Pavlik Jr., P. I., & Toth, J. (2010). How to Build Bridges between Intelligent Tutoring System Subfields of Research. In J. Kay & V. Aleven (Eds.), Proceedings of the 10th International Conference on Intelligent Tutoring Systems (pp. 534-543). Pittsburgh, PA.

How to Build Bridges between Intelligent Tutoring System Subfields of Research

Philip Pavlik Jr.[1] and Joe Toth[1]
[1] Human Computer Interaction Institute, Carnegie Mellon Univeristy, Pittsurgh, PA 15213
ppavlik@andrew.cmu.edu, jtoth@cs.cmu.edu

Abstract. The plethora of different subfields in intelligent tutoring systems (ITS) are often difficult to integrate theoretically when analyzing how to design an intelligent tutor. Important principles of design are claimed by many subfields, including but not limited to: design, human-computer interaction, perceptual psychology, cognitive psychology, affective and motivation psychology, statistics, artificial intelligence, cognitive neuroscience, constructivist and situated cognition theories. Because these theories and methods sometimes address the same grain size and sometimes different grain sizes they may or may not conflict or be compatible and this has implications for ITS design. These issues of theoretical synthesis also have implications for the experimentation that is used by our various subfields to establish principles. Because our proposal allows the combination of multiple perspectives, it becomes apparent that the current “forward selection” method of theoretical progress might be limited. An alternative “backward elimination” experimental method is explained. Finally, we provide examples to illustrate how to build the bridges we propose.

Keywords: Intelligent tutoring systems, situated cognition, cognitive psychology, design, perception.

1 Introduction

A very old story tells of 6 blind men who were asked to describe an elephant after feeling a part of it. One described it as a wicker basket (ear), one as a ploughshare (tusk), one as a plough (trunk), one as a granary (body), one as a pillar (leg) and one as a mortar (back) and one as a brush (tip of the tail). An ITS researcher trying to be faithful to basic theories of how people develop their capacity to think and act often feels likes he is reading research from these blind men describing an elephant. The ITS researcher turns to the fields of artificial intelligence, statistics, cognitive psychology, cognitive neuroscience, perceptual psychology, all fields of design, human computer interaction, computational modeling, constructivist learning theories, behaviorism, embodied cognition and situated cognition, to name several, and finds many alternative ways to characterize the development in proficiency that we commonly call learning. Unfortunately, these many different sources often result in more confusion than explanation because these multiple theories are difficult to integrate into a mental model of learning that is consistent. It seems the only alternative for the ITS researcher is to have several different inconsistent mental models of learning that must compete when it comes time to plan tutorial interventions to improve student learning.

2 The Problem

We are not the first to notice this problem with sciences of behavior as we know them. Donald Davidson in his influential paper "On the Very Idea of a Conceptual Scheme" [1] has addressed this problem by laying the blame at the feet of philosophers such as Thomas Kuhn. Kuhn has made the case that sciences progress by large paradigm shifts in which new ideas and worldviews replace old conceptual schemes [2], which are subsequently unintelligible from the perspective of the new paradigm. In criticizing the absurdity of this, Davidson says:

Suppose that in my office of Minister of Scientific Language I want the new man to stop using words that refer, say, to emotions, feelings, thoughts and intentions, and to talk instead of the physiological states and happenings that are assumed to be more or less identical with the mental riff and raff. How do I tell whether my advice has been heeded if the new man speaks a new language? For all I know, the shiny new phrases, though stolen from the old language in which they refer to physiological stirrings, may in his mouth play the role of the messy old mental concepts. The key phrase is: for all I know. What is clear is that retention of some or all of the old vocabulary in itself provides no basis for judging the new scheme to be the same as, or different from, the old. So what sounded at first like a thrilling discovery – that truth is relative to a conceptual scheme – has not so far been shown to be anything more than the pedestrian and familiar fact that the truth of a sentence is relative to (among other things) the language to which it belongs. Instead of living in different worlds, Kuhn's scientists may, like those who need Webster's dictionary, be only words apart. [1, p. 10-11]

Davidson's goes on to argue that the strong dualism of "scheme and content, of organizing system and something waiting to be organized", implied in Kuhn's theory of how ideas shift, is unintelligible. He makes the case that there is no duality of theory neutral reality and relative theory, but rather he argues sensory experience provide all our evidence for the acceptance of sentences as true or false. This being so he appeals to this basis of truth as a way of making meaningful disagreement possible between supposedly irreconcilable perspectives. He focuses on enlarging the basis of shared belief by using individuals' common beliefs and opinions as a starting point from which to improve translation between perspectives.

We argue here that similar bridges of translation should be built between the various education research subfields so that ITS will be more able to utilize multiple perspectives in creating educational interventions. To further this goal we offer here an integration of several of the major education research related subfields. We take a balanced approach in this integration, emphasizing how to achieve harmony between these sometimes combative, competitive subfields. This is not the same sort of project that Allen Newell advocated in finding a unified model of cognition because that project was a call for theoretical integration within a subfield [3], while this call is for a valid method to integrate disparate perspectives on learning for more practical reasons.

3 A Proposed Solution

Normally, a subfield of ITS can be partially defined by the units of analysis [4-5] it chooses to study. Thus, behaviorists study behaviors, cognitive psychologists focus on skills, situated cognition researchers study interactions, design specialists focus on form and function, and cognitive neuroscientists focus on the brain but attach cognitive labels to describe the actions of brain tissue. As we can see, one thing that these ITS-relevant behavioral scientists agree on is that a unit of analysis is necessary to analyze a scientific problem. This assumption helps in translating and aligning the subfields because, across each of the various perspectives, grain sizes are based on sensory evidence. It seems that we should be able to align the related sensory bases of each subfield (essentially agreeing that observation is a valid method) and thereby see how the disparate subfields' theoretical explanations map to each other. It seems in such cases you have three essential ways such subfield alignments may turn out.

3.1 Combination

The first possibility is that two perspectives do not explicitly make scientific conclusions about the same units of grain size. In this case we have two perspectives that are essentially orthogonal to one another, and it seems that the reasonable thing for an ITS designer to do is to incorporate both perspectives in the ITS being designed. At first his might not seem intuitive, but this incorporation may be crucial in both experimental and development ITS projects. This is because in ITS research, unlike research in a field like experimental cognitive psychology, we are ultimately interested in how educational research plays out in practical application. In other words, we are more interested in the ecological validity of our results (that they work in a real situation) because this utility within an ecological placement of the tutor is the most important result we wish to generalize in showing that our ITS will benefit real students in real classrooms. For example, in contrast, in cognitive psychology we are often looking for pure main effects and interactions unconfounded by manipulations of other variables. In cognitive psychology work, an assumption is that one wants to remove confounding variables so that main effects can show through and so that that the result that is found is not actually a hidden interaction with some unmanipulated variable. While this is a valid method given the goals, it is likely to be ineffective in establishing that our design has validity in ecological placements.

3.2 Integration

Our second situation is similar, where we have two perspectives, which, while not the same, nevertheless have clear isomorphism. The work of the Jilk, Lebiere, Anderson and O'Reilly illustrates how this is done in the case of aligning a cognitive theory with neuroscience theory [e.g. 6]. Essentially, while there may be some differences between the perspectives (see below on how to resolve conflict) in this situation the perspectives complement and strengthen each other and it becomes more clear that ITS research needs to pay attention to the learning related conclusions of such hybrid perspectives. In their work, these authors have attempted to identify how the ACT-R symbolic model can be mapped to connectionist architecture. They argue that this process is mutually informative and constraining and propose that no single perspective can capture the full richness of cognition. For instance, the authors note that their collaboration has led to a realization that neither theory answers the question of where symbols come from [6].

3.3 Resolution

Our final situation is when the ITS researcher encounters two perspectives that differ in what they predict is best for learning at a grain size. While this disagreement between perspectives may be hard fought in many cases, the ITS researcher might note that often times the common sense resolution admits some truth in both perspectives. Often this disagreement centers on issues of balance along a continuum. For example, constructivists often argue that learning is most effective when the student is able to participate in the building of understanding while direct instruction advocates argue that clear communications of information with some repetition are the most effective way of causing learning. In a case like this, most people's sensory experience probably supports some aspects of both theories, and this leads the ITS researcher to suspect a case where balancing the perspectives is most appropriate.

The assistance dilemma is one way to frame this competition between competing factors since the assistance dilemma describes considerations in balancing factors when one side of a continuum between factors can be described as more assistance and the other side of the continuum can be described as greater assistance [7-8]. For instance, the resolution between the dispute about the effectiveness of direct instruction compared to discovery learning [9] can be described as an assistance question because direct instruction is typically characterized as providing more assistance than is discovery learning. By characterizing the tradeoff along a continuum we seemingly convert what was a dichotomous theoretical question into a question about tradeoffs. Once characterized as a question about tradeoffs, it quickly becomes clear that the endpoints of the continuum (completely unassisted discovery and fully scaffolded problem solving) are both unlikely to be very useful. Rather it seems clear that some balance between telling information to and withholding information from the student will necessarily be optimal.

As we can see, by proposing that theoretical differences are defined by variation along a continuum, we avoid diametrical arguments that are insoluble. Of course, not all perspectives will hold up when their validity is inspected relative to other perspectives. We won't present examples of this here, but basically this admits that this method of combining perspectives does allow falsification. Falsification is a state where the ITS researcher has identified the predictions of 2 theories relative to some instructional decision and has good experimental evidence that one of the perspectives is incorrect in its conclusions relative the instructional design decision. If one accepts the original notion that it is plausible to suppose multiple educational factors moderate the influence of each other, then it is only in these cases where one perspective's predictions are probably incorrect that the predictions of that perspective are safe to exclude.

3.4 Implications for Experimentation

When considering this proposal to unify subfields, we can note that the combinatorics of exploring the space of the multiple principles becomes untenable if we wish to use a bottom up strategy of assuming the traditional null hypothesis (an unmanipulated, "vanilla" state of nature where everything except the tested principle is left out). However, when considering how to proceed with ITS research, it helps to think back to the full meaning of the null hypothesis. In many subfields the null hypothesis is typically considered to be a situation that is as "unmanipulated as possible" and taken to be the default state of nature. It is this typically barren backdrop against which we observe main effects when we manipulate a variable as our alternative hypothesis. However, this bias for simple experiments might simply be taken as a bias about our assumptions of the true default state of nature. If we instead assume a true state of nature that includes multiple "confounders" we may be creating a null hypothesis that is much closer to the default state of nature. Because of this, we might expect that any conclusions we might make will generalize more easily to other ecological valid states of nature (which would likely share many features).

Therefore, it seems that in ITS research we might validly consider the null hypothesis as using as many "confounding factors" as seem plausible from a naturalistic standpoint, determined from a review of the multiple ITS subfields. In this sort of situation, ITS research would begin with a many component ITS system well integrated into coursework and the classroom, and then each of the many factors could be varied experimentally, using typical experimental design to see if their removal reduces performance in the system. For example, while cognitive psychology experimentation answers questions like what is the effect of spaced practice given no other confounding variables using a simple memory task, this new experimental method would ask what is the effect of spacing and its interaction with the retention interval given its use in a complex context of learning items that will need to be used productively. In both examples then, our experiment might manipulate the spacing interval.

The key difference in this approach is that we assume that unexplored factors are present (and unmanipulated in our ecological null hypothesis) so that any conclusions we make about the investigated factors will have validity in a natural situation where we might expect some level of all plausible factors. This approach seems further justified in the case of education research specifically because the plausible threat is not that we will not include some highly effective components in each of our respective systems, but rather the threat is that the highly effective aspects of our systems will be nullified by factors that we have not included in our experiments (perhaps to reduce possible confounds). In such cases, it is possible that what we have considered confounding may actually be necessary to include so as to get valid results. Essentially, we are describing a likelihood of positive moderating variables that effectively gate the effect of the investigated variable. In such a situation, which we propose to be closer to the default classroom situation, it seems logical to do one's best to include all the possible moderators so as not to preclude an effect of the manipulated variable.

An analogy to the two most common ways statisticians have found to select terms in regression equations is useful to help see the structure of the argument we have made thus far. Unless you have time to search the entire combinatoric space of possible terms, forward selection and backward elimination are the main methods of trying to establish a best regression model. As we may recall, forward selection starts out with an empty model and tests the significance of adding each predictor to the model, adding the terms one at a time in the order of which is most significant. As we may note, this is very similar to how the experimental method is used in behavioral sciences. Like forward selection, behavioral science methods tend to begin with few variables being manipulated so that the other confounders or moderators should not have any effect. Obviously, like forward selection, our behavioral science methods produce many successes, and allow incremental additions to our progress toward ITS design that uses multiple principles to improve student learning in intelligent tutoring systems.

An alternative to forward selection is backward elimination. Backward elimination begins with all the variables and then begins to remove the ones that are discovered not to influence the predictions. Backward elimination may actually be a superior way to search for the best ITS because it allows us to retain all the potential moderators from the start, unlike forward selection. As was noted by in a review of variable selection, "It is often argued that forward selection is computationally more efficient than backward elimination to generate nested subsets of variables. However, the defenders of backward elimination argue that weaker subsets are found by forward selection because the importance of variables is not assessed in the context of other variables not included yet." [10]

4 Specific Resolutions

Our goal of providing specific examples of the combination of subfields is depicted in Table 1. Table one is organized with a partition of grain sizes in the left column and a subset of the relevant subfields along the top. For reasons of space, we have been forced to group subfields and ignore many findings to focus on some core principles in each subfield. Within each cell then, are the principles that apply at the grain size for the subfield. Using this chart we can then analyze cases where there is either no explicit dispute at the grain size, cases where the perspectives offer complementary approaches and finally cases where there is an explicit conflict to be resolved between the subfields.

4.1 Example Resolutions

Examination of table 1 can be used to help identify places where theories are making potentially conflicting predictions at matching grain sizes. From there, the theories can be broken down further to more closely identify any grain size isomorphisms. After this sensory process of object matching is completed, predictions relative to the objects can be compared for inconsistencies, which can then be resolved as described above.

However, this more clear analysis allows us to notice that many of the disputes between the theories seem to be issues where either side would admit some truth to the other. For instance, while situated/constructivist theorists specify the importance of authentic contexts, there is very little work in the cognitive literature which challenges this notion directly. While much cognitive work might be described as ignoring the importance of authentic contexts, this is very different than proposing that authentic contexts have a negative effect on student learning. Indeed, digging deeper into the cognitive literature allows us to unearth many specific findings that support the importance of context [e.g. 11]

Further, there appears to be no explicit reason why cognitive phenomenon would not be important in situated learning. For example, consider cognitive load [12] in real life situations. There seems to be no reason why the putatively cognitive mechanism of cognitive load would not affect students in authentic tasks with real world contexts. Indeed, because authentic contexts often include more details, it seems that an integration of situated theory and cognitive load theory offers advantages. By integrating these theories it would allow us to examine how much authentic context is useful and how much causes extraneous cognitive load. While this sort of synthetic approach is not always simple, it helps to explicitly reveal the best resolution to any contradiction when examining different perspectives on an issue important to ITS development.

5 Conclusions

What we are advocating here has much in common with design based research methods [13]. We are also advocating doing experiments in naturalistic settings and we also agree that “scientists must draw connections to theoretical assertions and claims that transcend the local context” [p.8]. However, unlike design based research advocates, we do not place a special emphasis on a theoretical starting point (e.g. situated cognition theory) for our design based analysis, but rather acknowledge that regardless of the perspective of the researcher (e.g. behaviorist, design theory, etc.) there are distinct advantages of using a backward elimination method that begins with a complex state of nature (the existing ITS system and its interrelations with students, peers and teachers) and then makes experimental variations. In contrast, we have argued that by beginning with a simple system and incrementally adding intervention characteristics the ITS designer will face challenges to ecological validity and have difficulty detecting effects that might be moderated by other factors. Since these moderating factors (e.g. student motivation controls learning even in a cognitively excellent ITS) may block effects that would show in a more natural design, it may even be necessary to use such a backwards procedure depending on the intervention being investigated.

Table 1. Map of subfield integration for each grain size.

	Cognitive, Social, Motivational Psychology	*Situated Cognition, Constructivism, Discovery Learning*	*Artificial Intelligence, Statistics, Computational Modeling*	*Neuroscience*	*Perceptual Psychology, HCI and Design*
Societal Learning Context		Interaction across levels[14]	SEM models[15]		Interaction design[16] Activity centered design[17]
Individual	Cognitive load[12] Self-explanation[18] Analogy[19] Worked examples[20-21] Scheduling[22] Testing[23] Theory of intelligence[24] Efficacy[25] Goal orientation[26]	Scaffolding[27] Affordances[28]	Models of individual differences[29]	Working memory training[30] School lunches[31]	
Peers	Vicarious Modeling[32]	Modeling[14]			
ITS software		Authentic contexts[33] Exploration contexts[34]	Question generation[35] User modeling[36]		Organization of gestalt[37] Pleasure[38]
Teachers		Modeling and Coaching[14]			

Acknowledgments. This research was supported by the U.S. Department of Education (IES-NCSER) #R305B070487 and was also made possible with the assistance and funding of Carnegie Learning Inc., the Pittsburgh Science of Learning Center, DataShop team (NSF-SBE) #0354420 and Ronald Zdrojkowski.

References

1. Davidson, D.: On the Very Idea of a Conceptual Scheme. Proceedings and Addresses of the American Philosophical Association: 47**,** 5-20. (1973)
2. Kuhn, T.: The structure of scientific revolutions. University of Chicago Press Chicago (1970)
3. Newell, A.: Unified theories of cognition. Harvard Univ Pr (1994)
4. Altman, I., Rogoff, B.: World views in psychology: Trait, interactional, organismic, and transactional perspectives. Handbook of environmental psychology: 1**,** 7-40. (1987)
5. Matusov, E.: Applying a Sociocultural Approach to Vygotskian Academia:Our Tsar Isn't Like Yours, and Yours Isn't Like Ours'. Culture & Psychology: 14**,** 5. (2008)
6. Jilk, D.J., Lebiere, C., O'Reilly, R.C., Anderson, J.R.: SAL: an explicitly pluralistic cognitive architecture. J. Exp. Theor. Artif. Intell.: 20**,** 197-218. (2008)
7. Koedinger, K.R., Pavlik Jr., P.I., McLaren, B.M., Aleven, V.: Is it Better to Give than to Receive? The Assistance Dilemma as a Fundamental Unsolved Problem in the Cognitive Science of Learning and Instruction. In: Sloutsky, V., Love, B., McRae, K. (eds.). Washington, D.C. (2008)
8. Koedinger, K.R., Aleven, V.: Exploring the assistance dilemma in experiments with cognitive tutors. Educational Psychology Review: 19**,** 239-264. (2007)
9. Kirschner, P.A., Sweller, J., Clark, R.E.: Why minimal guidance during instruction does not work: An analysis of the failure of constructivist, discovery, problem-based, experiential, and inquiry-based teaching. Educational Psychologist: 41**,** 75-86. (2006)
10. Guyon, I., Andre, E.: An introduction to variable and feature selection. Journal of Machine Learning Research: 3**,** 1157-1182. (2003)
11. Titone, D.: Contextual memory in a virtual world. Trends in Cognitive Sciences: 5**,** 376. (2001)
12. Kalyuga, S., Sweller, J.: Measuring Knowledge to Optimize Cognitive Load Factors During Instruction. Journal of Educational Psychology: 96**,** 558-568. (2004)
13. Barab, S., Squire, K.: Design-based research: Putting a stake in the ground. Journal of the Learning Sciences: 13**,** 1-14. (2004)
14. Collins, A.: Design issues for learning environments. International perspectives on the design of technology-supported learning environments**,** 347-361. (1996)
15. Ramayah, J.: A SEM investigation of e-learning: an illustrated perspective of a course website acceptance model among business students. International Journal of Information and Operations Management Education: 2. (2007)
16. Kolko, J.: Abductive Thinking and Sensemaking: The Drivers of Design Synthesis. Design Issues: 26**,** 15-28. (2010)
17. Gifford, B.R., Enyedy, N.D.: Activity centered design: towards a theoretical framework for CSCL. Proceedings of the 1999 conference on Computer support for collaborative learning, 22. International Society of the Learning Sciences, Palo Alto, California (1999)
18. Chi, M.T.H., de Leeuw, N., Chiu, M.-H., LaVancher, C.: Eliciting self-explanations improves understanding. Cognitive Science: A Multidisciplinary Journal: 18**,** 439-477. (1994)
19. Gentner, D.: Structure-mapping: A theoretical framework for analogy. Cognitive Science: A Multidisciplinary Journal: 7**,** 155-170. (1983)

20. Atkinson, R.K., Derry, S.J., Renkl, A., Wortham, D.: Learning from examples: Instructional principles from the worked examples research. Review of Educational Research: 70**,** 181-214. (2000)
21. Carroll, W.M.: Using worked examples as an instructional support in the algebra classroom. Journal of Educational Psychology: 86**,** 360-367. (1994)
22. Pavlik Jr., P.I., Anderson, J.R.: Using a model to compute the optimal schedule of practice. Journal of Experimental Psychology: Applied: 14**,** 101-117. (2008)
23. Karpicke, J.D., Roediger, H.L., III: The Critical Importance of Retrieval for Learning. Science: 319**,** 966-968. (2008)
24. Dweck, C.S.: Self-theories : their role in motivation, personality, and development. Psychology Press, Philadelphia, PA (2000)
25. Bandura, A.: Self-efficacy: Toward a unifying theory of behavioral change. Psychological Review: 84**,** 191-215. (1977)
26. Elliot, A., McGregor, H.: A 2 x 2 achievement goal framework. Journal of Personality and Social Psychology: 80**,** 501-519. (2001)
27. Wood, D., Wood, H.: Vygotsky, tutoring and learning. Oxford review of Education**,** 5-16. (1996)
28. Greeno, J.G.: Gibson's affordances. Psychological Review: 101**,** 336-342. (1994)
29. Navarro, D.J., Griffiths, T.L., Steyvers, M., Lee, M.D.: Modeling individual differences using Dirichlet processes. Journal of Mathematical Psychology: 50**,** 101-122. (2006)
30. Jaeggi, S.M., Buschkuehl, M., Jonides, J., Perrig, W.J.: Improving fluid intelligence with training on working memory. Proceedings of the National academy of Sciences of the United States of America: 105**,** 6829-6833. (2008)
31. Kleinman, R.E., Hall, S., Green, H., Korzec-Ramirez, D., Patton, K., Pagano, M.E., Murphy, J.M.: Diet, Breakfast, and Academic Performance in Children. Annals of Nutrition and Metabolism: 46**,** 24-30. (2002)
32. Bandura, A.: Principles of behavior modification. New York. (1969)
33. Lave, J.: Cognition in practice: Mind, mathematics, and culture in everyday life. Cambridge Univ Pr (1988)
34. Bicknell-Holmes, T., Hoffman, P.: Elicit, engage, experience, explore: discovery learning in library instruction. Reference Services Review: 28**,** 313-322. (2000)
35. Li, T., Sambasivam, S.: Question Difficulty Assessment in Intelligent Tutor Systems for Computer Architecture. Proceedings of ISECON 2003. (2003)
36. Desmarais, M.C., Maluf, A., Liu, J.: User-expertise modeling with empirically derived probabilistic implication networks. User Modeling and User-Adapted Interaction: 5**,** 283-315. (1996)
37. Chang, D., Dooley, L., Tuovinen, J.E.: Gestalt theory in visual screen design: a new look at an old subject. Proceedings of the Seventh world conference on computers in education conference on Computers in education: Australian topics - Volume 8, 5-12. Australian Computer Society, Inc., Copenhagen, Denmark (2002)
38. Green, W., Jordan, P.: Pleasure with products: Beyond usability. Taylor & Francis (2002)