Clinical Requirements Engineering

Stephen Fickas
Computer Science Department,
University of Oregon
Eugene, OR 97403
fickas@cs.uoregon.edu

ABSTRACT

In this paper, I make a case for integration of requirements engineering (RE) with clinical disciplines. To back my case, I look at two examples that employ a clinical RE approach, first, that of introducing email into the life of a brain-injured individual, and second, introducing digital darkroom tools into my life. The former uses a Brownfield approach by starting with an existing clinical process, cognitive rehabilitation, and then defining an RE process that fits. The latter uses a Greenfield approach that postulates a new clinical RE process that focuses on the problems some of us have using digital darkroom tools.

Categories and Subject Descriptors

D.2.1 **[Requirements/Specifications]:** *Elicitation methods (e.g., rapid prototyping, interviews, JAD), Languages*, *Methodologies (e.g., object-oriented, structured), Tools*

General Terms

Design, Human Factors

Keywords

Clinical Requirements Engineering, Goal Attainment Scale, Skill Assessment, Deferred Requirements/Goals, Requirements Monitoring

1. INTRODUCTION

I wish to convince the reader that the software engineering field might profit by looking to more "personal" disciplines for inspiration. Specifically, I am interested in clinical fields, where individuals are assessed, individual goals are acquired, each individual is given a tailored treatment package, the effectiveness of the deployed package is tracked for each individual and mid-course corrections can ensue. Currently, the packages delivered in clinical fields may have a software component, but this software is part of the clinical domain, targeted to professionals providing treatment support. In particular, the clinical software used has little or nothing to do with daily-living software *applications* that many of us take for granted, e.g., email, web browsers, digital photography tools, music-management tools, trip-planning tools. I am part of a multi-disciplinary team that is working towards making modern software applications available to the cognitively impaired population. This paper starts with a Brownfield approach, discussing the requirements engineering process that emerged from this multi-disciplinary, clinic-based effort. The case study I will use is that of making an email application available to the brain-injury population (a subset of the larger cognitively impaired population). I will also briefly mention some preliminary efforts to generalize our work to include those with a developmental disability (another subset of the cognitively impaired population), and in particular, children and young adults that are part of special education programs. The paper concludes with a more Greenfield discussion, considering how a clinic-based model of RE can be defined for new areas that might benefit from a personal approach. I will also discuss some of our attempts to automate pieces of the process. Section 2 begins the Brownfield discussion.

ICSE'05, May 15–21, 2005, St. Louis, Missouri, USA.

Mayo Interdisciplinary Program for Cognitive Rehabilitation (MIPCR)

People who enter the Mayo Interdisciplinary Program for Cognitive Rehabilitation may have problems with:

- Attention and concentration
- Memory
- Organization
- Problem solving
- Other cognitive skills that affect life management

Participants must be able to identify thinking problems and set treatment goals.

Before they can enter the program, candidates must go through an evaluation. A neuropsychologist reviews the candidate's medical history and referral. A neuropsychological evaluation may be conducted if one has not been completed recently.

Participants meet individually with an occupational therapist to set treatment goals and create a treatment plan. Additional services can come from a speech pathologist, recreational therapist, neuropsychologist, physical therapist, physiatrist and vocational and community resources.

These people and other professionals meet regularly to monitor the participant's progress. Family members also may take part in treatment sessions.

The length and number of treatment sessions is based on the participant's progress in meeting treatment goals.

http://www.mayoclinic.org/physicalmedicine-rst/mipcr.html

2. COGNITIVE REHABILITATION IS A CLINICAL DISCIPLINE

In the spring of 2000 I began a collaboration with a group of researchers who had an interest in assistive technology for the brain-injured population. These researchers had over 25 years total experience in working in the field of cognitive rehabilitation. One of my colleagues in the group, McKay Sohlberg, is co-author of a standard textbook in the field [11]. The problem the group wanted to tackle was the social isolation a person encounters after a brain injury. The proposal was to give the population access to email and community travel to help break this isolation. Before discussing this specific problem, I will provide some background on the brain-injury population and the discipline of cognitive rehabilitation.

The description of the Mayo Clinic's MIPCR program (see sidebar above) provides many of the high-level pieces of the cognitive rehabilitation field. In particular, it lists the general impairments that may be encountered with a brain-injured individual. It highlights the high-level, *individual* process of assessment, goal-setting, creation of treatment plan and then periodic monitoring. I would add to this the following background information.

- The term clinical centers on the *individual* diagnosis and treatment of *outpatients*. Cognitive rehabilitation fits this definition: it works with individuals and attempts to integrate them back into their communities.
- The size of the brain-injury population is roughly 6-10 million in the US [2].
- The impairments suffered can be broken into two broad classes, short-term memory and executive function. Impairments associated with executive function include issues with attention, problem solving and self monitoring.
- There is a range of severity on all impairments. The combination and severity of impairments varies so that it is not possible to build a "typical" profile of a member of the population.
- The life goals and interests of individuals vary widely, just as they do in the non brain-injured population.

What is missing from the MIPCR description is a reference to what is known as Assistive Technology (or AT). The use of AT will be found if one digs a little deeper into the activities of the cognitive rehabilitation staff listed in the MIPCR program. For instance, one or more of the speech pathologist, recreational therapist, neuropsychologist, physical therapist, or physiatrist may recommend some assistive device that can help an individual reach a goal. It is worth differentiating between non-computer AT and computer-based AT at this point. Non-computer AT has a successful track record, e.g., memory aids such as DayTimers are ubiquitous in the population and are viewed as highly effective. Computer-based AT devices have a less stellar record. Abandonment of such devices has been reported as high [6]. While there have been no clinical studies on this topic that I am aware of, informally professionals have pointed to two causes: (1) lack of initial fit of an AT device with an individual's needs, and (2) lack of device adaptation once an individual's needs change [6].

We are now ready to return to the problem: mitigating the social-isolation that comes with a brain injury. The interdisciplinary team that formed to tackle this problem included (a) computer scientists, (b) those doing research and practice in clinical cognitive rehabilitation, (c) those doing research into AT training, and (d) those focusing on qualitative research and field work. Two general solutions to the isolation problem were proposed: (1) giving an individual access to email, and (2) giving an individual access to social and entertainment venues within the community. I will focus on the email component in this paper: the team chose to look at email as a technology that could open lost channels of communication with family and friends, and potentially open new channels with some rudimentary pen pal services. The following plan of attack was created:

1. Conduct a set of focus groups to get a general view of the population's use of computers and their interest in using email.
2. Build a set of prototype email clients and conduct preliminary usability studies.
3. Put together an RE process that fit with the cognitive rehabilitation field. In essence, view email as a type of AT. Integrate it with current clinical practice.
4. Choose a set of brain-injured participants to work with. Because of the gradual change in the population, and because of the long-term abandonment issue, we choose a longitudinal approach, working with each participant for at least a year.

A summary of step 1, the use of focus groups, is that (a) participants were universally excited about the prospect of using email, and (b) highly pessimistic that they could use computers because of past failed attempts [8]. A summary of step 2, pilot usability studies, is that an email client with roughly three parameters of 15+ settings each could be made to work for the participants we studied [10]. This leaves us with the heart of the study, steps 3 and 4: defining a clinical process that delivers a software system (an email client) and measures its success over an interesting time frame.

3. THE DEFINITION OF A CLINICAL RE PROCESS FOR EMAIL

Our longitudinal participant group consists of nine brain-injured individuals. Three live in their own homes, four live in an assisted living facility with minimal supervision (they are free to come and go as they please), and two live in an assisted-living facility with high supervision. As I go through the steps in the RE process we have developed, I will take examples from our participant group to ground discussion.

3.1 A Goal Attainment Scale

The cognitive rehabilitation field uses what is known as a goal attainment scale to acquire the individual goals/desires of a person. Each goal is broken into a set of levels to provide a degree of attainment. For instance, a goal might be to be able to do shopping for meals, with this broken into degrees of satisfaction, e.g., can shop for all meals, can shop for special meals, etc.

Using this style, we attempted to capture the goals of an individual in terms of email use. We asked each participant to first list a goal and then five levels of attainment, from not attained to fully-attained. Several examples might help illustrate its use. (I will change the names of individuals for privacy.) First I will introduce

Mary. One of Mary's goals is to write letters to the editor of an online newsletter that runs articles on the brain-injury population. She divided this goal into the following categories:

Level 1 (not attained): will not be able to learn how to use email.

Level 2: will be able to write email opinions to friends.

Level 3: will be able to write an email opinion, that meets stylistic constraints, as a submission to the editor.

Level 4: will have a letter printed in the newsletter.

Level 5 (fully attained): will be invited to write a guest letter/opinion.

Another example is that of Don. One of his goals was to learn to email with no help. He divided this goal as follows:

Level 1 (not attained): will not be able to learn how to use email.

Level 2: can email, but only with lots of prompting and help.

Level 3: can email, with some prompting and help.

Level 4: can email with no prompting and help.

Level 5 (fully attained): can teach others how to email.

Each of the nine participants typically listed 4-5 goals. It is important to note the temporal nature of these goals, and in particular, the influence the division into attainment levels has on thinking of milestones. No participant expected to fully attain their goals from day one. Instead they were set as achievement targets. From a requirements engineering perspective, this led me to consider the notion of "deferred" requirements. I will return to this idea shortly.

3.2 An Assessment of Abilities

Members of the team felt strongly about the concept of functional assessment established in the cognitive rehabilitation field. The general idea is that you cannot get a true picture of an individual's abilities in using an AT device by giving them abstract tests. You need to test them using the real device in as realistic setting as possible. In our case, this translated into assessing their ability to write email. We developed an assessment tool built on top of our parameterized email client. Much like an optometrist will go through a series of differential tests to obtain a picture of someone's ocular ability, we developed an assessment process that "turned" the parameters of our email system to get the best fit with an individual's ability to click around an interface, read directions, write a composition, remember the task at hand, and stay focused. A snapshot of one setting of the interface is shown on the right above. This setting is on the power-user side of the scale, having minimal prompting and a split inbox that includes old and new mail. It also contains ten "email buddies" (all part of the research team), some of which can only be accessed by scrolling the inbox. None of our nine participants were able to effectively use this interface during assessment. However, two are currently using it now (after roughly 12 months of emailing).

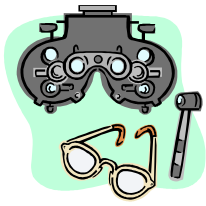

The outcome of assessment is an inventory of the abilities of an individual related to email. We developed an email-skill inventory that included more than 50 separate skills. Each was rated on a three-point scale of missing/sporadic/possesses [10].

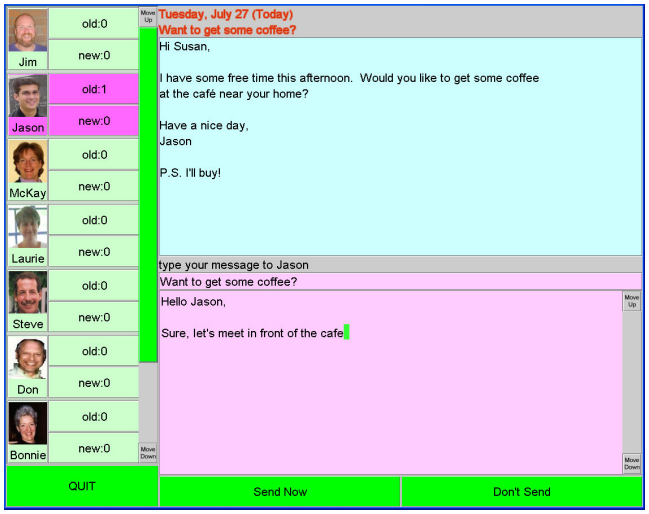

3.3 Providing a First System

At this point we are at an interesting juncture in the process. A decision is made on whether certain skills can be improved by training. If they can, then we may choose to delay delivery of a system until training has finished. This is based on the confusion-cost of changing a system in use. If we can improve skills through training, that may allow the individual to use a more powerful system. Of course, this brings up the question of "powerful". The majority of participants had no desire to become power users, Don from section 3.1 being an exception. The measure of the power of an email client is how well it supports an individual's goals. Perhaps an example will help. Manuel has stated (through his goals) that he is interested in emailing several times a week with a very small set of correspondents. He says he would like to write short notes to a few people – he likes to stay connected but is not a big writer. During the assessment, we find that he does not understand scrolling, i.e., he is missing this skill. We do have a training package for scrolling. Should we train Manuel in scrolling? Or should we skip the training and deliver a system to him immediately? His goals do not support the need for large buddy lists or long compositions; hence, we reason that he can start using email without the scrolling skill. Of course, we would like to do on-demand training if Manuel's goals change in the future, e.g., he does become more of a writer and composes long messages. (As a postscript, Manuel's goal has not changed over the 18 months he has been using the system. Ill-advised attempts by our team to encourage him to write longer messages and add new buddies have been firmly rebuffed. He meant what he said.) In summary, there is a three-way constraint satisfaction problem here: what skills are missing versus what skills are trainable versus what skills line up with goals.

The actual delivery of the system is a straightforward setting of parameters to instantiate a new email client. That is, we have developed an email client that allows us to add and remove features, and reconfigure the GUI, based on a set of XML specifications. Our parameter setting is actually a reconfiguration of the XML specification.

3.4 Monitoring and Adaptation

The clinic-based view of treatment in cognitive rehabilitation uses goal attainment scales, assessment, and then delivery of a treatment

package that is individualized. It also has a built-in step of continuous monitoring of progress, through regular follow-up sessions. In my view, this process has a direct tie with what I call goal-directed monitoring: use the individual's goals to decide what usage behavior is important to track. Using this idea, we developed a type of requirements monitoring, based on a mixture of the goal-attainment-scale (see section 3.1) and our knowledge of what abilities were prerequisites to the goals. We derived sub-goals from the different levels of each goal. We viewed these sub-goals as milestones (or stepping-stones) towards full attainment of a goal. An example might be useful here. Looking at Don's goal in section 3.1, he has a desire to be completely independent in his emailing, needing no help from the system (or humans). After completing an assessment (see section 3.3), it turns out Don will need a system with a full array of prompting and help. I view Don's sub-goals (stated as levels) as deferred requirements. We do not wish to abandon them. Instead, we will use them for monitoring. In particular, we will attach skill prerequisites to each sub-goal and monitor for those pre-requisites being met. When they are met (if ever), we can consider the adaptations we will need to make to Don's client to bring his deferred goal/requirement to the forefront. Getting back to Don's goal, once he shows proficiency with a portion of the emailing process, the actual adaptation taken will be to remove some piece of system scaffolding.

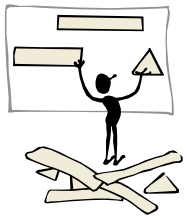

Mary illustrates somewhat the inverse of the adaptation done for Don: Don had something removed; Mary will have something added. In Mary's case, she wants to send opinion letters to the editor. We divided Mary's goal into sub-goals by using the attainment levels she provided. There are pre-requisites on each of these sub-goals, some tied to use of the email client and some tied to the constraints (placed by the newsletter editor) for accepted submissions. One example is that of length: a letter to the editor must not be longer than twenty lines. In monitoring Mary's correspondence with friends and family, she is initially observed to write very long emails, all well over 20 lines in length. However, with coaching from others, she is observed over time to write a wider variety of email, some short notes and some longer stories. The ability to write shorter size email triggers a prerequisite achievement, which is attached to her sub-goal of writing a legal submission. When all prerequisites are met, we can mark the sub-goal as active, moving it from the deferred state. This often triggers an adaptation to the system. In this case, the adaptation is to add the editor of the newsletter as a correspondent. This does not achieve her deferred goal of having one of her submissions accepted for publication, but does add another key prerequisite: the ability to submit online (the only form that is acceptable to the newsletter). She can now start submitting letters to the editor.

3.5 The Outcome

The process I describe in this section has been used successfully for all nine participants. While none has fully achieved all of their goals, no goals remain below level 3 for any individual. From the cognitive rehabilitation viewpoint, the project has been highly successful. As noted, virtually none of our focus group participants thought they would be able to use a computer, let alone be accomplished emailers. The clinicians now believe that email (as a tool to fight isolation) can find a home in the cognitive rehabilitation process [8].

From my point of view, I have moved the software development process to within the clinical framework that the team is most comfortable with. This took modifications. An individual assessment process had to be established. A means of capturing individual goals, and linking deferred goals to monitoring and adaptation had to be defined. As a manual process, this has been honed over several years to a routine methodology, allowing new participants to be brought into the project and moved smoothly through the steps. More details on the project can be found at www.think-and-link.org.

4. EDUCATION IS ANOTHER INTERESTING CLINICAL DISCIPLINE

I have used the word clinical up until now in the rather narrow sense of injury clinics, largely from a medical view. However, there are other fields that I would classify as clinical in the sense that they take a very personalized view of a problem. Special Education is a good example. There are well established (and legally mandated) procedures for working with special-needs kids in the K-12 setting. I was brought into this discipline by the suggestion that a usable email client was out of reach by many of the kids that are in special needs programs. While I do not have the experience here that I have with the brain-injury population, I have found several processes in place that fit with the personalized view of the problem. In the US, the most highly accepted and standardized process is the IEP: Individualized Education Plan. Currently the IEP focuses on the educational components of a child's school experience. I have begun to explore, with the director of a high school special needs program, the possibility of working internet communication tools into IEP content. We are interested in extending the IEP assessment to include the possible introduction of email and chat, both desktop and mobile versions, into a child's or teenager's life. My conjecture is that the process described in section 3 can be overlaid with the IEP process, extending assessment, goal acquisition and tracking. The challenge is another Brownfield one, finding the integration points within the well-entrenched IEP process.

5. A HYPOTHETICAL GREENFIELD EXAMPLE

To close out and summarize the previous sections, I have used a case study, delivering email technology to the brain-injured population, as an example where a clinical view of software development has success. I have also hinted that the same approach has potential to work with other clinical fields, e.g., special education, which identify personal goals and help individuals work towards them. Along these lines, our team has started a project to introduce a travel assistant (PDA-based) into the life of brain-injured individuals, allowing them to achieve their goals to make trips into the community – many are housebound at the moment [9]. We are attempting to use the process discussed in section 3 as a basis for this work. To summarize, there are many areas where segments of our population are currently unable to achieve their goals. In some of these areas, clinical practice has grown up to help. In some of these clinical practice areas, individuals' goals line up with delivery of software-based systems. If one believes that the RE process must conform to existing clinical practice (which I do), then the modifications I discuss in section 3 should be a guideline to what is needed in doing clinical RE.

At this point, I will admit to being selfish. I would like software on my computing devices that conforms to *my* personal goals. In essence, I would like quid pro quo: I have been spending effort working software engineering into existing clinical processes; I would like someone to invent a clinical process to help me with some of my goals that rely on software solutions. Let me ground this with an example. I am an amateur photographer. I have a goal of showing some of my work. I am reasonably proficient in the darkroom, and have submitted some of my photographs to local art shows. However, I have found that transferring my physical darkroom skills to a digital darkroom has been a challenge. In essence, I am Photoshop impaired. As far as I know, there is no clinical practice that has grown up around this impairment. But perhaps we can consider what it would require to define one. Using section 3 as a guide, we will need a means of capturing the various goals of darkroomers. We will need to establish an abilities-assessment process that can produce a skills inventory. We will need to define an adaptable form of a digital darkroom tool (perhaps by overlaying an existing tool with an adaptation framework) that supports change as needed. We will need to build a monitoring process that can track usage and link this to prerequisite skills, and eventually to adaptations and goal achievement. In this section, I will go through a paper-design of the associated RE process. Along the way, I will introduce some first steps my colleagues and I have been making on providing automation.

5.1 A Goal Attainment Scale

For brevity, I will look at just one of my goals for doing digital darkroom work: working with B&W photos. I have come up with the following levels under this goal:

> *Level 1* (not attained): will never become proficient in any part of running a digital darkroom.
>
> *Level 2*: will be able to translate *exposure time* in a physical darkroom into *grayscale percentage* in a digital darkroom.
>
> *Level 3*: will be able to translate *light intensity* in a physical darkroom into *grayscale percentage* in a digital darkroom.
>
> *Level 4*: will be able to translate *print area* in a physical darkroom into *pixelized, bounded areas* in a digital darkroom.
>
> *Level 5* (fully attained): using a digital darkroom, will be able to reproduce prints achieved in my physical darkroom.

One might argue that these goals are not particularly unique to me, but are shared by the majority of digital darkroom users. However, after discussing the problem with a digital darkroom expert (MFA in Photography), it appears that only a small percentage of those using digital darkroom tools work on B&W, and even a smaller set are attempting to transfer skills from a physical darkroom setting. And most of them are professionals (as opposed to an amateur like me). He believes my goal is becoming more and more niche, as new generations of photographers grow up with purely digital medium. I am becoming a dinosaur.

I'd like to pause a moment at this point, and discuss potential ways to formalize goals. This can pave the way for later automation. The RE field has at least two projects/tools that focus on the representation of requirements as goals, the OME project and the KAOS project [3]. I have worked informally with the OME group to evaluate their tool's ability to represent the type of individual goals I have discussed in section 3. The result of the OME group's work on this problem is reported in [5]. My group, more recently, has been working closely with the KAOS tool [4]. I have several reasons for using KAOS for this problem. First, it supports a goal-hierarchy that is a natural fit with our current representation of general goals and their refinement into levels/sub-goals. It also supports the quantification of goals into ones that have observable success measures. Second, it provides temporal logic operators that support reasoning about deferred goals. For example, the operators sometime in the future (◊), the next state (o), always in the future (□), always in the future unless (*W*), and always in the future until (*U*) allow us to reason about the timeline of goals. Looking at the top portion of the figure below, you can see my first attempts to represent the letters-to-the-editor goal of Mary, using the KAOS Objectiver tool (www.objectiver.com). There is a general goal pattern that we have recognized in the email project of wishing to participate in an online forum (e.g., newsletter, discussion group). We use this pattern here to capture Mary's goal. To personalize the pattern, we must fill-in details on the particular newsletter Mary wishes to target.

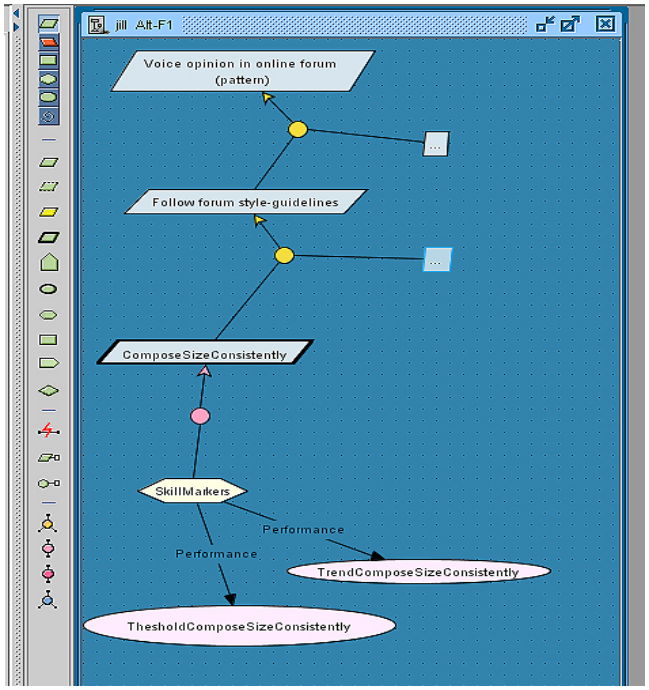

5.2 An Assessment of Abilities

In our work on email, my team mapped out 50 specific skills that impinged on the use of an email tool. We then used these skills to build an assessment tool and associated assessment procedure. I'd like to say that I have mapped out the same skill set needed for digital darkroom work, and then built an assessment tool around that set. I have not. However, I conjecture that it is possible; the skill set is not unknowable. I suspect a good first step would be to collect together all of the published training material on doing digital darkroom work, and use this to draw up a first cut at a skill set. Another candidate technique is to use task analysis from HCI, possibly working backwards from an existing tool, e.g., Photoshop. We used a bit of both techniques on the email problem [10]. At the same time, I must note that the full definition of this skill set is likely to be effortful: it took our team 6 months to hammer out the skill set for the email application. It is one of the key challenges in taking up a clinical approach to a problem.

Turning for a moment back to formalization, I have begun working with William Robinson on a means to connect KAOS goals with skills, and then eventually to monitors, using his tool ReqMon [7]. We have chosen to draw distinctions among skills, goals, and measures. A *goal* is a desired property of the user; it describes what "ought" to occur. A *skill* supports a goal by enabling it: it is a prerequisite to be successful in achieving the goal. A skill can support multiple goals. A *measure* is a means of linking runtime events with skills. Measures operationalize skill observation by specifying the events necessary in the usage of the runtime system to mark a skill as acquired.

I have included a bit of detail below to illustrate the idea. Mary's goal of meeting the newsletter constraint of writing legal size submissions is linked to a skill, ComposeSizeConsistently. Assume that the goal is instantiated with a size of 1-20 lines.

Skill Maximize ComposeSizeConsistently(EmailSize size)
Def
ThesholdComposeSizeConsistently(size) ∧
TrendComposeSizeConsistently(size)

Mary's goal is linked to this skill. We have chosen to put two prerequisites on this skill: (1) that Mary can consistently compose messages of the desired size, and (2) that she is not starting to backslide on the skill.

Measure ThesholdComposeSizeConsistently(EmailSize size)
Def
Average(PercentSize(size, ComposedEmail(*),24h),-14d))
≥ 75%

Measure TrendComposeSizeConsistently(EmailSize size)
Def
Slope(PercentSize(size,ComposedEmail(*),24h),-14d)) ≥ 0

The measures compute the average and trend of success for the last 14 days, where each 24-hour period is a data point. Can we define similar skill and measure components for my goals (see 5.1)? The question is what is observable? The example above, of measuring Mary's skills of writing the right size emails, is straightforward to observe – it is quantitative. Doing digital imaging has some quantitative components, but some qualitative ones as well. I will return to this issue in section 5.4.

5.3 Providing a First System

I would expect a clinical, RE analyst to now look at my goals and abilities and make a determination on how much to train me (if at all) before delivering my first dark room system. There is a plethora of training material on using digital darkrooms, so that is not an issue. Assume that I am basically at level 1, unable to use a digital darkroom tool effectively, i.e., during assessment I was unable to transfer any of my physical darkroom knowledge over consistently. Let's assume my clinician foregoes the traditional route of giving me a stack of books to read, courses to take and training videos to watch. Instead, she suggests a first system that is tailored to one or more of my goals. Where does this system come from? We built the adaptable email application from scratch, focusing on its fit with a clinical process. Will someone do the same for digital imaging tools? Construct a framework that can be used in a clinical RE setting? There are approaches in the software engineering field that appear relevant, e.g., architectures for adaptable systems, self-healing systems, product-line development. None that I am aware of have attempted to link to individual goals and abilities. There is also the economic issue. Even if an existing approach could be used to provide my first system, would vendors be interested in refactoring their software to meet my goals?

5.4 Monitoring and Adaptation

I have foreshadowed the problem here in previous sections. To support deferred goals, we need a means of (1) monitoring user activity, and (2) adapting the system when called for. In the email case study, we have done some preliminary work to close the deferred-goal loop in an automated fashion. We have been able to use definitions like those in section 5.2 to raise adaptation alerts, and then tie those alerts to actual changes to the email client. The foundation of this work is an email client that makes activity observable and is adaptable. As noted, we built our email client from the ground-up with these two issues in mind.

Turning back to my goal of doing digital darkroom work, what are the options? I could attempt to convince someone to build a digital darkroom system from scratch, paying attention to a clinical process that supports different goals and abilities (e.g., mine), and allows observation and adaptation; there are several open source projects that would be a starting point, e.g., imageJ, JIU. Or, I could attempt to convince a vendor of an existing digital darkroom procuct to refactor or wrap their product in a way that supports a clinical process. Or, I could attempt to use what software is on the market, attempting to bend their current personalization features to my needs. My attempts at the latter do not make me optimistic that it is a good option; the peronsalization features of today's software applications have little to do with individual goals and abilities.

Another problem is that few applications support the types of detailed external observations that are necessary to support deferred goals and adaptation. (At least two execptions are Sun's JVM, which has spawned a set of third party tools based on its observable events, and the Smart Connectors project, which has found a means of tapping into DLLs to observe low level events and support runtime operations [1].) I'll end my Greenfield example here, and discuss its feasibility in the next section.

6. TAKING STOCK

I have attempted to convince the reader that there are an interesting set of software development problems rooted in existing clinical fields. I've called these Brownfield problems to reflect my belief that we, as software engineers, must conform if we expect to fit in. I've made a start in this paper by looking at a specific clinical RE process that is tied to the introduction of email into the life of a brain-injured individual. The resulting process was one that fit with the established clinical procedures I found in place. It has proven quite successful, both in terms of meeting goals of the individuals and being accepted by the clinicians.

The email project barely covers the possibilities. There are a host of applications that would be of interest to this population (and to the larger cognitively impaired population), but remain out of their ability range, e.g., iTunes, photo sharing, graphic design. I can only hope that other teams might form in the future to tackle one or more of these challenging problems.

I have also attempted to convince the reader that the clinical viewpoint, one that puts individual goals and abilities first, might have new fields to conquer. Perhaps a bit fancifully, I defined a new clinical area around those with a Photoshop impairment. Using

myself as an example, I sketched the steps that a clinical RE process could provide to allow me to achieve my goals.

In my current view, I believe that the Brownfield approach is the more tractable. One is attempting to integrate into an already existing process. There are roles and practices established. Funding models are already in place through federal, state and local agencies. Even *partial* automation can help reduce the burden on the human staff to take on another treatment component.

On the other hand, who will staff the clinical (Greenfield) process I describe in section 5? Perhaps there are economics that could make such a clinical process viable. I, personally, would pay to be part of such a clinic: to have an assessment process lead to software delivered that I can use and will adapt when needed. Would others pay that have the same clinical needs? Along the same lines, where does the application software come from that will be at the heart of the clinic? This is a question for both Brownfield and Greenfield approaches. My team built the email application from scratch to conform to the concepts described in section 3. We did not see an option of taking an existing email application and bending it to our needs. More optimistically, there are a strong set of open-source projects that would appear to be natural starting points for other applications that we might consider, e.g., web-browsing, chat. In the end, the viability of taking a clinical approach to meeting individual software needs remains uncertain, bound up in the economics that can make it sustainable. I believe automation can play a key role, but there are health-care politics and policy, and open-market forces that must also play a part.

As a final note, and an attempt to close the loop on my Brownfield/Greenfield distinctions, it turns out that Manuel and I share a Photoshop impairment. Manuel, prior to his injury, had the digital darkroom concepts well in hand; he now has problems with the complexity of the tool interfaces. I am the inverse. I do not have the concepts in hand, but (mostly) understand the tool interfaces. Perhaps an enterprising software engineer can meet both our needs.

7. ACKNOWLEDGMENTS

The National Institute on Disability and Rehabilitation Research has been a generous contributor to the work reported here. I would also like to acknowledge the National Science Foundation's support, under grant IIS-0313324, and Intel's support of the community access project.

This paper reports on a team effort. I acknowledge the contribution of the unsung students on the team over the last four years: Jim Allen, Moira Burke, Andrew Fortier, Harrison Green-Fishback, Jason Prideau.

I would like to thank Harold Velline for giving me valuable advice on digital darkroom tools on the market today.

Lastly, I would like to acknowledge that my inability to use Photoshop should in no way impugn the reputation of that product. I freely admit to being challenged in making the switch from physical darkroom to digital darkroom. This basic problem would likely render any of the many digital photography tools on the market unusable by me.

8. REFERENCES

[1] R. Balzer and N. Goldman. Mediating Connectors. In *Proceedings of the 19th IEEE Conference on Distributed Computing Systems*, Austin, TX, May 1999, 73-77.

[2] Brain Injury Association, Inc., (2000). BIA Community Awareness Presentation. CD ROM. Available from the Brain Injury Association, Inc.

[3] Dardenne, A., A van Lamsweerde, S Fickas. Goal-directed requirements acquisition. *Science of Computer Programming*, vol. 20, pp. 3-50, Elsevier 1993

[4] Feather, M., S. Fickas, Axel van Lamsweerde, and C. Ponsard, Reconciling System Requirements and Runtime Behavior. In *Proceedings of the International Workshop on Software Specification and Design (IWSSD'98)*, Isobe, 1998.

[5] Hui,B., Sotirios Liaskos, John Mylopoulos, Requirements Analysis for Customizable Software: A Goals-Skills-Preferences Framework. In *Proceedings of 11th IEEE International Conference on Requirements Engineering*, Monterey Bay, CA

[6] LoPresti, E., A. Mihailidis, and N. Kirsch, Assistive technology for cognitive rehabilitation: State of the art. *Neuropsychological Rehabilitation* 2004

[7] Robinson, W. Monitoring Web Service Requirements. In *Proceedings of 11th IEEE International Conference on Requirements Engineering*, Monterey Bay, CA, pp. 65-74, 2003.

[8] Sohlberg, M.M., Fickas, S., Ehlhardt, L. & Todis, B. (in press). Case Study Report: The Longitudinal Effects of Accessible Email for Four Participants with Severe Cognitive Impairments. *Journal of Aphasiology*, 2005

[9] Sohlberg, M. M., Fickas, S., Hung, P., Lemoncello, R. (2004). Community Navigation Profiles for Six Individuals with Severe Cognitive Impairments. Poster session presented at *the annual meeting of the National Academy of Neuropsychology*, Seattle, WA.

[10] Sohlberg, M. M., Ehlhardt, E., Fickas, S., & Sutcliff, A. (2003). A pilot study exploring electronic mail in users with acquired cognitive-linguistic impairments. *Brain Injury*, 17(7), 609-629.

[11] Sohlberg, M. M., & Mateer, C. A. (2001). *Cognitive Rehabilitation: An integrative neuropsychological approach*. New York: Guildford Press