Problem solving and mathematical knowledge

Joseph Corneli*

December 15, 2010

Abstract

This report describes the research goals, and intermediate milestones related to an investigation of the relationship between problem solving and mathematical knowledge in an online mathematics community.

The proposal is to build a problem-solving layer over the encyclopedia layer that comprises the central feature of the current PlanetMath.org. Research will proceed by examining the activities of people in this space (e.g. connecting, discussing, working, recording, sharing, learning, etc.) and analysis of these activities in context, pursuant to creating useful adaptive recommendations for learners. The investigation will include a qualitative component, based on participant observation of peer learning in this space.

We propose to add three core features to the software system that underlies PlanetMath: (1) building blocks of an open platform for testing and evaluating various learning and instruction methods, (2) algorithms for recommendations following the cognitive tutoring approach, and (3) support for end-user participation in developing problem sets and relevant analytics.

Outcomes will include a statistical study of how various activity patterns correlate with indicators of learning, and a set of narratives that assemble these key factors into problem-solving and teaching/learning strategies.

Keywords: mathematical problem solving, ecological approach, activity streams, cognitive tutoring, coherence, peer-based learning.

*`j.a.corneli@open.ac.uk`, Knowledge Media Institute, The Open University, Milton Keynes, UK. The research work described in this paper is partially funded through the ROLE Integrated Project, part of the Seventh Framework Programme for Research and Technological Development (FP7) of the European Union in Information and Communication Technologies.

Contents

Outsiders see mathematics as a cold, formal, logical, mechanical, monolithic process of sheer intellection; we argue that insofar as it is successful, mathematics is a social, informal, intuitive, organic, human process, a community project. – Social Processes and Proofs of Theorems and Programs, DeMillo, Lipton, and Perlis.

1 The Topic: mathematical activities

Mathematics isn't just made of Definitions, Theorems, Proofs, Corollaries, Lemmas, and Examples – there are also Conversations, Questions, Intuitions, Experiments, Diagrams; and so on. Given the central role of problem solving in mathematical thinking ([1], [2]), we can be justified in restricting ourselves to considering the relationship between mathematical problem solving and its noetic and social context.

To be clear, we intend to focus on textbook-style problems at the undergraduate and beginning graduate level. This space is characterized by both practical requirements[1] and the need for pedagogically appropriate challenges.[2] Furthermore, there is a huge difference between learning methods for solving "non-problematic problems" by rote, and learning how to deal with the challenging, possibly open-ended, problems that require the most creative aspects of mathematical thinking.[3]

As a background to these interesting issues, there are certain important foundational facts. Most problems are impossible to solve without knowing what the terms in the problems mean. Knowing the definitions of terms isn't enough to solve many problems: one might have to hop out 2 or 3 more links through hypertextual chains of relationships to get the relevant material. The corresponding "search space" could become large if search was done in a blind way.

Accordingly, it is often useful to supplement mechanical paper-chasing with well-thought-out questions addressed to an appropriate expert. Another tack is to look for relationships, analogies, or indeed any intuitive idea that offers to simplify the problem solving situation.

Considering the "problem" of problem solving, we find that it is always going to be underspecified: it concerns real humans and real-world applications. The challenge taken up in this project is to bridge these interesting but informal

[1] "*Textbooks are much more explicit in enunciating the laws of mathematics or of nature than in saying anything about when these laws may be useful in solving problems.*" [3], cited in [4].

[2] "*If students are learning, they should be making fewer errors over time. However, the activities given to the students over time should also increase in difficulty. In a well constructed curriculum, these two forces should cancel each other out, leading to a fairly constant error rate over time.*" [5]

[3] "*One consequence of experiencing the curriculum in bite-size pieces is that students learn that answers and methods to problems will be provided to them; the students are not expected to figure out the methods by themselves.*" [6]

aspects of problem solving into a system of quasi-formal Definitions, Theorems, Proofs and Expositions. When compared with the classical investigations of the subject of problem solving [7], we have quite a number of assets, including large collections of encyclopedia-style mathematics articles in the quasi-formal style just mentioned, copious quantities of additional data in the form of current research articles, and always-on systems for discussing and adding to these corpora.

2 The Research Question: how do people solve problems?

What are the activities involved in mathematical problem solving, and how do these activities relate to mathematical knowledge?

Some relevant activities are likely to be: *connecting*, *discussing*, *working*, *recording*, *sharing*, and *learning*. Unlike possible "deeper" categories, like intuition – and unlike the three-step, four-step, and five-step frameworks of Bruner, Pólya, and of Kilpatrick *et al.*, respectively, (see Section 4.1.1) – these terms can be explicitly mapped to actions that take place within a learning environment.

These six dimensions will receive further discussion throughout Section 4.1. Of course, they are not yet the definitive answer to our question, since they have not been given a precise interpretation, nor tested, nor refined, but they do give a sense of what an answer might look like in its initial stages. The next section describes, in overview, a plan for operationalizing the investigation. Section 5.6 focuses on some of the critical questions that will be useful in analyzing the results of this investigation.

3 Sketch of a plan: problems over a mathematics encyclopedia

The basic plan is to build a problem-solving layer over the encyclopedia layer that makes up the current PlanetMath. Research can then proceed by examining the activities of people in this space. In general it will be interesting to look at the ways activity in the space connects with encodings of knowledge in the base layer.

Adding facilities to PlanetMath for *submitting*, *solving*, *hyperlinking*, *marking*, *discussing*, and *keeping track of* problems would be a satisfying beginning. This approach becomes scientifically interesting because it will allow us to gather new information about problem solving behavior that goes beyond these basics. Each of the activities above can linked e.g. to Activity Streams[4], and relevant user actions can be counted and correlated with one another, as well as with data garnered from textual analysis or participant observation. We can also vary the system itself to bring more or less detail into focus, i.e. to work in

[4] `http://activitystrea.ms/`

perspectives ranging from bulky observations (like number of problems solved versus number of articles written) to much more detailed, fine-grained, observations related to responses to "cognitive tutoring" or other recommendation systems.

3.1 A scenario

To be clear, the "main actor" would be a mathematics learner – and most likely a learner at the university level, since that's where most of PlanetMath's current content and user base comes from. Other people would surely be involved (e.g. teachers or course designers) but this wouldn't be the focus of the research project.[5] Using the upgraded PlanetMath system, our hypothetical learner would find it possible to use such a system to *share* problems of interest (either globally or with members of a specific group), to *link* them to relevant material in PlanetMath's encyclopedia and to other related problems, to *work* through them (again, either individually or as part of a group), to *ask* for help from other people on the site, to get or give *advice* about solution techniques, and to *record* what they have mastered or what they are having trouble with. At any given point in time, the learner can get a fairly clear picture of what they know and of what they are trying to *learn*.

3.2 Some analogies

This project will be what similar to a project that creates a radio drama and soundtrack layer over the collection of field recordings in The Freesound Project (`http://www.freesound.org/`), or a literary criticism layer over Project Gutenberg (`http://gutenberg.org`).[6] This extra layer comprises, among other things, a new social dimension in such spaces. We can ask how activities focusing on the "application layer" relate to activities focusing on "resource layer". These activities will in general be quite different: solving problems *vs* writing an encyclopedia; making and listening to radio plays *vs* creating re-usable audio samples; discussing literary texts *vs* scanning and proofreading literary texts. However, even when the activities are very different, they may correlate in thought-provoking ways (e.g. a much-discussed text also tends to be well-proofread).

We should also be on the look out for inhomogeneities within the application layer and user population, which point to different requirements in the resource layer. For example, one might guess that "connecting" is important for dealing with advanced problems, whereas "discussing" is more important to people

[5]Problems can come from many sources. For example, the 1912 book "Plane Trigonometry" by Ms. Sidney Loney (famously useful to the young Srinivasa Ramanujan) passed into the public domain in the UK in 2009. It has been reprinted many times, most recently in 2005. Clearly the problems it contains continue to be relevant. There are also several textbooks about more contemporary mathematical topics available under free licenses; for a partial list, see `http://www.opentextbook.org/category/maths/`. In addition, ArXiv now asks people uploading research papers to choose a license (`http://arxiv.org/help/license`), and many of these papers will contain material that contributors could use when creating new problems.

[6]Cf. `http://openshakespeare.org/`, `https://github.com/nickstenning/annotator`.

working on basic problems. Such differences would point to different knowledge needs, indeed, the need for different ways of knowing, and this would inform future design and development work.

4 Literature review

4.1 General background

***We develop an overview of some of the factors that bear on problem solving, and discuss what it would mean to add support for problem-solving in the PlanetMath context. So far, the factors we've identified include* sharing, working, discussing, connecting, learning, *and* recording.**

Without seeking to (re)define these terms, nor yet to say precisely how they will be modeled, we discuss the ideas that inform the way we think about these terms, in Sections 4.1.1–4.1.6, respectively.

4.1.1 Wisdom and strategy

Jerome Bruner's book "Toward a theory of instruction" [9] is very succinctly summarized as follows:

> "In his research on the development of children (1966), Bruner proposed three modes of representation: enactive representation (action-based), iconic representation (image-based), and symbolic representation (language-based)."[7]

One implication of the above is that intuition may often come from "tasks" involving simulations, games, or puzzles – certainly not always from "problems" in the usual sense, or from mathematical knowledge in the quasi-formal sense. We certainly intend to be open to contribution of open-ended tasks, but they will not be focus of this study.

Pólya's "How to solve it" [8] considers the case of self-guided learners, but focuses its attention on classroom instructors who will guide learners through the well-known four-step process of understanding the problem, devising a plan, carrying out the plan, and looking back over the solution. (Some further social history surrounding Pólya's work on problem solving is presented in a rather celebratory spirit in [11], but the same paper notes criticism from Alan H. Schoenfeld describing the work as "scientifically problematic" and "epiphenomenal rather than real".) For the moment, it suffices to say that in the PlanetMath context, the instructional function would in general be embedded in contributed learning materials, and/or distributed among various peer mentors. Thus, again, we intend to be open to various instructional offerings without making them the focal point of the study.

[7] `http://en.wikipedia.org/w/index.php?title=Jerome_Bruner&oldid=394387000`

According to Kilpatrick et al. [10] (also discussed with regard to testing in [12]), mathematical proficiency has the following components:

- conceptual understanding: the ability to understand mathematical concepts, operations and relations,
- procedural fluency: the ability to use procedures appropriately, accurately and flexibly,
- strategic competence: the ability to represent, formulate and solve problems in mathematics,
- adaptive reasoning: the capacity for logical thought, explanation and justification,
- productive disposition: the perception that mathematics is a sensible, useful, and effective tool.

These dimensions, particularly the last one, point to another important contextualizing feature, namely "application". Connections to real-world or scientific applications are again something that we intend to be open to, without focusing on.

In short, the key for this study will be people sharing mathematical problems and their own ideas about how to solve them. Our initial choice of activities to focus on is inspired by the earlier discussions reviewed in this section. These discussions will continue to be an important part of the background as we think about how problem solving works in practice: we aspire to add the next level of detail and precision.

4.1.2 Information processing theories of problem solving

John R. Anderson has led the development of a computational-psychological ("information processing") model, ACT-R theory [16] which has been applied in mathematics education contexts (cf. [5], [13]). In brief, ACT-R models the cognitive behavior of problem solvers (both effective and non-effective), and recommends the patterns of effective problem solvers to those who run into trouble. Some of the relevant contemporary applications of ACT-R appear to lie in the domain of natural language understanding [14], [15], pursuant to developing a "synthetic teammate". This project is especially useful because it comes with working source code.[8] ACT-R is not the only cognitive tutor available: see Section 4.3.2 for further discussion.

4.1.3 Sociological factors and peer support

At present we plan to leave some of the most commonly examined sociological factors (e.g. cultural background or economic status) as hidden variables in the

[8] `http://sourceforge.net/projects/synavo/develop`

study. However, we may see wide variability in the user population, e.g. when considering the way different learners approach social interaction.

We are less interested in the fact that in some cultures, "socialization to work collaboratively starts at an early age" [17], than in the possible role of peer support and encouragement in learning [18]. Socializing in forums and by direct messaging is possible in PlanetMath at present, and a key facet of the study is to look in detail at the data on these interactions.

4.1.4 Coherence

Correctly answering a mathematical problem, requires mastery of the underlying concepts, and this, in turn, comes from a coherent understanding of related facts. Thus coherent patterns of hypertextual reference (e.g. starting with undefined terms, but moving into related material and strategies) will generally help learners [19]. Encouraging third party curatorial activities that maintain coherence, and keeping track of indicators of coherence or incoherence that come from people working on solving problems will be of considerable importance in the study. There is ample room for discussion about what "coherence" means, though some simple guidelines have been found widely useful [20].

4.1.5 Motivation and economic aspects

In general we can assume that people will participate in PlanetMath because of the direct benefits of doing so. In this study we are interested in forms of participation that result in the specific benefit of learning.

Why might people go to PlanetMath to learn? The motivations may be similar to the motivations that lie behind "Science 2.0" [21]. Some of the same sorts of benefits that can accrue to corporations or researchers who adopt peer-production behaviors can presumably apply to students who collaborate well.

Just as in the world of free software, there is ample room to consider mixed production models for "free mathematics" and the "Learning 2.0" activities we will be examining in this study. In other words, money can change hands within a collaborative environment: for example, it would be possible to build an infrastructure to support paid tutoring in the PlanetMath context.

However, all else equal, the existence of payments shouldn't change the way learners interact with knowledge very much. In fact, there is even some evidence that artificially incentivizing "helping" can change the way peer tutors interact *with one another* in ways that are of a strictly negative value in a peer tutoring context [22].

Thus, at least at the outset, it seems suitable to give recognition for both learning and helping, but not to go further with an incentive scheme to support either at the present time.

This seems a suitable place to note that the benefits of engaging in peer tutoring in mathematics is the subject of current research, and that the overall benefits are suspected to be highest when the peers are not too far separated in mathematical maturity.[9]

4.1.6 Personalization – what about "me"?

As indicated in Section 4.1.2, one of the key issues will be to keep track of the appropriate range of activities related to the problem solving space, in order to turn it into a useful research environment. There are various protocols for tracking activity and attention; one light-weight protocol currently in use on sites ranging from Github to Facebook is known as Activity Streams [23], and as a representation format it seems to offer everything we might ask for (at least at present). As for selecting what activities to keep track of (and, more fundamentally, what activities to support), we have now established a list of candidate activities (Sections 3,3.1) and we will endeavor to create opportunities to mine for more using standard techniques [24].

Combining these data with data about user customizations (e.g. input method for mathematical text) may lead to various useful adaptive personalizations of the resource.

4.2 A brief critique of related work outside the literature

There are quite a few websites (including some of growing popularity) that are devoted in whole or in part to support for mathematical problem solving. However, these sites are not always subject to the research-level scrutiny, and those which have received a degree scrutiny do not show very convincing results (in contrast, for example, with the documented successes obtained by Carnegie Learning[10]).

Kahn Academy[11,12] and Alcumus[13] are two relatively recent websites that offer interactive problems and keep track of user success. The second of these is built to correspond with a series of grade-school level texts, while the former is being built in concert with online video tutorials in elementary mathematics through Calculus, elementary Differential Equations, Linear Algebra, and similar. Correspondences between the videos and the problems is not made clear at present, but presumably can be expected to arrive soon since all of this content is being developed in-house. The Kahn Academy's interactive problem collection currently goes, step by step, from Addition through to the Chain and Product rules.

There are various complete mathematics curricula that feature an integrated online component or that are meant to be worked with entirely online, e.g. Time4Learning.com.[14] These services generally address home-schooled children or otherwise replace in-person private tutoring. Third party evaluation of one

[9] `http://www.york.ac.uk/iee/research/t_peer_learning_paired_maths.htm`
[10] `http://www.carnegielearning.com/research/references/`
[11] `http://www.khanacademy.org/exercisedashboard`
[12] `http://bjk5.com/post/1664635835/constellation-knowledge`
[13] `http://www.artofproblemsolving.com/Alcumus/Introduction.php`
[14] `http://www.time4learning.com/math.htm`

such program, First in Math, noted "small but significant effects" on a standardized test.[15]

There are other websites that offer occasional help with mathematical problems, including `http://mathproblems.info/`, `http://www.webmath.com/`, and `http://calc101.com/`.

None of these websites appear to offer much support for university-level mathematics (little beyond basic calculus as yet). There *is* support for Calculus I, II, and III, as well as Linear Algebra, Number Theory, and Abstract Algebra in the "Calculus on the Web" (COW) project.[16,17] While COW provides a very nice collection of interactive problems, it does not seem to be set up as a research tool (in any case, there doesn't seem to be any research published about it). From the point of view of the current research proposal, it has the following drawbacks: there is currently no crowdsourcing component or open source release; and indeed, the project seems firmly Web 1.0, in that it has no significant opening for social interaction whatsoever. Similar, though less extreme, remarks would seem to apply to authoring content for the ActiveMath system, which is difficult but at least possible [25].

Indeed, support for social interaction and end-user contribution does not appear to be a strong point for any the projects mentioned in this section. This suggests that these systems are not particularly good at supporting, nor can they be usefully used to study, any connecting, discussing, or sharing activities that can be relevant to problem solving. Furthermore, they are not, at present, particularly useful environments in which to look for as-yet-unidentified factors in problem solving.

4.3 Review of contemporary methods

Two contemporary research programmes seem especially relevant the project: the "ecological approach" to peer collaboration championed by Gord McCalla [26], and cognitive tutoring methods applied in a social context. Taken together these approaches give a succinct but coherent picture of the uses of social interactions in computer-assisted learning.

4.3.1 Peer-to-peer intelligent tutoring

In [30], the authors discuss a novel and apparently quite general approach to sequencing educational content, namely a sort of collaborative filtering based on knowledge acquisition instead of ratings. Using this technique, a student is presented with content that has provided the greatest advantage to students in a "similar" situation, i.e. with respect to their learning profiles, and relative to metadata about the system's learning objects. The system's metadata is updated in a dynamic way as learners interact in the system in order to reflect

[15] `http://www.firstinmath.com/pdfs/FIM_WestEDstudy.pdf`
[16] `http://www.math.temple.edu/~cow/`
[17] `http://www.math.temple.edu/~gmendoza/psPapers/cow.pdf`

the results of these interactions. We note that a similar filtering process is part of a recent EU funded project, "étoile cascades ideas" [33], which aims to deliver education in a scalable way.

We note that the work of Kurhila and Miettinen on "the role of the learning platform in learning" ([27], [28]) is related, though this work focuses more on user-facing features of the platform than on algorithms to exploit user contributions.

The tradition appears to go back to [34], where the authors advocate for "restructuring schools as knowledge-building communities". Although it is not a school, PlanetMath does seem to already be a knowledge building community. Nevertheless, there remains room for considerable expansion of the range of meaningful end-user participation at PlanetMath (notably including the notion of opening up discussion around problems and solutions).

4.3.2 Cognitive tutors for collaborative learning

While there is no shortage of references associated with cognitive tutors in general, especially the ones developed at Carnegie Learning[18], the code for the Carnegie Learning systems isn't publicly available. Erin Walker's work [32], [31] describes experiments adapting Carnegie Learning's systems to offer support for peer tutoring. In particular, instead of coaching a student through mathematics problems, the system she studied coached students in tutoring other students. The essential outcomes of these studies were neither positive nor negative impact on learning outcome (in comparison to students who used the standard, non-social, algebra tutoring software), and the authors suggest that "further research on how to optimize collaboration support for particular interaction conditions may be necessary" [32]. The commentary on this article by Elisabeth Paus and Ina Jucks emphasizes the notion that fully computer-mediated communication in peer tutoring will make assessment easier (which we agree with). They also suggest that multiparty interaction should be considered in the future, with the computer serving as an "additional communication partner".

A related study with the Andes tutor [36] showed that peer interaction can keep learners from making deep errors. This is an encouraging result which suggests that relatively unstructured peer interactions may be a useful complement to the structured information presented in intelligent tutoring systems (e.g. useful because ITS's don't yet have full natural language interaction modes, but peer learners do). This paper is used an interesting methodology, the "microgenetic" approach to measuring learning "as it occurs" [37]. This approach, which is succinctly reviewed in [38], seems particularly useful for a distributed learning context like PlanetMath, where it would be infeasible to administer before-and-after tests in a routine way.

[18] http://www.carnegielearning.com/research/references/

5 The Research Design: a quarterly plan

We propose three core features: "building blocks", "algorithms for recommendations" and "support for end-user participation". We plan to do the study with two development phases and one assessment phase. Phase 1 is about building a problem solving environment. Phase 2 adds a cognitive tutor. Phase 3 focuses on analysis and narrative. These phases will be complemented throughout by a "participant observer" process designed to capture qualitative feedback.

5.1 Overview of implementation goals

These are the main features planned for the project:

Building blocks Build a platform that tracks actions, and that can flag up interesting bits of data using (hyper-)textual analysis and data mining of actions.

Algorithms for making recommendations Use data about previous experiences (not just atomic encounters, but learning pathways) to create recommendations (not just content, but also relevant activities and *strategies*). I.e., generate heuristics.

Support for end-user participation Do all of this in a way that is transparent and that encourages end-user participation in the process!

We will make two passes through this list, once without and once with cognitive tutoring support. Thus, for example, the recommendation of similar articles in Section 5.4.3 will not be personalized, but recommendation of learning pathways in Section 5.5.3 will be.

5.2 Participant observer cycles at P2PU

Mathematics courses are in demand at the Peer-2-Peer university, where the author has already ran one pilot course[19] and posted outlines of several new short courses`http://en.wikiversity.org/wiki/User:Arided/MathCourses`.

P2PU courses run once a quarter, and are populated by self-selectedly self-motivated learners. Thus P2PU seems an ideal setting in which to get feedback on the software we plan to develop. The value proposition for students is that this can help structure emerge in the by-default unstructured peer learning experience.

As a mater of risk management: P2PU learners are likely to have a somewhat different set of motivations from those held by "average" PlanetMath users, which should help us explore a wider range of use conditions. Additionally, any software components that may be tricky to scale up can be tried first with this smaller population.

[19] `http://p2pu.org/general/diy-math`

Perhaps most usefully, participant observation in this context will provide a qualitative dimension that will complement the data analysis planned for PlanetMath as a whole (Section 5.6). To this end, the set of questions outlined in Section 5.4.1 should be reiterated each quarter with a group of P2PU participants. A version of these a questions will be adapted into a questionnaire to be completed by participants at the end of each cycle.

5.3 A note on priorities and timing

Perhaps it goes without saying that Phase 2 builds on Phase 1, in as much as we need a problem solving space and various activities in this space in order to model what works for learners. In addition to being logically primary, Phase 1 is also likely to be somewhat simpler than Phase 2 from a design and implementation point of view. This will allow some work towards the Phase 2 milestones to take place during Phase 1: in general the dates associated with the milestones described in the outline which follows are dates for delivery. That is to say, development efforts will be "frontloaded" to the extent possible.

Each quarter corresponds to a set of questions, and each can be considered to be a small focused research project. In general the only thing that we require is usage data, which we have reason to believe will be readily available (Section 5.7).

Both positive and negative outcomes about the effects of cognitive tutoring are interesting (as we saw in Section 4.3.2). That said, the project should be able to address the questions from Section 5.5.3 in particular, even if the cognitive tutoring techniques do not prove completely successful. To this end, these questions will also be "frontloaded" into a post-mortem analysis for each quarter.

5.4 Phase 1: the problem-solving environment

The first phase will build on a collaborative project to rewrite and expand PlanetMath's software that is already quite well along[20], and that is already being used in classes at Jacobs University, Bremen [39].

5.4.1 Q1 2011: A research-ready version of PlanetMath

PlanetMath already exists, but we're updating the software to make it look better and make it extensible; in particular, setting it up so that people can upload "problems" and "solutions" – just blocks of text at this stage.

Development goals:

- Finish pluginifying PlanetMath.
- Incorporate Activity Streams.

[20] `http://trac.mathweb.org/planetary`

- Add problem and solution objects.

Questions:

- Do people use it?
- Do they like it?
- Do they suggest any improvements?
- Compare P2PU users cohort with the PlanetMath user population as a whole.
- Count links versus solutions (are problems that are well-connected to the encyclopedia solved more?)

5.4.2 Q2 2011: Support for sequencing and fragmenting problems

Add support for building "problem sets" and taking problems and solutions apart into steps. This will be useful for looking at dependencies between individual steps and specific pieces of background or procedural knowledge (Section 5.4.3). This in turn will form the locus for cognitive tutoring style recommendations in Phase 2 (Section 5.5.1), working along the now-classic lines of [40][21]

Development goals:

- Support for structured problem sets.
- Support for presenting a solution as a sequence of steps.

Questions:

- Do people use these features?
- Do people preferentially use these features?

If people do not care to use these features, either we will have to work with less detailed connections between problems and background material, or do more intensive linguistic analysis to take problems and solutions apart into steps automatically in the next phase.

5.4.3 Q3 2011: Textual and hypertextual analysis

Textual and hypertextual analysis to find "similar" documents (useful e.g. for retrieving the relevant background information associated with a problem set). The Concept Forest [41] technique is attractive because it can do similarity analysis *quickly* on the fly.

Development goals:

[21]"*The basic idea in learning by example is to induce the production rules used by the expert who generated the example. Each pair of lines, or states, in the example leads to thelearning of one production rule, with some part of the input (the first line) as a condition and the operation performed on the first line the action.*"

- Find similar documents using "Concept Forest" analysis.
- Extend the WordNet-based analysis with terms from PlanetMath's existing "thesaurus" to enhance the Concept Forest approach.
- Use these techniques to make background information available on demand.
- Include ways to indicate whether a given link or background item is helpful or unhelpful, and reweight the algorithm based on this feedback.

Questions:

- Can we identify patterns and dependencies between problems and/or steps so as to get a sense of when people are progressing to solve "harder" problems?
- Do these features impact solvability of problems?
- Can we say convincingly say that fewer bogus links are generated as time goes by?

5.5 Phase 2: intelligent tutoring in the large

The second phase will build on available open source cognitive tutoring software, Andes, and/or the ACT-R based Synavo.

Although there is no example of a cognitive tutor hooked into a large-scale social environment, there are several existing cognitive tutoring systems that could be adapted to this purpose. Note in addition to the (open source) ACT-R system[22] and the (publicly available for research purposes, but not open source) CTAT[23], the is the Andes system (see [35]), which is both open source and under active development at present.[24] Synavo seems to be the most complete working example of an open source ACT-R based system.[25]

5.5.1 Q4 2011: Support for hints

People should be able to ask for and receive hints, in the presence of a cognitive tutor. E.g. I could work on a problem up to a given point and then say "I'm stuck", or I could look at a problem and say "I have no idea what sort of technique to use here, I'm stuck". At this point, help could come from another person or possibly from the cognitive tutor, on a limited set of problems.

One shouldn't assume that the best result is to walk the student through the problem in question: another perfectly good option is to find some simpler related problems and make sure the student understands those.

[22] `http://act-r.psy.cmu.edu/actr6/`
[23] `http://ctat.pact.cs.cmu.edu/index.php?id=download`
[24] `https://github.com/bvds/andes`
[25] Cf. Footnote 8

I like the idea of improving the tutor based on the effectiveness of previous help. It seems simple enough to add a little feedback button (was this suggestion helpful? yes/no) attached to every message from the system. Certainly trickier to "learn" from that feedback, but no doubt doable.

Development goals:

- Add a feature for saying "I need help here".
- Add a cognitive tutor that listens for these events and responses to them, and that can itself suggest. either (a) tips for working with this problem itself; (b) places to look for easier background material.

Questions:

- Do people who use the cognitive tutor exhibit changes in their problem solving rate or problem solving ability?
- Do they see improvements in mastery of the material?

5.5.2 Q1 2012: A platform for open analytics

A way for people to easily create little widgets that will display things like "give me all of the problems that my friends have solved" or "give me all of the problems in algebra that no one has looked at for 6 weeks or more" or "give me all the problems similar to this one that have been successfully solved", etc., with the idea that some of these queries could be used by the tutor if the student doesn't think to use them.

Development goals:

- Create a way to make and share semantic queries.
- Hook the collection of queries into the tutor so that it can use them as a subroutine of hint generation.

Questions:

- Does this improve the efficacy of the tutor by the metrics mentioned in the previous section?

5.5.3 Q2 2012: Support for learning pathways

What sequence of steps has been effective for other learners whose background is something like mine? This question could be answered either applied at the macro level, recommending a sequence of problems to look at, or at the micro level, recommending steps to take when examining a given problem, or some mix, as in "You could attempt this series of problems, but be prepared to ask for help."

This might include a way for people to "look inside" the cognitive tutor and tweak its behavior based on what they find helpful. In any case, obviously as part of the design it will be important for me to be able to make tweaks as the investigator, and this could be done in an open way (e.g. as a project on github).

Development goals:

- Suggest and evaluate possible learning pathways based on learner profiles.
- Share information about how these suggestions came into being and allow users to tweak the suggestions at the outset or after trying them.

Questions:

- How are sequences identified as "successful"?
- When is success of a learning pathway transferable between learners?

5.6 Phase 3: assessment, transition, and writeup

What factors can we analyze? How do these various factors correlate with indicators of learning, like "number of problems solved"? How can we "narrativize" those key factors in a set of new problem-solving heuristics and teaching/learning strategies?

5.6.1 Q3 2012: Critical analysis

We now have a complete working system. Can we follow the methods of John Champaign [30] to measure total learning among subpopulations, and make a clear statement about the heuristics they appear to use? Can we follow the methods of Erin Walker *et al.* [32] and make some definitive statement about the effects of the social aspects of our tutoring system on a given population?

Data analysis goals:

- Assess the factors to use in a statistical analysis.
- Correlate these factors with various indicators of learning.

Questions:

- What are the most appropriate indicators of learning to use? (This could be as simple as saying "those people who use the cognitive tutor appear to solve more problems and achieve better mastery of the subject than those who don't.") In the PlanetMath context, following the microgenetic approach, which focuses on the four-plus-one dimensions of "path", "rate", "breadth", "source", and "variability" (cf. [38]).

As remarked in [12], "An important key factor for successful proficiency testing is to assess the testing method: which abilities will be tested, in which way the test will be performed, what the tested skills are etc. Performing the proficiency test with the web-based survey lays restrictions on the assessment of results. For example, how to assess student's ability to formulate and solve problems." Treating solutions as a sequence of steps (Section 5.4.2) should give us adequate material to make relatively rich assessments, but we will need to further clarify the precise criteria to use in these assessments.

5.6.2 Q4 2012: Distilling new strategies

Having made six cycles of development, testing, and analysis and one more cycle of statistical analysis, what remains is to put all of these things into one coherent story. It is to be expected that some of our efforts will have been more useful than others, and this cycle is about capturing this information in a useful narrative: what works under what conditions?

Writing goal:

- "Narrativize" the most important factors in a set of strategies for problem-solving and learning/instruction.

Questions:

- Have we been able to identify *new* heuristics for problem solving in this study?

5.7 Overall risk management

We should note that the theory of commons-based peer production [42] says that, in order to work:

1. the potential objects of peer production must be modular;
2. the modules must be small in size (noting that heterogeneous granularity will allow people with different levels of motivation to collaborate by contributing smaller or larger grained contributions);
3. the integration mechanism must run at a fairly low cost (either through automation or enforced social norms).

The first two conditions are easily met by the efforts in Phase 1 towards sequencing and fragmenting solutions (Section 5.4.2). Assistance from a cognitive tutor in the process of identifying useful learning pathways should especially help with otherwise costly integration (Section 5.5.3). That said, PlanetMath has a track record of strictly enforced social norms surrounding integration[26], and the introduction of new tools that make it easy to analyze contributed content will be a further help (Section 5.5.2).

Although these are likely to be necessary conditions for success, an even more critical condition is that people should be interested in the system. One reason for working with P2PU is that both organizers and students have expressed

[26] http://wiki.planetmath.org/AsteroidMeta/one_week_in_october

significant interest in mathematics courses. As P2PU continues to grow in popularity, there should be a continual flux of people interested in learning standard mathematical subjects available to test out the system (i.e. in addition to the many PlanetMath visitors, some of whom are of course likely to take an interest in the system; some may even become interested in using it in their own courses).

6 Conclusion and Future Work

We have described a research project that can form the basis of an open platform for testing and evaluating various learning and instruction methods.

In [13], the progenitors of ACT-R and Cognitive Tutoring call for the creation of something like a "Federal Education Administration", with reference to the "Food and Drug Administration", the agency that approves or disapproves foodstuffs and medicines for sale or distribution in the US. From an international perspective, it would likely be preferable to have a non-governmental, non-partisan organization involved in running an open platform for testing and evaluating various learning and instruction methods.

The work we propose will take several concrete steps in that direction. The core of the project is to critically examine methods associated with the "ecological approach" to computer-mediated learning, and the cognitive tutoring approach to recommendations, as they apply to problem solving in an open online mathematics community.

The same infrastructure can be reused in further studies of factors we have chosen to leave out at present (such as cultural or sociological factors). The results of our analysis should also be useful for further improvements and extensions to the basic system, and, indeed, we can reasonably hope that this platform will be a site of ongoing improvements to the infrastructure for learning mathematics. As indicated in the paragraphs above, the same basic infrastructure (Section 5.4.1) can be used to test future innovative educational strategies that can meet the fairly minimal requirement of plugging into the Planetary system.

For example, although we have targeted students in this project, future rounds of work could target teachers and researchers in a similar fashion. Further, as we mentioned in Section 3.2, there are some abstract parallels between the work proposed in this project and similar projects that could be undertaken in a variety of Web 2.0 contexts. Although the methods used here will not carry across all such parallels, certainly they can directly bear on similar applications in other STEM disciplines, and may well be of interest to anyone working at the nexus of online education and online communities.

In sum, the major contribution of the research project will be new methods for exploiting large quantities of hypertextual and social interaction data associated with problem solving in a knowledge-rich peer-produced online environment to extract effective learning and problem solving strategies in mathematics.

It is the thesis of this book that society can only be understood through a study of the messages and the communication facilities which belong to it; and that in the future development of these messages and communication facilities, messages between man and machines, between machines and man, and between machine and machine, are destined to play an ever-increasing part. – The Human Use of Human Beings, Norbert Wiener.

References

[1] Newell, A. (1980). Reasoning, problem solving and decision processes: The problem space as a fundamental category. In R. Nickerson (Eds.), Attention and performance VIII. (pp. 693-718). Hillsdale, NJ: Lawrence Erlbaum. http://shelf1.library.cmu.edu/cgi-bin/tiff2pdf/newell/box00042/fld03567/bdl0003/doc0001/newell.pdf

[2] Halmos, P. (1980). The heart of mathematics. American Mathematical Monthly, 87, 519-524.

[3] Simon, H.A. (1980). Problem solving and education, in D.T. Tuma and R. Reif (Eds.), Problem Solving and Education: Issues in Teaching and Research (pp. 81–96). Hillsdale, NJ: Erlbaum.

[4] John D. Bransford, Ann L. Brown, and Rodney R. Cocking, (Eds.) (2004). How People Learn: Brain, Mind, Experience, and School, Washington, D.C.: National Academy Press

[5] Ritter, S., Anderson, J. R., Koedinger, K. R., & Corbett, A. (2007). Cognitive Tutor: Applied research in mathematics education. Psychonomic Bulletin & Review, 14, 249-255.

[6] Schoenfeld, A. H. (1992). Learning to think mathematically: Problem solving, metacognition, and sense-making in mathematics. In D. Grouws (Ed.), Handbook for Research on Mathematics Teaching and Learning (pp. 334-370). New York: MacMillan.

[7] Newell, A., & Simon, H. A. (1972). Human problem solving. Englewood Cliffs, NJ: Prentice-Hall.

[8] George Pólya (1945). How to Solve It, Princeton University Press

[9] Bruner, J. S. (1966). Toward a theory of instruction. Cambridge Mass: Harvard University Press.

[10] Kilpatrick, J., Swafford, J., & Findell, B. (Eds.), (2001). Adding it up: Helping children learn mathematics. Washington, DC: National Academy Press.

[11] Frank, Tibor (2001). George Pólya and the Heuristic Tradition: Fascination with Genius in Central Europe. The Polanyiana, 6(2). http://www.kfki.hu/~cheminfo/polanyi/9702/frank.html

[12] Hanna Kinnari (2010). A study of mathematics proficiency. In Ion Mierlus-Mazilu, ed., 1st International Workshop on Mathematics and ICT: Education, Research and Applications, Bucharest, Romania. http://buletinstiintific.utcb.ro/bs/arhiva2010/modelling_nr3_2010_special_issue.pdf

[13] Anderson, J.R., Reder, L.M., & Simon, H.A. (2000). Applications and Misapplications of Cognitive Psychology to Mathematics Education. Texas Educational Review (Summer).

[14] Ball, J. (2010). Context Accommodation in Human Language Processing. Proceedings of the Natural Language Processing and Cognitive Science Workshop. Lisbon: INSTICC Press. (http://www.doublertheory.com/ContextAccommodation.pdf)

[15] Ball, J., Freiman, M., Rodgers, S., & Myers, C. (2010). Toward a Functional Model of Human Language Processing. Poster presented at the 32nd Annual Conference of the Cognitive Science Society. Portland, OR. (http://www.doublertheory.com/HumanLanguageProcessing.pdf)

[16] J. R. Anderson, D. Bothell, M. D. Byrne, S. Douglass, C. Lebiere, and Y. Qin (2004). An integrated theory of the mind. Psychol Rev, 111(4), pp. 1036–1060. http://act-r.psy.cmu.edu/papers/403/IntegratedTheory.pdf

[17] Rose Asera (2001). Calculus and Community: A history of the emerging scholars program. The College Board http://professionals.collegeboard.com/profdownload/pdf/calcandcomm_3947.pdf

[18] Steinberg, L., Dornbusch, S., and Brown, B. (1992). "Ethnic Differences in Adolescent Achievement: An Ecological Perspective," American Psychologist, Vol. 47.

[19] Foltz, P.W. (1996). Comprehension, Coherence and Strategies in Hypertext and Linear Text, In Rouet, J.-F., Levonen, J.J., Dillon, A.P. & Spiro, R.J. (Eds.) Hypertext and Cognition. Hillsdale, NJ: Lawrence Erlbaum Associates. http://www-psych.nmsu.edu/~pfoltz/reprints/Ht-Cognition.html

[20] Stephen Neale (1992). Paul Grice and the Philosophy of Language, Linguistics and Philosophy 15, pp 509–559. http://ncs.ruhosting.nl/bart/talks/paris2010/Neale1992.pdf

[21] Don Tapscott and Anthony D. Williams (2006). Wikinomics: How Mass Collaboration Changes Everything. New York: Portfolio.

[22] James Garbarino (1975). The impact of anticipated reward upon cross-age tutoring, Journal of Personality and Social Psychology, 32(3), pp. 421–428.

[23] Mango Spring, Inc. (2010). White Paper: Why microblogging and activity streams will overtake email for enterprise communications and collaboration. http://securedocs.mangospring.com/mkt_docs/whitepapers/HowActitityStreamsWillReplaceEmail.pdf

[24] Matthijs den Besten, Alessandro Rossi, Loris Gaio, Max Loubser and Jean-Michel Dalle (2008). Mining for Practices in Community Collections: Finds From Simple Wikipedia, Open Source Development, Communities and

Quality, IFIP International Federation for Information Processing, Volume 275/2008, pp. 105–120.

[25] Paul Libbrecht and Christian Gross (2006). Authoring LeActiveMath Calculus Content, Mathematical Knowledge Management, MKM'06, LNAI http://www.activemath.org/pubs/Libbrecht-Gross-Experience-Report-Authoring-LeAM-MKM-2006.pdf

[26] McCalla, G. (2004). The Ecological Approach to the Design of E-Learning Environments: Purpose-based Capture and Use of Information About Learners. Journal of Interactive Media in Education, 7. Special Issue on the Educational Semantic Web.

[27] Jaakko Kurhila, Miikka Miettinen, Petri Nokelainen, Henry Tirri (2004). The Role of the Learning Platform in Student-Centered E-Learning, ICALT, Fourth IEEE International Conference on Advanced Learning Technologies (ICALT'04), pp.540–544.

[28] M. Miettinen, J. Kurhila, H. Tirri, On the Prospects of Intelligent Collaborative E-learning Systems (2005), In the Proceedings of the 12th International Conference on Artificial Intelligence in Education (pp. 483-490). IOS Press.

[29] John Champaign and Robin Cohen (2010). Peer-Based Intelligent Tutoring Systems: A Corpus-Oriented Approach (Young Researchers Track), In Proceedings of Intelligent Tutoring Systems (ITS), Pittsburgh, USA, June, 2010.

[30] John Champaign and Robin Cohen (2010). A Model for Content Sequencing in Intelligent Tutoring Systems Based on the Ecological Approach and Its Validation Through Simulated Students (Intelligent Tutoring Track), In Proceedings of FLAIRS Conference, Daytona Beech, Florida, May, 2010

[31] Harrer, A., McLaren, B. M., Walker, E., Bollen, L., and Sewall, J. (2006). Creating Cognitive Tutors for Collaborative Learning: Steps Toward Realization. User Modeling and User-Adapted Interaction, 16(3-4), pp. 175–209.

[32] Diziol, D., Walker, E., Rummel, N., & Koedinger, K. R. (2010). Using Intelligent Tutor Technology to Implement Adaptive Support for Student Collaboration. Educational Psychology Review, 22(1), pp. 89–102.

[33] P. Bourgine and J. Johnson. Etoile Project: Social Intelligent ICT-System for very large scale education in complex systems. Geophysical Research Abstracts, Vol. 11, EGU2009-13965-1, 2009 EGU General Assembly 2009 http://meetingorganizer.copernicus.org/EGU2009/EGU2009-13965-1.pdf

[34] Scardamalia, M. and Bereiter, C. (1994): Computer Support for Knowledge-Building Communities. The Journal of the Learning Sciences, 3(3), pp. 265–283.

[35] Kurt VanLehn, Collin Lynch, Kay Schulze, Joel A. Shapiro, Robert Shelby, Linwood Taylor, Don Treacy, Anders Weinstein, and Mary Wintersgill (2005). The Andes Physics Tutoring System: Lessons Learned. International Journal of Artificial Intelligence in Education, 15(3), pp.147–204 http://www.andestutor.org/Pages/AndesLessonsLearnedForWeb.pdf

[36] Robert G.M. Hausmann, Brett van de Sande, and Kurt VanLehn (2008). Trialog: How Peer Collaboration Helps Remediate Errors in an ITS. Proceedings of the Twenty-First International FLAIRS Conference http://www.learnlab.org/uploads/mypslc/publications/flairs08-099.pdf

[37] Siegler, R. S., & Crowley, K. (1991). The microgenetic method: A direct means for studying cognitive development. American Psychologist, 46, pp. 606–602 http://www.psy.cmu.edu/~siegler/sieglercrowley91.pdf

[38] Gerald J. Calais (2008). Microgenetic Analysis of Learning: Measuring Change as It Occurs. National Forum Of Applied Educational Research Journal 21(3). http://www.nationalforum.com/Electronic%20Journal%20Volumes/Calais,%20Gerald%20J%20Microgenetic%20Analysis%20of%20Learning.pdf

[39] Catalin David, Deyan Ginev, Michael Kohlhase and Joseph Corneli (2010). eMath 3.0: Building Blocks for a Social and Semantic Web for Online Mathematics & eLearning, In Ion Mierlus-Mazilu, ed., 1st International Workshop on Mathematics and ICT: Education, Research and Applications, Bucharest, Romania http://buletinstiintific.utcb.ro/bs/arhiva2010/modelling_nr3_2010_special_issue.pdf

[40] David M. Neves (1978). A Computer program that learns algebraic procedures by examining examples and by worknig test problems in a textbook. Proceedings of the Second National Conference of the Canadian Society (or Computational Studies ol Intelligence, pp. 191–195

[41] J. Z. Wang and W. Taylor (2007). Concept forest: A new ontology-assisted text document similarity measurement method. In WI '07: Proceedings of the IEEE/WIC/ACM International Conference on Web Intelligence, pp. 395–401, Washington, DC, USA, 2007. IEEE Computer Society

[42] Y. Benkler (2005): Common wisdom: Peer production of educational materials. Center for Open and Sustainable Learning at Utah State University. http://www.benkler.org/Common_Wisdom.pdf