8 LA is the learner, TA is the mentor/teacher, p is a proposition, know and suspects are attitudes (that do not make a commitment to the actual truth of p), and intends is an imminent commitment to action. By taking the session 1 example, we can put the Appropriateness Condition (AC) in brackets to illustrate the Move Function (MF) of partner modelling. ACs are the conditions under which a tree branch in Figure 8 is taken. MFs were described in the section on 'Theoretical Approach'; they define the purpose of a communicative act in a specific context.

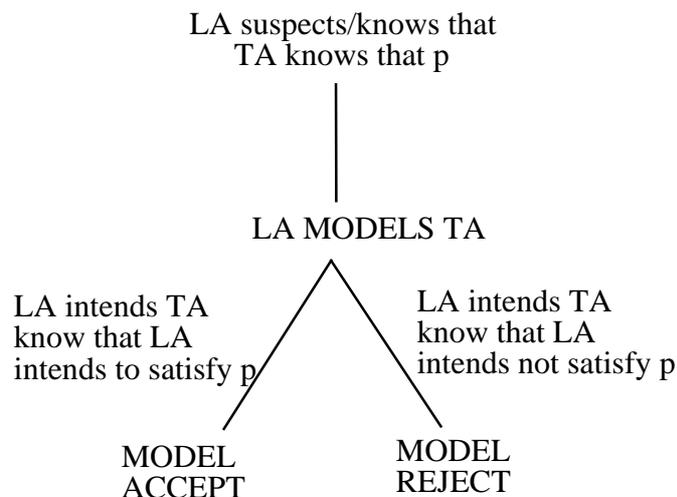


Figure 8. Appropriateness conditions of partner modelling

Partner modelling example 1

1. I thought at first, when he said tempo, I thought he's thinking that

AC[LA suspects]

2. that I would like it slower',

AC = [that TA knows that p]

MF = LA MODELS TA

3. would I like it slower? And thought about it, and my answer was that I actually think that it would be better faster.

AC = [LA intends TA to know that LA intends not satisfy p]

MF = MODEL REJECT

Step 3 above eventually led to the communicative acts of 'question' at LA7 and the 'request' shown at LA8 in Table 1. For the session 4 example, the partner modelling is as follows:

Partner modelling example 2

1. So now when [the mentor] says: 'Make it more irregular',

AC = [LA knows that TA knows p]

MF = LA MODELS TA

2. I go out of my way to make it more irregular.

AC = [LA intends TA to know that LA intends to satisfy p]

MF = MODEL ACCEPT

The above approach to understanding of an utterance, both for its content and its function, has been used by Bunt in the TENDUM dialogue system (Bunt, Beun et al., 1985). In TENDUM move functions (Bunt et al. call them communicative functions) play a role in two places: in the interpretation of natural language inputs, and in the planning of a continuation of the dialogue. What we are proposing here is a combination of Bull and Broady's and Bunt's work. A teaching agent could be used to interact with the learner to develop the a set of partner models. The teaching agent could model the learner, i.e. (TA believes (model of LA)). The

teaching agent could also develop a model of what it suspects the learner's model is of itself, i.e. (TA believes (LA believes (model of TA))). There are two reasons for doing this modelling. First, both models could be used to influence TA's next turn and utterance choice (Ndiaye and Jameson, 1996). Second, the agent could allow pairs or groups of students to come together to discuss each others' model set. The first proposed reason for modelling is discussed further below. One problem with the second proposal is that learners are not always good at explaining their internal processes. However, this is exactly the type of interaction that we wish to encourage with agent-based ILEs. Getting students together to collaboratively explain their partner models may encourage reflective engagement. The idea that a teaching agent is viewed not as a stand-alone resource, but as an assistant to be used in the classroom is discussed towards the end of this sub-section.

Another interesting link, between the result about the occurrence of reciprocal modelling and the student modelling that can be pursued is the AI research, mentioned above, on nested (look-ahead) modelling (Ndiaye and Jameson, 1994; Ndiaye and Jameson, 1996). Ndiaye and Jameson (1994) address the issue of making their natural language processing system, called PRACMA, 'transmutable', i.e. where the system is able to take on either of two possible roles in a dialogue. Transmutability would, the authors claim, enhance the system's ability to anticipate and interpret a dialogue partner's reasoning and behaviour. PRACMA models noncooperative dialogues between a buyer and a seller and includes a module that makes use of Bayesian meta-networks for reasoning about the dialogue partner's beliefs and evaluations. Clearly, such an approach would be useful because, as we saw above, a student may decide to be cooperative or noncooperative, depending on the way in which they decide to model the teacher (the learner may decide to accept or reject what they think the teacher wants from them).

The partner modelling given in example 1 above could be represented as a simple Bayesian beliefs network in a PRACMA-like teaching agent, i.e. (TA believes (model of LA)). The learner's impression of what tempo to have at step 1 in example 1, has been influenced by the teacher's comments, i.e. "AC[LA suspects]". This dependency could be represented as a parent-child link in a Bayesian beliefs network from TEMPO to TA's COMMENT, because from the point of view of LA, TA's comment is probabilistically "caused" by the truth about TEMPO. When TA asks a question about TEMPO, LA would begin its updating of this part of its network, i.e. its evaluation of what it thinks TA wants the tempo to be, using propagation. LA would then go on, at other nodes, to use the outcome of its evaluation of what TA expects to reject or accept this expectation, and to then formulate a verbal response (utterances composed of Communicative Acts that serve a particular Move Functions).

If the PRACMA-like teaching agent is then to take on the role of the teacher (i.e. of itself) it must reason about these types of probabilistic inferences by LA in order to anticipate LA's reaction to what TA decides to say, and possibly also to interpret LA's behaviour as reflecting LA's underlying beliefs and intentions. Ndiaye and Jameson's (1994) Bayesian meta-network could be used to represent TA's changing beliefs about the state of LA's network. Conditional properties are used to link nodes in the meta-network, which in the case of Ndiaye and Jameson's (1994) work are derived from an unpublished empirical study involving the role-playing of sales dialogues. Our own empirical data, presented in this paper, could similarly be used to derive conditional properties for a meta-network. Furthermore, we would suggest that the post-experimental cue approach used in our study (to capture beliefs and intentions of dialogue participants) has a large degree of generality. This 'cue' approach could be used in a similar manner in many domains to generate empirical data that can form the basis of conditional properties (which are similar to the appropriateness conditions in Bunt's approach) for a meta-network.

Ndiaye and Jameson (1996, p. 143) point out that predicting the possible responses of a dialogue partner has some large computational overheads. Such an approach requires the teaching agent to reason in both roles (TA and LA) about the goals and potential communicative acts of both roles. There is no reason in principle why the teaching agent cannot build a succession of models, representing an imagined dialogue between itself and a learner, thus allowing it to look farther into the future. This would be like a chess program looking several moves ahead. Thus, when TA is predicting LA's next response, TA could consider how LA will anticipate TA's subsequent move, etc. Building a decision tree to represent the various

combinations of possible moves is computationally costly (as the trees get deeper, additional expansions of the tree become less worthwhile as they are becoming further removed from what is likely to happen). Ndiaye and Jameson (1996, p. 143) suggest that look-ahead should be dependent on (a) the resources available to the system and (b) the assessed importance of correct prediction.

We can envisage two instances when priority should be given to look-ahead modelling. First, the predicting of possible dialogue turns and utterances is similar to our notion of creative reflection, where LA is being asked to predict how a phrase will sound before it is played back. A computer-based teaching agent that was able to develop its own partner models of creative reflection a few moves ahead, based on a learner's musical input, and that was able to share this model with a learner would be a convincing agent. The second priority to look-ahead would involve diagnosis of 'entrenched' partner modelling. The first session in the empirical study was deemed by the mentor as more successful, possibly because he valued surprise. Of course, surprises have to be the right kind of surprises! One implication for look-ahead modelling is that learners should not be allowed to become too 'entrenched' in reacting on the basis of one model (i.e. the learner's model of the teacher). We have already pointed out above that, although the student in session 4 in the empirical study appeared very capable and indeed outspoken, his stated goal of going "out of my way to make it more irregular" may have been counter-productive. In this case it would appear that the learner's modelling of the teacher's expectation influenced the student's decisions in a negative way. Session 4 accounted for all 13 occurrences of 'not accept yet' shown in Table 3, which is where the learner does not appear to accept the validity of the tutorial 'task' or 'offers' made by the mentor on how to proceed in the session. Of course, this finding is tentative. If entrenched modelling was detected by the accumulation of utterances like the category 'not accept yet', then the agent could perform deep 'look-ahead' modelling to find the best combination of turns and utterances to encourage the learner to develop a new model.

Implications for instructional planning

Instructional planning has recently been defined as an attempt to orchestrate the activities of the learning environment (Wasson, 1996, p. 28). Specifically, instructional planning is now perceived as technique to support the individualisation of the learning activity. It may involve a process of

mapping out of a global sequence of instructional goals and actions that enables the system to provide consistency, coherence, and continuity throughout an instructional session and enables this global sequence to be interspersed with local goals generated when instructional opportunities arise ... (Wasson, 1996, p. 24)

The goals and plans are used by the system to indicate what instructional strategy is required given the current state of the learning environment. An interesting link exists between the state-transition networks resulting from the empirical study described in this paper and the approach for instructional planning used in MENO-tutor (Woolf and McDonald, 1984). MENO-tutor uses a Discourse Management Network (DMN) to plan instruction. The DMN (or tutoring component) "contains 40 states similar to the state of augmented transition network, or ATN" (Woolf and McDonald, 1984, p. 68). The nodes in the ATN correspond to tutorial actions that constitute the basic components of a theory of tutorial dialogues. The DMN makes decisions about what discourse transitions to make and what information to convey. The DMN is a "set of decision units organised into three planning levels that successively refine the actions of the tutor" (Woolf and McDonald, 1984, p. 67). These three levels have hierarchical dependencies (actions at one level may be a refinement of actions at the level above). The arcs of the DMN define the sequence of states normally traversed by the tutor. State transitions thus correspond to default tutorial paths. A set of metarules are also provided that can, if triggered, move the focus to any state in the network.

MENO-tutor made a useful attempt at the development of domain independent discourse planning. Woolf and some co-workers (Woolf, Murray et al., 1988) have used Tutoring Action Transition Networks that are similar to the DMN as a control tool for facilitating the

specification and modification of prototypical patterns of tutorial behaviour. However, the following issues have been raised regarding the DMN. First, Wenger (1987, p. 256) has noted that in MENO-tutor the articulation of the teaching principles upon which decisions are based are not explicitly represented. These principles are implicitly embodied in the arcs and metarules of the DMN. Douglas (1991) has also pointed out that the DMN, and related later work by Woolf, proposes a structure for discourse that is largely independent of the pragmatics of the particular context: "Thus, the history of the discourse, the student (hearer) model, and the tutor's (speaker's) intentions are informally implied" (Douglas, 1991, p. 128). Douglas in fact makes a similar point to Wenger: that the DMN can say whether a particular set of state transitions should occur, but is unable to explain why they should occur at that point in the dialogue. If an approach to using transition networks is to be generalisable to other domains, then the question of 'why' a particular exit from a node is taken needs addressing. Furthermore, the requirement to know 'why' a particular intervention (i.e. choice of transition) is appropriate is particularly relevant in the case of an Intelligent Learning Environment, where an agent needs to reason dynamically about appropriate actions to manage the learning process in a particular situation.

One result of our empirical study was State Transition Networks (STNs) for the seven mentoring stages. In our approach the arcs represent subgoals, which may lead to action or communicative acts. The nodes represent a state at which a decision is made about which transition should be selected next. These STNs can be used to structure interactions and to embed those interactions in computer systems (Winograd, 1987-1988). The networks could be used in an ILE to provide means-ends beliefs about which goals satisfy a particular mentoring stage. Often, more than one exit was possible from a state node. A way of structuring interactions would be to offer these options to a learner as a menu (the options would vary from one node to the next). Preferred options could be highlighted and some mentoring and negotiation (Baker, 1994) could take place if a learner did not accept one of these options (one approach to TA decisions about preferences is described in the next sub-section). Thus, the STNs derived from our own empirical study could be used to give learners options for choice and would give the ILE expectations of answers and a principled position from which to mentor (principled in that the STNs represent descriptive abstraction of what one teacher actually did with four students).

Implications for models of decision making

This section addresses the issue of whether it is in principle appropriate to apply descriptive frameworks (like the typical mentoring STNs and the script identified as results in this paper) in prescriptive way to guide system design. This discussion is related to the distinction in the AI literature between descriptive and prescriptive models (especially for decision making). Slade (1994) has pointed out that

there is a basic dichotomy in the decision-making literature between prescriptive and descriptive models. Prescriptive or normative models focus on how people *should* make decisions, while descriptive theories explore how people *do* make decisions ... descriptive theories can be viewed as bottom-up models, in which the data define the significant features and dimensions of the model. The resulting theory is derived to match the data. (Slade, 1994, p. 194)

Prescriptive methods of decision analysis are mathematically precise and have as a requirement a method of estimating probabilities and payoffs of outcomes. The Bayesian meta-network approach proposed above in the section on student modelling is one possible prescriptive approach. However, many aspects of learning and teaching lack principled ways of making these mathematically precise decisions. Descriptive models of decision making recognise that an agent does not optimise, but rather satisfices (i.e. the teaching agent selects options that are good enough) by examine the 'likelihood' of an outcome being achieved. The position taken in this paper, which is in agreement with Slade's position (1994, p. 202), is that goal-based models of decision making in agent-based ILEs must move beyond a strict

adherence to mathematical formulations of decision analysis. However, each approach to decision making can gain from the lessons learnt by the other approach.

This issue of taking descriptive basis for system design can be restated as the question: What is the nature of the argumentative link between the analysis-description of what a human teacher did and the design of a system? The relation can not be one of direct transfer of expertise, for a number of reasons. On the purely dialogue side, you have open-ended spoken dialogue versus constrained human-computer dialogue. And then, artificial agents are not meant to be copies of human ones. As was pointed out in the introduction, the interaction analysis framework and the study described in this paper are part of a teaching agent design approach, described in detail elsewhere (Cook, in preparation), that aims to make practical use of empirical research in teaching agent development. We argue, therefore, that because very little work has yet been done on how to develop an artificial agent of this kind (i.e. to systematically link empirical data to agent design), the best starting point is to look at what human teachers do, and to then implement descriptive models of that. Refinements to the agent and to guiding theories or frameworks (e.g. the KMf) can then take place on the basis of what happens in the real target dialogue environment when students use the system. Any refinement would thus take place as a result of formative evaluations.

Teaching agents like Blandford's WOMBAT (Blandford, 1991; Blandford, 1994) tend to draw upon the agent designer's intuitions in order to assign preference weightings in mechanisms that make decisions about what teaching intervention to make. In WOMBAT, for any goal that can be satisfied by more than one subgoal or action, there is a list of possible subgoals or actions. For each possible subgoal or action there is a list of means-ends beliefs about what 'values' (an affective attitude held by the agent, e.g. promote monitoring) that action satisfies, plus there are related functions that compute under what conditions a value is relevant (i.e. when a relevance function will return true). There is a separate list of all the values, with a numerical weight attached to each value, to reflect its relative importance to the system. In Blandford's system all preference weightings are set, by the designer of the agent, on the basis of the designer's intuitions. However, by drawing on the study described in this paper, an agent would be provided with empirically based measures of the degree (a preference weighting) to which a subgoal or action satisfies a value. Data indicating the most popular combinations of learner intervention and teacher response (analysis two and three in our study) would provide an empirically based resource for assigning preference weights to a teaching agent interactive response mechanisms. Empirically based measures of the degree (i.e. preference weighting) to which a subgoal or action satisfies an agent's goals can be calculated (the mentoring script shown in Figure 7 illustrates this approach). Thus, for example, if a learner asks a question, the mentor could have a preference weighting on the basis of our results, to making an assertion response, of 6.6 (out of 10), which is a medium-strong preference. This is easily calculated using: $(\text{intervention-response score} \times 10) / \text{total score for that intervention type}$.

Whilst the preference weightings method only gives a numerical weighting, it could be used with the STNs described above (which provide means-ends beliefs about which goals satisfy a particular mentoring stage). The preference weightings approach described above could be used to deal with any conflict that may arise through multiple goals (i.e. how to select a goal) or the need to negotiate new goals because a learner rejects the currently proposed mentoring plan (as defined by the STN).

CONCLUSIONS

The need for dialogue is especially relevant in open-ended, problem-seeking domains such as musical composition learning. Given the open-endedness in music, of both what problems could be addressed, and the space of possible 'solutions', any educational intervention must be similarly open. Teaching interventions can not be restricted to a simple correct or incorrect response. Thus, problem-seeking domains usually require some form of open dialogue between teacher and student. This shifts the emphasis away from the assertion of facts and towards interactions that encourage the types of creative, metacognitive and critical thinking that has been reported in this paper. Like knowledge in the humanities, for example, in the domain of

musical composition knowledge is essentially problematic: it is not just a question of solving a problem, it is more a question of seeking out the nature of the problem and then devising an approach to solving it. This paper has introduced many definitions of metacognitive processes that are relevant to problem-seeking and has described how a mentoring approach could be used to promote learner creative, metacognitive and critical thinking. However, it is not immediately obvious as to how one should go about designing artificial teaching agents that are capable of engaging learners in such problem-seeking interactions.

The approach followed in the research described in this paper is to base system design on the empirical study of interactions between human teachers and students, in the problem domain. However, if a teaching agent's design is to be based on the study of human behaviour, we require an appropriate analysis technique. Therefore, we claim that one contribution of this paper is an original user-centred framework (the KMf), for investigating interactions in problem-seeking domains, that is based on a principled and systematic relationship between interaction data and system design. The KMf provides a taxonomy and definitions of the pedagogical goals involved in a mentoring style of teaching and a convenient tool for the analysis of protocol data in terms of speech acts and associated goals. The KMf provides teacher and learner goal trees that contain goals and subgoals at the top two levels and an utterance level (of communicative acts) at the bottom level. This framework was used to interpret experimental data. However, the interpretation of the results went further than proving the adequacy of the framework; it generated several new results which, we propose, could be useful in the design of ILE user models, instructional planning and decision models for interaction. A prototype teaching agent has been developed that utilises, in the design process, some of the results from our empirical study (the agent is not the focus of this paper, see Cook, in preparation).

For the KMf to be readily generalisable to other areas, like the teaching of social sciences, then we would probably need to exclude the subgoal, in our framework, relating to "creative imagine opportunity" and replace it with subgoals relating to, for example, inter-subjective understanding (see Goodyear and Stone, 1992). Goals relating to metacognitive and critical thinking would, however, appear to be particularly relevant to the domain of social sciences. The KMf may, therefore, have the potential for application in other domains that rely on aspects of creative, metacognitive and critical thinking. Future work will focus on both evaluating the implemented prototype teaching agent with users (Cook, in preparation) and on testing the generality of the KMf approach to interaction analysis and ILE design by applying it to other problem-seeking domains.

Acknowledgements

Many thanks to the following people: to Michael Baker for (i) making various insightful comments on some of the work presented in this paper, and (ii) commenting in detail on various drafts of this paper; to Simon Holland for providing detailed comments on a draft; to Nigel Morgan, who was the teacher-composer in the study and co-worker on Coleridge; to the students who took part in the study; and to the IJAIED reviewers, who made many useful corrections, suggestions and in one case provided additional helpful references.

This work was supervised by Michael Baker and Simon Holland at the Computing Department, The Open University, UK. Funding to support this research has been provided by Thames Valley University. This paper is an extended version of a paper that appears in the AI-ED 97 proceedings.

References

- Adey, P. and Shayer, M. (1994). *Really Raising Standards: Cognitive intervention and academic achievement*. Routledge.
- Ashman, A. F. and Conway, R. N. F. (1997). *An Introduction to Cognitive Education*. Routledge.
- Austin, J. L. (1962). *How to do things with words*. New York: Oxford University Press.

- Baker, M. (1989). Negotiated Tutoring: An Approach to Interaction in Intelligent Tutoring Systems. Centre for Information Technology in Education, Unpublished Ph.D. Thesis No. 7, The Open University.
- Baker, M. (1994). A model for negotiation in teaching-learning dialogues. *Journal of Artificial Intelligence in Education*, 5(2), 199-254.
- Blandford, A. E. (1991). Design, Decisions and Dialogue. Unpublished Ph.D., IET, The Open University, UK.
- Blandford, A. E. (1994). Teaching Through Collaborative Problem Solving. *Journal of Artificial Intelligence in Education*, 5(1), 51-84.
- Boden, M. A. (1990). The Creative mind: Myths and Mechanisms. Weidenfeld and Nicolson.
- Brown, A. L. (1987). Metacognition, Executive Control, Self-Regulation, and Other Mysterious Mechanisms. In Weinert, F. and Kluwe, R. (Eds.), *Metacognition, Motivation, and Understanding* (pp. 65-166). Hillsdale, NJ: Lawrence Erlbaum Associates.
- Brown, G. and Yule, G. (1983). *Discourse Analysis*. Cambridge Textbooks.
- Bull, S. and Broady, E. (1997). Spontaneous Peer Tutoring from Sharing Student Models. In Du Boulay, B. and Mizoguchi, R. (Eds.), *Artificial Intelligence in Education: Knowledge and Media in Learning Systems*. Proceedings of AI-ED 97 World Conference on Artificial Intelligence in Education, Kobe, Japan (pp. 143-150). Amsterdam: IOS Press.
- Bunt, H. C. (1989). Information dialogues as communicative action in relation to partner modelling and information processing. In Taylor, M., Neel, F. and Bouwhuis, D. (Eds.), *The structure of multimodal dialogue* (pp. 47-74). North-Holland: Elsevier Sciences Publishers.
- Bunt, H. C., Beun, R. J., Dols, F. J. H., Linden, J. A. and Schwartzberg, G. O. (1985). The TENDUM dialogue system and its theoretical basis. IPO Annual Progress Report, 19. Institute for Perception Research. Eindhoven, The Netherlands.
- Burrows, K. (1989). Three Divergent Approaches to "Critical Thinking": Introduction and Annotated Bibliography. Kansas City Region Teachers of Philosophy (available from ggale@vax1.umkc.edu).
- Chi, M. T. H., Bassok, M., Lewis, M., Reimann, P. and Glaser, R. (1989). Self-Explanation: How Students Study and Use Examples in Learning to Solve Problems. *Cognitive Science*, 13, 145-182.
- Chi, M. T. H., de Leeuw, N., Chiu, M.-H. and LaVancher, C. (1994). Eliciting self-explanations improves understanding. *Cognitive Science*, 18, 439-477.
- Cohen, P. R. and Perrault, C. R. (1979). Elements of a Plan-Based Theory of Speech Acts. *Cognitive Science*, 3, 177-212.
- Conati, C., Larkin, J. H. and VanLehn, K. (1997). A Computer Framework to Support Self-Explanation. In Du Boulay, B. and Mizoguchi, R. (Eds.), *Artificial Intelligence in Education: Knowledge and Media in Learning Systems*. Proceedings of AI-ED 97 World Conference on Artificial Intelligence in Education, Kobe, Japan (pp. 279-286). Amsterdam: IOS Press.
- Cook, J. (1994). A small study of the three reflective R's: a framework for categorising reflective composer-learner-teacher interactions. Computing Department Technical Report No. 94/10, The Open University, UK.
- Cook, J. (1996). Knowledge Mentoring: towards a framework for linking dialogue analysis to teaching agent design. In Brna, P., Paiva, A. and Self, J. (Eds.), *European Conference on Artificial Intelligence in Education (EuroAIED)* (pp. 199-205), held at Fundacao Calouste Gulbenkian, Lisbon, Portugal, September 30 - October 2, 1996. Edicoes Colibri.
- Cook, J. (1997). Mentoring, Metacognition and Music: An analysis of the interactions used to support creative reflection. Computing Department Technical Report No. 97/01, The Open University, UK.
- Cook, J. (in preparation). Knowledge Mentoring as a Framework for Designing Agents for Supporting Musical Composition Learning. Ph.D. thesis, Computing Department, The Open University.
- Cook, J. and Morgan, N. (1998). Coleridge: A computer tool for assisting musical reflection and self-explanation. *Association for Learning Technology Journal*, 6(1), 102-108.
- Croll, P. (1986). *Systematic Classroom Observation*. London: The Falmer Press.

- Daloz, L. (1990). Mentorship. In Galbraith, M. (Ed.), *Adult Learning Methods* (pp. 205-224). Krieger.
- Dewey, J. (1933). *How we think: A restatement of the relation of reflective thinking to the educative process*. Boston, Mass.: Heath.
- Dillenbourg, P. (1993). Analysis of the protocol of JG-22.9.93. Paper presented at the SMILE Workshop, September 29th - October 1st. Department of Computing, Lancaster University.
- Douglas, S. (1991). Tutoring as interaction: detecting and repairing tutoring failures. In Goodyear, P. (Ed.), *Teaching Knowledge and Intelligent Tutoring* (pp. 123-147). Ablex.
- Draper, S. W. (1994). Constructivism, other theories of the teaching and learning process, and their relationships. Department of Psychology, University of Glasgow, Glasgow G12 8QQ.
- Ericsson, K. A. and Simon, H. A. (1993). *Protocol Analysis (Revised Edition)*. Cambridge, Massachusetts: MIT Press.
- Foss, C. (1987). Productive thrashing in a computerized tutoring system: the acquisition of error management skills. *Proceedings of AI-ED 87*.
- Fox, B. A. (1993). *The human tutorial dialogue project: Issues in the design of instructional systems*. Hillside NJ: Lawrence Erlbaum.
- Freedman, M. (1993). *The Kindness of Strangers: Adult Mentors, Urban Youth, and the New Voluntarism*. Jossey-Bass Inc.
- Goodyear, P. and Stone, C. (1992). Domain Knowledge, Epistemology and Intelligent Tutoring. In Moyses, R. and Elsom-Cook, M. T. (Eds.), *Knowledge Negotiation* (pp. 69-95). Academic Press.
- Holland, S. (1989). Artificial Intelligence, Education and Music. Ph.D., CITE Report No. 88, Unpublished Ph.D. thesis, The Open University, UK.
- Kiss, G. (1986). High-level Dialogue in Man-Machine Interaction. HCRL Technical Report No. 44, KMi, The Open University.
- Kiss, G., Clark, M., Hopkins, C., Cliffe, N. and McTear, M. (1988). Axiological Attitudes in Dialogue Agents. HCRL Technical Report No. 45, KMi, The Open University, UK.
- Kiss, G., Domingue, J. and Hopkins, C. (1991). A Multi-Agent System for Telephone Network Management. HCRL Technical Report No. 72, KMi, The Open University.
- Lerdahl, F. (1988). Cognitive constraints on compositional systems. In Sloboda, J. (Ed.), *Generative Processes in Music: The Psychology of Performance, Improvisation, and Composition*. Oxford Science Publications.
- Levinson, S. C. (1983). *Pragmatics*. Cambridge: Cambridge University Press.
- Lipman, M. (1991). *Thinking in Education*. Cambridge University Press.
- Morgan, N. (1992). Transcript of tutorial sessions with undergraduate composers at Dartington College of Arts. Research Report, Dartington College of Arts Library, Dartington, Devon, UK.
- Morgan, N. and Tolonen, P. (1995). *Symbolic Composer Professional - a software application for Apple Macintosh computers*. Tonality Systems, Amsterdam.
- Ndiaye, A. and Jameson, A. (1994). Supporting Flexibility and Transmutability: Multi-Agent Processing and Role-Switching in a Pragmatically Oriented Dialog System. In Jorrand, P. and Sgurev, V. (Eds.), *Proceedings of the Sixth International Conference on Artificial Intelligence: Methodology, Systems, Applications* (pp. 381-390). Singapore: World Scientific Publishing.
- Ndiaye, A. and Jameson, A. (1996). Predictive Role Taking in Dialog: Global Anticipation Feedback Based on Transmutability. In Carberry, S. and Zukerman, I. (Eds.), *Proceedings of the Fifth International Conference on User Modeling* (pp. 137-144). Boston, MA: User Modeling, Inc. Also available from <http://zaphod.cs.uni-sb.de/KI>.
- Ng, E. and Bereiter, C. (1995). Three levels of Goal Orientation in Learning. In Ram, A. and Leake, D. B. (Eds.), *Goal-Driven Learning* (pp. 355-380). MIT Press. This paper is reprinted from *The Journal of Learning Sciences*, 1(3&4), 243-271, 1991.
- Pilkington, R. and Parker-Jones, C. (1996). Interacting with computer-based simulation: The role of dialogue. *Computers and Education*, 27(1), 1-14.
- Power, R. (1979). The organisation of purposeful dialogues. *Linguistics*, 17, 107-152.

- Schraw, G. (1994). The Effect of Metacognitive Knowledge on Local and Global Monitoring. *Contemporary Educational Psychology*, 19(2), 143-54.
- Schraw, G. and Moshman, D. (1995). Metacognitive theories. *Educational Psychology Review*, 7(4), 351-371.
- Searle, J. R. (1969). *Speech Acts, an essay in the philosophy of language*. Cambridge: Cambridge University Press.
- Searle, J. R. (1976). A classification of illocutionary acts. *Language in Society*, 5, 1-23.
- Self, J. (1992). BRM - A framework for Addressing Metacognitive Issues in Intelligent Learning Environments. In Brahn, J. and Lasker, G. (Eds.), *Artificial Intelligence - Theory and Application* (pp. 85-90). Windsor, Ontario: IAS Press.
- Self, J. (1993). Dormobile: a vehicle for metacognition. Department of Computing, Lancaster University. Also published in Chan, T. and Self, J. (Eds.) (1995), *Emerging computer technologies in education*. AACE.
- Sinclair, J. and Coulthard, R. (1975). *Towards an analysis of discourse: the English used by teachers and pupils*. OUP.
- Slade, S. (1994). *Goal-Based Decision Making: An Interpersonal Model*. Hillsdale, NJ: Lawrence Erlbaum Associates.
- Sloboda, J. (1985). *The Musical Mind: The cognitive psychology of music*. Oxford University Press.
- Slonimsky, N. (1947). *Thesaurus of Scales and Melodic Patterns*. Collier MacMillan.
- Sternberg, R. J. (1985). How to Teach Intelligence. *The Education Digest*, April, 21-25.
- Vygotsky, L. S. (1978). Mind in Society. *The Development of Higher Psychological Processes*. Edited by Cole, M., John-Steiner, V., Scribner, S. and Souberman, E. Harvard University Press.
- Walton, D. N. (1989). *Informal Logic: A Handbook for Critical Argumentation*. Cambridge University Press.
- Wasson, B. (1996). Instructional Planning and Contemporary Theories of Learning: Is this a Self-Contradiction? In Brna, P., Paiva, A. and Self, J. (Eds.), *European Conference on Artificial Intelligence in Education (EuroAIED)* (pp. 23-30), held at Fundacao Calouste Gulbenkian, Lisbon, Portugal, September 30 - October 2, 1996. Edicoes Colibri.
- Wenger, E. (1987). *Artificial Intelligence and Tutoring System*. Morgan Kaufmann.
- Winkels, R. (1992). *Explorations in Intelligent Tutoring and Help*. Amsterdam: IOS Press.
- Winograd, T. (1987-1988). A Language/Action Perspective on the Design of Cooperative Work. *Human-Computer Interaction*, 3, 3-30.
- Winograd, T. and Flores, F., Eds. (1986). *Understanding Computers and Cognition: A New Foundation for Design*. Norwood, New Jersey: Ablex.
- Woolf, B. and McDonald, D. (1984). Building a Computer Tutor: Design Issues. *IEEE Computer*, September, pp. 61-73.
- Woolf, B., Murray, T., Suthers, D. and Schultz, K. (1988). Knowledge Primitives for Tutoring Systems. *Proceedings of the International Conference on Intelligent Tutoring Systems (ITS-88)* (pp. 491-498), held in Montréal.

Appendix

Tables A and B give a brief summary of the goals, subgoal and utterance categories used in the interaction analysis. This summary is provided as a quick reference to assist understanding of the interaction analyses provided in this paper.

Table A. Summary of goals and subgoals

<u>Goal/subgoal</u>	<u>Brief definition</u>
Creative thinking goal	Constrained to creative reflection: the ability of a learner to imagine opportunities in novel situations and to then make accurate predictions about these opportunities.
Creative imagine opportunity subgoal	Interaction (verbal and/or musical and/or actions) by either teacher or learner that concerns the mental imaging of a creative idea in a novel context.
Creative make prediction subgoal	Interaction by either (i) the teacher to elicit a prediction, or (ii) the learner that indicates that a prediction has been made about how an imagined novel opportunity will sound when played.
Creative accurate prediction subgoal	Interaction by either teacher or learner that indicates that a successful prediction has been made about how an imagined novel opportunity will sound when played.
Metacognitive intervention goal	Has an implicit intention of adding a goal to the learner's head to help them monitor their own progress or to reflect.
target M or Ref subgoal	Usually an open-ended question that involve the teacher's attempts to elicit verbal self-explanations from the learner about their own attempts at creative reflection. M = Monitoring, Ref = Reflection.
Metacognitive thinking goal	An understanding of knowledge (i.e. monitoring and reflection) that can be reflected in either effective use or overt description of the knowledge in question.
Monitoring evaluate subgoal	Interaction that involves some evaluative comment about the match between a prediction and an outcome is indicative of monitoring.
Monitoring diagnose subgoal	An attempt to diagnose why something did or did not work.
Reflect predict subgoal	A pause (.9 seconds upwards) where the context shows that the learner reflects about a prediction before actually using a mentoring goal to make a prediction.
Reflect imagine opportunity subgoal	Pauses (.9 seconds upwards), or utterances like 'umm', where the learner reflects about an opportunity before actually using an mentoring goal to state what that opportunity is.

Table A (continued). Summary of goals and subgoals

<u>Goal/subgoal</u>	<u>Brief definition</u>
Critical thinking goal	Thinking that (1) facilitates judgement because it (2) relies on criteria, (3) is self-correcting, and (4) is sensitive to context. To paraphrase Lipman (1991) on critical thinking: 1. Where knowledge and experience are not merely possessed but applied to practice, we are likely to see clear examples of judgement. Good judgement takes everything relevant into account and are the products of skilfully performed acts. 2. Criteria are reasons; a particularly reliable kind. Criteria, and particularly standards, are among the most valuable instruments of rational procedure. 3. Insofar as each participant in the community of inquiry is able to internalise the methodology of the community (where members begin looking for and correcting each other's methods and procedures) as a whole, each is able to become self-correcting in his or her own thinking. 4. Thinking that is sensitive to context involves recognition of exceptional or irregular circumstances (the character of a witness in a trial may govern our view of a statements truth); special limitations, contingencies, or constraints; and overall configuration (remarks taken out of a statement and used elsewhere may change context).
Critical judgement subgoal	Reaching a conclusion about a complex situation or phenomenon, generally without algorithmic deduction or calculation.
Critical probing subgoal	Probing can be a focused used of observations of a student's phrase. Probing provides guidance on how to structure new knowledge.
Critical challenging subgoal	This is adversarial in that a proposition, for example, may not be accepted as true without further evidence being provided.
Critical clarification subgoal	Further elaboration of some point may be requested or given because (i) a previous attempt was unclear or (ii) a response was required.
Critical give reasons subgoal	Any interaction that involves the giving of criteria as a reason.
Critical give evidence subgoal	Similar to a court of law, if some claim or judgement is made then evidence to back up a the claim or judgement may be requested.
Motivation goal	The sense of a willingness to pursue activities.
Motivation intrinsic subgoal	This usually involves the teaching agent giving an account of its approach to composition. Such a description would be intended to motivate the learner to develop their own approach to composition, perhaps on similar lines.
Motivation extrinsic subgoal	Providing a concrete reason for pursuing a task.

Motivation encouragement subgoal	Utterances like repeating another agent's utterance, or "Right" are viewed here as often having the intention of giving positive feedback.
Task goal	This is a description of what the teaching and learning interactions will aim to achieve, it provides a statement against which learning outcomes can be measured.

Table B. Summary of acts and relations

<u>Act</u>	<u>Brief definition</u>
Communicative Acts (a main term)	Used to realise, with different contents, many of the upper level KMF intermediate and pedagogical goals.
assertion	Communicates that the speaker believes/accepts the content to be a fact.
assertion confirmation	A earlier assertion by another speaker has been accepted by the current speaker.
question	There are different types of question. A wh-question communicates that speaker wants hearer to provide certain information, more or less specific, and that speaker believes hearer can provide that information.
request	Like questions, except concerned with action, getting other to do something, with exception of information providing which would make it a question: speaker wants hearer to perform action A, speaker believes hearer can perform action A. Or announcing that something needs to be done and then doing it.
offer & accept	A suggestion of how to proceed, 'who should do what next'. If the teachers accepts the learner's offer, or if the student accepts the teacher's, then the teacher and learner set about getting 'to work'. If an offer is spread over a number of turns then code as 'offer continue'. If confirming a previously accepted offer code as 'accept confirm'.
reject	An offer is not taken up, or an assertion is not accepted.
transform	Where an offer is changed in some way based on negotiation but remains similar to the original offer. Here, "transform" means a communicative act has been made, e.g. an offer, and it (transform) bears a special <i>relation</i> to a previous one of the other speaker to the extent that the contents of the current offer "transforms" that of the previous.
complete	This is where, for example, the teacher or student leaves a sufficiently long gap ($\geq .7$ seconds) for the other to add the correction or continuation to the end of a sentence. Here, "complete" means a communicative act has been made, e.g. an assertion, and it bears a special <i>relation</i> to a previous one of the other speaker to the extent that the content of the current assertion "completes" that of the previous act (e.g. an offer).
action	Non-linguistic forms of communication (e.g. musical actions that have communicative intent or pointing at the computer screen).

Note: transform and complete are not dealt with in this paper. See the appendices in Cook (in preparation) for details.