

Concepts, Structures, and Goals: Redefining Ill-Definedness

Collin Lynch, *Learning Research and Development Center & Intelligent Systems Program, University of Pittsburgh, Pittsburgh, PA, USA.*
(collinl@cs.pitt.edu)

Kevin D. Ashley, *School of Law; Learning Research and Development Center; & Intelligent Systems Program, University of Pittsburgh, Pittsburgh, PA, USA.*
(ashley@pitt.edu)

Niels Pinkwart, *Institut für Informatik, Technische Universität Clausthal, Clausthal-Zellerfeld, Germany.*
(niels.pinkwart@tu-clausthal.de)

Vincent Aleven, *Human-Computer Interaction Institute, School of Computer Science, Carnegie Mellon University, Pittsburgh, PA, USA.*
(aleven@cs.cmu.edu)

Abstract. In this paper we consider prior definitions of the terms “ill-defined domain” and “ill-defined problem.” We then present alternate definitions that better support research at the intersection of Artificial Intelligence and Education. In our view both problems and domains are ill-defined when essential concepts, relations, or criteria are un- or underspecified, open-textured, or intractable requiring a solver to *recharacterize* them. This definition focuses on the core structural and pedagogical features that make problems and domains ill-defined while providing a consistent and functional frame of reference for this special issue and for future work in this area. The concept of ill-definedness is an *open-textured concept* where no single static definition exists. We present the most suitable definition for the present goals of facilitating research in AI and Education, and addressing the pedagogical need to focus learners on addressing this ambiguity.

Keywords. ill-defined problems, ill-defined domains, argumentation, metacognition.

INTRODUCTION

In his 1961 paper, “Steps Towards Artificial Intelligence” (Minsky, 1995), Marvin Minsky noted that artificial intelligence, as he conceived it, is largely concerned with problems that are *well-defined*. For this class of problems, first delineated by John McCarthy (1956), there exist systematic methods to validate correct solutions. Thus, they are *Turing-recognizable* languages (Sipser, 1997). All other problems fall into the residual category: *ill-defined*. This definition set up a “frontier” view of ill-definedness. Under this view, problems are ill-defined until they are surveyed, cleared, and strung with barbed wire. According to this view, using terms like “ill-defined” suggests that one simply has not tried hard enough to

develop general methods for problem solving or objectively evaluating the validity of proposed solutions. We disagree.

In our opinion, domains like law, ethics, history, public policy, and architecture are *inherently* ill-defined, as are most of the problems in them. This differentiates them from domains such as Newtonian mechanics. Far from needing to eliminate this ill-definedness, researchers in AI and education should focus on developing systems that address it directly for both functional and pedagogical reasons. In our opinion, the fact that a domain or problem is ill-defined should not be an impediment to useful work in AI and education; rather, as the work in this volume demonstrates, it is an opportunity.

In this work we present a definition for the terms *ill-defined problem* and *ill-defined domain* intended to provide a concise basis for framing and motivating both the discussion in this special issue and future work in this area. We will highlight some of the definitions that have followed McCarthy's before presenting our own. We then conclude with a survey of related work along these lines, including new work published in this volume.

PREVIOUS DEFINITIONS OF ILL-DEFINEDNESS

McCarthy's original characterization has been reexamined and extended by a number of researchers including Walter Reitman (1964; 1965), Alan Newell (1969), Herbert A. Simon (1973), James F. Voss (2006; Voss, Greene, Post, & Penner 1983), Namsoo Shin, David H. Jonassen, and Steven McGee (2003; 1997), Antonija Mitrovic and Amali Weerasinghe (2009) and ourselves (Lynch, Ashley, Alevan, & Pinkwart, 2006). In addition to these general reviews other authors have focused on specific classes of ill-defined problems such as Design Spaces (Goel & Pirolli, 1992), Medical Diagnosis (Pople, 1982), and Wicked Problems (Rittel & Webber, 1973; Conklin, 2006).

Each author approached the topic with a different purpose. Minsky sought to differentiate classical search problems from other tasks. Newell sought to address the relationship between a problem's definition and the abilities of the problem solver, Simon to illustrate extensions of the General Problem Solver, and Voss to examine the process of problem solving in these domains. The resulting *goal-driven* definitions cover many but not all of the same concepts, and vary in their structure and emphasis. Each served to focus attention on the aspects of ill-definedness of interest to the author. Shin, Jonassen & McGee focused on the educational role of ill-definedness but did not address the role of AI in that process. Mitrovic and Weerasinghe, by contrast, focused primarily on the structural aspects of constraint-based tutoring systems in ill-defined domains.

In this work we take into account both the cognitive and structural aspects of ill-defined domains and problems. Rather than attempting a complete survey of prior definitions, we will focus in detail on the work of Reitman, Newell, Simon, Voss, Shin, Jonassen and McGee, citing other authors where appropriate. Their work is illustrative of the range of definitions and has been frequently cited by others.

Clarification of Terms

Before proceeding with this discussion, we will clarify a few terms. The words *ill-structured* and *ill-defined* are often used synonymously in the literature. We adopt the latter here and, when necessary, have substituted for clarity. Additionally, some of the work surveyed, including our own, focuses on ill-defined *problems* while other work focuses on ill-defined *domains*. In our view, a problem is defined

by one or more goals that the solver must achieve, such as winning a court case, designing a house, or calculating the speed of a moving car. In order to do so, they must apply relevant knowledge to achieve their goal given their initial position or knowledge. Domains, by contrast, are conceptual spaces or fields of study such as law and Newtonian mechanics. They are characterized by relevant declarative, pedagogical, or structural knowledge such as legal statutes, rhetorical techniques, or the principles of linear kinematics and the acceleration constant.

In educational settings, domains are structured pedagogically, in terms of the declarative or procedural knowledge that students should understand or the professional practice for which they study. Problems are used functionally to teach students about the domain. The domain delineates, both explicitly and implicitly, the knowledge required for problem solving. In addressing a problem and obtaining feedback, students exercise, expand, and refine their command of that knowledge. We will return to this relationship between domains and problems below.

Reitman, Newell and Simon: Components as Constraints

Minsky's original characterization was first adapted by Reitman (1964, 1965). Reitman noted that most of the problems addressed by human problem solvers, from design to hiring, are ill-defined. In Reitman's view, problems are defined by a set of initial characteristics and a *problem requirement* that specifies the task to be performed and the solution criteria. Suppose, for example, a composer is given a sheet of paper and a pianoforte, and then commanded to "compose a fugue." The assigned materials constrain both the problem solving process and the solution. Providing the composer with a piano for example, restricts him, perhaps unintentionally, to producing fugues that are "pianistic." Reitman illustrated his discussion with a think-aloud protocol of music composition collected from an experienced composer.

In Reitman's formulation a problem is well-defined if the problem solver has access to a complete description of the problem requirements and to all relevant concepts and terms or *problem components*. It is ill-defined if one or more of these components is left unspecified or is *open-textured*. An open-textured concept is one whose application is not automatic but which requires judgment and is context-dependent (Berman & Hafner, 1985). Open-textured concepts are common in ill-defined domains such as law where the definitions of concepts such as "negligence," "income," or "vehicle" are a source of dispute, or in music where pieces are structured in terms of "themes" and "voices."

Solvers of ill-defined problems must decide how to "close" the open constraints by filling in the missing problem components. A solver proceeds through a steady process of transformation or *recharacterization* of the initially ill-defined problem specification to a better defined and more solvable one. Although Reitman cautions against taking it too far, he analogizes this process to linguistic transformations where the term "fugue" transforms to "exposition plus development plus conclusion" (Reitman, 1965). During this process, the original problem is mapped to many overlapping subproblems where the solution to one subproblem in turn constrains the others. For example, selecting a theme for a fugue in turn constrains the choice of counter-themes and the number of voices.

This process of recharacterization, as described by Reitman (1965), is performed both to break the problem down and to tease out other relevant constraints. In some cases the problem solver will make an a priori design decision, such as when Reitman's composer declared his intention at the outset to produce a three-part fugue. In other cases, the solver will reassess a partial solution to identify implicit constraints within it. Half-way through the process Reitman's composer noticed a fourth "voice" in

his fugue and continued with that in mind. While this process is often characterized by monotonically increasing commitments, this is not always the case. Reitman's composer, for example, abandoned his original theme when it came into conflict with a counter-theme that he developed later on. Simon (1973) described a similar process in the design domain noting that it occurs both when problems are solved by a single solver (e.g., an individual architect designing a home) and when a large-scale project is distributed across separate solvers (e.g., the design of a warship). In the former case, an architect might come upon a particularly creative or attractive top-floor layout that requires additional supporting walls below. In the latter case, the addition of a new weapon by the Director of Naval Ordinance might necessitate major structural changes by the Engineer-In-Chief's office or concessions by the Director of Torpedoes.

According to Reitman (1964), the presence of unspecified components in a problem guarantees that no solution will be universally accepted. While a given solution may be acceptable to some, the problem-solving context, the characterization decisions made by the solver, and the judges involved guarantee a range of viable alternatives. For instance, musical legend holds that Frederick the Great, King of Prussia, once commanded Bach to improvise a fugue on the spot with six obbligato voices and a theme of his choosing. Bach wisely chose Frederick's own "Royal Theme," and the resulting fugue made its way into Bach's "Musical Offering," a set of pieces dedicated to the Emperor (Hofstadter, 1979). In this case the judge had an implicit preference that was cleverly divined by the great composer. In Reitman's example, on the other hand, the composer was free to compose on any theme within the conventions of the piano and the fugue.

Experienced composers are better able to compose novel fugues than novices or non-musicians. A master composer of fugues like Bach would find the task far better-defined than a novice with no rules or experience to fall back on. This concept of *solver power* was noted by Newell (1969) who focused on the relationship between the power or potential of a given problem solver and the problem itself. He specifically characterized ill-defined problems as those for which only weak problem-solving methods are available. Simon (1973, 1957) in turn noted that many existing models of problem-solving behavior ignored the important role that resource limitations play in guiding, or bounding, solver behavior.

Voss: Rhetorical Problem Solving

Like Reitman, Voss (Voss, 2006; Voss, Greene, Post, & Penner, 1983) focused on human problem solvers. Voss was concerned with the differences between expert and novice problem solvers and the role of argumentation in ill-defined problems, particularly large-scale social problems. Voss and his colleagues analyzed a series of expert and novice problem solvers each tasked with addressing the then Soviet Union's ongoing agricultural problems. He provided a more detailed, and solver-oriented list of criteria than Simon, identifying the following characteristics of ill-defined problems, solutions to those problems, and the problem-solving process:

1. Ill-defined problems:
 - (a) have vaguely stated goals;
 - (b) no unambiguously right or wrong answers;
 - (c) unstated or assumed problem constraints; and
 - (d) require a large database of relevant information that is often difficult to access.

2. Solutions to ill-defined problems:
 - (a) are rarely correct or incorrect but fall on a range of acceptability; and
 - (b) cannot be judged on their own but require some implementation and evaluation to test.
3. Solvers of ill-defined problems:
 - (a) divide their work into “problem representation” and “problem solving” phases; and
 - (b) justify their solutions by means of argument.

Voss argued that the process of solving ill-defined problems is “rhetorical in nature.” He characterizes well-defined problems as being particularly amenable to computer simulation. His focus, however, on the process by which solution decisions are *justified* is novel. This process was present to some degree in the examples presented by other authors but was not a primary factor in their analyses. Much of his analysis agrees with that of Reitman, Newell, and Simon. This is particularly true of his database criterion (1d). In presenting this, Voss was addressing some of the same issues highlighted by Newell (1969) in his discussion of solver power, and Simon (1957) in his critique of the idealized “economic man.” Simon criticized many existing economic theories for assuming that human decision-makers always make optimal decisions. Simon countered this by proposing a model of “satisficing” where solvers choose the “best” action they can given a limited amount of processing power and background information.

Voss, like Reitman, described problem decomposition in terms of subgoals and subproblems. When faced with the problem of poor Soviet agriculture, for instance, Voss’ experts focused on identifying relevant factors such as the lack of arable land, the need to convince ideologically-motivated decision-makers, and overcome the legacy of previous failures. This is similar to Reitman’s description, which dealt explicitly with decomposing a problem into subproblems, for example, “finding a theme.”

Voss, Greene, Post, & Penner (1983) separated the problem-solving process into distinct recharacterization and problem-solving phases. In his view the experts first frame a problem and then solve it, sometimes repeating the cycle as needed. Goel & Pirolli (1992) described a similar process in their discussion of design spaces. This was a more structured form of the process than Reitman described, but it is similar to Simon’s (1973) example of warship design in an organizational context where, for example, the need for a new design is identified by the First Sea Lord (head of the British Navy) and then passes through a general requirements phase before being passed to each department from Naval Engineering to the Director of Torpedoes. Each group recharacterizes the design according to their needs and makes modifications by imposing new constraints on the problem, setting new goals, and generating new conflicts that must be reconciled.

Shin, Jonassen & McGee: Educational ill-definedness

Voss’ work was continued by Shin, Jonassen, & McGee (Shin, Jonassen, & McGee, 2003; Jonassen, 1997), who focused on ill-defined problems in education, specifically on the cognitive and metacognitive skills necessary to solve ill-defined problems and on the process of problem solving. Their study centered on astronomy problems, such as designing an experiment to verify claims made by another researcher or selecting the best site for a new telescope.

Jonassen (1997) characterized the solving of ill-structured problems as a design process and highlighted the role of metacognitive strategies such as self-explanation play in problem solving. In that paper he articulated a multi-step process. The first step was identifying an appropriate “problem space”

in which to work and articulating the problem's contextual constraints. He also highlighted the role of cases in his list of problem criteria, noting that ill-defined problems provide no general rules for deciding cases and very often exhibit inconsistent relationships among concepts, principles, and rules. Unlike previous authors, Jonassen (1997) drew a distinction between *ill-structured* problems for which multiple solutions exist and *ill-defined* problems where the definition of the problem or problem components are unclear. The two are intertwined in his subsequent discussion so we treat them collectively.

Subsequently Shin, Jonassen, & McGee (2003) enumerated a list of characteristics for ill-defined problems including: 1) the absence of one or more problem elements; 2) the presence of vague or unclear goals; 3) the existence of multiple solutions and multiple solution criteria; 4) uncertainty about what concepts, rules and principles are relevant to a solution or their organization; and 5) the requirement that learners to make and defend judgments about the problem. They then extended the problem-solving process described by Jonassen (1997).

The authors departed from earlier researchers, notably Simon, in arguing that solvers of ill-defined problems must construct a representation of the problem containing all possible states in order to solve it. We take Simon's view that a complete space, even in traditionally "well-defined" problems, can be intractably large making explicit representation impractical. In articulating their definition, Shin and his coauthors focused on the pedagogical characteristics of ill-defined problems. Their goal was to highlight the role that metacognition and justification play in the solving of these problems and the fact that solutions to ill-defined problems differ procedurally and pedagogically from well-defined ones.

PROPOSED DEFINITIONS OF ILL-DEFINEDNESS

As the foregoing discussion illustrates, ill-definedness is an open-textured concept. While each of the provided definitions covers the same or similar problems, its application varies depending upon the authors' goals, and the characterization they choose. No single definition can be considered "definitive." Rather, they all serve as functional tools to frame the discussion at hand. With that in mind:

[1] A problem is *ill-defined* when essential concepts, relations, or solution criteria are un- or under-specified, open-textured, or intractable, requiring a solver to frame or *recharacterize* it. This recharacterization, and the resulting solution, are subject to debate.

[2] *Ill-defined domains* lack a single strong domain theory uniquely specifying the essential concepts, relationships, and procedures for the domain and providing a means to validate problem solutions or cases. A solver is thus required to structure or *recharacterize* the domain when working in it. This recharacterization is subject to debate.

For ill-defined problems this framing/recharacterization is an essential part of the problem-solving process. It may include redefining aspects of the problem to relate it to relevant domain rules and concepts; identifying clear solution criteria; reinterpreting essential rules and concepts according to the present goal; and analogizing or distinguishing the current problem from prior cases. This kind of reasoning is common in law where an advocate will choose precedents (prior cases) and legal rules that

are advantageous for their client, and seek to describe the facts of their case accordingly. In a search and seizure case, for example, where officers had executed an unwarranted search of a mobile home the advocates choose whether to describe it as a ‘home’ – thus emphasizing the privacy implications – or to describe it as a ‘vehicle’ – thus emphasizing mobility and the risk of flight. Each characterization aligns the case with different legal principles and precedents, and carries with it different implications for the warrant requirement.

In addition to aligning the problem with relevant knowledge, recharacterization may also make the problem more tractable by moving from impractical standards such as “the *optimal* chess move” to “a *strategic* move.” Or it may focus the solver on a bounded space of relevant concepts, rules, and relations. The importance of tractability was discussed by both Simon (1973) and Voss (2006). Diplomats and mediators, for example, expend a great deal of effort seeking to clarify and refine their actual goals so as to move from the need to achieve “a lasting peace” to the specific positioning of military forces. The problems of negotiation are discussed in detail by Kim, Hill Jr., Durlach, Lane, Forbell *et al.* (2009). Due to the presence of unspecified, open-textured, and intractable concepts, and the need for recharacterization, ill-defined problems and domains typically:

1. involve open-textured concepts and competing domain principles that are subject to debate;
2. lack widely accepted domain theories identifying relevant concepts and functional relations;
3. cannot be readily partitioned into independent subproblems;
4. have prior cases that are facially inconsistent;
5. involve the need to reason analogically with cases and examples;
6. have a large or complex solution space that prohibits one from enumerating all possible characterizations or solutions;
7. lack formal or well-accepted methods to verify solutions;
8. lack clear criteria by which solutions are judged;
9. are not considered to be “solved” when one solution is presented but may be readdressed by multiple, often distinct, solutions;
10. involve disagreements among domain experts regarding the adequacy of the solutions; and
11. require solvers to justify their solutions through argument.

These characteristics have been discussed by other authors on par with absent information and the need for recharacterization. In our view, there is a causal relationship among them. Features such as the lack of general domain theories, expert disagreement, or facially inconsistent cases, arise out of the presence of open-textured, unspecified and intractable criteria and the consequential need for recharacterization. A formal algebraic theory, for example, requires clearly defined concepts and terms that may be automatically applied to the facts of a given problem and to the background information. Absent such formality in the problem or domain, a single strong and complete theory cannot be readily defined. This is particularly true of policy debates and other “wicked problems” (Rittel & Webber, 1973; Conklin, 2006) where solvers must recharacterize a problem in order to transform a mandate to “promote the general welfare” into programs for highway construction and public education. It is this variability and characterization choice that prompts expert disagreement.

It is not the case that all ill-defined problems exhibit every one of these characteristics. Doctors and diplomats argue about the application of open-textured concepts or past cases to the problem at hand

and justify most if not all of their decisions, sometimes well after the fact. Architects and other artists, by contrast, are rarely asked to justify or even to articulate their design decisions. Nevertheless, all ill-defined problems share the core need to *recharacterize* the problem in order to solve it.

As we noted above, ill-defined domains lack a centrally organizing domain theory describing the relationship among concepts and providing a means to validate cases. When working in such domains it is frequently necessary to enumerate competing concepts and theories as well as conflicting precedents, to select among them, and to defend the choices made. As a consequence, problems in ill-defined domains tend to be ill-defined, since debating the relevance of and relationships among concepts means recharacterizing the problems that involve them.

This is not to say that such domains are free of theories, rather that the theories are not well-accepted, are not formally specified, or are incomplete. Architecture, for example, has been systematized by theorists such as Alexander (1977, 1979) who delineated both a functional language for describing “spaces” and rules for its application. These conventions are often employed by practitioners to characterize their problems and justify their solutions or to scaffold novice students. They are not, however, treated as complete and final even by their adherents.

Ill-defined domains are often defined in terms of prior cases just as movements or styles of art like Dadaism, Modernism, and Punk are defined extensionally by the art or artists within them (e.g., Man Ray and NOFX), rather than logically by formal rules stating what *isn't* Modern and *is* Punk. When making an argument in an ill-defined domain, or describing some feature of it, it is often necessary to explicitly analogize or distinguish the present work to or from prior cases and competing theories.

Many domains have a well-defined core characterized by a single, strong theory, but exhibit ill-definedness at the “fringes” where the theory begins to break down and otherwise well-defined concepts fail. Lakatos (1976) provides an illustration of this process in the domain of geometry. In his dialogue students discuss candidate definitions for the concept of a polyhedron. The proposed definitions are challenged with open questions or selected counterexamples that prompt redefinitions of this seemingly basic concept.

We do not view the distinction between well-defined and ill-defined domains as a binary one but as poles on a weak continuum. At one end lie *strictly well-defined problems* such as the puzzle problems described by Jonassen (1997) (e.g., “Missionaries and Cannibals”), while on the other extreme lie artistic problems such as the need to “express fear in your work.” As one moves along this continuum, the role that theories and cases play in structuring the domain and evaluating solutions changes. Teaching students to work in these domains means teaching them to recognize the varying roles that theories and cases play, and to compare, analogize, and critique alternatives while defining their own.

We believe that these definitions are amenable to work in AI and education. They focus attention on the core features that make problems and domains ill-defined. They also focus on recharacterization, the central skill for ill-defined problem solving, and a pedagogically important aspect of ill-defined problems and domains. The definitions also frame the problem for future research: How can we design educational technologies that support flexible recharacterization of ill-defined problems and domains?

Borderline Cases

Ill-defined and well-defined problems are often facially similar. Both can involve recharacterization and, to novice problem solvers, both can appear intractable or underspecified. Consider the following

problems:

Car A 1000kg car rolls down a 30 degree hill. What is the net force acting on the car?

Checkerboard Given a checkerboard with two opposing corners removed, is it possible to tile the board with dominoes each one covering two adjacent tiles?

Fraction Given a proper fraction $\frac{X}{Y}$, if you add 1 to both the numerator and denominator, will the resulting fraction be larger, smaller, or the same?

The first is a standard problem in introductory physics. It may be solved using kinematics, the Work-Energy Theorem, or a combination of the two. Problem solvers must recharacterize the problem in terms of a particular set of equations as part of their solution process. This characterization is not, however, driven by the presence of open-textured, underspecified, or intractable concepts, nor will it change the outcome of the problem. Both sets of equations fit into the same theory of classical mechanics, both are equivalent and, if applied correctly, will produce the same result. Neither characterization is, in the context of physics, subject to reasonable debate. Thus, although the problem necessitates recharacterization it is not ill-defined.

The Checkerboard problem has a lengthy history in AI. The common solution recharacterizes the problem in terms of paired tiles and color matching. However, the problem includes no open-textured or unspecified concepts. If one considers all possible board sizes, then the space of possible tilings grows large but is still regular. Thus the problem is tractable. No recharacterization of the problem will change the answer nor will it be subject to reasonable debate. If, however, one seeks the *least-creative* solution as in McCarthy (1999), then it becomes ill-defined as creativity is an open-textured concept, subject to ongoing debate (Buchanan, 2001).

This is also the case for the fraction problem. The problem itself is well-defined with a clear logical solution.¹ For a student who is unfamiliar with logic and algebra, the problem may appear ill-defined as they may not know where to begin. This, however is categorically different from the problems of architecture or ethics where no amount of expertise can provide *the* indisputable answer. Similarly, while there exist a number of viable representations for the fraction problem ranging from pieces of pie to glasses of water, all valid representations are equivalent unless one changes the criterion to seek the best or *most pedagogically effective* representation. Thus, in making the distinction between ill-defined and well-defined problems the framing of the problem, and assumptions regarding its solution, make a crucial difference.

¹In algebraic terms, where \bowtie designates the unknown relation, the solution is:

$$\left(\frac{X}{Y} \bowtie \frac{(X+1)}{(Y+1)}\right) \Rightarrow ((X(Y+1)) \bowtie (Y(X+1))) \Rightarrow (XY+X) \bowtie (YX+Y) \Rightarrow X \bowtie Y \quad (1)$$

$$\text{Thus: } (X < Y) \Rightarrow \left(\frac{X}{Y} < \frac{(X+1)}{(Y+1)}\right) \quad \& \quad (X = Y) \Rightarrow \left(\frac{X}{Y} = \frac{(X+1)}{(Y+1)}\right) \quad \& \quad (X > Y) \Rightarrow \left(\frac{X}{Y} > \frac{(X+1)}{(Y+1)}\right) \quad (2)$$

ILL-DEFINEDNESS, AI AND EDUCATION

Pedagogical Roles of Ill-defined Problems

As Reitman (1964) noted many, if not most, of the problems addressed by human problem solvers are ill-defined, a point that was echoed by Jonassen (1997) who described them as “the kinds of problems that are typically encountered in everyday practice.” Some domains such as law, music, policy, and design are ill-defined and are characterized by ill-defined problems. In these domains, the job of a practitioner is to deal with this ambiguity, identify the open-textured concepts, apply them, and in some cases exploit the ambiguity to develop creative solutions. The pedagogical role of ill-defined problems is to train students in these processes, giving them an opportunity to practice these skills in an authentic way.

Jonassen (1997) argued that ill-defined problems are more interesting, meaningful, and thus motivating to learners. Shin, Jonassen, and McGee (2003) further argued that solvers of ill-defined problems employ different skills, including metacognitive skills such as goal monitoring and reflection, than do solvers of well-defined problems. Jonassen (1997) also argued that the domain dependent nature of ill-defined problems compels problem solvers to treat them as “realistic situations” where the solution requires access to a large amount of relevant and well-organized domain knowledge. Finally, in Shin, Jonassen, and McGee (2003), they further argued that solving ill-defined problems requires domain knowledge that is structured around experiences rather than fundamental principles, as these problems rely more on case-based reasoning than formal derivations.

In their analyses Reitman, Simon, Voss, Shin, Jonassen and McGee all highlight the need to explore alternate problem characterizations. Voss, Greene, Post, and Penner (1983), in his analysis of the Soviet Agriculture problem noted that the expert problem solvers spent a great deal of time analyzing their characterizations of the problem and considering alternate characterizations. If, for example, the task was seen as a political problem then the solver would need to address issues like land reform. If, however, it was primarily seen as a technical issue then they would address it through education and investment. Novices, by contrast, focused exclusively on “low level” factors such as the lack of arable land or the need for more tractors, and proceeded to address them separately with little consideration for the interaction among factors or the consequences of their decisions.

Pople (1982) identified a similar issue of early commitment in his work on medical diagnosis. He argued for adopting educational strategies designed to support early exploration of diagnoses. This need for guidance was strongly echoed by Shin, Jonassen, and McGee (2003) who noted that, while problem-solving strategies were not a significant predictor of success on ill-defined problems, metacognitive knowledge such as monitoring skills and ability to consider alternate goal formulations was predictive.

Voss, Greene, Post, and Penner (1983) tested the same body of novices both before and after a course on Soviet policy and noted little change in their behavior, suggesting that the skills required to solve ill-defined problems were not being taught in classrooms. In addition to experts in policy formation and novice students, Voss also examined the behavior of “nonexpert-experts,” that is, individuals with expert-level training in scientific reasoning but with no policy experience. This group behaved very much like novices, highlighting the importance of domain knowledge in the problem-solving process.

This brings us back to Minsky’s frontier view of ill-defined problems. Under this view researchers in AI and education should focus on defining and structuring ill-defined domains, or subsets of them, so

as to provide well-defined tutoring environments for instruction. While such an approach has a place, particularly in scaffolding lower-performing students, this is not a complete solution. Restricting students to the drained and tilled land will not help them to traverse the swamp. Instead, researchers in AI and education should develop systems that focus on the ambiguity, support students in the process of recognizing underspecified or open-textured concepts, help them to frame and recharacterize problems, teach them to explore their ramifications, and then to assess and respond to the consequences of their decisions.

Implementations

Developers of tutoring systems or other instructional materials may explore a number of avenues in helping students to address ill-defined problems. As noted above, Pople (1982) endorsed the development of systems that support the exploration of alternate problem hypotheses to prevent early student commitment. Jonassen (1997) went further, articulating a problem-solving process and a series of pedagogical recommendations including the development of case bases and supporting the construction of knowledge bases that reflect real-world knowledge.

Tutoring systems can be and have been constructed that support these kinds of analyses. In a prior discussion (Lynch, Ashley, Alevén, & Pinkwart, 2006), we identified four human tutoring strategies that were ideal for incorporation into intelligent tutoring systems: case studies; collaboration; weak theory scaffolding; and expert review. We also described a number of systems that employed one or more of these strategies to provide useful instruction. Some of these systems were constructed around “classical” well-defined models but employed them to facilitate more open ended analysis (e.g., CATO (Alevén, 2003)). Others used constraints to express an open but guided domain for exploration (Holland, 2000; Mitrovic & Weerasinghe, 2009) or focused on the support of open discovery in the target domain (Suthers, Toth, & Weiner, 1997). Still others took a step back from modeling the *solutions* to a given problem and instead supported the analysis of prior examples (Pinkwart, Alevén, Ashley, & Lynch, 2007) or in leveraging peer guidance (Cho & Schunn, 2007).

As the papers in this volume demonstrate, research in the design and use of tutoring systems for ill-defined domains is ongoing. Ogan, Alevén, and Jones developed a tutoring system for intercultural competence (Ogan, Alevén, & Jones, 2009). The goal of this instruction is to help students in learning to analyze and explain cultural differences, a basic recharacterization task. Cultural learning depends heavily on students’ interpretations both of events and language, and the domain lacks strong rules for case comparison. Their system, ICCAT, takes a *process-support* approach to tutoring. Students are given an authentic problem in cultural translation, in this case a film clip from the target culture, and the system guides them in a clear analysis process consisting of noticing key features, making analyses, receiving feedback, and then reflecting upon their choices.

By contrast, Kim, Hill Jr., Durlach, Lane, Forbell *et al.* (2009) chose a more model-driven approach to the task of tutoring cross-cultural negotiation. In their analysis of the problem space, they note that negotiation practice involves characterizing more abstract interests (i.e., what one really wants) and then, in the course of negotiation, assessing whether a given proposal satisfies those desires. Students need to both reify their personal goals and to recharacterize the debate in terms of them, thus making the problems ill-defined. Their system, BiLAT, is constructed around a strong domain model that governs system behavior. Students using the system engage in a simulated negotiation in a game-like setting

based upon real-world situations. They receive feedback from the model both as explicit guidance and simulated responses from the in-game characters. Here the decision to construct an explicit though partially-hidden model enables the system to directly scaffold novices while still keeping a range of open behaviors.

Unlike Ogan et al. and Kim et al., Kazi, Haddaway and Suebnukarn chose a more explicit focus on the role of recharacterization (Kazi, Haddawy, & Suebnukarn, 2009). Their task domain is problem-based learning in medicine and the need to support creative student solutions to diagnosis problems. They describe extensions to the COMET system that allow it to use preexisting expert solutions and a medical ontology to support novel student solutions. Students using the system enter diagnostic hypotheses using a graphical markup language, explicitly recharacterizing the problem in functional medical terms. The system employs an ontology and stored expert solutions to recognize when the students employ more general, or more specific variations of the experts' reasoning. This allows the students to explore a range of viable diagnoses while still receiving system guidance.

CONCLUSION

In this article we have examined prior definitions of ill-defined problems and ill-defined domains with the goal of framing discussion of ill-defined problems in an AI and education context. Ill-definedness is an open-textured concept, and many researchers have approached it with diverse goals. The resulting definitions share many of the same essential concepts but focus on aspects of ill-definedness relevant to the authors' underlying purposes. We then presented our own definitions for ill-defined problems and ill-defined domains with the goal of setting a workable stage for future work.

The definitions that we presented here focus on the presence of un- or underspecified, open-textured, or intractable concepts in ill-defined problems and domains, and on the role of recharacterization in addressing them. In our view these are the root features from which other salient but optional aspects of ill-definedness spring, such as a lack of expert agreement and the need to justify solutions. Instructing students in this 'real-world' process of recharacterization is a central pedagogical approach for ill-defined problems that researchers in this field have begun to address.

REFERENCES

- Aleven, V. (2003). Using background knowledge in case-based legal reasoning: A computational model and an intelligent learning environment. *Artificial Intelligence*, 150, 183–237.
- Alexander, C. (1977). *A Pattern Language*. New York, New York: Oxford University Press.
- Alexander, C. (1979). *The Timeless Way of Building*. New York, New York: Oxford University Press.
- Berman, D. M. & Hafner, C. (1985). Obstacles to the development of logic-based models of legal reasoning. In C. Walter & L. E. Allen (Eds.) *Computing Power and Legal Reasoning*, (pp. 183–214). Saint Paul, Minnesota: West Publishing Co.
- Buchanan, B. G. (2001). Creativity at the metalevel: AAAI-2000 presidential address. *AI Magazine*, 22(3), 13–28.
- Cho, K. & Schunn, C. (2007). Scaffolded writing and rewriting in the discipline: A web-based reciprocal peer review system. *Computers and Education*, 48(3), 409–426.

- Conklin, J. (2006). *Dialogue Mapping: Building Shared Understanding of Wicked Problems*. Chichester, England: Wiley.
- Goel, V. & Pirolli, P. (1992). The structure of design problem spaces. *Cognitive Science*, 16, 345–429.
- Hofstadter, D. R. (1979). *Gödel, Escher, Bach: An Eternal Golden Braid*. New York, New York: Vintage Books.
- Holland, S. (2000). Artificial intelligence in music education: A critical review. In E. R. Miranda (Ed.) *Readings in Music and Artificial Intelligence*, (pp. 239–274). Amsterdam: Harwood Academic publishers.
- Jonassen, D. H. (1997). Instructional design models for well-structured and ill-structured learning outcomes. *Educational Technology: Research and Development*, 45(1), 65–94.
- Kazi, H., Haddawy, P., & Suebnukarn, S. (2009). Expanding the space of plausible solutions in a medical tutoring system for problem-based learning. *International Journal of Artificial Intelligence in Education*, 19(3), 309–333.
- Kim, J. M., Hill Jr., R. W., Durlach, P. J., Lane, H. C., Forbell, E., Core, M., Marsella, S., Pynadath, D., & Hart, J. (2009). BiLAT: A game-based environment for practicing negotiation in a cultural context. *International Journal of Artificial Intelligence in Education*, 19(3), 289–308.
- Lakatos, I. (1976). *Proofs and Refutations*. Cambridge, England: Cambridge University Press.
- Lynch, C., Ashley, K. D., Alevan, V., & Pinkwart, N. (2006). Defining ill-defined domains; a literature survey. In V. Alevan, K. D. Ashley, C. Lynch, & N. Pinkwart (Eds.) *Proceedings of the First International Workshop on Intelligent Tutoring Systems for Ill-Defined Domains*, (pp. 1–10). Jhongli Taiwan: 8th International conference on Intelligent Tutoring Systems.
- McCarthy, J. (1956). The inversion of functions defined by turing machines. In C. E. Shannon & J. McCarthy (Eds.) *Automata Studies, Annals of Mathematical Studies*, (pp. 177–181). Princeton, New Jersey: Princeton University Press.
- McCarthy, J. (1999). Creative solutions to problems. In *Proceedings of the AISB Workshop on Artificial Intelligence and Creativity*. Edinburgh, Scotland, United Kingdom: AISB.
- Minsky, M. (1995). Steps to artificial intelligence. In G. F. Luger (Ed.) *Computation & Intelligence: Collected Readings*, (pp. 47–90). Menlo Park, California: AAAI/MIT Press.
- Mitrovic, A. & Weerasinghe, A. (2009). Revisiting ill-definedness and the consequences for ITSs. In V. Dimitrova, R. Mizoguchi, B. du Boulay, & A. C. Graesser (Eds.) *Artificial Intelligence in Education: Building Learning Systems that Care: From Knowledge Representation to Affective Modelling, Proceedings of the 14th International Conference on Artificial Intelligence in Education, AIED 2009*, (pp. 375–382). Amsterdam: IOS Press.
- Newell, A. (1969). Heuristic programming: Ill-structured problems. *Progress in Operations Research*, 3, 361–413.
- Ogan, A., Alevan, V., & Jones, C. (2009). Advancing development of intercultural competence through supporting predictions in narrative video. *International Journal of Artificial Intelligence in Education*, 19(3), 267–288.
- Pinkwart, N., Alevan, V., Ashley, K. D., & Lynch, C. (2007). Evaluating legal argument instruction with graphical representations using largo. In R. Luckin, K. R. Koedinger, & J. E. Greer (Eds.) *Artificial Intelligence in Education, Building Technology Rich Learning Contexts That Work, Proceedings of the 13th International Conference on Artificial Intelligence in Education, AIED 2007*, (pp. 101–108). Amsterdam: IOS Press.

- Pople, H. E. (1982). Heuristic methods for imposing structure on ill-structured problems: The structuring of medical diagnostics. In P. Szolovits (Ed.) *Artificial Intelligence in Medicine*, (pp. 119–190). Boulder Colorado: Westview Press.
- Reitman, W. R. (1964). Heuristic decision procedures, open constraints and the structure of ill-defined problems. In M. Shelly, II & G. L. Bryan (Eds.) *Human Judgments and Optimality*, (pp. 282–315). New York, New York: John Wiley & Sons Inc.
- Reitman, W. R. (1965). *Cognition and Thought: an Information Processing Approach*. New York, New York: John Wiley & Sons Inc.
- Rittel, H. & Webber, M. (1973). Dilemmas in a general theory of planning. *Policy Sciences*, 4, 155–169.
- Shin, N., Jonassen, D. H., & McGee, S. (2003). Predictors of well-structured and ill-structured problem solving in an astronomy simulation. *Journal of Research in Science Teaching*, 40(1), 6–33.
- Simon, H. A. (1957). *Models of Man Social and Rational: Mathematical Essays on Rational Human Behavior in a Social Setting*. New York, New York: John Wiley & Sons Inc.
- Simon, H. A. (1973). The structure of ill-structured problems. *Artificial Intelligence*, 4, 181–201.
- Sipser, M. (1997). *Introduction to the Theory of Computation*. Boston, Massachusetts: PWS Publishing Company.
- Suthers, D., Toth, E., & Weiner, A. (1997). An integrated approach to implementing collaborative inquiry in the classroom. In R. Hall, N. Miyake, & N. Enyedy (Eds.) *CSCL '97: Proceedings of the 2nd International Conference on Computer Support for Collaborative Learning*, (pp. 272–279). Georgia: International Society of the Learning Sciences.
- Voss, J. F. (2006). Toulmin's model and the solving of ill-structured problems. In D. Hitchcock & B. Verheij (Eds.) *Arguing on the Toulmin Model: New Essays in Argument Analysis and Evaluation*, (pp. 303–311). Berlin: Springer.
- Voss, J. F., Greene, T. R., Post, T. A., & Penner, B. C. (1983). Problem solving skill in the social sciences. *The Psychology of Learning and Motivation*, 17, 165 – 215.