Submission to the 2002 International Conference on Intelligent Tutoring Systems

Toward Measuring and Maintaining the Zone of Proximal Development in Adaptive Instructional Systems

Tom Murray (1) & Ivon Arroyo (2)
(1) Cognitive Science, Hampshire College, Amherst, MA 01002
(2) Computer Science, University of Massachusetts, Amherst, MA 01003
tmurray@Hampshire.edu, ivon@cs.umass.edu

Abstract. Intelligent tutoring Systems (ITSs) adapt content and activities with the goals of being both effective and efficient instructional environments. They have goals for students to be challenged and guided in an optimal way--without being too overwhelmed with difficult material or too bored with easy or repetitive material. We propose a particular definition of the zone of proximal development (ZPD) as a general way to describe what all ITSs try to do, and we propose a foundational analysis of instructional adaptivity, student modeling, and system evaluation in terms of the ZPD. We give an operational definition of the ZPD and give an example of its use, and summarize how instructional methods such as scaffolding can be used to maintain ZPD-learning. We also explain how our definition of the ZPD can lead to a more complete model for efficient and effective instruction than common mastery learning criteria.

Keywords: Intelligent Tutoring Systems, instructional strategies, zone of proximal development, scaffolding

Introduction

Many exposés in educational technology build upon educational constructs such as scaffolding, apprenticeship learning, and the zone of proximal development (ZPD), that can be traced back to Vygotsky. These concepts are rarely clearly defined or operationalized [Wertsch 1984]. The field of intelligent tutoring systems is based on building models [Clancey 1982] and it behooves us to try to specify our principles in a computational form. Our goal in this paper is to give a specific and concrete model showing one possible implementation of the concepts of zone of proximal development and scaffolding. We propose that the model is general, not in reflecting a general version of the diverse interpretations of Vygotskian constructs, but in its applicability to ITS design issues.

Our exposition centers around the idea of adapting instruction to keep students within a "zone" where they are neither too frustrated nor too bored. We propose a particular method for measuring this zone and we elaborate on the variety of means available to an ITS for trying to keep students in this zone if it is inferred that they have slipped out of it (by using different types of scaffolding for instance).

We could say that the primary goals of all instruction are for learning to be both efficient and effective. I.E. we want the learner to learn as much as possible in as little time as possible . The primary method ITS's use to attain these goals is to adapt instruction to the needs of the learner and to the pedagogical properties of the content. To be both efficient and effective we want to provide just the right number and content of expository interactions (content and feedback) and interrogative interactions (exercises or activities). Vygotsky describes the zone of proximal development as "the distance between the actual development level as determined by independent problem solving and the level of potential development as determined through problem solving under adult guidance or collaboration of more capable peers" [Vygotsky 1978 pg 86]. The ZPD is commonly used to articulate apprenticeship learning approaches [Collins et al. 1989]. In the prototypical description of learning in the ZPD the learner is involved in a task that is realistic in terms of its complexity and context, and is apprenticing with an expert mentor. Instruction progresses from simply observing the expert perform the task to taking on increasingly difficult components of the task (individually and in combination) until the apprentice can do the entire task without assistance. The assistance is called "scaffolding" and the removal of assistance is called "fading." Such a scenario of learning in the ZPD has implications for the design of learning environments and educational activities, but has limited implications for adaptive instruction. The description reminds us that authentic tasks involve the use of multiple concepts and skills in a concerted way, and that learners must be engaged ("situated") in an integrated task context to learn the sub-skills in their proper context and relationships (Lajoie & Lesgold [1992] calls this "holistic" as apposed to didactic instruction). It articulates a zone within which tasks are too difficult to accomplish without assistance, but which can be accomplished with some help. However, it does not give much guidance on how to determine that zone, what and when to scaffold, and when and what to fade. Our goal in this paper is to give a concise operational definition of the ZPD, and to discuss how to maintain learning in the ZPD.

Our sense of the term ZPD is compatible with how it is used in many contexts. We want to give assistance in order to keep the learner at their leading edge--challenging but not overwhelming them. The ZPD can be characterized from both cognitive and affective perspectives. From the cognitive perspective we say that material should not be too difficult or easy. From the affective perspective we say that the learner should avoid the extremes of being bored and being confused and frustrated. (But *some* cognitive dissonance is usually necessary). Both boredom and confusion can lead to distraction, frustration, and lack of motivation. Of course the optimal conditions differ for each learner and differ for the same learner in different contexts.

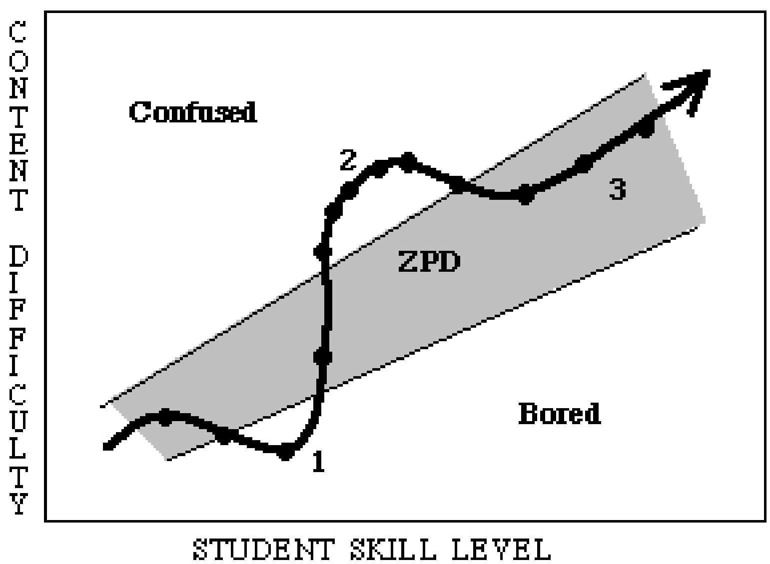

Figure 1: ZPD Illustration

Figure 1 illustrates our meaning of the ZPD. It shows a "state space" (or "phase plane") diagram illustrating a student's trajectory through time in the space of tutorial content difficulty versus the student's evolving skill level. The dots on the trajectory indicate either unit time or lesson topics, and are included to illustrate that progression along the trajectory is not necessarily linear with trajectory length. For example, the dots are bunched up in some placed and spread out in others. In practice, each tutor (human or machine) has limited resources and possibilities of assisting the student, so the "effective ZPD" is defined by the difficulty of tasks possible if the student is given the *available* help (Luckin & du Boulay [1999] call this the "zone of available assistance"). We are only concerned with the effective ZPD for a particular learning environment. This zone will also differ according to each student's tolerance for boredom and confusion. The ZPD is neither a property of the learning environment nor of the student; it is a property of the interaction between the two. We say that the student or learning is "in the ZPD" when the student demonstrates efficient and effective learning. The delineation of the exact zone that is the goal for instruction (shaded area in the figure) is defined by the instructional strategy, and is not a property of the student per-se. This is similar to saying that the criterion for mastery learning is defined by the instructional strategy. We can not directly control the cognitive properties of the student, so it is the *tutor* that must adapt to keep the student in the ZPD.

Though this concept of ZPD makes intuitive sense and provides an attractive metaphor for designing instruction and analyzing learning, it is not operationalized. What is this zone and how do we determine 1) if the student is in it or 2) how to adapt instruction to keep the learner in it? It might seem that the goal of measuring when a student is in the ZPD is intractable and impossible to operationalize. But a human tutor has a workable estimate of when the student is in the zone. Alos, students have a great deal of flexibility and tolerance for non-optimal instruction, so we can reasonably aim to just put them in the "ball park." We can tell when a student is clearly *not* in the ZPD at the two extremes. If a learner indicates that they are bored or if they consecutively answer many similar problems correctly, we can infer that they are in the bored-zone. If a learner is unable to solve a problem using the available means of assistance (i.e. have reached a non-constructive impasse) then we can infer that they are in the confused-zone.

Mastery vs. ZPD learning

Before describing our operational definition of the ZPD we need to articulate a simple but general framework for adaptive instruction upon which to base our discussion. At a course level we can say that all instruction is comprised of three elements: sequencing content, providing opportunities for practice, and giving feedback. Intelligent tutors can adapt at all three levels. For most of our discussion we will make the simplifying assumption that all content is organized into "topics" and topic "difficulty levels." Difficulty levels are used to describe the different levels of performance or understanding within a topic. We need to include difficulty levels so that we can define a problem "equivalency set" (or simply "problem set") as a set of problems, activities, interactions, etc. associated with a particular topic and difficulty level. We assume that a "mastery learning" instructional method is used. That is, the learner is repeatedly given problems (and feedback) for the same learning goal until they demonstrate its mastery in some way. The repetition of this material does not need to be sequential, i.e. the tutor can put the learning goal on hold and treat different topics, but it does not forget the goal to come back and treat the unmastered topic. Note that the number of "tries" (problem solution attempts) is the same as the number of hints plus one (a "hint" can be a simple as "wrong, try again"), which is equivalent to the number of mistakes made on a problem attempt.

Our assumption that mastery learning is used can be almost universally applicable to ITSs if we define mastery learning broadly. Mastery does not imply perfection, but satisfactory performance. Learners can do better than the "mastery" level for a content unit, but the instructional goal of the system is for the student to achieve mastery on a content unit (or units) before moving on. To *not* use mastery learning implies skipping a content unit when it is not mastered. However, missing prerequisites will eventually have to be dealt with when they are needed. Thus for ITSs the goal of effectiveness is met. But mastery learning does not address efficiency. It only assumes that there is *enough* instructional material, practice opportunities, and help (feedback and scaffolding) available to the learner for them to eventually master each learning goal. The mastery criterion does not measure or detect whether a student achieves mastery in an inefficient way, for example through a slow and tedious process or through an inconsistent and frustrating one. Incorporating the ZPD into systems goes "beyond mastery" in a sense, by introducing the goal of efficiency along with effectiveness (mastery). We can say that ITSs have two goals: content mastery (primarily effectiveness) and keeping the student in the ZPD (primarily efficiency). Of course these goals are not independent. If the student is confused or bored this inefficient learning will usually lead to loss of effectiveness.

Next we make the point that, in the ZPD paradigm, it is not desirable to have a student get all of the answers correct. If a student consistently gets the first 3 problems correct on every topic, and then moves on to the next topic, the student is probably not being challenged enough. So the best case is for the student to fail at some items and to have multiple attempts at problems.

An operational definition of the ZPD

Wertsch [1984] and others have attempted to give a clearer definition of the ZPD than is available from Vygotsky's sketches of the construct, but even more precise operational definitions are required for machine tutors.

Hadgaard [1991 p. 350] notes that the ZPD is "an analytic tool necessary to plan instruction and to explain its results". Our operational definitions of tutoring within the ZPD can be used for both system evaluation and for computer tutoring strategies, as we shall demonstrate in subsequent sections.

Mastery criterion. First we will operationalize mastery learning in a common fashion. The mastery learning criterion determines when the student can move on to the next content unit, while the ZPD measurement will determine whether the student learning was efficient for the previous (or current) problem set. It is not practical to infer mastery (or ZPD) based on one task. We will call P the minimum number of times a learner should be given a problem exercise on a particular topic. Here is an example problem sequence showing the number of hints given on problems in a problem set: (3, 1, 0, 0). The student needed 3 hints on the first problem, 1 hint on the second, and then got two correct without hints. There are several possible methods for defining mastery criterion for a problem set. We choose the criterion of getting M out of the previous N correct. This method has the property that problems done before the moving window of N problems do not affect the score (ancient errors are forgiven). It also tolerates guesses and slips as noted below. Further, we will allow N to be equivalent to P (the minimum number of problems allowed), because this streamlines the method and we can think of no good reason to complicate the scheme by making them different. Let us look at a number of possible sequences ("hint vectors") to illustrate this "M out of P" mastery criterion.

1. (3, 1, 0, 0). A prototypical sequence. The learner gets better and reaches mastery.
2. (0, 4, 3, 1, 0, 0). The first problem seems to have been a lucky guess.
3. (4, 4, 0, 0, 3, 0). The fifth problem seems to have been a "slip" or random error.
4. (2, 2, 2, 1, 1, 0, 0). Illustrates very gradual learning or improvement.
5. (4, 4, 0, 0). Illustrates sudden learning or improvement---an "aha" experience.

All of these sequences achieve mastery under the criterion of getting two correct out of the last three problems (in #3 the final 2 problems would not have been needed).

ZPD criterion. Our criterion for ZPD assumes that there is *some* mastery criterion in effect, but any reasonable criterion for mastery could be used in place of the "M out of P" method above and the following analysis of ZPD would still hold. Remember that our goal is to challenge the student just the right amount: not too much, not too little. Challenge level can be inferred from the number of failed attempts or hints needed to solve a problem. It does not make sense to measure whether the learner is in the zone for one problem. Being in the zone is determined for a problem *set* (or more generally for some sequence of problems). We want just the right number of hints in a problem set, which we will call "H." We do not want "less than or equal to H" (this contrasts with mastery criterion where we could accept M "or more" correct answers). As with mastery learning, there are alternative ways to implement this general idea. Compare problem sequences #4 and #5 above, showing gradual and "aha" learning. For these two sequences any criterion that counts the total hints in the *last N* attempts would give differential results for what we think are equivalent examples in ZPD terms. So our goal is to have the student get *exactly* H hints in the problem set, regardless of the number of problems seen. Or, to be more practical, we want the student to stay in the zone of getting H +/- DH (delta H) hints.

We say that the learner is in the bored zone and the problems are too easy if they need less than H-DH hints. We say that they are in the confused zone and the situation is too difficult if they need more than H+DH hints. The exact values for P and H will strongly depend on the content and pedagogical goals and style. These values may even differ for different topics or activity types in a tutorial. As we will note later, they may also be adapted based on the student model.

Let us give our operationalized definition of the ZPD a name: "Specific ZPD" or **SZPD** (like "specific heat" in material science). The SZPD has three parameters, **H**, the goal number of hints in each problem set, **DH**, the allowed variation in H to consider the situation within the ZPD, and **P**, the minimum number of problems the student is guaranteed to see (under normal circumstances). The SZPD is a property of the system's instructional strategy. The measurable learner property is **H***, the actual number of hints given in a problem set. We define **Z**, a measurement of how close the student's performance in a problem set is to the goal: $Z = H^* - H$. The area $Z > DH$ delineates the confused zone, and the area $Z < DH$ delineates the bored zone. It is left to the instructional designer to specify the default SZPD parameters for a particular system, and how the system will respond to the learner being in the bored or confused zones. Ultimately it is a matter of empirical testing to determine the values of H, DH, and P that correspond to the teaching and learning style desired.

Finally we will mention another method for determining whether the learner is in the ZPD---asking them. The system can provide buttons for indicating that the learner is confused or bored, or that the material is too difficult or easy. This information can be used instead of or in combination with the more analytic SZPD method in determining how the system reacts to non-ZPD learning.

An example of using the SZPD to evaluate an ITS

In this section we describe how we are using the SZPD to evaluate a pool of data obtained from use of a mathematics ITS, Animalwatch [Beal and Arroyo, in press]. The analysis looks at a set of trials in which different hinting strategies were used. The goal is to determine which hinting (or "help") strategies are best suited for particular types of individual differences (such as gender and cognitive developmental level). The data analysis is not yet complete, and results will be published in a future paper. In this paper we simply use the study to illustrate use of the SZPD for system evaluation.

Animalwatch

Animalwatch is an Intelligent Tutoring System for basic arithmetic and fractions that offers word problems about endangered species. It takes the student through a series of word problems dynamically chosen from a large database of word problem templates, which are instantiated with appropriate operands, depending on the student's current proficiency. When the student enters an incorrect answer, Animalwatch provides help through progressive hints.

The study

We are in the process of using our operational definition of the ZPD in the analysis of a pool of data obtained from the use of Animalwatch over several years. We have done some analysis of hint effectiveness before on this system and concluded that students behaved differently with different hint styles [Arroyo, 00]. This time, we will integrate data from uses of Animalwatch over a period of 3 years (including 300 subjects). This is a post-hoc formative evaluation–the system was designed and the data was gathered prior to the formulation of our SZPD theory.

Our study focuses on how students with different cognitive abilities benefit from hints with different levels of abstraction ("concrete" and "formal" hints). The problem selection mechanism is the same for all versions of the ITS, only the hinting methods vary. We measured students' cognitive development level with a Piagetian cognitive development pre-test instrument [Arroyo, 01]. Students were randomly assigned to two different versions of Animalwatch: one providing concrete help, another providing formal numeric help.

Analysis of Animalwatch data

We said before that the ZPD is a property of the interaction between the learner and learning environment. Thus, we expect similar students to have a different average number of hints (H*) for different hint methods (i.e. they take longer or shorter to reach mastery). In our analysis we assume a goal value for H and compare students' Z values (how far they are from H) for sequences of equivalent problems. A problem equivalency set in Animalwatch consists of problems for a particular topic (e.g. subtraction) and a specific difficulty (e.g. three-digit subtraction involving borrowing). Due to the problem and topic selection strategies (described in [Arroyo, 01]) the problems in an equivalency set are not always seen sequentially (i.e. problems from other sets may be interspersed).

We have hypotheses about how Z values should change for students of similar cognitive ability who have been provided concrete or formal hints. Table 1 shows fictitious data that illustrates our hypothesis: that we expect low cognitive development student to perform better when given concrete help and high cognitive development students to do better with formal help. The table assumes a value of H=5 (for example, in the top left data cell Z=H*-H=7-5=2). We also expect high cognitive development students to have lower Z values than low cognitive development students, as we expect them to need fewer hints in order to reach mastery. Statistically significant findings in the data would enable us to create adaptive hinting strategies in future versions of the system, where we could change the hinting style based on cognitive development level.

	High cognitive development	Low cognitive development
Concrete help	(3,2,1,1,0,0) Z=2	(3,2,3,1,0,0) Z=4
Formal help	(3,1,1,0,0) Z=0	(3,3,2,1,2,1,0,0) Z=7

Table 1. Fictional Z values for low and high cognitive development students vs. help type

The example in Table 1 compares individual students, but statistical analysis requires aggregating over each of the four groups of students. Rather than use the hint sequence for a single student and problem set, we average the hints for all students in a group for a particular problem set. We compare student Z values by averaging the first problem seen by all students in a group, then the second problem, etc. So that the hint vectors in Table 1 becomes a sequence of the average number of hints received for the Nth problem in a problem set. For example if two students' hint vectors were (4,2,1,0,0) and (3,2,1,1,0) the vector for the average hints is (3.5, 2, 1, 0.5, 0) and the Z value for H=5 would be (7 - 5 =) 2. We plan to compare these "composite Z values" for the four groups indicated in Table 1, and across the various topic levels (problem equivalency sets) in the tutorial.

The Z value gives us an overall look at ZPD behavior aggregated over a problem set, but we may also want to investigate at a smaller grain size what is happening within the problem sets. For example, this would help us distinguish the "gradual" vs "sudden" learning behaviors in example hint vectors #4 and #5 above. To do this we will be producing and inspecting graphs such as the one in Figure 3, which illustrates four hint vectors averaged over for each group in Table 1: A. is for High-Cdevel/Concrete, B. is Low-Cdeve/Concrete, C. is High-Cdevel/Formal, and D. is Low-Cdevel/Formal. We intend to use such graphs to visually compare (and use the associated hint vectors to analytically compare) ZPD-learning among subject types and across problem sets.

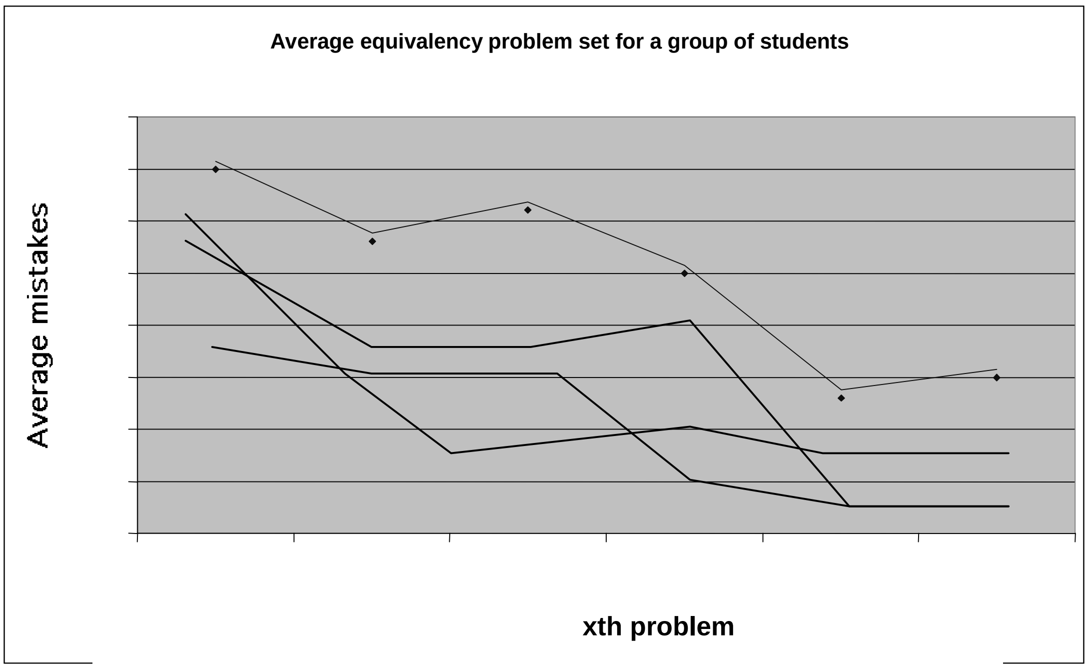

Figure 3. Average hint vectors vs subject group

The ZPD for planning, scaffolding, & adaptation

In this section we look at the types of adaptations that an ITS can make when it is determined that the tutoring session has drifted outside of the ZPD.

Keeping the student in the ZPD involves maintaining an optimal degree of new material, and/or level of challenge. The term "scaffolding" is used to describe tutorial interventions or decisions serving this goal. It is not our goal to give a definition of scaffolding that most accurately reflects either Vygotsky's theories or modern theories of apprenticeship learning [Collins et al]. Rather, we make a reasonable appropriation of the term to refer to any instructional decision or method that has the goal of keeping the learner in the ZPD. The general methods commonly used include: adaptive content sequencing, providing cognitive tools, and hinting and related forms of feedback and help including partial problem solutions.

"Hints" are problem solving assistance that gives information or focuses attention in ways that improve the chances that the learner will be able to solve a problem. The major other form of help is problem decomposition scaffolding. In this method the problem is broken into components which are individually easier to solve. Some forms of scaffolding will assist with problem decomposition and other higher order skills to allow the learner to focus on domain specific schema formation, while other forms of scaffolding will assist the learner in domain-specific skills or answers so that they can practice the overall approach to the problem.

It is also possible to "scaffold" both known and unknown skills. The case of scaffolding unknown skills is clear. But the tutor can also automate a task the student already knows how to do or that is inconsequential to the current learning goals. For example, providing a graphing tool for students that already know how to create graphs, or where using and understanding graphs is important but learning how to construct them is not an important goal.

We have already mentioned the common method of providing progressive levels of hints. Examples include Lajoie & Lesgold [1992]. Arroyo [2000] and Luckin [1999] have found that different types of help differentially benefit learners with different characteristics. As mentioned above, this argues for providing not only different levels of hints but different types of hints and help.

Conclusion

In this paper we offer an analysis of the Vygotskian constructs of scaffolding and the zone of proximal development within the context of adaptive instructional system design and evaluation. We say that scaffolding is any instructional assistance, intervention, or planning that assists the learner in maintaining a learning experience that is within the ZPD. We define the ZPD as a zone of instructional interaction wherein the material given to the learner is neither too difficult nor too easy, or, to phrase it in affective terms, wherein the learner is neither too bored nor too confused, as they progressively master instructional objectives. In posing this definition we extend the traditional concerns of instructional systems (and extend the traditional interpretation of the ZPD) from one concerning only effectiveness to also explicitly include the goal of efficiency. We assume that in all individually paced instruction there is some masterycriterion, so that learning effectiveness is guaranteed for competed topics. When we add the goal that just the right amounts and types of information, practice tasks and help is given, we address efficiency as well as effectiveness. We maintain that the primary goals of all adaptive instructional systems are effectiveness and efficiency.

The learner's state and progress can be determined by analysis of pretests, task performance, direct communication (as in "I'm confused") and other actions (such as tool use and navigation). We focus here on task performance and measure ZPD learning by counting the number of problem attempts (or hints given) for a set of equivalent problems. We define the "specific ZPD" (SZPD) as having three parameters: H (the goal number of hints in a problem set), DH (the allowable deviation from this goal), and P (the minimum number of problems in a problem set, as determined by the mastery criterion). The SZPD parameters are properties of the instructional strategy.

The SZPD can be used for post-hoc analysis and formative evaluation of systems, or it can be used for in dynamic adaptation in tutoring strategy rules. To illustrate its use, we give an example of our in-progress use of the SZPD in a post-hoc evaluation of a mathematics tutor. The goal of the evaluation is to determine the relative effectiveness of different hinting styles vs. learner gender and cognitive development level.

Our proposal for an operational definition of the ZPD (the SZPD) is a specific and practical method for measuring ZPD-learning. In contrast, the question of how adapt or respond to non-ZPD learning (i.e. how to scaffold and fade) is an under-constrained problem with many degrees of freedom. We have described the purposes and characteristic of various forms of scaffolding as a step towards articulating and discovering the appropriate applicability conditions for these various forms.

References

Arroyo, I.; Beck, J.; Woolf, B.; Beal, C.; Schultz, K. (2000) Macroadapting Animalwatch to gender and cognitive differences with respect to hint interactivity and symbolism. Proceedings of the Fifth International Conference on Intelligent Tutoring Systems. Montreal, Canada. June 2000. pp. 574-583.

Beal, C. R., & Arroyo, I. (in press). The Animal Watch Project: Creating an intelligent computer math tutor. In S. Calvert, A. Jordan, & R. Cocking (Eds.) Children in the digital age. New York: Praeger.

Beck, J.; Stern, M.; Woolf, B. (1997) Using the Student Model to Control Problem Difficulty. In the Proceedings of Sixth International Conference on User Modeling. pp. 277-289.

Brown, A., Ellery, S. & Campione, J. (1994). Creating Zones of Proximal Development Electronically. In J.G. Greeni & S. Goldman (Eds.) Thinking practices: A symposium in mathematics and science education. Hillsdale, NJ: Erlbaum.

Collins, A., Brown, J.S., Newman, S. (1989). "Cognitive Apprenticeship: Teaching the Craft of Reading, Writing, and Mathematics. In (L.B. Resnick (Ed.) *Knowing, Learning, and Instruction.* Hillsdale, NJ: Erlbaum.

Lajoie, S., & Lesgold, A. (1992). "Apprenticeship Training in the Workplace: Computer-Coached Practice Environment as a New Form of Appre3nticeship." In Farr & Psotka (Eds) *Intelligent Instruction by Computer: Theory and Practice.* Washington D.C.: Taylor & Francis, pp. 15-36.

Luckin, R; du Boulay, B. (1999) Ecolab: the development and evaluation of a Vygotskian design framework. International Journal of Artificial Intelligence in Education. Volume 10. Number 2. pp. 198-220

Tudge, J. (1991). "Vygotsky, the zone of proximal development, and peer collaboration: Implications for classroom practice." In L. C. Moll (Ed) Vygotsky and education : instructional implications and applications of sociohistorical psychology. New York: Cambridge University Press, 1990. pp. 155-172.

Vygotsky, L.S. (1978). *Mind in Society: The development of higher psychology processes.* Cambridge MA: Harvard University press.

Wertsch, J. (1984). "The zone of proximal development: Some conceptual issues." In B. Rogoff, J. Wertsch (Eds.) *Children's Learning in the "Zone of Proximal Development"* pp. 7-18. San Francisco: Jossey-Bass Inc.