Studying Problem Solving Through the Lens of Complex Systems Science: A Novel Methodological Framework for Analyzing Problem-Solving Processes

Manu Kapur (manu_kapur@hotmail.com)
Communication, Computing, and Technology in Education
Teachers College, Columbia University, New York, NY10027

John Voiklis (jv37@columbia.edu)
Communication, Computing, and Technology in Education
Teachers College, Columbia University, New York, NY10027

Charles Kinzer (kinzer@tc.columbia.edu)
Communication, Computing, and Technology in Education
Teachers College, Columbia University, New York, NY10027

Abstract

Viewed through the lens of complex systems science, one may conceptualize problem-solving as a complex adaptive activity. Theories of biological evolution point to an analogical equivalence between problem solving and evolutionary processes and, thus, introduce innovative methodological tools to the analysis of problem-solving processes. In this paper, we present a methodological framework for characterizing and analyzing these processes. We describe two measures that characterize genetic evolution—*convergence* and *persistence*—to characterize the problem-solving process, and instantiate them in a study of problem-solving interactions of collaborative groups in an online, synchronous environment. We conclude with a discussion of issues relating to reliability, validity, usefulness, and limitations of the proposed methodology.

Introduction

This proposal springs from a shared, situative, epistemological belief that learning and, particularly, problem solving, function as continuous, dynamic processes distributed in space and time over multiple actors, actions and artifacts, influencing and being influenced by the environment in a complex, adaptive, and iterative manner. Understanding this process ranks among the most important challenges facing cognitive (particularly educational) research (Akhras & Self, 2000; Barab, Hay, & Yamagata-Lynch, 2001). While researchers have accounted for several cognitive, metacognitive, and socio-cultural factors, theories and models of problem solving remain unable to account for multiple, interrelated, and dynamically changing actors, contexts, processes, and outcomes. Thus, measures and methods for tracking the evolution of problem-solving processes in ways that account for changing conditions within a problem space are needed (Barab et al., 2001; Derry, Gance, Gance, & Schlager, 2000). The proposed methodology is a step in that direction.

The Nature and Processes of Problem Solving

Humans engage in a myriad of purposeful, goal-directed behaviors (Anderson, 2000): from negotiating traffic, to negotiating contracts, to implementing disaster relief. Even a seemingly mundane task may involve numerous variables connected in ways that are indirect, recursive, and time-dependent (Simon, 1978; Funke, 1991). Further, the problem solver can combine and recombine these variables to form multiple trajectories to multiple outcomes (Quesada, Kintsch, & Gomez, 2001). In such an intractable problem space, a solution evolves through processes akin to organic evolution: instead of exhaustive searches, the problem solver iteratively charts, follows, and modifies a trajectory to an ever-shifting destination (Newell & Simon, 1972), ignoring or setting aside other trajectories and destinations. Sometimes the iterative operations perform optimally; more often, a problem solver on an equivalent trajectory must settle for satisfactory results (Simon, 1978). How, where, and why do certain processes succeed and others fail?

Latent Problem Solving Analysis (LPSA) approaches the modeling of problem-solving processes without *a priori* assumptions about the contours of the problem space, the trajectories through that space, nor about the final destination (Quesada et al., 2001). LPSA tries to discern each of these from the data; comparing the problem-solving processes—operations and/or problem states—from multiple trials (Quesada et al., 2001). While LPSA avoids the teleological fallacy, a full account of evolution may require additional methodological tools. Evolution implies not only change but change over time; Quesada et al. (2001) indicate that LPSA does not discern the order in which operations and/or problem states occur and, thus, it cannot discern the effect of time. To this end, some researchers have pointed to Complex Systems Science as a framework for understanding the evolutionary dynamics of human problem solving (Mitchell, 1998; Port & van Gelder, 1995).

Complex Systems Science (CSS)

CSS provides a both a theoretical and methodological framework for studying how interactions among the parts of any given system change over time and culminate in the behaviors of the system as a whole (Bar-Yam, 1997; Crutchfield, 1994). CSS identifies adaptation as one macroscopic behavior shared across systems—biological, physical, and cognitive. A complex adaptive system (CAS) changes its behavior in response to environmental and selfgenerated feedback, often in an attempt to achieve a goal; goal-seeking adaptations that occur on a collective scale and/or over multiple iterations emerge as evolution (Bar-Yam, 1997). Problem solving, too, involves iterative goal-seeking adaptations (or operations) through which an individual or a collective tries to reduce discrepancies between an initial state and goal state (Newell & Simon, 1972). Thus, one can view the problem solver (individual or collective) as a complex adaptive system that evolves as problem operations develop and change over time. CSS, therefore, facilitates a strong two-way analogy between adaptive problem solving and evolutionary process.

Measures for Characterizing the Problem-Solving Process

CSS has aggregated several mathematical methods for measuring evolution, including measures for *convergence*, *persistence*, and *phase transition*. Convergence and persistence measures have been used to study cultural and conceptual evolution (Cavalli-Sforza & Feldman, 1981); they may prove informative for describing and explaining problem-solving processes.

Convergence, here, involves three component measures—*number*, *function*, and *fitness* of problem states (Heylighen, 1988). When one imagines the problem-solving process as a sequence of problem states, the *number* of states from initial to goal state serves both as a temporal and spatial measure: a *tick* on the evolutionary clock or *step* along the evolutionary path (Heylighen, 1988). In biological evolution, each mutation reconfigures the gene. Similarly, in problem solving, each operation reconfigures the problem state. This reconfiguration may increase or decrease the difference between the reconfigured and goal states and, thus, the distance (number of *ticks* or *steps*) required to reach the goal state. Each operation, then, has a positive or negative *impact* on the problem-solving process; it increases or decreases the difference between the current, problematic state and a specified goal state. However, with the exception of the initial operation, the configuration of a problem-state and its distance from the goal state reflects the cumulative impact of all the operations up to that particular state. If problem solving means minimizing this distance, then cumulative impact reflects the *fitness* of the problem-solving process at the given state (Heylighen, 1988).

Persistence measures are based on the intuitive idea that a component (e.g. a gene, genetic trait, concept, or strategy) that gets used, again and again, over time has proven itself useful. Bedau & Packard's (1992) measure of persistence, *evolutionary activity*, proves informative on this componential level—measuring the extent and intensity with which each component gets used over time–as well as systemic level—the extent and intensity with which the system can generate persistent components. *Evolutionary activity* emerged from *Artificial Life*, where it was used to study the *vitality* of an artificial biosphere, but can apply to any adaptive system in which one can identify and isolate components and their usage statistics, including natural biospheres, chemical systems, computational systems, and mental systems (Bedau & Packard, 1992). For example, in biological evolution, the components are genes or genetic traits. In problem solving, concepts, strategies, or functional categories can serve as components. As with convergence, each operation and each configuration of the problem state will affect the dynamics of these problem-solving components (Bedau, Snyder, & Packard, 1998); conversely, the dynamics of these components inform the evolutionary activity structures of the problem-solving process. When added to convergence analyses, persistence may reveal how multiple evolutionary processes converge on similar paths without implying a single best path.

Methodology

To illustrate the proposed methodology, we describe and demonstrate how it was used in a study of computer-supported, collaborative, problem-solving interactions.

Research Context and Data Collection

Participants included sixty 11th grade students (46 male, 14 female; 16-17 years old) from the science stream of a co-educational, English-medium high school in Ghaziabad, India. They were randomized into 20 groups of three and instructed to collaborate with their group members to solve two problem scenarios. Both presented an authentic car accident scenario that required the application of Newtonian kinematics to solve. The study was carried out in the school's computer laboratory, where group members communicated with one another only through synchronous, text-only chat. The chat archived the transcript of each discussion as a text file. These 20 transcripts, containing the problem-solving interactions and solutions produced by the groups, formed the data used in our analyses.

Coding Problem-solving Interactions

Quantitative Content Analysis (QCA) (Chi, 1997) was used to segment and code interactions using an interaction coding scheme developed by Poole and Holmes (1995), namely the Functional Category System (FCS) (see Table 1). Two trained doctoral students independently coded the interactions with an inter-rater reliability of .85.

Table 1: Functional Category System (FCS)
(Adapted from Poole & Holmes (1995), p. 104)

1. Problem Definition (PD)
1a. <u>Problem Analysis</u>: Statements that define or state the causes behind a problem
1b. <u>Problem Critique</u>: Statements that evaluate problem analysis statements

2. Orientation (OO)
2a. <u>Orientation</u>: Statements that attempt to orient or guide the group's process.
2b. <u>Process Reflection</u>: Statements that reflect on or evaluate the group's process or progress

3. Solution Development (SD)
3a. <u>Solution Analysis</u>: Statements that concern criteria for decision making or general parameters for solutions
3b. <u>Solution Suggestion</u>: Suggestions of alternatives
3c. <u>Solution Elaboration</u>: Statements that provide detail or elaborate on a previously stated alternative.
3d. <u>Solution Evaluation</u>: Statements that evaluate alternatives and give reasons, explicit or implicit, for the evaluations.

3e. <u>*Solution Confirmation*</u>: Statements that state the decision in its final form or ask for final group confirmation of the decision.

4. Non-Task (NT)

Statements that do not have anything to do with the decision task. They include off-topic jokes and tangents

5. Simple Agreement (SA)

6. Simple Disagreement (SDA)

The *unit of analysis* was semantically defined as the function(s) that an intentional statement served in the problem-solving process. Therefore, each statement was segmented into one or more interaction unit(s) and coded into the functional categories of the FCS. Once coded, a time-ordered sequence of functional categories or codes represented each problem-solving discussion.

Measuring Convergence & Fitness

Convergence of problem-solving interactions may be broadly defined as the extent to which the group discussion leads to a solution as perceived by the group. To model the telic aim of problem-solving interactions and develop a measure for convergence, we used a two-state Markov model (Ross, 1996). An *a posteriori* impact value of 1, -1, or 0 was assigned to each interaction unit depending upon whether it pushed the group discussion towards (impact = 1) or away (impact = -1) from the goal, or maintained the status quo (impact = 0). This was done with an inter-rater reliability of .93.

More formally, let the problem space be defined by n interaction units; each assigned an impact value of 1, -1, or 0. Further, let n_1, n_{-1}, and n_0 denote the number of interaction units assigned the impact values 1, -1, and 0 respectively such that $n_1 + n_{-1} + n_0 = n$. Then convergence, $C(n)$, may be defined as $C(n) = \frac{n_1 - n_{-1}}{n_1 + n_{-1}}$. The number of zeros is not factored into the calculation of convergence because interaction units assigned a zero impact, by definition, maintain the convergence level of the discussion. It is easy to see that the convergence value will always lie between -1 and 1.

Note that the numerator in the formulation of $C(n)$ is a measure of position, $P(n) = n_1 - n_{-1}$. In other words, if the problem-solving process is a sequence of steps along a straight line - some forward (impact = 1) and others backward (impact = -1) - then the difference between the total number of forward and backward steps gives the position relative to (or distance from) the starting point, i.e., the start of the discussion. Convergence then is the mean distance from the starting point.

Convergence can also be conceptualized as measure of *fitness* of the entire discussion. The higher the convergence, the higher the fitness of the discussion. Extending this conceptualization to all problem states and not just the final one, we can define fitness as the temporal measure of convergence, i.e., at any point in time in the discussion, how close a group is to reaching the goal state – an ideal solution to the problem. Therefore, the fitness statistic at an arbitrary point in time in the problem-solving process is defined as the convergence value up to the interaction unit at that point in time, with the final fitness level of the entire problem-solving process being the convergence value itself. Recalling that time refers to *ticks* on the evolutionary clock (i.e. an arbitrary time t corresponds to, say, the i^{th} interaction unit), the fitness $F(t)$ at time t in the discussion may be defined as $F(t) = C(t) = \frac{n_1(t) - n_{-1}(t)}{n_1(t) + n_{-1}(t)}$ where $n_1(t)$ and $n_{-1}(t)$ represents the number of interaction units coded as 1 and -1 respectively, up to and including the i^{th} interaction unit. Plotting the fitness value on the vertical axis and time (as defined above) on the horizontal axis, provides a representation (also called the fitness curve) of the problem-solving process as it evolves in time. Figures 1 and 2, drawn to the same scale, present four major types of fitness curves that emerged from the 20 problem-solving discussions in our study.

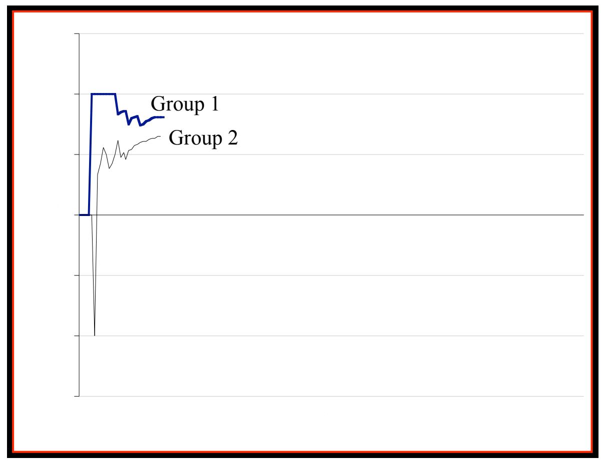

Figure 1: Fitness curves of two short discussions

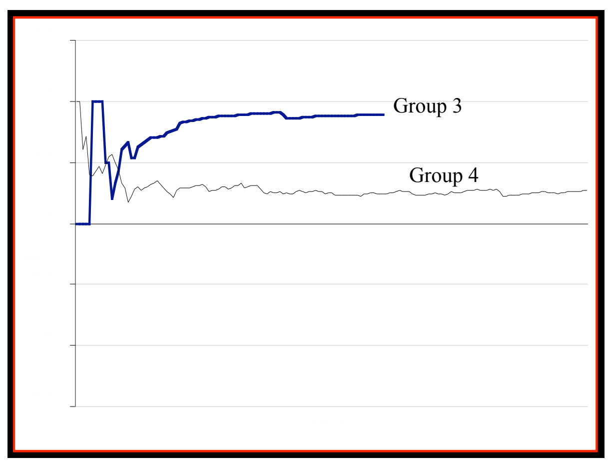

Figure 2: Fitness curves of two long discussions

Interpreting Fitness

In our view, there are five aspects to interpreting the fitness analysis. First, because the fitness value at a given time indicates proximity to an ideal solution (with higher values indicating greater proximity), fitness curves that trend upwards indicate problem-solving processes that are getting closer to an ideal solution (fitness = 1), and vice versa. Hence, fitness curves provide a quick snapshot of the entire problem-solving process in terms of how short or long it was as well as how close or far the discussion was from an ideal solution at any given point in time.

Second, the shape of the fitness curve is informative about the paths respective groups take toward problem solution. For example, groups 1 and 2 converged at approximately the same fitness levels (about 0.65, indicating positive movement toward an ideal solution), but their paths to this point were quite different. Group 1's discussion moved toward an ideal solution immediately when compared to group 2, whose initial approach seemingly took them away from the goal (indicated by the negative fitness initially) only to recover later. Similarly, comparing groups 3 and 4, we can see them settling into different plateaus of fitness albeit after some chaos (fluctuations in fitness levels) initially. Further, comparing groups 1 and 2 with groups 3 and 4, we can see that the discussions of groups 1 and 2 ended quickly whereas those of groups 3 and 4 settled into an "equilibrium" after the initial fluctuations. What is most interesting is that this interpretation of fitness curves provides a view of paths to a solution that are lost in analysis systems that consider only a given point in the solution process, thus assuming that similar behaviors or states at a given point are arrived at in similar ways. As different paths can lead to similar results, uni-dimensional analyses that consider only single points in time (often only the solution state) are not consistent with what we know about problem-solving processes and are not informative about movement toward a goal.

Third, the fitness curve of groups 3 and 4 also highlight the notion of "fitness inertia," i.e., having settled into fitness equilibrium, these groups found it difficult to move in new directions. Of course, group 3 did not have a need to do so, as their high fitness value indicates movement toward an ideal solution. But implications of fitness inertia for groups that equilibrate at low fitness levels indicating no or very little movement toward higher fitness levels, such as what occurred with group 4, are grave. It follows from this that the eventual performance of groups exhibiting fitness inertia can be predicted early on in the discussion. Because our analyses showed convergence (and not the position) to be a significant predictor of group solution quality ($F = 50.245$, $p < .0001$), it preliminarily suggests that the net number of positive steps (the position) is not as critical to the success of a discussion as convergence is. This can be explained by the fact that convergence, being a ratio, is designed to be more sensitive to initial steps, both positive and negative, than steps that are taken later on in the process. Hence, convergence takes into account not only the number of positive and negative steps, but also the order in which they are taken. Many studies of problem solving typically focus on the number of positive steps (such as instances of higher-order thinking, questioning, etc.) as an indicator of the quality of the discussion and learning. Our methodology reveals that a simple frequency count needs to be combined with a measure that takes into account the temporality and order of the steps as well.

Fourth, the end-point of the fitness curve represents the final fitness level or convergence of the discussion. From this, the extent to which of a group was able to solve the problem can be deduced. In other words, we can deduce that, comparatively, group 3 did the best followed by group 1, group 2, and finally group 4. Furthermore, the final fitness levels can also be compared with the maximum fitness level of 1. One might imagine that an ideal fitness curve is one that has all the pushes in the right direction, i.e., a horizontal straight line with fitness equaling 1. However, the data suggests that, in reality, some level of divergence of ideas may in fact be a good thing. Note that, at present, one can only extract a comparison either between groups or with the upper and lower bounds of fitness (1 or -1). But, with repeated application in other research contexts and settings and over multiple studies, norms for absolute values of convergence and fitness will begin to emerge.

Finally, based on the above analysis of the characteristics of fitness curves and what they tell us about the problem-solving process, we can begin to conceptualize how problem-solving processes (individual or collective) may be scaffolded to achieve optimal outcomes. For example, the fitness curves of groups 2 and 4 suggest a need for scaffolding early on in the discussion.

Persistence

In addition to looking at the fitness characteristics of a discussion as a whole, one can also examine how ideas or families of ideas emerge and persist during the course of the problem-solving discussion. In our study, these families of ideas are represented by the 6 major functional categories - problem definition, orientation, solution development, non-task, simple agreement, and simple disagreement—into which all interactions were categorized. Treating each functional category as a component of the problem-solving system, its *usage* (or *persistence*) can be tracked as a measure of evolutionary activity. The central assumption is that components of a complex system that persist and continue to be used make greater contribution to the system. Equivalently, functional categories that persist and get used repeatedly make a greater contribution to the problem-solving activity. Therefore, by examining the persistence of functional categories, we can gain insights into the problem-solving process that would otherwise remain elusive.

More formally, let $f_k(t)$ denote whether the k^{th} functional category exists in the problem-solving system at time t:

$$f_k(t) = \begin{cases} 1 & \text{if component } k \text{ exists at time } t \\ 0 & \text{otherwise} \end{cases}$$

$f_k(t)$ is simply an *activity indicator function* that "switches on" each time an interaction unit belonging to a particular functional category exists in the discussion. In order to measure the usage of a functional category, we can define a corresponding function – an *activity incrementation function* – that increases by 1 each time the indicator function "switches on." Then, the value of the incrementation function for the k^{th} functional category at time t, say $a_k(t)$, reflects its cumulative *usage* up until time

t, i.e., the persistence of the functional category up until time t. Formally, $a_k(t) = \sum_0^t f_k(t)$.

Figure 3 shows the persistence curve of the problem definition and solution development functional categories for two groups. We decided to illustrate persistence using these two categories because they had the most manifest interactions compared to the other four categories.

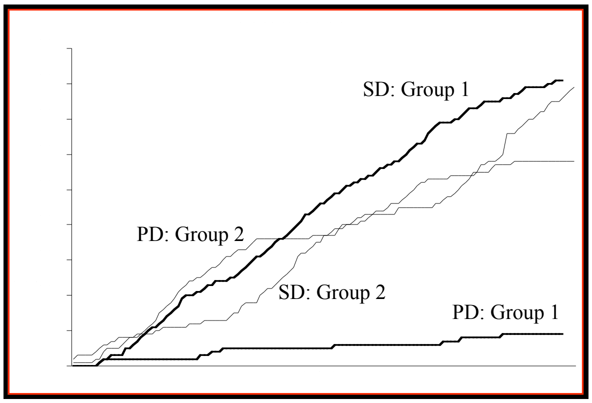

Figure 3: Persistence curves of Problem Definition (PD) & Solution Development (SD) functional categories

Interpreting Persistence

First, being a cumulative function, it is a non-decreasing curve whose end-point indicates the total activity in a given functional category, i.e., the number of interaction units in that functional category. Often, it is this number that is used as a measure in quantitative content analysis. However, the number alone does not indicate anything about the evolutionary activity of the functional category it represents. Persistence curves provide that trajectory from which meaningful insights may be drawn.

Second, a plateau on the persistence curve of a functional category indicates a period in a discussion where no interactions of the type that the functional category represents take place. Therefore, persistence curves that plateau often and for long periods are indicative of a passive functional category. Similarly, a persistence curve that does not plateau is indicative of an active functional category. For example, the problem definition (PD) functional category for group 1 is an example of a passive functional category whereas the PD functional category for group 2 is an active one. In other words, this suggests that group 1 either did not see the need to define the problem or was able to define it quickly and move on, whereas group 2 seemed to need much more time and discussion for problem definition. Note that, in either case, this does not indicate whether or not the problem definition was correct, which can be revealed by cross-validating persistence curves with fitness curves.

Third, persistence curves bring out the notion of competition among functional categories. For group 1, only the SD functional category is active whereas both PD and SD functional categories are active for group 2. This suggests that the problem-solving process was by and large linear for group 1: they defined the problem early on and then worked on developing a solution. There was little or no competition between the PD and SD functional categories. However, the process was quite the opposite for group 2: their attempts to define the problem and develop a solution were iterative and intermingled making the process non-linear and chaotic. There was high competition between the two functional categories. At this point, it is difficult to use the level of competition to make inferences about the quality of the discussion or the resulting solution. However, repeated application in other research contexts and settings and over multiple studies will provide greater validity for the inferences.

Usefulness and Limitations

Reliability and Validity

The inferences that one can draw from the new measures are strong in so far as the coding scheme is reliable and valid. In this study, we opted to use an existing coding scheme, namely the functional category system (FCS) developed by Poole and Holmes (1995). The reasons for choosing the FCS as the interaction coding protocol include:

i. The FCS was developed specifically for the purpose of studying small-group, collaborative interactions in problem-solving contexts,
ii. The FCS categories are theoretically well-grounded in the cognitive and educational theories of problem solving thereby increasing their content validity, and
iii. The FCS has been tried and tested in several research studies (e.g. Jonassen & Kwon, 2001) making it inherently more reliable than developing an entirely new coding scheme (Rourke & Anderson, 2004).

Limitations

As with any new methodology, its repeated application and modification over multiple data sets is needed before strong and valid inferences about the underlying cognitive processes can be made (Rourke & Andersen, 2004). Another limitation includes the requirement of capturing rich and meaningful data in which there is ample opportunity for evolutionary structures and goal-seeking adaptations to occur. In our study, we ensured this by making the objects of the activity—the problems—rich in context. While capturing the data was made easy due to the technology, data analysis proved time consuming. As such, this approach is a useful analytical framework for researchers but not for classroom teachers. However, inferences drawn by researchers using our methodology may have implications for the classroom, work-group, or organization especially with regard to the design and scaffolding of instruction and learning environments for problem-solving tasks.

Usefulness of the Methodology

A major strength of the methodology is in its potential application to other problem-solving settings and contexts. We argue that the proposed methodology would be applicable to the analysis of any process that is a) goal-directed, b) complex and adaptive, and c) well-manifested through rich and meaningful artifacts (which we broadly define to include not only physical behaviors, actions, and products but also conceptual artifacts such as concepts and ideas). As such, the methodology may be applied to individual or collaborative problem-solving, in domains other than physics, with other populations, in a modality other than online, synchronous chat, and using other categorization coding schemes.

Future Directions

In an extension of this research, we are developing and testing new measures, especially at a macroscopic level of analysis. In particular, we are focusing on isolating the phases through which the problem-solving process moves. Such sequences of phases often alternate between stable phases, interspersed with chaotic phases. One can then calculate and predict the probabilities of moving from one phase to another using Hidden Markov Models (HMM). As a result, one may begin to understand when and why phase transitions, cascades and catastrophes (sudden, mass change), as well as stable phases emerge; more importantly, one may begin to understand how the configuration of one phase may influence the likelihood of moving to any other phase. Whether one can control or temper these phases, or whether such control or temperance might prove an unwise practice remains an open question which, even if only partially answered, will be a major breakthrough in characterizing and modeling the problem solving process.

Through such an endeavor, cognitive and education researchers who wish to study the problem-solving process will find choices among several lenses at several resolutions. With measures to analyze number, function, fitness, sequencing, and transition of states, as well as the evolutionary activity of components (concepts, strategies, or functional categories), one can zoom from the micro- to macroscopic properties and behaviors of the problem-solving process.

References

Akhras, F. N., & Self, J. A. (2000). Modeling the process, not the product, of learning. In S. P. Lajoie, *Computers as cognitive tools, volume two: No more walls* (pp. 3-28). Mahwah, NJ: Lawrence Erlbaum Associates.

Anderson, J. R. (2000). *Cognitive psychology and its implications*. New York: Worth Publishers.

Barab, S. A., Hay, K. E., & Yamagata-Lynch, L. C. (2001). Constructing networks of action-relevant episodes: An in-situ research methodology. *The Journal of the Learning Sciences*, *10*(1&2), 63-112.

Bar-Yam, Y. (1997). *Dynamics of complex systems*. Perseus Publishing: New York.

Bedau, M. A., & Packard, N. H. (1992). Measurement of Evolutionary Activity, Teleology, and Life. In C. Langton, C. Taylor , D. Farmer, and S. Rasmussen, eds., *Artificial Life II*, Santa Fe Institute Studies in the Sciences of Complexity, Vol. X, Redwood City, CA: Addison-Wesley, pp. 431-461.

Bedau, M. A., Snyder, E., & Packard, N. H. (1998). A classification of long-term evolutionary dynamics. In C. Adami, R. Belew, H. Kitano, & C. Taylor (Eds.), *Artificial Life VI* (pp.189-198), Cambridge: MIT Press.

Cavalli-Sforza, L. L., & Feldman, M. W. (1981). Cultural transmission and evolution: a quantitative approach. Princeton, N.J.: Princeton University Press.

Chi, M. T. H. (1997). Quantifying qualitative analyses of verbal data: A practical guide. *The Journal of the Learning Sciences*, *6*(3), 271-315.

Crutchfield, J. P. (1994). The calculi of emergence: Computation, dynamics, and induction. *Physica D*, *75*, 11-54.

Derry, S. J., Gance, S., Gance, L. L., & Schlager, M. (2000). Toward assessment of knowledge-building practices in technology-mediated work group interactions. In S. P. Lajoie (Eds.), *Computers as cognitive tools, volume two: No more walls* (pp. 29-68). Mahwah, NJ: Lawrence Erlbaum Associates.

Funke, J. (1991). *Solving complex problems: Human identification and control of complex systems*. In R. J. Sternberg & P. A. Frensch (Eds.), Complex problem solving: Principles and mechanisms (pp. 185-222). Hillsdale, NJ: Lawrence Erlbaum.

Heylighen, F. (1988). Formulating the problem of problem-formulation. In R. Trappl (Ed.), *Cybernetics and Systems '88* (pp. 949-957). Dordrecht, Netherlands: Kluwer Academic Publishers.

Jonassen, D. H. & Kwon, H. I. (2001). Communication patterns in computer-mediated vs. face-to-face group problem solving. *Educational Technology, Research and Development*, *49*(1), 35-52.

Mitchell, M. (1998). A complex-systems perspective on the "computation vs. dynamics" debate in cognitive science. In M. A. Gernsbacher and S. J. Derry (eds.), Proceedings of the 20th Annual Conference of the Cognitive Science Society---Cogsci98, 710-715.

Newell, A., & Simon, H. (1972). *Human problem solving*. Englewood Cliffs, NJ: Prentice Hall.

Poole, M. S., & Holmes, M. E. (1995). Decision development in computer-assisted group decision making. *Human Communications Research*, *22*(1), 90-127.

Port, R. F., & van Gelder, T. (1995). *Mind as motion*. Cambridge, MA: MIT Press.

Quesada, J.F, Kintsch, W. and Gomez, E. (2001) A Computational Theory of Complex Problem Solving Using the Vector Space Model (part I): Latent Semantic Analysis, Through the Path of Thousands of Ants. In Cañas (Ed.) Cognitive research with Microworlds, pp. 117-131.

Ross, S. M. (1996). *Stochastic processes*. New York: John Wiley & Sons, Inc.

Rourke, L., & Anderson, T. (2004). Validity in quantitative content analysis. *Educational Technology, Research and Development*, *52*(1), 5-18.

Simon, H. A. (1978). *The science of the artificial*. Cambridge, MA: MIT Press.