

TWOFOLD CONTINUITY IN IMMERSIVE LANGUAGE LEARNING

Henry Hamburger & Tucker Maney

*George Mason University
and
Naval Research Laboratory*

Overview: An immersive language learning environment undertakes to engage the student in a two-medium communication process: a conversation supplemented by graphical interaction in an ordinary scene on the computer screen. The fundamental rationale for such a system is that it promotes language learning by enabling the student to use the new language, not analyse or translate it. In this paper, we examine two constellations of issues that arise in trying to provide computer-based language immersion, issues concerning discourse and issues of tutorial strategy, and consider how to deal with their apparently conflicting demands.

An Immersion System: The dual continuity requirement can be grasped best in the context of a specific design and implementation, in particular, the FLUENT framework (Hamburger & Hashim, in press) and its first instantiation, FLUENT-1 (Hamburger and Maney, 1990). The most distinctive characteristic of FLUENT, even in its current prototype form, is its fine-grain two-medium (spatial as well as linguistic) interaction between student and tutor. Within realistic everyday partially animated scenes, co-ordinated by everyday goal structures, both student and system can both discuss and manipulate objects in meaningful ways. For example, in a kind of dialogue interchange we call Movecaster, the student is empowered to move or alter the state and relationships of an object, sometimes in indirect ways. The tutor then makes a relevant comment. Specifically, in one of the FLUENT-1 microworlds, the student can do things like make a screen character's hand turn a faucet on, pick up soap, get it wet, and wash, rinse and dry things, etc. These screen actions alter an internal model, and can lead the tutor to describe an action or its resulting state or to critique an action in terms of a common sense goal structure.

Capabilities like these, together with appropriate tutorial decisions, make possible the immersion approach, in which a student can extend her/his understanding to include a new language aspect that is introduced in a context that makes its meaning clear. The above action examples come from the Washroom World, one of several FLUENT-1 microworlds each devised with particular language learning objectives in view. Relationships among microworlds turn out to play an important role in both kinds of continuity.

Rationale: Before diving into the issues that the method raises, one might well ask if the whole idea has merit. Cognitive and motivational arguments for the immersion approach are provided at some length elsewhere (Hamburger and Hashim, in press), but a brief orientation is appropriate here. We draw inspiration from the immersive foreign language classroom methods of 'total physical response' and 'the natural approach',

for which success is reported by Krashen & Terrell (1983). These human-based teaching methods employ conversation about simple concrete subject matter. The visible objects that they use are directly adaptable to the computer screen. Moreover, the use of simple subject matter presumably keeps down cognitive load, giving language learning a chance (cf. Sweller, 1988). FLUENT's approach is in harmony with the current 'communicative' pedagogy, based on the insight that to know a natural language is not just to produce grammatical sequences of words, but to produce output that expresses what you mean, and not just to parse or recognise grammatical input, but to understand it (cf. Richards & Rodgers, 1986). In other words, the learner must associate language with conceptual material. FLUENT therefore aims to make it possible for the learner to associate new words and phrases with ideas already in the context or picked up from visual clues.

The Two Continuities: An immersive system is committed to two kinds of continuity, conversational and educational. First, in co-operation with the student, it must maintain a coherent dialogue, so that the student can continue to have an accurate context in which to interpret and formulate new sentences. This requirement is present for conversation in general, but is crucial for immersive language learning in light of the student's need for a partially language-independent source of information about what is being said, to support the learning of new aspects of language without translation. Second, since it is actually the new language itself, not the conversational subject matter, that is the educational objective, it follows that the need for educational continuity will influence the choice of language aspects to be introduced at a particular point in a student's progress, and how frequently those aspects should be reused. This second kind of continuity is the province of the syllabus in conjunction with cognitive and pedagogical principles. The two kinds of continuity correspond roughly to two aspects of natural language generation, to wit, what to say and how to say it.

Conversational Continuity: Our discussion of conversational continuity begins with a look at some known discourse issues, which we elaborate and deploy to address the needs and opportunities of FLUENT.

Discourse Knowledge Enables the Immersive Approach: Discourse knowledge and its use help a converser to make sense of a conversation, in part by imposing an intentional structure. Such a structure is directly concerned with having a conversation achieve practical purposes, such as informing one party that the other wishes him to do something. Intentional structures also provide a basis for key aspects of co-ordinating the conversers' attention. See Grosz & Sidner (1986) on this point and for a comprehensive and insightful discussion of discourse issues. For a language receiver, discourse structures arise from language understanding within a dialogue and in turn permit further understanding later in the dialogue. The co-operative language emitter, on the other hand, will somehow have to take this receiver task into account, so that what is emitted will be comprehensible to the receiver.

It is plausible to suspect that a significant portion of discourse knowledge transcends individual languages. Such knowledge might be based, for

example, on general reasoning or on language universals, or even on cultural knowledge in the case of similar cultures with different languages. An immersive language-learning environment should be able to take advantage of such knowledge. Putting it inversely, lack of such knowledge might be a serious handicap for a language beginner in an immersion situation, where co-ordination of attention is crucial. Any non-transferring discourse knowledge would therefore be a high priority area for instruction. Even if all the discourse knowledge itself carries over, there will still be an important need for the learning environment to convey the language-specific tools that let the student make early use of this knowledge.

In order to put discourse knowledge to use, the immersion student needs to learn early how the new language expresses linguistic boundary markers like G & S's examples 'First...' and 'Incidentally...' as well as other means of structuring discourse. Discourse segmentation crucially affects the determination of the referents of pronouns and reduced definite noun phrases. The immersion student needs both to learn the pronouns in the new language, since that is part of the curriculum, and then also to use them with (possibly trans-language) discourse knowledge to deal correctly with situations in which other language phenomena are to be learned.

Even with a full definite noun phrase (like 'the noun' in English) as opposed to a pronoun, it is presumably true across many languages that the referent is chosen from those that are instances of the type associated with the head noun. Here we are interested in the discourse rules that may be carried over from the student's L1 that can enable the student as receiver to select the right referent from among those present and of the appropriate type. A discourse rule might, for example, give preference to the object that has been tapped to participate in satisfying the current subgoal.

Goal Shifts and Continuity: In a microworld that makes significant use of a goal structure, as does the Washroom World of FLUENT-1, it is important for the student and tutor to have a shared view of what the current goal or subgoal is. Successful maintenance of goal agreement places a constraint on what the tutor can possibly be saying and thereby helps the student make an unambiguous connection between language and meaning. It is to foster goal agreement that we use simple everyday domains with familiar goal structures. For example, we assume that it will be clear enough that washing involves soaping up, rinsing and drying, in that order, and that each of these in turn has subgoals, like getting a towel, using it and putting it back on the towel rack.

Suppose that the student is carrying out a spatial task under the supervision of the tutor, in the new language. After the student has completed a subgoal, one thing the FLUENT-1 tutor can say is something that means, 'Now that you have done X, do Y', where X and Y are the just completed subgoal and the next one to be done according to the familiar goal structure. If the student fully understands this remark, it has the effect of confirming that there is agreement on what subgoal is under consideration. If the student's understanding of the sentence lacks some aspect, s/he may actually infer and learn that aspect from the combination of knowing the rest of the sentence, knowing the relevant subgoals

and having the discourse knowledge that such a transitional remark is appropriate.

Different kinds of remarks with different consequences for language education are appropriate in related but different goal-related circumstances. For example, to introduce going on to the first subgoal of a new goal, the tutor can use a sentence that means 'To do X, first do Y'. Another occasion for such comment arises upon completion of a subgoal that also completes its supergoal. Even failures in subgoal sequencing can create educational opportunities, by giving rise to remarks like 'You did X. Before doing X, you need to do Y'. Thus the curious, obstreperous or absent-minded student can continue to receive exposure to the new language.

Goals bear an important relation to intent. An intentional structure expresses the language emitter's intent about things that have to do with the receiver, like the receiver's intent or a potential receiver act. Such an act might in turn be a goal or subgoal in a FLUENT goal structure, whose predicates have to do with changes of state in the physical world. For example, 'wash', changes the surface state of an object or a part from dirty to clean (and dry). Intentions about the physical world in FLUENT are analogous to the discourse purposes (or discourse segment purposes) of G & S.

Intentions that concern language learning are distinct from the world intentions just discussed. Language-learning intentions are also important to FLUENT, since it is a language-tutoring system, but these intentions, unlike those directly involved in the conversational subject matter, can remain a hidden agenda of the tutor, rather than materialising explicitly in the conversation. The student can, however, be given some control over what is taught through the selection of a microworld, which can in turn determine significant aspects of the linguistic content of the interaction.

Microworld Shifts and Continuity: Conversations in FLUENT can be structured not only among goals and subgoals but also among microworlds and submicroworlds. Microworld shifts, like goal shifts, come in different varieties with differing affects on language use. In the case of microworlds, it will be seen that an important issue is whether the new one shares objects with the old one. Generally speaking, shifts in both goals and microworlds have consequences for attention, which, in turn, can affect the potential referents of noun phrases with pronouns or demonstratives.

A specific linguistic consequence of shifts across microworld structures has to do with the choice of definite versus indefinite article. We first show how the type of shift affects the 'previous-mention' property, an ephemeral, conversation-based property of objects. We then show that previous mention affects definiteness. First consider a student who in the course of normal progress satisfactorily completes the work offered by one microworld and moves on to another. The new microworld may be unrelated to the previous one, but it may, on the other hand, be a superworld of the old one, sharing specific objects with it. In the latter case, the shared objects constitute a form of conversational continuity, and the important implementational point is that these objects should not be re-initialised with respect to the property of whether or not they have

been previously mentioned in the conversation, as they would upon a shift from an unrelated microworld. Similar considerations apply if the student is floundering and the tutor or the student decides to move to a simpler subworld with a subset of the objects.

Having seen the effect of microworld structure and shifting on the previous-mention property, we shall now show how previous mention in turn has crucial effects in the choice of a definite versus indefinite article. The key is that the definiteness choice is made differently on entry to a situation than it is after various entities have been introduced. The various aspects of definiteness in English are notoriously difficult for learners whose native language lacks articles, and presents problems even for speakers of French (despite its being close to English and making a similar distinction). Therefore let us look at the choice of articles in detail, in relation to FLUENT, on entry to a microworld. To do this, we write out a hypothetical scenario introducing a learner to Receptacle World, in which small, movable objects can be placed on large, stationary ones, whose top surfaces have the receptacle property. The scenario comprises mostly descriptive sentences, especially at the outset, but there is also a command and a question. The dialogue is conversationally continuous and deliberately a little repetitive. The reader can imagine the screen actions that would accompany this dialogue.

- (1) Here is a box.
- (2) It is small and grey.
- (3) It is a small box.
- (4) The small box is on a table.
- (5) Here is another box.
- (6) This box is big.
- (7) It is also grey.
- (8) Both boxes are grey.
- (9) Put a box on the floor.
- (10) That's the small box.
- (11) Put the other box on the floor.
- (12) Where are the boxes?

Notice the use of 'a' with the noun 'box' in (1), (3) and (9). In the last of these, no particular box is intended, so one may say that the intent is semantically indefinite, which we define as reference to an arbitrary member of a set of cardinality at least two that is specified by the remainder of the noun phrase. It seems reasonable enough to have the syntactically 'indefinite' article used for semantic indefiniteness, if our linguistic terminology has been well devised. Turning to (1) and (3), however, we see that they do appear to refer to a particular box. Stating that previously unmentioned individual objects are introduced with the indefinite takes care of (1). However, in (3) we have already mentioned the box and might therefore expect a definite. Indeed that is just what we get, in that the subject, 'it', is indeed definite. The second noun phrase in (3) is descriptive, as opposed to referential, thereby giving us yet a third reason to use the article 'a'. That (1) and (3) differ in their reasons for using 'a' can be confirmed by recasting them in the plural, as 'Here are some boxes' and 'They are small boxes', respectively. Notice that the first of them has

the indefinite plural 'some' and the other is a bare plural, having no overt article.

The remaining sentences in the scenario display many important related phenomena, but the above analysis suffices to make the general point that syntactic choices are affected by situational aspects. Another situation-syntax tie arises in the use of the restrictive relative clause, a syntactic construction whose use is infelicitous, and hence potentially confusing, unless the situation demands the distinction that the particular clause makes; see Hamburger & Crain (1987). We conclude that it is a complex but useful enterprise to instill a knowledge of the relationship between syntax and situation into both foreign language-learning environments and foreign language learners.

Interchange Types and Dialogue Schemas: A session in FLUENT is a dialogue organised into three levels. At its top level of organisational structure, the session is composed of dialogue schemas, each realised at the next level as a sequence of interchanges between the student and the system tutor. Each interchange comprises a small number of turns, often just one by each party. A turn is simply the output, linguistic, spatial or both, of either the student or the tutor, from the time one of them gets control of the floor until that party relinquishes control to the other. We first introduce several educationally useful interchange types, then use them to compose dialogue schemas.

To fix ideas, we give several specific interchange types. Each of them is named after the role of the tutor, so 'quizmaster' consists of a question by the tutor followed by an appropriate student answer. Reversing roles gives 'oracle', a more demanding type, since here it is the student who produces the question. Questions tend to be linguistically more complex than their answers, and production tends to trail comprehension. Another kind of interchange, called 'servant', consists of a student command followed by the tutor carrying out that command spatially on the screen. Reversing roles yields 'commander'. Yet another interchange type is 'tourguide', a very undemanding type for the student. Here, the tutor makes both an action and a relevant comment, which the student need only acknowledge. Also undemanding is 'movecaster', in which the student gets to choose an action and the tutor then describes it or its effect. Both here and in 'commander', the student carries out spatial actions by controlling the hand movements of a character in the scenario.

Tutorial planning for FLUENT-1 was in the province of the system designer, but will be carried out by the system tutor in subsequent versions. In either case it is useful to do this planning at the level of the dialogue schema, which can be regarded as a skeletal plan that guides the selection of several successive interchanges of specified types, independently of subject matter. By providing integrated chunks larger than just a single turn or interchange, the dialogue schema allows tutorial decisions to be made at a higher level, making the job more manageable. We are devising a modest-sized set of dialogue schemas that make sense in conversational terms and can fulfill various tutorial objectives. Several of them are built into FLUENT-1 and will be abstracted into a generalised form for flexible use in subsequent work.

A particularly simple dialogue schema that is useful upon entry into a microworld calls for several rounds of tourguide, one to introduce a common noun for each object in that microworld. If a microworld has prerequisite microworlds, then only the new objects, those not found in any of the prerequisite worlds, need introducing (unless a student model indicates otherwise). From this description, it is clear that a dialogue schema is not just a specified finite sequence of interchange types, but rather an abstract procedure expressible in terms of loops and other simple control structures in a dialogue schema language.

Besides introducing new objects, it is also appropriate to introduce new actions. For this purpose a more complex dialogue schema is appropriate, one using a round of tourguide followed by one or more rounds of commander with the same action used in tourguide but with different objects fulfilling the roles of the action. The purpose of the commander interchanges is to permit the tutor to confirm that the student has at least temporarily attained comprehension of the relevant action description. Notice that a simple commander interchange does not include specification of a follow-up reaction by the tutor. Such a reaction is desirable and should depend on whether or not the student has given a correct response. To achieve this effect, it is necessary to include in the dialogue schema a loop control specification that repeats tourguide with the same action if the student errs, but that after, say, two consecutive successes with different objects moves on to the next action.

The notion of dialogue structure advanced here for FLUENT can be used as a framework for discussion of the styles of teachers, tutors, textbooks and other kinds of CALL systems. In particular, consider the common textbook chapter structure that begins with a story or multi-paragraph description, along with relevant word translations, each restricted to its sense in the accompanying textual material. Such a chapter typically continues with some comments on new grammatical forms and their semantics, and then goes on to exercises, such as fill-ins, questions, translations from the L1, and transformations within L2. The exercises are typically closely tied to the textual material and the new grammar. These materials constitute, in the best cases, an integrable assemblage, but the job of integrating them is left to the student.

Looking at such a textbook structure from the viewpoint of the above kind of dialogue analysis, one of course notices that the types of interchange, if they can even be called that, are severely limited by the intrinsic lack of responsiveness of a book. Also note that the dialogue (monologue, really) schemas of the textbook structure as described are apparently highly repetitious, each consisting of just a single interchange type, in that all the textual material comes at once, and so do all the exercises of a particular type. Many CALL programs share this propensity for repetition of interchange type, especially those that only have one or two interchange types. Whatever advantages such repetition may have, it does not lead to the natural use of language in context that was argued above (with an example concerning definiteness in English articles) to be essential for picking up subtleties in the relationship of syntax to situation. Another benefit of FLUENT's interleaving of description with commands

and questions in both directions is that it safeguards against grammatical inflexibility.

Educational Continuity: By educational continuity, we mean things like staying on a topic long enough for it to be learned, putting prerequisite topics ahead of topics that rely on them, ordering roughly by difficulty, and putting similar topics together to facilitate useful analogies. The general point is that the student should be working on something s/he is in a position to learn.

Obstacles: Two key constructs in tutoring systems are the syllabus and the student model. The former specifies the longer-term goals and the latter charts progress toward those goals, thereby supporting educational continuity. Unfortunately, two major obstacles – one for the syllabus and one for the student model – arise as soon as we begin to think about educational continuity in the context of immersive language learning environments.

As for syllabus, there is no insurmountable difficulty constructing one; the trouble is to adhere to it, given that the tutor must not say any old thing that uses a particular syntactic aspect, but is obliged to meet the demand of conversational continuity, that is, to say something coherent in the context of the dialogue. To be a just a bit more precise, suppose that conversational continuity dictates what the system should say, in the form of a meaning representation that is then simply passed to a how-to-say-it module of a natural language generation system; see for example McKeown (1985). With such an arrangement, there is no reason to believe that the result will employ syntactic aspects and lexical items that are appropriate for the current student. How, then, do we get educational continuity? One strategy is to provide for at least some flexibility either in the process of selecting what to say or in the process of language generation. For this flexibility to result in language choices helpful to the student, it would then be necessary for the responsible module to access and be influenced by some form of student model. This brings us to the second obstacle to educational continuity, the one related to the student model.

Even if the system had a response selector or a natural language generator that could succeed part of the time in acting on requests to use particular aspects of syntax, it is not clear that it is practical to devise a student model of sufficient precision to provide relevant information. In particular, that information might consist, in effect, of statements that the student has mastered some grammatical rules, has had trouble on others and perhaps has had limited exposure, if any, to others, and so on. To obtain information at this level of granularity, the system would have to engage in some version of grammatical inference on the student's grammar, a notoriously difficult enterprise, indeed one that is probably impossible under certain conditions. Further complications would arise if there were a systematic formal mismatch between the tutor's computational kinds of rules and students' cognitive ones; see Chanier (1991).

Using Microworlds to Structure the Syllabus: A straightforward approach to dual continuity can be formulated directly in terms of grammar, if one is willing to commit to grammatical induction on student responses to form a fine-grained, grammar-based student model. Given such a stu-

dent model, one would try to adapt the language generator to assume the burden of reasoning about the deployment of grammatical rules deemed pedagogically timely. Since there is good reason to believe that each part of this approach would be exceedingly difficult to accomplish, we seek a simpler, yet effective approach, mindful of Self's (1988) cautionary words about excessive complexity in the student model.

Our alternative approach allows for tutorial decisions about three kinds of options, the choice of a microworld, a dialogue schema, and a situational aspect. We comment briefly on each of the three, before proceeding to consider the choice of microworld at greater length. Of the three kinds of choices, the microworld is the one at the highest level of organisation, and is consequently made least frequently. Each microworld is deliberately constructed to introduce particular language aspects or provide practice on them. If it can be determined from the student model that the featured aspects of some microworld are currently appropriate, then choosing that microworld will foster educational continuity.

At the next level, the choice of dialogue schema can promote educational continuity in a different way. Rather than influence what language aspects are used, this choice allows the tutor to adjust the degree of difficulty, by allowing the tutor to invoke easy interchange types, like *Tourguide* and *Movecaster*, on entry to a microworld, when the student is encountering unfamiliar language material, and later to move to forms of interchange that are more challenging and also that yield information on the student's progress that can be used to update a student model.

At the third level, the choice of situational aspect takes advantage of a certain limited flexibility that does exist even at the level of a single turn, within the constraint of conversational continuity. This choice is the one that distinguishes between describing an action ('The girl took the book off the shelf'), describing its most direct result ('(Now) the girl has the book'), or commenting on some other aspect of the situation ('(Now) the shelf is empty', or 'Both books have been taken'). A corresponding distinction can typically be made if the dialogue mode calls for a question or a command. A note of caution is essential here. The immersion approach assumes that a learner sometimes uses the situation to complete her/his understanding of a sentence that would otherwise not be understood. This will work most smoothly for the student if the tutor has chosen the most obvious or prominent situational aspect as the basis for comment. Commenting on less obvious aspects of the situation can provide important language exposure, but the tutor must avoid doing this when an upward shift has just been made in the difficulty of the microworld or the dialogue schema.

We now focus on the construction of microworlds and their relationship to the difficulty of the language that is used with them. The student model associated with this approach will, like the tutorial decision-making algorithm itself, be expressed in terms of coarse information, namely the microworld prerequisite relations and the level of difficulty of the dialogue schemas. This approach should in principle be combinable with any progress that can be made on the approaches, since they will involve tutorial decision-making at different levels.

Our search for a simpler way to gain control of educational continuity

starts with a reconsideration our own efforts in constructing microworlds for FLUENT-1. Our experience there indicates that while it is no trivial job to design a model world in which a pre-specified aspect of language will find natural use, nevertheless once a model is constructed, it is easy enough to determine what aspects of language occur naturally in it. It is also easy to strip away facets of it to create subworlds with smaller ranges of things that can be said in a reasonably natural conversation. These observations suggest that at least a portion of a workable strategy could be to construct many microworlds by subworlding, on the assumption that the consequently narrowed subject matter in those subworlds will in turn constrain the aspects of language employed in them. This approach was used in FLUENT-1 to create review environments. Also, a beginning student can be provided simple language by conversing in simple worlds, especially if attention is restricted to simple, prominent situational aspects.

More generally, this strategy takes advantages of existing human knowledge of microworld design to provide a way to structure the language syllabus. If the choice of (sub)microworld does have the desired effects, then the problem of causing a fine-grained student model to suitably influence the language generation process at a particular turn in a dialogue will have been replaced by the easier problem of making a course student model influence tutorial decisions only at the higher level of choosing what microworld to be in.

In support of this proposal, we clarify the claim that simpler microworlds result in simpler language. Consider Washroom World introduced earlier, which has movable objects that can be placed in various relative positions on the surfaces of stationary objects and also has a goal structure. Dropping the goal structure leaves Receptacle World, in which objects can be moved around in arbitrary order and there is no rationale for saying anything about such things as the completion of subgoals ('You have gotten the soap off your hands') or the appropriateness of particular actions ('Now you need the towel'), though it still makes sense to talk about position ('push the soap to the left'). Alternatively, to retain the notion of goal structure but drop the flexibility of positional relationships, one can go to Face World, in which one assembles a face by moving face parts from a palette to preset positions within a face outline. Here the presence of a goal, the completion of a face, means that it does make sense to talk about needs ('The face needs a nose'), but there is less to say about positional relationships. This is partly because this world is deliberately designed for inflexible positioning of the facial features, so that pushing something to the left, as in the above example, is not possible.

In this discussion it is important to keep in mind that what can be said in a particular world hinges more crucially on its underlying model than on its visual presentation on the computer screen. For example, upon seeing a cup resting on a table, a human being, being possessed of much ordinary knowledge, could conceivably think it appropriate to say, 'The cup can be lifted unless it is stuck to the table'. The world model underlying this scene, however, may well lack any representation of the relation of stuckness. Thus when we say that a world is simple and that certain complex things cannot be said about it, we do not deny that a

human being can think of complex concepts to express in complex ways in terms of what is depicted on the screen, but rather that the world model does not support those concepts.

The preceding observation can be inverted, to good effect. That is, suppose we do, as human designers, think of something to say about a scene in some microworld that uses an important linguistic construction but is not supported by the world model. Then it may be possible to add the needed concept to the world model to create a new and useful microworld that will be a follow-on to the original one. For example, in Receptacle World, it looks like you might be able to say 'move the soap closer to the brush', but the model does not have a concept of distance. By incorporating a distance concept and tying it suitably to the language processing capability, we create a new world in which the comparative adjectives 'closer' and 'further' can be used meaningfully.

For the moment, let us think of a prerequisite relationship as something that holds between two microworlds. The foregoing discussion indicates that reasonable prerequisites for Washroom World would be both Receptacle World and Face World, that is, a conjunction (and-structure) of prerequisites. There will also be cases of a need for disjunctions (or-structures). The need for a disjunction arises if there are several alternative microworlds that can provide the same language education. We shall therefore allow for expressing both conjunctions and disjunctions of prerequisites and more generally for a conjunction of disjunctions (though not an arbitrary and/or tree), in a manner we now describe.

What each microworld prerequires can be expressed abstractly in terms of the need for certain kinds of experience. We put these experience types into a list, which is to be regarded as a conjunction. In the case of Washroom World, for example, one can envision a conjunction of two prerequisites, one being experience with goals, and the other being experience with position and movement. Turning to what might be called the result side of a microworld session, we note that successful exposure to a microworld is intended to provide certain experiences or benefits. These benefits, also expressed as a list, serve as preparation for other microworlds. By associating with each microworld two lists, which we will call its prerequisites and benefits, we obtain, as promised above, an implicit prerequisite structure in the form of conjunction of disjunctions. This is because the list of prerequisites is a conjunction of experiences, each of which is an implicit disjunction. To obtain in explicit form the disjunction of microworlds for an experience, one finds all the microworlds that have a benefit list containing that experience.

Organising things in this way yields a certain amount of flexibility. The designer can simply state that some particular microworld prerequires experience with a topic without having to worry about what particular topics provide that experience. Notice that on this view a particular microworld, say Washroom World, may appear not to provide experience with any new material, but only to provide the challenge of combining a variety of material in new ways. Another use for this kind of structure can be to allow the student to pick a target microworld whose subject matter is of special interest to him or her. A tutorial planning algorithm would then

conduct a search, taking into account the student's known achievements, if any, as represented in the student model, to find a minimal succession of microworlds that the particular student would need to cover to be prepared for the desired world.

In sum, we have begun to explore how microworlds and discourse knowledge can contribute to the design and effectiveness of an immersive language learning environment. We have pointed to some potential pitfalls, but also to a course that appears to steer clear of them.

References

- Chanier, T., Pengelly, M., Twidale, M. & Self, J. (in press) 'Conceptual modelling in error analysis in computer-assisted language learning systems', In M. Swartz and M. Yazdani (eds), *The bridge to international communication: Intelligent Tutoring Systems for Foreign Language Learning*. New York: Springer-Verlag
- Grosz, B. & Sidner, C. (1987) 'Attention, intentions, and the structure of discourse', *Computational Linguistics* 12,3: 175-204.
- Hamburger, H. & Crain, S. (1987) 'Plans and semantics for human processing of language', *Cognitive Science* 11, 1.
- Hamburger, H. & Hashim, R. (1991) 'Foreign language tutoring and learning system', in M. Swartz and M. Yazdani (eds), *The bridge to international communication: Intelligent Tutoring Systems for Foreign Language Learning*, New York: Springer-Verlag
- Hamburger, H. & Maney, T. (1990) 'TPR in a box'. Paper presented at Georgetown University Roundtable on Languages and Linguistics, pre-session on computational linguistics, Washington, DC. March 9, 1990.
- Krashen, S.D. & Terrell, T.D. (1983) *The natural approach*, Oxford: Pergamon Press.
- McKeown, K. (1985) *Text Generation*, New York: Cambridge University Press.
- Richards, J.C. & Rodgers, T.S. (1986) *Approaches and Methods in Language Teaching: A Description and Analysis*, Cambridge: Cambridge Language Teaching Library.
- Self, J. A. 'Bypassing the intractable problem of student modelling'. Proceedings of the conference on intelligent tutoring systems, Montreal, 1988.
- Sweller, J. (1988) 'Cognitive load during problem-solving: Effects on learning', *Cognitive Science* 12, 2: 257-286.
- White, L. (1987) 'Against comprehensible input: The input hypotheses and the development of second-language competence', *Applied Linguistics* 8: 95-110.