
Explainable Recommendation for Self-Regulated Learning

Michael Freed
Melinda Gervasio
Aaron Spaulding
Louise Yarnall

MICHAEL.FREED@SRI.COM
MELINDA.GERVASIO@SRI.COM
AARON.SPAULDING@SRI.COM
LOUISE.YARNALL@SRI.COM

SRI International, 333 Ravenswood Ave., Menlo Park, CA 94205 USA

Abstract

Recent years have seen rapid advances in intelligent technology to support online learning, but these have primarily targeted formal educational contexts such as classrooms and e-courses. In contrast, the predominant form of adult learning in the workplace is informal and self-directed. Learners self-assess competency, set goals, find relevant learning resources, and initiate learning activities covering many topics at different depths at different points in time. Our approach to supporting self-regulated learning is embodied in PERLS, a mobile personal assistant application that serves as a virtual mentor for informal learning. A key component of PERLS is its recommendation system, designed to adaptively co-construct a path towards desired learning outcomes with the learner. Recommendations are made largely on the basis of value propositions, each a persuasive explanation for taking a particular learning action to advance along a particular learning path. In this paper, we present a process model of self-regulated learning used in PERLS and our approach to generating explained recommendations and using them to co-construct learning paths.

1. Introduction

Rapid innovation in mobile computing and intelligent personal assistant technology presents an opportunity to better support *self-directed learning*, an activity that is widely seen as critical to individual job success and overall workforce adaptability. Self-learning is the predominant form of learning in the workplace (Marsick & Watkins, 1990; Livingstone, 1999). Adults routinely learn job-related knowledge for which little or no formal instruction is available, doing so through a self-assembled mixture of resources at times, places, and pace of their own choosing. Informal learning presents challenges such as discovering needs, identifying resources to address those needs, self-assessing progress, and coping with conditions that strain a learner's determination and meta-cognitive faculties. Fortunate individuals find mentors to help meet these challenges. For those not so fortunate, technology able to provide mentor-like support can improve learning outcomes.

We have implemented and field tested PERLS, a personal assistant aimed at reducing several practical and motivational difficulties of learning outside formal training contexts. One such difficulty is finding content for activities such as pre-study exploration and post-study sustainment which are important for self-learning but often absent from formal instruction. PERLS infers whether the learner is currently in a pre-study, study, or post-study phase, and recommends phase-

appropriate content. A second difficulty is accommodating adult patterns of learning which, in contrast to school age children, typically involve short, intermittent bursts of activity at convenient times. PERLS mainly recommends short form “microcontent,” intelligently mixing content meant to advance learning goals with content meant to promote engagement, and prioritizing recommendations to minimize time spent finding desirable content, making brief windows of available time useful for learning. A third difficulty is persuading users that recommendations are worth their time. In keeping with accepted best practices for in-person training, PERLS explains the rationale for each recommendation, letting the user to decide if the rationale is compelling.

This paper presents our approach to virtual mentorship in PERLS, with a particular focus on describing how PERLS generates explainable recommendations appropriate for self-directed learning in informal learning settings. We first describe the PERLS application, and the process model of self-regulated learning (SRL) it uses to represent and track individual learning trajectories. We then describe how recommendation candidates are generated, ranked, and tagged with explanations. We conclude with a brief overview of field studies evaluating user response to PERLS recommendations.

2. PERLS Overview

The PERLS mobile app uses a card-based user interface, where each card displays a recommendation and a swipe gesture advances to the next recommendation (Figure 1). In general, recommendations that come earlier in the sequence are predicted to be most appropriate for the learner, given the user’s inferred interests, stage of learning, and current attitude. The app recommends content on a mixture of topics, including new ones meant to help the user discover new needs and interests. And it mixes appealing content meant to promote engagement with challenging content that advances learning goals. The rationale behind a recommendation is surfaced through a “sell point,” a pithy statement that explains the recommendation in motivational terms (e.g., “Your peers find this to be important” and “This will help you retain what you learned last month”).

The card-based interface also lets PERLS serve up specialized content types such as event notifications, quiz cards (single, multiple-choice questions), and action cards that suggest or offer to carry out actions that improve learning outcomes. For example, action cards can be used to encourage learners to set topic-specific objectives, make plans, and reflect on progress. Such actions, along with the deliberate combination of exploratory learning with goal-driven learning, and




Figure 1. PERLS card-based UI.

engaging content with challenging content, help replicate strategies used by strong self-learners.

PERLS prioritizes recommendations in a three-stage process. First, it generates a candidate set of recommendations based on inferred topic interests. Next it ranks candidates based on contextual factors, including estimates of the user’s current phase in the SRL trajectory for each topic, the appropriateness of candidate content for that phase, and the relevance of different motivations at that phase. Finally, PERLS adjusts the actual sequencing of recommendations to more closely align with the user’s immediate situation. This corresponds to steps a good human mentor might take to help a learner advance in their learning trajectories – i.e. to gain an understanding of the user’s level of interest and learning goals, determine relevant situational factors, make content recommendations that fit both long term learning goals and short term interests, attitudes, and situational context. The learner retains control, select among content and action options presented by the system. This is crucial, as both PERLS and the learner have different, equally valuable information for deciding the most effective next step in the overall learning path. Whereas PERLS has access to indices of learning resources and models of effective pedagogy, learners have privileged knowledge of their own preferences and other internal states.

3. PERLS Self-Regulated Learning Model

The foundation for PERLS’s recommendations is a process model of Self-Regulated Learning (SRL). The model describes self-learning as a set of activities grouped into three main phases—*Explore*, *Study*, and *Sharpen*. Most current technology to support learning focuses on *Study*, in which the learner makes a concentrated effort to achieve a learning goal, usually in a formal, structured context such as a classroom or e-course. Self-learning, which occurs primarily in informal learning contexts, encompasses earlier activities that sets the stage for successful *Study*, and later activities that maintain and enhance outcomes.

The PERLS SRL model divides each of the three phases into a set of key activities or subphases (Figure 2). These define intermediate self-regulated learning transitions and objectives. For example, establishing motivation, orientation, and confidence while *Dabbling* improves knowledge intake during *Familiarization*. Content appropriateness varies systematically by sub-phase. For example, dabble content should ideally be brief, engaging, and not very difficult because learners in the Dabble phase will not typically have committed to studying the topic. The model captures the range of possible paths, with individuals varying in where they

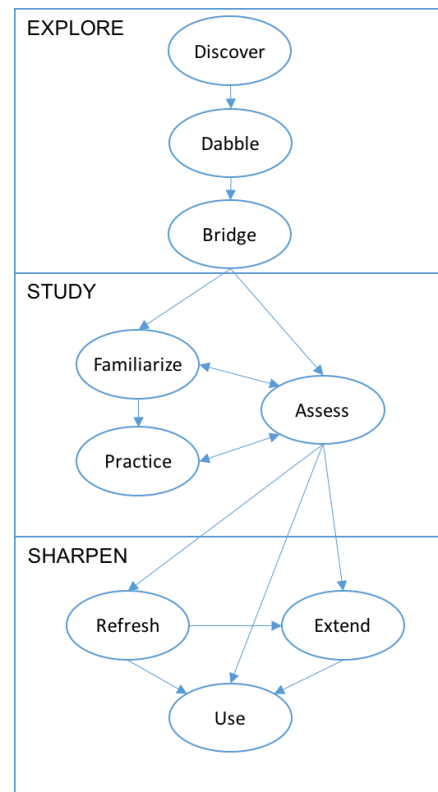


Figure 2. The PERLS Self-Regulated Learning (SRL) Model.

enter and exit, and whether they skip or repeat a given phase.

The *Explore* phase starts with *Discovery*, where the learner first becomes aware of a topic. In some cases, discovery results from intentional exploration—for example, a learner discovering potentially interesting topics while browsing news content. In others, discovery happens incidentally while doing other things or in response to colleagues drawing attention to the topic. Learners who are naturally curious, perceptive, and social will tend to become aware of important topics in a timely way; others benefit from technology that supports discovery.

After learners have discovered a topic, they may begin *Dabbling*—engaging in light interaction with topic materials, consistent with a low level of commitment to long-term learning. Dabbling establishes motivation, confidence, and conceptual orientation to a topic in preparation for higher-commitment learning. Dabbling content will be relatively lightweight and engaging—e.g., brief narratively structured readings, engagingly edited video, and brief games or game-like experiences.

Bridging is the process of preparing for more intensive, high-commitment learning. Learners self-assess their level of competence in the topic and begin to formulate their learning plan, setting expectations, and gauging the time and effort it will take to achieve their learning goals. During bridging, learners check their current understanding and confidence, connect with more advanced learners and mentors, and preview learning materials.

Learners who make the transition to the *Study* phase have made the commitment to gain some level of competence, to complete a course, or to achieve some other goal. Formal instruction is a special case of *Study* since learners may achieve goals through either formal or informal methods. Study begins with *Familiarization*, where the learner goes through introductory study materials to build a foundation for learning (e.g. knowledge of terms, concepts, procedures, and principles).

During *Practice*, learners build fluency in retrieving and applying acquired knowledge, including development of procedural skills. Practice typically involves building expertise by focusing on specific gaps and elements in one’s knowledge and skills.

Throughout the *Study* phase, learners will also typically be in *Assessment*—feedback, either from others or self-generated, both formal and informal, on their proficiency level in the topic and progress toward learning objectives. Early formative assessment provides the learner with feedback on performance, allowing course corrections and adoption of new strategies to improve results. Later, summative assessment evaluate performance against criteria-based standards or group norms, and determine when the learner is ready for more advanced learning.

Once learners have achieved the desired level of competence in a topic, they enter the *Sharpen* phase for sustaining acquired knowledge. In *Use*, learners build fluency by applying their newly acquired knowledge to real-life problems or situations and understand of the range of situations where the knowledge is most commonly applicable.

In the *Refresh* subphase, learners seek to check or strengthen knowledge or skill. This may be driven by some external circumstance where their knowledge or skill will be on display (for example, a meeting or a presentation) or by an innate drive for increased proficiency.

A learner in the *Extend* subphase is seeking to enrich basic knowledge or to deepen a basic skill beyond the requirements for basic proficiency. Learners extending their knowledge build on existing competence foundation—for example, by learning about unusual/corner cases or by transitioning into subtopics.

4. Content Recommendation for Self-Regulated Learning

Technology-enhanced learning presents challenges not typically addressed in the recommender systems community (Manouselis et al., 2011). For example, conventional recommenders, including those designed to be sensitive to context (Verbert et al., 2012), typically focus on optimizing user response to a single metric such as ad click-throughs or monetary value of supplementary purchases (Adomavicius & Tuzhilin, 2005). In contrast, learning-oriented recommenders should use different metrics for different phases of learning. For example, when in a discovery phase on a particular topic, the recommender should optimize for obtaining *positive or negative* evidence of learner interest in the topic. In other cases, it might need to optimize for having the user meet an externally imposed Study completion deadline or acting to reinforce prior learning during the critical period for retention.

Technology focused on adult learning must deal with additional challenges. Whereas most recommendation methods rely on statistical machine learning approaches requiring a great deal of training data, workplace-oriented learning content is often interesting to a small number of people and of no interest to most others. As a result, SRL recommendation needs to rely heavily on techniques used for recommendation bootstrapping such as rules and decision-theoretic methods. And, in accordance with established principles for adult learning, self-learning recommendations need to be explained in convincing detail (Knowles, 1984). Thus, the mentor must have some mechanism for representing and reasoning about learner motivation. As described below, the recommendation approach incorporated into PERLS meets these adult learning requirements by estimation topic interest, generating a set of recommendation candidates consisting mainly of learning objects on high interest topics, identifying potential reasons to accept each candidate, predicting the strength of user motivation for each such reason, and then ordering the presentation of recommendations to prioritize the strongest.

4.1 Topic Interest

All PERLS recommendations will be on topics of interest to the learner, although the learner's level of interest, learning stage, and particular learning goals may vary from topic to topic. PERLS thus tracks user interest in *topics*¹, monitoring for direct evidence of user interest and propagating this to nearby topics in the corpus structure. Direct evidence comes in various forms, such as the user starring or subscribing to a topic, setting a learning goal, or starting some learning content. In determining the *base interest* in a topic from direct evidence, we use the following heuristic: explicit intention > explicit interest > demonstrated interest. Thus, setting a goal to complete a course (explicit intention) is stronger evidence of interest than starring a topic (explicit interest), which is stronger than completing a *learning object* (LO) under a topic (demonstrated interest).

Multiple instances of the same type of evidence within a short time period are indicative of greater interest while evidence that occurs further in the past indicates lower interest. Thus, we increase the base interest value for *intensity* (number of instances of the same evidence type within

¹ The PERLS corpus is organized through four kinds of content groups: *topics* (sets of learning objects (LOs) or groups corresponding to semantically meaningful topics), *collections* (arbitrary sets of LOs), *courses* (ordered sets of LOs), and *feeds* (short-lived LOs that arrive over time). Since these distinctions are irrelevant for the purposes of this paper, we simply use *topics* to refer to all the groups.

a given time period) and discount it for *recency* (time since the most recent event in the cluster). For the intensity adjustment, we want a function that gradually asymptotes to the maximum interest level. For example, for a base interest of b , a maximum interest level of 10, and n instances of the evidence type within a designated period of time, we can calculate the additional interest for user u in a topic t due to intensity as:

$$\Delta_{intensity}(u, t) = (10 - b) * (1 - \frac{1}{n})$$

For the recency adjustment, we want the inverse—a smooth discounting toward the minimum interest value (0). For example, if we allow a grace period of d days (i.e., within d days, an event is still considered ‘recent’), we can calculate the discount for an evidence type whose most recent event was m days ago as:

$$\Delta_{recency}(u, t) = \max(\log_2(m - d + 1), b + \Delta_{intensity}(u, t))$$

The *direct interest* of a user in a topic t is then

$$interest_{direct}(u, t) = b + \Delta_{intensity}(u, t) - \Delta_{recency}(u, t)$$

After computing direct interest for all topics, we can calculate *indirect interest*—i.e., interest inferred from the direct interest in neighboring topics. Intuitively, interest in a subtopic translates to strong interest in a parent topic, so we distribute uniformly the direct interest in a topic to its parent topics. For example, if a child topic c with direct interest $interest_{child}$ has n_p parents, then the indirect interest contributed by the child to a parent can be calculated as:

$$\Delta_{intp}(u, c) = interest_{child}(u, c) - \log_{10} n_p$$

In contrast, indirect interest from a parent topic decreases with the number of child LOs of that topic. A topic with many immediate child LOs is likely to be a relatively self-contained focus of interest and thus interest in that topic is unlikely to transfer to its subtopics. On the other hand, interest in a topic with no child LOs strongly implies interest in at least some of its subtopics. Thus, for computing indirect interest from a parent to its children, if a parent topic p has n_{clo} child LOs and n_{ctopic} child subtopics, then we might calculate the indirect interest contributed by the parent to a child subtopic as:

$$\Delta_{intc}(u, p) = interest_{parent}(u, p) - \log_4 \frac{n_{clo} + 1}{n_{ctopic}}$$

The final interest level of a topic t is the sum of the direct interest in the topic and the maximum of the indirect interests from its parents and children, i.e.,

$$interest(u, t) = interest_{direct}(u, t) + \max(\max_{c \in Children(p)} \Delta_{intp}(u, c), \max_{p \in Parents(c)} \Delta_{intc}(u, p))$$

The use of *max* to aggregate indirect interest greatly simplifies the calculation of topic interest but, more importantly, supports a straightforward explanation of why PERLS believes a learner is interested in a topic. We note that the specific functions used for calculating direct and indirect interest are less important than their characteristics (e.g., the effects of intensity gradually

diminishing with the number of events) and we anticipate refining these functions as we obtain additional PERLS usage data. Given a user’s interest levels in the different topics, PERLS considers as candidates for recommendation all the LOs under topics meeting a minimum interest threshold.

4.2 Value Propositions

PERLS uses diverse features of users, content, and situation factors as inputs. With enough training data, statistical machine learning approaches could be used to acquire accurate ranking functions for the recommender. However, the application of recommending work-related micro-content is a small data problem: unlike general education domains (e.g., algebra) where many learners may consume each bit of content and content rarely becomes obsolete, workplace content is typically fragmented, low-circulation, and perishable. Standard approaches are also not amenable to generating the explainable recommendations that are critical to adult learning.

We thus adopt an approach with a significant domain modeling and knowledge engineering component. In particular, PERLS recommendations are based on the notion of a *value proposition* (VP), that represents a potentially compelling explanation or “reason” for a user to accept a recommendation. There is a large scientific literature on human motivation, although none are completely suited to self-directed learning. Some are too broad, covering all of human experience while providing little insight on learning in particular (e.g., Maslow, 1943; Reiss, 2004). Others apply specifically to learning but focus on particular learning behaviors and contexts (e.g., Lepper & Malone, 1987) or on goal setting (e.g., Locke, 1968).

Perhaps the best known scientific idea regarding motivation is the distinction between intrinsic and extrinsic motivation—internal desires such as curiosity vs. external rewards or punishments. However, ongoing research in Self-Determination Theory shows that motivations are inherently neither one nor the other (Ryan & Deci, 2000). We thus take a simpler approach of categorizing VPs by whether they involve *endogenous* or *exogenous* motives. Endogenous motives involve an innate connection between learning activity and outcome (e.g., learning to play a guitar for fun or because you want to make music), while exogenous motives focus on rewards and punishments (e.g., learning to play a guitar to impress people or make money).

We are developing an ontology of VPs based on these and other high level distinctions grounded in relevant psychological literature. VPs based on specific motivational types tend to be more persuasive but apply in a narrowed range of conditions. So our goal is an extensible library containing numerous and diverse VPs. The ontology is important for guiding the expansion of this library and assessing its coverage.

4.3 Recommendation Strength

Intuitively, PERLS looks for the most compelling reason to recommend a particular LO to a learner. In general, multiple VPs will apply to any one LO and there may be significant, unknown semantic overlap between them. Thus, PERLS calculates as the recommendation value or *strength* of a LO for a user to be the *maximum* strength of any VP for that LO. Let u be a user, o be a candidate LO, and V be the set of all VPs. Then the recommendation *strength* of o for u is:

$$strength(o, u) = \max_{v \in V} strength(u, v, o)$$

where $strength(u, v, o)$ is the product of the VP's *importance* and *fitness* to u according to o :

$$strength(u, v, o) = importance(u, v, o) * fitness(u, v, o)$$

4.3.1 Importance

Importance is a sum of three values: *topic importance*, *urgency importance*, and *VP importance*. *Topic importance* can be interpreted as the value of a LO due solely to it being on a topic of interest while *urgency importance* is its value due solely to the existence of a pressing need to consume the content sooner rather than later. While one could ostensibly create VPs that capture motivations along these lines (and we do), PERLS factors them separately because they capture fairly universal motivations. That is, all learners will naturally be more interested in topics of interest and they will be more inclined to attend to urgent content.

Topic importance is a function of the learner's interest in the topic—at its simplest, topic interest itself, i.e.,

$$importance_{topic}(u, t) = interest(u, t)$$

For urgency, PERLS tracks a number of *urgency factors*, each representing some immediate need to consume the LO—for example, because of an approaching deadline (whether self-imposed or set by an external authority), or because the content is associated with some location the learner is currently near. *Urgency importance* is the maximum value of any of these factors—i.e., the most urgent need. Let U be the set of urgency factors and $urgency(u, o)$ be the value of the factor u for a LO in the user's current context. Then the urgency importance of a LO for a user is:

$$importance_{urgency}(u, o) = \max_{u \in U} urgency(u, o)$$

The primary component of importance is *VP importance*, which captures the rationale behind a recommendation—i.e., how motivating a particular VP will be to a learner, given the learner's current state of mind and stage of learning on a given topic. There are two main components to VP importance: *attitude* and *phase*. *Attitude* represents the learner's current disposition toward learning. For example, a learner with an *achievement attitude* in the *Familiarize* subphase of *Study* is looking to make significant learning progress and is thus likely to be receptive to more challenging content. In contrast, a learner in the same subphase with a *discovery attitude* is still looking to get the lay of the land and will probably be more amenable to lighter content. Given a set of possible attitudes, we attach to each VP a baseline VP importance for each subphase for that attitude that captures how important that VP is generally to a learner with that attitude. For example, VPs with high baseline values for the discovery attitude in the Explore subphases might include “This is trending” or “You might be interested in this because of your interest in X.” Meanwhile, VPs with high baseline values for the achievement attitude in the Study subphases might include “This will raise your competency in X” or “This will complete the course.”

The second component of VP importance is *phase*, which represents the user's current stage within the SRL model described earlier. More specifically, within any subphase, the user may be in one of four states: *Not Ready*, *Ready*, *Progressing*, and *Done*. For a VP to apply to a LO given the user's learning status for a topic, the user must be either *Ready* or *Progressing* in the subphase, the content must be appropriate for that subphase, and the VP must have a nonzero baseline value

for the subphase. PERLS currently constructs a discrete probability distribution over these states for each subphase to provide an estimate of the user being in a particular state within a subphase. The probabilities are inferred through a *Markov Logic Network* (MLN) (Richardson & Domingos, 2006), which we chose because it provides a principled way to integrate probabilities with logical rules. Phase estimation is a task for which we have significant knowledge about different pieces of evidence that a user is in a particular phase/subphase/state but where there is some uncertainty in that knowledge and in the evidence itself. For example, if a user starts a few *Dabble* LOs in a topic, the user is likely to be in the *Ready* or *Progressing* states for *Dabble*. After completing several such LOs, the learner is likely to be *Done* with *Dabble*. If learners are *Ready*, *Progressing*, or *Done* in a subphase, they are likely to be *Done* with any precursor subphase. By encoding such rules within a MLN, PERLS can use the learner’s activities to infer the probabilities over the different states. In addition, when learner data becomes available, MLNs naturally lend themselves to automatic adjustment of rule weights through machine learning techniques.

Given the phase estimates, PERLS can determine the learner’s current state within a subphase to know whether a VP applies. A straightforward approach would be to set the highest-probability state in a subphase as the learner’s current state in that subphase. However, this would lead PERLS to always recommend content that matches its best guess as to the user’s current learning stage. While reasonable, this approach relies heavily on accurate estimates and prevents fortuitous exposure to topics of potential but unverified interest. To address this *exploration/exploitation tradeoff*, PERLS employs an ϵ -greedy exploration strategy: with probability $1 - \epsilon$, PERLS chooses a high-probability state and with probability ϵ , a low-probability one. We set the high/low threshold at 0.25 (uniform probability over four states). So, for example, if the states $\{Not\ Ready, Ready, Progressing, Done\}$ have probabilities $\{0.2, 0.4, 0.3, 0.1\}$ respectively, then exploit will pick *Ready* with probability $\frac{0.4}{0.7}$ and *Progressing* with probability $\frac{0.3}{0.7}$, whereas explore will pick *Not Ready* with probability $\frac{0.2}{0.3}$ and *Done* with probability $\frac{0.1}{0.3}$.

The final factor for determining VP importance is *fitness-to-phase*—i.e., how well the LO suits a phase/subphase/state. There are different possible interpretations of suitability—for example, that an educator has determined the content to be appropriate or that the learner will agree that the content is desirable in their current context.

Given that the user may be *Ready* or *Progressing* in multiple subphases, that the VPs have different baseline importance values for different subphase-states, and that the LOs are suitable for different phases, we want to choose the combination of VP and subphase-state that leads to the highest VP importance value.

Let $S(u,t)$ be the subset of subphase-states selected for the user u for a topic t according to the exploration/exploitation strategy described above. Further let $importance_{base}(v,p,s,a)$ be the baseline importance of the VP v for the subphase-state (p,s) when the user has attitude a , and $fittophase(o,p,s)$ be the LO fitness to the subphase-state. Currently, PERLS relies on learner attitude being given (e.g., by the user indicating their own attitude) but eventually, we anticipate estimating the learner’s attitude based on observed behavior (e.g., the rate at which the user is flipping through content). Similarly, $fittophase(o,p,s)$ is currently provided by corpus contributors but could potentially be learned from learner data in the future. The value of a VP v for a LO o in the subphase-state (p,s) when the learner has attitude a is:

$$importance_{VP}(u, v, a, t, o) = \max_{(p,s) \in S(u,t)} importance_{base}(v, p, s, a) * fittophase(o, p, s)$$

Finally, importance is a weighted sum of topic importance, urgency importance, and VP importance. As with the formulas presented for calculating interest, the calculation of importance values and fitness values are meant to provide prescriptive criteria, and we expect that some parameter adjustment or formula modification will be needed to better match actual experience.

4.3.2 Fitness

The second factor in $strength(u, v, o)$ is *fitness*—i.e., the likelihood that the LO will deliver the value defined by the VP. A LO provides strong support for a VP if the VP is true of that LO with high certainty, the content is suited to a selected subphase-state for its topic, and it is high quality (e.g., it is appealing, effective, up-to-date). We have already discussed LO suitability for a phase state so we focus on VP truth value and content quality here.

Rather than attempting to define *VP truth value* as an absolute measure, we take a more practical approach of defining it by whether a user is likely to agree that the VP is true of the LO. The criteria for determining the probability that a VP is true differ for each VP, so each VP is associated with a unique estimation function within some normalized range (e.g., [0,1]). VP truth value estimation can vary widely. Any information represented in PERLS is potentially relevant, including situational, learning, social, and interaction context; and user profile, population, and corpus data as well as trends or patterns computed over them. While some VP values may be easily determined—for example, by checking for specific annotations or metadata on a LO, other calculations may be more involved. For example, determining that “This is hot with your peers” requires identifying the peer group and calculating statistics on their learning; while determining that “This is the sort of thing you like first thing in the morning” could involve learning a classifier based on features of the content, time of LO interaction, user feedback, and so on. Regardless of the information required to for the VP truth value estimate, computing the value itself should be reasonably efficient as the computation will have to be made for all potential recommendations.

Content quality addresses the desire to prefer recommending higher quality content over lower quality ones. In line with the objective of extensibility, PERLS does not subscribe to a well-defined set of quality criteria nor does it require information about their relative importance. Instead, PERLS allows quality criteria to be defined over a wide range of factors such as production value, instructional effectiveness, enjoyability, and author popularity. We cast content quality as a multiplicative discount factor ranging from [0,1] so that the discount in importance is inversely proportional to the quality of the content.

Finally, we have

$$fitness(u, v, o) = truth(v, o) * fittophase(o, p, s) * quality(o)$$

where (p,s) is the subphase-state determined earlier to yield the highest importance value.

5. Summary

Our approach to content recommendation has been implemented in the PERLS mobile app. In addition to providing a mobile-based app for recommendations to support self-directed learning,

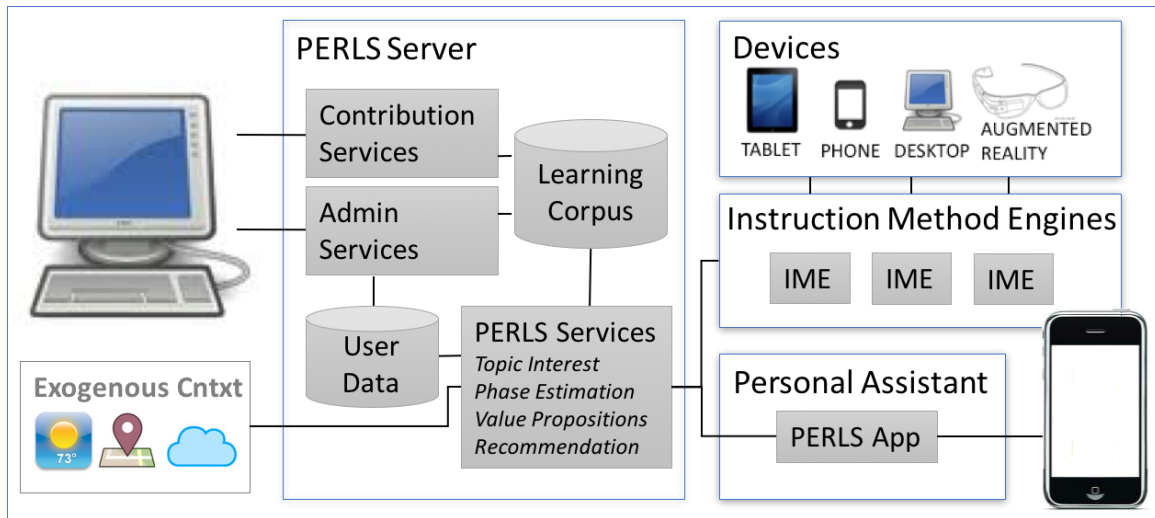


Figure 3. PERLS serves as an integration point for a variety of content providers.

PERLS also serves as an integration point for different content providers and content delivery methods (Figure 3). Selecting a card in PERLS initiates a learning activity, with some activities supported natively in the PERLS mobile phone app and others linking to external applications integrated with PERLS (Freed et al., 2017). This enhances PERLS ability to support lifelong (or at least employment-long) learning for members of large organizations where learning activities are typically supported by diverse systems.

PERLS has been deployed in limited contexts within our organization for the purpose of conducting user studies to explore and validate design concepts (Freed et al., 2014), and has recently been tested in collaboration with several Department of Defense organizations. First, it served as the primary user interface to a variety of content providers and backend services embedded within the Total Learning Architecture (TLA) (Regan et al., 2013), a learning system integration standard under development by DoD's Advanced Distributed Learning (ADL) Initiative. Second, it was used to support disaster response-related training in collaboration with USNORTHCOM and the Joint Knowledge Online (JKO) office, with a specific focus on learner response to anytime/anywhere microlearning (short, frequent bursts of learning) as a potential alternative or enhancement to traditional course-based learning. Learner response was strongly positive (Freed et al, 2017b), and efforts are underway to carry out larger scale tests leading to eventual adoption. Finally, in collaboration with National Security Office (DLNSEO), PERLS was tested as a means to support advanced Mandarin language learners, also with strong positive user response.

Self-directed, informal learning is the predominant form of adult learning in the workplace and yet remains largely unaddressed by existing work in intelligent learning assistants. Much of this has to do with the ill-defined, partially structured, dynamic, and uncertain nature of informal learning, which makes traditional approaches designed for formal, classroom-based learning ill-suited to the task. In our work on PERLS, we have developed a model of self-regulated learning that captures a wide variety of informal learning trajectories. The PERLS approach to content

recommendation is centered on the notion of value propositions, which provide the motivational rationale behind recommendations. By tracking user activity, PERLS can estimate users' level of interest in different topics, their learning goals, and their progress through the SRL model, enabling PERLS to make recommendations that best suit the user's current learning context.

Acknowledgments

This material is supported by the ADL Initiative under Contract Numbers W911QY-12-C0171 and W911QY-16-C-0019. Any opinions, findings, conclusions, or recommendations expressed in this material are those of the authors and do not necessarily reflect the official views of the U.S. Government or Department of Defense. The authors would also like to thank Kenneth Wingerden, Brian Blonski, and Nicholas Boorman for helping bring PERLS to life.

References

- Adomavicius, G., & Tuzhilin, A. (2005). Toward the next generation of recommender systems: A survey of the state-of-the-art and possible extensions. *IEEE Transactions on Knowledge and Data Engineering*, *17*, 734–749.
- Freed, M., Folsom-Kovarik, J.T., & Schatz, S. (2017a). More than the sum of their parts: Case study and general approach for integrating learning applications. *Proceedings of the 2017 Modeling and Simulation Conference*. Paris, France: WASET.
- Freed, M., Yarnall, L., Spaulding, A., & Gervasio, M. (2017). A mobile strategy for self-directed learning in the workplace. *Proceedings of the 2017 Interservice/Industry Training, Simulation, and Education Conference*. Orlando, FL: NTSA.
- Freed, M., Yarnall, L., Dinger, J., Gervasio, M., Overholtzer, A., Pérez-Sanagustin, M., Rochelle, J., & Spaulding, A. (2014). PERLS: An approach to pervasive personal assistance in adult learning. *Proceedings of the 2014 Interservice/Industry Training, Simulation, and Education Conference*. Orlando, FL: NTSA.
- Knowles, M. (1984). *The adult learner: A neglected species*. Houston, TX: Gulf Publishing.
- Livingstone, D. W. (1999). Exploring the icebergs of adult learning: Findings of the first Canadian survey of informal learning practices. *Canadian Journal for the Study of Adult Education*, *13*, 49–72.
- Lepper, M. R., & Malone, T. W. (1987). Intrinsic motivation and instructional effectiveness in computer-based education. *Aptitude, Learning, and Instruction*, *3*, 255–286.
- Locke, E. A., Cartledge, N., & Koeppel, J. (1968). Motivational effects of knowledge of results: A goal-setting phenomenon? *Psychological Bulletin*, *70*, 474–485.
- Manouselis, N., Drachsler, H., Vuorikari, R., Hummel, H., & Koper, R. (2011). Recommender systems in technology enhanced learning. In F. Ricci, L. Rokach, B. Shapira, and P. Kantor (Eds.), *Recommender systems handbook: A complete guide for research scientists and practitioners*, 387–409. Boston, MA: Springer.
- Marsick, V. J., & Watkins, K. E. (1990). *Informal and incidental learning in the workplace*. London, England: Routledge.
- Maslow, A. H. (1943). A theory of human motivation. *Psychological Review*, *50*, 370–396.

- Regan, D., Raybourn, E. M., & Durlach, P. J. (2013). Learner modeling considerations for a personalized assistant for learning (PAL). In R. A. Sottolare, A. Graesser, X. Hu, and H. Holden (Eds.), *Design recommendations for intelligent tutoring systems: Learner modeling*. Adelphi, MD: U.S. Army Research Laboratory.
- Reiss, S. (2004). Multifaceted nature of intrinsic motivation: The theory of 16 basic desires. *Review of General Psychology*, 8, 179–193.
- Richardson, M., & Domingos, P. (2006). Markov logic networks. *Machine Learning*, 62, 107–136.
- Ryan, R. M., & Deci, E. L. (2000). Self-determination theory and the facilitation of intrinsic motivation, social development, and well-being. *American Psychologist*, 55, 68–78.
- Verbert, K., Manouselis, N., Ochoa, X., Drachsler, H., Bosnic, I., & Duval, E. (2012). Context-aware recommender systems for learning: A survey and future challenges. *IEEE Transactions on Learning Technologies*, 5, 318–335.