Borderline Cases of Ill-definedness – and How Different Definitions Deal with Them

Kevin Ashley[1], Collin Lynch[2], Niels Pinkwart[3], and Vincent Aleven[4]

[1]University of Pittsburgh, Learning Research and Development Center,
[2]University of Pittsburgh, Intelligent Systems Program, Pittsburgh, PA, USA
Pittsburgh, PA, USA
[3]Clausthal University of Technology, Department of Informatics, Germany
[4]Carnegie Mellon University, HCI Institute, Pittsburgh, PA, USA
{ashley,lynch}@pitt.edu

Abstract. In prior work, we have proposed definitions of ill-defined problems and domains that focus on the need for framing or recharacterizing the problem in order to generate plausible solutions, where the recharacterizations and solutions are themselves subject to debate. In this paper we elaborate on the framing or recharacterization component of this definition in light of some plausible counterexamples and in contrast with some recent alternative definitions. While the question of defining ill-defined problems and domains is itself ill-defined, we conclude that our formerly proposed definition withstands the counterarguments.

Keywords: ill-defined problems; ill-defined domains, ITSs for ill-defined domains

1 Introduction

In educating students, many problems assigned to them are well-defined. That is, the problem as stated is fairly constrained, the relevant knowledge for solving it is more-or-less apparent and available to the solver, and the answer is unambiguously right or wrong. For instance,

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

By contrast, here is an example of an ill-defined problem of a type that instructors, instructional designers, and ITS designers encounter on a daily basis:

> `**Best-Pedagogical-Solution-Example:**` `What is the most pedagogically effective solution of the problem, "A 1000kg car rolls down a 30 degree hill. What is the net force acting on the car?"`

In solving this latter problem, various questions necessarily arise that make it ill-defined, questions such as,

- "What does "pedagogically effective" mean?",
- "Effective in what pedagogical context, with what type of students, and for what purpose?", and

- "What are the criteria for assessing a proposed solution's pedagogical effectiveness, and how will solutions be compared?"

Arguments will follow proposing definitions of concepts, criteria for assessing solutions, methods for comparing the solutions in relation to the criteria, etc.

In nearly every domain of human endeavor and in nearly every classroom in higher education, the interesting questions are ill-defined. Certainly, in teaching students ITS design, the goal is learning to solve ill-defined problems like the latter example, not just well-defined problems like the former. Arguably, however, ITS technology is better suited to teaching well-defined problems. The Car Example is typical of questions assigned to high school physics students class. This is the sort of problem that ITS technology has proven it can handle effectively. Students' solutions can readily be compared with expert solutions and used to model the level of students' understanding of the problem and to trigger progressively more explicit hints.

By contrast, in solving an ill-defined problem, the relevant constraints on the problem are not readily apparent from the mere statement of the problem, but the solver must uncover the relevant constraints through a process of exploration. For instance, the solver must explore plausible answers to the questions about what "pedagogically effective" means, and what the relevant context, goals, and criteria are for addressing the problem. Often, the answers are arguments in free-form text justifying one solution over others. Different solvers may frame the ill-defined problem differently according to their knowledge, beliefs, and attitudes, and thus may generate different arguments, many of which may be quite defensible. Since an ITS cannot readily interpret a student's natural language arguments, updating a student model and using it to select guidance, feedback and new problems to tutor students, are especially problematic for a computer tutor teaching ill-defined problem solving.

Nevertheless, in recent years, there has been a growing appreciation in the ITS community of the pedagogical importance of developing methods for addressing ill-defined problems. See, e.g., (Aleven, et al. 2006; 2007; 2008; Pinkwart, et al. 2010). This may reflect the fact that, as students progress through the educational system, the problems at the center of their education become increasingly ill-defined, perhaps because ill-defined problem-solving lies at the core of professional practice in fields not only like law, business, government, and ethics, but even in scientific fields as one moves to the research frontiers of theoretical development and experimental design. In addition, many decisions one makes when *writing* an essay or research article tend to be ill-defined.

Given these considerations, the question of how one should define ill-defined problems and domains, both generally and in the ITS context, becomes increasingly pressing. This question is not likely to be resolved easily or for all time; the question of how to define ill-definedness is itself ill-defined and subject to argument exactly as described above. Indeed, it is useful and enlightening for the community of scholars at this Workshop on Intelligent Tutoring Technologies for Ill-Defined Problems and Ill-Defined Domains to debate this issue and to consider alternative plausible definitions and the arguments for and against them.

To that end, in the remainder of this paper, we present in Section 2 alternative recently proposed definitions of ill-definedness including our own definitions. In Section 3, we consider some plausible counterexamples to our definition and use them in Section 4 to elaborate on the meaning of crucial concepts in our definition as well as to compare and contrast the various recent definitions. Our goal is to encourage a discussion of the definitional issues at the workshop that will help all of us in this community of interest to consider which definitions work best, generally and in the ITS context, in light of the underlying issues.

2 Definitions of ill-definedness

In this section, we present four alternative recently proposed definitions of ill-definedness (for a history of sources of definitions of ill-defined problems and domains, see Lynch et al. 2010). In prior work, we have proposed the following definitions of ill-defined problems and ill-defined domains (Lynch et al. 2006; 2010):

> [1.1] 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.
>
> [1.2] 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.

A central feature of our definitions is the requirement that a problem solver frame or recharacterize the problem some or all of whose essential concepts, relations or solution criteria are under-specified, open-textured, or intractable. By "framing" or "recharacterizing" the problem, we mean restating or refining aspects of the problem in order to align it with specific domain concepts; redefining existing rules according to the present goals; clarifying the solution criteria; or analogizing the problem to and distinguishing it from prior examples.

Another definition that has received some currency in recent ITS work summarizes a set of criteria proposed by Simon (1978, p. 286).

> [2] An ill-structured problem is defined by Simon (citing Simon 1973) as one that is complex, with indefinite starting points, multiple and arguable solutions, or unclear strategies for finding solutions (citing Fields 2006). (Nkambou et al. 2008)

A further possible strategy is to define ill-defined problems as the complement, in effect, of a definition of well-defined problems as in Le and Menzel (2008).

> [3] [R]equirements … have been proposed as criteria a problem must satisfy in order to be regarded as well-defined: 1) a start state is available; 2) there exist a limited number of relatively easily formalized transformation rules; 3) evaluation functions are specified and 4) the goal state is unambiguous (citing Jonassen, et al. 1999). If one or several of these conditions is violated, the problem is considered

ill-defined (citing Ormerod 2006). … However, “the boundary between well-defined and ill-defined problems is vague, fluid and not susceptible to formalization” (citing Simon 1973).

A very similar “complementary” definition, also in terms of a state space view, may be found in Mitrovic and Weerasinghe (2009). These authors define as well-defined, tasks or problems having clear start states, goal states, transformations (i.e., problem-solving procedures) that are known to the decision-maker and a correct solution. By negative implication, an ill-defined task or problem has underspecified start states, transformations, or solution criteria. Specifically, the authors emphasize two dimensions for discussing ill-definedness, the definedness of the task and of the domain, which can order the domains/tasks along a continuum, from well- to ill-defined. Their implicit definition is apparent in the following:

[4] Although the ER model [Entity-Relationship data model] itself is well-defined, the task of developing an ER schema for a particular database (i.e. conceptual database design) itself is ill-defined: the initial state (i.e. the set of requirements) is usually underspecified and ambiguous, there is no algorithm to use to come up with the solution, and finally the goal state is also underspecified, as there is no simple way of evaluating the solution for correctness.

3 Possible counterexamples of the “framing” definition

While the authors of any of the above definitions might well agree that framing or recharacterization of a problem may be useful in dealing with ill-definedness, only definitions [1.1] and [1.2] incorporate it formally into the definitions of ill-defined problems or domains.

One challenge to these definitions is that framing or recharacterization is useful in solving problems generally, including well-defined problems. Consider the following examples (Lynch, et al. 2010):

Checkerboard Example: Given a checkerboard with the two opposing corners removed, is it possible to tile the board (i.e., fully cover it) with dominoes each one covering two adjacent tiles?

Fraction Example: Given a proper fraction X/Y, if you add 1 to both the numerator and denominator, will the resulting fraction be larger, smaller, or the same?

Each of these examples illustrates the utility of framing or recharacterizing for problem solving, but the problems are not ill-defined. For instance, the definition of well-defined problems that appears in the complementary definition of ill-definedness in [3] makes that clear. The Checkerboard and Fraction examples each have a clear start state and unambiguously right-or-wrong goal state for which the evaluation criteria are clear and uncontroversial, and there are a few relatively easily formalized transformation rules by which the problem can be framed of recharacterized into a readily solvable form.

The Checkerboard problem is commonly posed in introductory AI classes to illustrate the utility of an appropriate recharacterization of the problem. If one recharacterizes the problem in terms of paired tiles and color matching, it becomes apparent that removing two opposing corners, each of the same color, necessarily implies that the answer is “no”. The Fraction Example also has one right answer that can be derived in multiple ways, each involving a recharacterization of the problem in terms of logical and algebraic expressions, or of interpretation of appropriate graphical representations of fractions.

Even the Car Example from the introduction illustrates the utility of framing or recharacterizing. It may be solved using linear kinematics or the Work-Energy Theorem, or a combination of these, but recharacterizing in terms of a set of equations is required.

Thus, based on these counterexamples, one might argue that framing or recharacterization is not unique to ill-defined problem-solving and is not a defining characteristic contrary to definitions [1.1] and [1.2].

4 Response to counterexamples

The counterexamples, we submit, are not fatal to proposed definitions [1.1] and [1.2], and the response to them underscores the importance of a proviso of each definition: “This recharacterization, and the resulting solution, are subject to debate.” We argue that framing to bring a problem within a method that will generate a verifiably correct answer is different in kind from that required to bring a problem toward a defensible solution that will not be, indeed, cannot be, verifiably correct.

In solving a well-defined problem, the framing is a search for a solution template whose concepts, relationships, and procedures, when applied, will enable the solution to “click into place.” It involves mapping the problem’s facts into the template, often, a small set of alternative formulae or procedures appropriate to that type of problem, using the procedures associated with the template to generate an answer, and validating that the answer is correct. The templates are “representational tricks” of the trade; knowledge of which templates are appropriate to the problem, how to perform the mappings from the problem to the template, and how to validate the answers are key. While the answer to a well-defined problem may need to be justified with an argument, the criteria for validating an answer usually are accepted by the relevant audience (e.g., physics or math instructors) and are applied in a straightforward way. The generally accepted criteria for validating an answer help constrain the framing.

By contrast, in solving an ill-defined problem, a decision-maker is not framing the problem to apply a template or algorithm; rather he frames the problem to identify solutions and to make arguments for and against the solutions. The framing serves to explore plausible solutions and arguments. It involves not just a search for useful concepts to apply to the problem facts that may lead to a plausible solution. Since the concepts frequently are open-textured (e.g., “pedagogically effective”), one must search both for concepts and ways to define them that lead to a plausible solution. Framing is needed to search even for criteria to assess the proposed solution that are acceptable to the relevant audience, and arguments that justify the solution as a good in terms of the concepts and criteria.

In the Car, Fractions, and Checkerboard Examples, the characterization of the problems for solution is not driven by the need to specify meanings for concepts that are open-textured, underspecified, or intractable, nor will it change the outcome of the problem.

For the Car Example, linear kinematics and the Work-Energy Theorem are equivalent and each fits into a common theory of classical mechanics. If applied correctly, each produces the same verifiable result. Within the context of classical physics, neither approach is controversial or the subject of reasonable debate. Thus, although the Car Example may require a kind of recharacterization to apply the relevant formulas, it is not ill-defined. This is also the case for the Fraction Example. There are multiple ways to recharacterize the problem for solution, including recharacterizing it in terms of pieces of a pie or glasses of water, but each different representation yields the same solution. The problem does not include any open-textured or unspecified concepts whose alternative plausible meanings need to be explored; the clear, logical solution and the criteria for evaluating it are similarly beyond debate.

Although a different kind of problem, the Checkerboard Example also lacks any open-textured or unspecified concepts. One might recharacterize it as a search problem (with a huge search space), but no recharacterizing of the problem will change the answer nor will the answer be subject to reasonable debate. If the problem were to find the least-creative solution as in (McCarthy, 1999), it would introduce an open-textured concept and raise questions about what "creative" means, in what context, and for what purpose, and how will degrees of creativity be compared, etc. The new problem would be ill-defined; the recharacterization of the problem adding constraints to answer those questions makes it a matter of debate and argument (Buchanan, 2001).

5 Discussion

Apart from challenging our own definitions, the above examples also underscore some differences among the various definitions of ill-defined problems and domains introduced above.

For one thing, according to definitions [3] and [4], the lack of "a limited number of relatively easily formalized transformation rules" or of an "algorithm to use to come up with the solution" renders a problem ill-defined. Obviously, solving the Car and Fraction Examples involves transformation rules and an algorithm; if one knows which rules or algorithm to apply the solution method is straightforward. Is there, however, a set of easily formalized transformation rules or an algorithm for *recognizing* which solution method to apply? If not, would definitions [3] and [4] characterize even these examples as ill-defined?

According to our definition [1.1], however, the problems are well-defined because no essential concepts, relations, or solution criteria are un- or under-specified, open-textured, or intractable, requiring a solver to frame or recharacterize them. Moreover the recharacterizations, and the resulting solutions, are not subject to debate; there is one correct solution regardless of the representation. The problem solver does have to recharacterize these well-defined problems to solve them, but as argued above, that is not the kind of recharacterization that contributes to making the problems ill-defined.

Presumably, on Definition [2] the Car and Fraction Examples are well-defined, too. They are not complex, have definite starting points, only one non-controversial correct solution, and clear strategies for finding solutions.

In addition, (Mitrovic and Weerasinghe 2009) invite taking into account a learner's knowledge or ability (i.e., or power in Simon's sense), but we think this consideration only adds to the difficulties of defining ill-definedness. For instance, the Fraction Example may appear ill-defined to a student unfamiliar with logic or algebra who does not know where to begin. As a result, he may be unable to recharacterize the problem into a solvable form. Or the solver may simply not see the trick that leads to solving the Checkerboard example. That, however, should not make these problems ill-defined. Considerations of learner power are categorically different from truly ill-defined problems where, as recognized by [1.1], because the recharacterizations are subject to debate, no amount of expertise can provide an indisputable answer.

Similarly, a state space view, as in (Mitrovic and Weerasinghe 2009) (and Simon), adds to the problems of defining ill-definedness. If one considers the generalized Checkerboard Example, with all possible board sizes, then the space of possible tilings grows as large as one likes but the problem is still regular and subject to the representational trick that is not subject to debate, and thus remains well-defined.

On the other hand a virtue of the complementary definitions [3] and [4], and of the disjunctive definition [2] is that they accommodate the intuition that problems lie along a spectrum of well- and ill-definedness. On these definitions, a problem becomes more or less ill-defined as more or fewer of the disjunctive criteria are satisfied. In particular, Mitrovic and Weerasinghe (2009) present a two dimensional view representing the ill-definedness of a task or a domain along which problems may vary. While definitions [1.1] and [1.2] do relate the ill-definedness of problems and domains, they do not appear to accommodate the spectra in any obvious way. A problem of the complementary definitions, [3] and [4], however, is that these or'd (i.e., disjunctive) criteria seem too liberally to ascribe ill-definedness to problems (as argued above).

6 Conclusion

A final consideration is the extent to which the definitions take into account the educational context of the problems, and if so, how they do it. Expert instructors in some domains explicitly employ the terminology of ill-definedness. For instance, in domains like law, professional ethics, medicine, business, and public policy decision-making, debate and argumentation about the application of open-textured terms are explicitly incorporated into the pedagogy, and viewing the problem from the (multiple) viewpoints of those affected by the decision is seen as one driver of how to frame the problem. Something like this appears to be the case in design fields, too, where the arguments about whether and how alternative designs satisfy constraints and what those constraints really mean would likely be important in pedagogy. Definitions [1.1] and [1.2], and to a lesser extent [2], employ this language. Interestingly, although definition [4] is applied to a design task, developing an ER scheme for a database, there is no focus on debating characterizations of the problem. Definition [3] is the most divorced from a particular pedagogical context.

Of course, whether this is a good thing or not is itself, a matter of debate, part of the ill-defined nature of the task of defining ill-definedness. The presentation of alternative current definitions of ill-definedness, the consideration of some possible counterexamples, and their use to underscore the differences among the definitions, it is hoped, will continue the debate on this important issue at the workshop.

References

Aleven, V., Ashley, K., Lynch, C., & Pinkwart, N. (Eds.) (2006). Proceedings of the Workshop on Intelligent Tutoring Systems for Ill-Defined Domains at the *8th International Conference on Intelligent Tutoring Systems.* Jhongli (Taiwan), National Central University.

Aleven, V., Ashley, K., Lynch, C., & Pinkwart, N. (Eds.) (2007). Proceedings of the Workshop on AIED Applications for Ill-Defined Domains at the *13th International Conference on Artificial Intelligence in Education.* Marina Del Rey California 2007.

Buchanan, B. G. (2001). Creativity at the metalevel: Aaai-2000 presidential address. *AI Magazine*, 22(3):13-28.

Fields, A.M. (2006) Ill-structured problems and the reference consultation: The librarian's role in developing student expertise. Reference services review, Emerald 34(3), 405–420.

Jonassen, D.H., Tessmer, M. and Hannum, W.H (1999). *Task analysis methods for instructional design*, Erlbaum 1999.

Le, N.T. and Menzel, W. (2008) Using Constraint-Based Modelling to Describe the Solution Space of Ill-defined Problems in Logic Programming. Advances in Web Based Learning (ICSL 2007) Lecture Notes in Computer Science. V. 4823. pp. 367-379.

Lynch, C., Ashley, K. D., Aleven, V., and Pinkwart, N. (2006). Defining ill-defined domains; a literature survey. In Aleven et al. (2006), pages 1-10.

Lynch, C., Ashley, K. D., Aleven, V., and Pinkwart, N. (2010). Concepts, structures, and goals: Redefining ill-definedness. *International Journal of Artificial Intelligence in Education*: Special issue on Ill-Defined Domains. (in press).

McCarthy, J. (1999) Creative solutions to problems. In *AISB'99 Symposium on AI and Scientific Creativity*, pages 44–48, Edinburgh, Scotland, April.

Mitrovic, A. and Weerasinghe, A. (2009). Revisiting ill-defiedness and the consequences for ITSs. In Dimitrova, V., Mizoguchi, R., du Boulay, B., and Graesser, A. C., editors, AIED, volume 200 of Frontiers in Artificial Intelligence and Applications, pages 375-382. IOS Press.

Nkambou, R., Nguifo, E.M., and Fournier-Viger, P. (2008) Using Knowledge Discovery Techniques to Support Tutoring in an Ill-Defined Domain. Intelligent Tutoring Systems. Lecture Notes in Computer Science V. 5091. Pp. 395-405.

Ormerod, T. C. (2006) Planning and ill-defined problems, in R. Morris & G. Ward (Eds.): *The Cognitive Psychology of Planning*, London: Psychology Press.

Pinkwart, N., C. Lynch, K. D. Ashley, V. Aleven, eds. (to appear 2010). *International Journal of Artificial Intelligence in Education*, Special Issue on AIED applications for ill-defined domains.

Simon, H. (1973) The structure of ill-structured problems, *Artificial Intelligence*, No. 4, pp. 181-201.

Simon, H. (1978) Information-processing theory of human problem solving. In: Estes, W.K. (ed.) *Handbook of learning and cognitive processes*. Human information, vol. 5.