

Semantic and Episodic Learning to Integrate Different Opportunities for Life-Long

J. T. Folsom-Kovarik, **Randolph M. Jones**, Dylan Schmorrow

Soar Technology

Ann Arbor, MI

jeremiah@soartech.com, rjones@soartech.com, dylan.schmorrow@soartech.com

ABSTRACT

The Advanced Distributed Learning (ADL) Initiative has developed the Training and Learning Architecture (TLA) with the goal of using information technology to change the paradigm of education from occasional classroom study and training to pervasive and lifelong activity. TLA views education content providers as services that produce educational content relevant to particular learning needs and contexts. A content brokering service assesses an individual learner's current learning needs, and recommends content that is suitable to those needs but also appropriate to the situation the learner is in (e.g., recommending listening to a particular podcast if the learner is driving somewhere). Key technologies for the TLA vision are a Learning Record Store (LRS), which stores a continuously updated record of learning activity and outcomes, as well as a content meta-tagging language that enables the mapping of particular educational tools and content to specific situations. One challenge associated with content meta-tagging is that it requires significant manual effort, especially as content and technologies change. TLA also does not yet have a capability to identify new meta-tags or relationships between existing tags, implying that it may miss some opportunities for effective instruction. We describe a new research effort called FLUENT (Fast Learning from Unlabeled Episodes for Next-generation Tailoring), which will learn new tags and relationships to improve the overall coverage and effectiveness of content delivery in TLA. FLUENT will use a hybrid machine-learning approach that includes episodic learning, heuristic search based on analogical mapping, and an explanation-based learning capability that uses a background knowledge base of causation in instruction to discover relationships from examples. The knowledge-based learning approach will allow effective learning in a domain where statistical learning methods would suffer from the sparse data available.

ABOUT THE AUTHORS

Dr. Folsom-Kovarik is a Lead Scientist with Soar Technology, Inc. (SoarTech). His research interests include computational approaches to implement and advance instructional strategies, methods, and best practices that accelerate or deepen learning and education through adaptive, intelligent computer systems. His research has contributed to cognitive ergonomics, authoring, and instructor control in intelligent tutoring and training. A second research area is advancing learner model accuracy, robustness, and computer understanding to enable planning ahead and deliver more effective teaching and training. He earned a PhD from the University of Central Florida.

Dr. Jones, Senior Artificial Intelligence Engineer at SoarTech, is a leading developer of knowledge-rich intelligent agent software. He has been principal investigator for a variety of advanced R&D projects funded by ONR, ARI, DMSO, DARPA and other DOD agencies. He has previously held teaching and research positions at Colby College, the University of Michigan, the University of Pittsburgh, and Carnegie Mellon University. His areas of research include computational models of human learning and problem solving, executable psychological models, and full-spectrum intelligent behavior models. He earned a BS in Mathematics and Computer Science at UCLA, and MS and PhD degrees from the Department of Information and Computer Science at the University of California, Irvine.

Dr. Schmorrow, Chief Scientist at SoarTech, leads the advancement of research and technology tracks to build intelligent systems for defense, government, and commercial applications that emulate human decision-making. He also serves as a Potomac Institute for Policy Studies Senior Fellow, Editor of the Theoretical Issues in Ergonomics Journal, and the Technical Advisor for the Applied Human Factors and Ergonomics Conference Series. He is a retired US Navy Captain and his past service includes OSD, DARPA, NAWC, NRL, ONR, Naval Postgraduate School, and Executive Assistant to the Chief of Naval Research. Dr. Schmorrow holds a PhD in Experimental Psychology from Western Michigan University, as well as MS degrees in Psychology and Philosophy.

Semantic and Episodic Learning to Integrate Different Opportunities for Life-Long Learning

J. T. Folsom-Kovarik, Randolph M. Jones, Dylan Schmorow

Soar Technology

Ann Arbor, MI

jeremiah@soartech.com, rjones@soartech.com, dylan.schmorow@soartech.com

Modern technology creates an opportunity for effective, efficient, and pervasive lifelong learning. While lifelong learning can benefit anybody, the potential for pervasive learning is particularly important for skills that require high rates of rehearsal or that involve technology and techniques that change often. Because such skills are often also associated with tasks that are in high demand, it is difficult for learners to find the time to keep their skills up to date.

These challenges are being addressed by the Advanced Distributed Learning (ADL) Initiative. This initiative began in 1999, with the goal of using information technology to improve the state of the art and opportunities for learning outside of classrooms. A major emphasis of ADL is the development of educational technology organized around online information services, integrated into an architecture called the Training and Learning Architecture (TLA). The TLA has several goals, including matching training content to individual learner needs, but also to provide any-time learning in any situation where the learner is willing and able to take advantage of it.

TLA's success hinges in part on the Learning Record Store (LRS), which encodes a language to describe educational material, educational services, skill competencies, individual learner attributes and assessments, and relevant situational descriptors. LRS uses this language to allow *content brokering*. LRS defines a set of metatags that can be used to discover and provide opportunistic mapping of educational materials to particular individuals in particular situations. This content brokering capability is intended to allow TLA to implement a form of *pervasive* or *any-time* learning, in which learners can interact with educational content at any time they are willing and able, instead of restricting learning opportunities to traditional classroom and training settings.

A challenge to the cost effectiveness of TLA is the need to manually evaluate and tag educational resources with the educational situations they may be relevant to. This paper describes an approach to automating this tagging process through an integration of episodic learning and explanation-based learning (EBL), which combines a generalized knowledge base for potential relationships between educational materials, experiences, and outcomes with traces of observed experiences stored in the form of episodic memory traces within a cognitive architecture. The generalized knowledge base is used to explain observed episodes, producing LRS mappings from successful explanations and generating new metatags and relationships from failed explanations. These learning mechanisms will be incorporated into TLA as a set of interacting modules called FLUENT (Fast Learning from Unlabeled Episodes for Next-generation Tailoring).

RESEARCH CHALLENGES

The TLA encapsulates the view that educational materials are service brokers that can match their services to individualized learner needs, situations, and opportunities. The ability to match services to opportunities is key to the architecture. Matching is enabled by the LRS, which defines a language to represent these services and opportunities. A primary challenge in TLA involves how to manage the complex library of LRS tags that need to be associated with each service, learner, and situation, in order to generate the most effective mappings to provide opportunistic, pervasive learning.

An associated challenge is the need to create and manage the LRS tags in efficient ways, without requiring large amounts of data to create the tags and relationships. Because individual learners and learning opportunities are often different from each other in a variety of ways (i.e., each one has properties that make it unique), it would take numerous case-study observations to generate good statistically motivated tags. In the absence of any additional knowledge about tags and relationships, there is no alternative to using slow statistical approaches to learning the

relationships between learning content and learning opportunities. Individual groups of learners taking courses with the same content are often small, perhaps on the order of 30-40 people, especially for specific technical skills (Folsom-Kovarik, Wray, & Hamel, 2013). The small numbers, combined with the uniqueness of each learner, creates a situation where there is a dearth of data to collect for statistical learning. The challenge is to enable fast learning by introducing appropriate amounts and types of educational knowledge into the system, without simply shifting the burden of the effort to engineering this required knowledge.

To address this challenge, we propose a form of knowledge-rich explanation-based learning (EBL). The term EBL has been used to describe a family of different types of learning systems for the past three decades (see, for example, Mitchell, Keller, & Kedar-Cabelli, 1986; and DeJong & Mooney, 1986). Some of this work dictates a very specific definition of what EBL is. In this paper, we use the term EBL in the broad sense, meaning a learning approach that is based on using knowledge to explain (or attempt to explain) observed relationships between learning services and learning opportunities, in order to shortcut the need to collect numerous observations and analyze them statistically. In particular, we exploit a specific hybrid approach to EBL, in which successful explanations tune knowledge by providing search control for better explanations, and the search control in turn focuses a knowledge acquisition method that is able to fill gaps in the explanatory knowledge. This approach to EBL has been successfully demonstrated to generate fast learning with modest knowledge-engineering demands in cognitive modeling and educational applications (Jones & Flesichman, 2001; VanLehn, Jones, & Chi, 1992; VanLehn & Jones, 1993).

EXAMPLE USE CASE

We present an example to illustrate the FLUENT vision, although many more use cases are possible. In this example, an Army Warfighter is using TLA to access an augmented-reality tool for gunnery targeting practice. The gunnery tool records the Warfighter's progress via detailed xAPI events in the LRS. At a certain point, the Warfighter makes gunnery mistakes that indicate a persistent error in applying sine and cosine to compute target geometry. We assume that a quick remediation with a trigonometry tutor tool that exists in TLA will keep the gunnery practice on track. But the gunnery tool does not recognize the need and does not record explicit metatags about the potential need for trigonometry instruction. The gunnery tool only records the fact that some failures occurred on this particular simulated training scenario. Because there is no way to draw the link, the coarse skill estimate for trigonometry remains high in this Warfighter's learning record.

This situation illustrates an opportunity for FLUENT to assist TLA to improve the individual tailoring of its instructional recommendations. In this case, FLUENT first observes the *episode* representing the gunnery tool failures. The present episode is made up of three consecutive failures on this specific scenario, recorded to the LRS by this specific gunnery tool. The episode also contains learning context information such as skill estimates, demographics, and past activity. From this present episode, FLUENT recalls a closely matching episode where a Marine had a high trigonometry skill estimate and also failed three times in a row on this particular module. The recalled episode is prominent in memory because the Marine next chose to review a trigonometry tutor, and immediately became able to pass the gunnery module. The retrieved episode is integrated with an explanation that FLUENT generated, suggesting that the cause of the improved performance was the trigonometry tutorial. As a result of the match, FLUENT infers that the same trigonometry review may help in the current situation. Furthermore, FLUENT is able to offer the past explanation as an easily understood justification for its recommendation.

If trigonometry remediation does not help the Army Warfighter, FLUENT could reduce the strength of retrieval cues for this particular episodic memory and explanation, which could also trigger an attempt to re-explain the past episode. For the purposes of this example, we assume the Warfighter accepts the recommendation and quickly succeeds at gunnery. FLUENT has discovered and used fine-grained knowledge about one of the many interactions between tools in the TLA ecosystem. As a result, FLUENT's recommendations enhance TLA with a new source of data to enable meta-adaptation and improved learning.

DISTRIBUTED AND INDIVIDUALIZED EDUCATIONAL MATERIAL

In the current state of the art (see Figure 1), TLA *tailors* learning content to better help individual learners through macro-adaptation and micro-adaptation of learning content. An example of macro-adaptation (adaptation between learning experiences) might be selecting the next lesson or module to recommend. An example of micro-adaptation (adaptation within a learning experience) might be causing a tool to offer more hints during a test question or increase the speed of a simulation.

The high-level mechanism behind TLA tailoring is suggested by Figure 1. A Persistent Learner Profile in TLA maintains summary skill estimates and the Learning Record Store (LRS) records learner performance in a detailed event stream. However, TLA must summarize the rich details of the LRS event stream to provide a coarse skill estimate for common use by TLA tools – different tests, training systems, and so on. The summarized coarse skill estimate is the primary method for TLA tools to interchange learner information, but it loses much detail compared to the rich store of LRS events. As a result, individual TLA tools require engineering effort in order to interpret LRS events recorded by other tools (Figure 2). Therefore, TLA tools do not often take advantage of all the detail in the LRS.

Metadata and Tagging

One of the key enablers for the TLA vision is the ability for each educational “service provider” to be tagged with indices that indicate when the services of that provider may be useful. Such tags include the particular skills, subskills, or knowledge units that the provider addresses, the types of resources that enable the provider, the types of environments in which the provider is most effective, and other metadata that allows TLA to determine when a particular educational service should be considered.

The current approach to managing this metadata is to tag all educational service providers and educational content with indices defined by the LRS. Unfortunately, this is a time-consuming and error-prone task. It requires careful attention to managing the metadata, and it requires ensuring that the tags for each TLA component continue to match the actual capabilities of the component as it evolves, has content added, receives software updates, etc. Additionally, a manual tagging process makes it difficult to identify synergistic effects between different TLA component services. This in turn produces a situation where the maximum potential of TLA is not being realized, because without sufficient metadata tags, TLA will miss opportunities to provide effective educational materials at effective times for particular learners and situations.

A need exists for automated tailoring that can benefit all TLA tools with the detail in the LRS. *Next-generation tailoring*, using adaptation that carries detailed knowledge across the boundaries between tools, supplements coarse skill estimates with adaptive recommendations inferred from the detailed data in the LRS event stream. This kind of adaptation finds interactions between tools, environmental factors, and other mediating or moderating effects that can be discovered because of the detailed LRS event stream. Putting these together yields improved tailoring that integrates all available tools, that continues to improve over time, and that can be quickly authored or learned from relatively small data.

AUTOMATED MARKUP

The purpose of the FLUENT effort is to create a new capability in

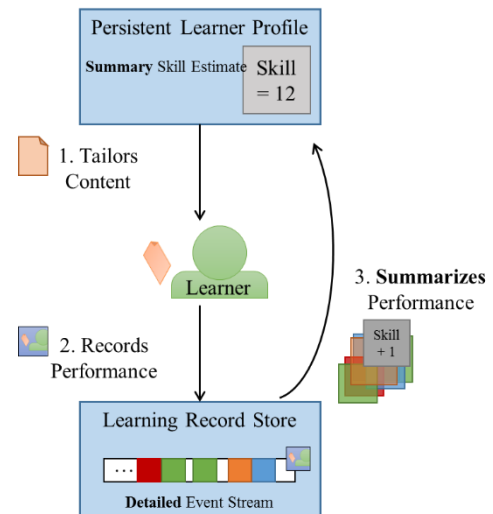


Figure 1. TLA tailors content for a learner's needs and records performance to continue tailoring recommendations. TLA summarizes the detailed LRS event stream to give a coarse skill estimate.




Figure 2. An individual TLA tool such as a tailored tutor or test requires engineering effort to interpret detailed events created by other TLA tools (green and blue blocks). As a result, purposeful interaction between tools is reduced.

the TLA, so that users learn target skills in greater depth, apply skills across multiple training systems, and learn faster, with less time wasted on poorly chosen material. The new capability is *meta-adaptation*, or adaptation across the learning experience that takes full advantage of the multiple systems and contexts available through TLA now and in the future. Our research effort leverages the unique possibilities of the TLA technical environment and improves on the state of the art in TLA's efficacy and efficiency.

The FLUENT research effort will create a meta-adaptation capability that lets TLA direct a single, tailored learning trajectory and unify multiple TLA tools such as individual assessments, simulation environments, intelligent tutoring systems, and transmedia learning opportunities. The research will enhance the persistent learner profile within TLA with the ability to reason across TLA tools, improving TLA recommendations by automating meta-adaptation creation tasks without requiring technical expertise or large amounts of data. TLA will accomplish this through fast learning implemented in its hybrid explanation-based learning algorithms.

Figure 3 illustrates a high-level organization for FLUENT. Several TLA components are sources of information for FLUENT: the Learning Record Store, content repositories and discovery devices such as SCORM (Sharable Content Object Reference Model), and supplemental data sources such as any sensor inputs that might be available and learner characteristics or demographics. FLUENT processing passes through memory segmentation and labeling, memory retrieval, and explanation. The output of explanation is returned to TLA in the form of tailoring recommendations.

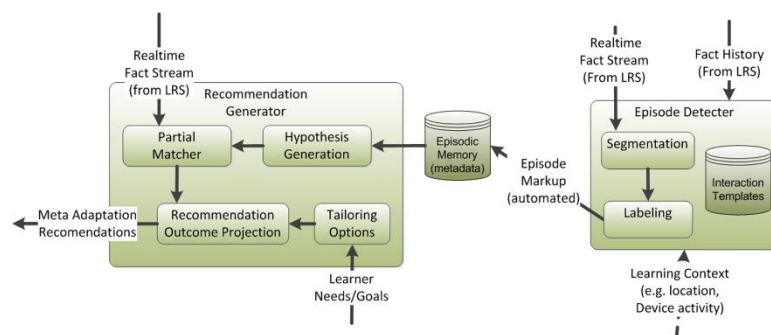


Figure 3. High-level architecture for FLUENT.

Memory labeling is a process that works on anonymized learner histories recorded as activity streams in the LRS. Given a sequence of LRS events, FLUENT must detect and store episodes that it deems memorable. To achieve this, memory labeling compares the learner histories to *interaction templates* that define the expected patterns to look for, such as unexpected changes in proficiency, instances of tools working together that might lead to an interaction effect, or simply temporal overlap. Using these templates as background knowledge

FLUENT generates a hypothesized explanation of how events of certain types may have caused events of other types in the event stream. Using knowledge of which types of events are general causal and which are effects, the explanation process segments the continuous event stream input into discrete episodes and labels the episodes with the learning outcomes it detects, together with explanations of the outcomes. Then the episodes are stored in the *episodic memory*, which enables efficient storage and retrieval with data structures designed for incremental updates as episodes change over time and hierarchical memory elements that can be easily generalized.

If FLUENT is unable to generate an appropriate explanation for a series of events, it attempts to find past explanations for similar series of events and work from them. A sufficiently matching explanation may only need one or two adaptations to apply to the new episode. If an adapted explanation can be found, FLUENT will not only (tentatively) mark the explanation as successful, it will also remember the relational connections it had to create to “patch the gaps” in the retrieved explanation. This new knowledge will be verified or disconfirmed through future recommendations and outcomes. This process allows FLUENT to improve its explanatory knowledge base over time, which in turn will cause FLUENT to improve its recommendations.

Improving markup through successful and failed explanations

This approach to explanation and learning is a hybrid of learning mechanisms for search control and explanation completion, which was first created and robustly investigated by VanLehn and Jones (1993) in the Cascade system. The Cascade work modeled the ability of college students to learn physics knowledge by explaining worked physics example problems. Example problems are analogous to the event streams provided by the LRS. That is, they explicitly represent significant, observable pieces of a causal process, but they are missing the causal knowledge that

links the observable pieces together. Numerous experiments have shown that human learners are able to improve their knowledge and skill by careful explanation generation, linking the observable portions of worked examples. We will adapt the algorithmic implementation of these learning processes from Cascade into the FLUENT research, replacing explanatory physics knowledge with explanatory knowledge of causation in educational domains.

When an individual learner uses TLA, the resulting LRS events become cues for memory retrieval. As the event stream updates over time, it is matched against episodic memory to find similar past episodes. A partial matcher lets FLUENT find episodes that are similar but need not be exactly the same; for example, under certain conditions the location where training takes place might vary without changing the nature of the learning episode. The partial matcher defines the ways episodes may deviate and still be considered relevant. Next, since matches to past episodes are partial, an episode mapping process determines what the past episodes might tell about the current learner situation by drawing analogies to “fill in the blanks” of the parts that differ. Finally, the outcomes associated with the matched episodes are output as *interaction forecasts* to the explanation process.




Figure 4. FLUENT will use an explanatory knowledge base to explain connections between observed episodes describing learning activities, situations, and outcomes. Explanations created for past episodes will provide search control for new explanations. In cases where new explanations cannot be completed, search-control knowledge will focus the knowledge-acquisition method to fill gaps in the explanatory knowledge.

Explanation in the FLUENT system refers to the process of making sense of the current situation, using past explanations for assistance and search control, when possible. Figure 4 shows an abstract view of the search FLUENT will engage in to explain an observed sequence of episodes. Explanation is guided by search-control knowledge by analogy to past successful explanations. In cases where FLUENT does not have sufficient knowledge to explain a learning outcome, the technical challenge is that the system does not necessarily know if the explanation failed because it has insufficient explanatory knowledge, or if it failed because there really is no interesting causal connection in the current analysis. Past explanations assist in focusing the search to areas of the incomplete explanation search that are most likely to represent *knowledge gaps*. Heuristic background knowledge, in the form of *overly general rules* about causation, suggest potential explanatory steps that could complete the explanation. If a plausible explanation is discovered, there is no guarantee that it is actually correct. However, FLUENT can then make recommendations based on the newly acquired plausible explanation, and use feedback from future outcomes to determine whether the plausible explanation is indeed accurate. This hybrid learning process will use interaction forecasts as retrieval cues to find analogical search-control knowledge in episodic memory. The various interaction forecasts are synthesized according to an algorithm that can balance multiple factors such as memory prominence, recency, and closeness of match. In addition, successful explanations attached to stored episodes serve as search-

control knowledge in performing the episode mapping by generating new explanations. Finally, the available tailoring options are evaluated to determine what meta-adaptation FLUENT should recommend.

Episodic learning

A goal of FLUENT is to be able to learn from and generalize few or even single observed learning experiences. This form of *fast learning* is enabled (in part) in human cognition through the using of *episodic memory*. In human experience, episodic memory is commonly described as remembering rather than knowing (Tulving, 1985). A person might *remember* the time when they visited a zoo as a child, with all the sights and smells. From this remembered episode they might *know* facts about animals, if they encoded such semantic knowledge or if they think back to what they experienced after the fact. The remembered episode is located in time and can have positive or negative emotions attached to it that lead a person to seek to recreate the episode or avoid it. The key features of human episodic memory for our purposes are that it is automatically created or stored, it captures events in their context including time, and it can be recollected or retrieved for purposeful consideration. From a functional perspective, episodic memories provide a heuristic function for accessing semantic knowledge that may relate a particular learning situation to learning outcomes. Some modern cognitive architectures provide direct support for computational implementations of episodic memory, to assist in this type of heuristic learning and retrieval (e.g., Anderson & Lebiere, 1998; Laird, 2012).

Episodic memory and learning can be viewed as a form of analogical, or case-based reasoning (e.g., Kolodner, 1993). However, there are some distinctions between how computational models of analogy and episodic memory are typically employed. Cases in case-based reasoning are usually fairly large structures that record long sequences of decision making. While episodic memory is also certainly large, the episodes retrieved at any one time are typically small. This allows systems that use episodic memory to compose different memories into novel chains of inference, when appropriate. In addition, episodic learning is implemented as a pervasive and mostly automatic cognitive process. The filtering and semantic encoding of episodes is implemented through separate cognitive processes. In case-based reasoning or deliberative analogical reasoning, the selection, encoding, and elaboration of the case library is typically a knowledge-intensive process, in an attempt to anticipate the best future retrieval and adaptation of cases. VanLehn and Jones (1993) demonstrated that this finer-grained and semi-automatic approach to storing, retrieving, and adapting episodes is well suited to the hybrid explanation-based learning mechanism that we plan to adapt for FLUENT. In addition, Jones and Langley (2005) demonstrated that this approach to episodic memory provides a qualitative account for several varieties of human skill learning.

Heuristic search through the effective markup space

An important capability of episodic memory is to fluidly redefine which features are contained in each episode, as well as what the boundaries of an episode should be. This enables episodic memory to link events that may be an important part of the learning context, or to discover interactions between a demographic feature and the learning effectiveness of a particular module. The same capability enables FLUENT to make use of unlabeled episodes, or episodes whose label is determined automatically rather than set by a human user.

The goal is for this capability to automate tasks that currently require technical ability from a content creator or an instructor who controls instructional content. *Instructor-mediated* design is a pattern the authors use to help ensure training systems fit practitioner needs without introducing a barrier of technical ability (Folsom-Kovarik et al., 2013). The pattern increases authors' or instructors' control over the content and operation of a training system through focus on the content those users can most easily control. Our approach assists in cases where instructional or content changes require technical expertise or formal specification.

Giving instructors and authors control over how their training systems work can improve system acceptance and effectiveness in real-world use. It can reduce costs, turnaround time for changes, and the errors introduced during communication between end users and developers. However, when adaptive elements are complex, direct control can place technical burdens on users or, more likely, result in incompleteness and incorrectness in the authored system. Task automation in FLUENT helps enable complex meta-adaptation and even give users control, without requiring they have great technical skill.

FLUENT also makes possible increased control of TLA recommendations for nontechnical instructors and other users. Two example use cases are episode insertion and an episodic open learner model. We envision that FLUENT could communicate with end users by describing episodes that it used as the basis of its internal reasoning. For example, FLUENT could support a specific recommendation to a learner by describing a past episode when a similar action had a specific positive outcome, together with the explanation that FLUENT generated to understand the outcome. As another example, we envision individual instructors could control FLUENT recommendations by teaching it episodes that describe which conditions they want to trigger a particular outcome. Rather than requiring the instructor to author many specific cases or understand a programming language to create the generalization explicitly, FLUENT would be able to generalize from a few inserted episodes and present to the user the effects they are creating in an easily understood manner.

The labels associated with retrieved episodes determine (among other things) whether the episode represents a positive experience to be recreated or negative experience to be avoided. Labels in FLUENT can be automatically ascribed to an episode based on an observed impact on learning. The computation requires that we consider large or small improvements differently for learners with a low or high initial skill estimate. Furthermore, FLUENT will not consider the coarse-grained skill estimates as entirely reliable but will instead base its deductions on fine-grained activity records. As a result of all these, the automation of episode labeling will be a topic of research and refinement.

An additional area of research will be the utility of aggregating multiple episodes for improved recommendation. In the current state of the art, individual retrieved episodes match against the current situation context, and the best match determines the outcome of the process in a meta-adaptation recommendation. However, with very large numbers of episodes collected from many individual learner experiences, we anticipate that aggregating episodes for efficiency or for better generalization will improve FLUENT's recommendations.

THE ROAD AHEAD

The anticipated end results of the FLUENT research effort are published algorithms and open-source software that make meta-adaptation possible and practical in an ecosystem of many TLA tools and learning contexts. The research is also anticipated to produce a well-defined and interoperable architecture for that ecosystem that we refine over the course of several spiral iterations, making it increasingly capable until it reaches transition readiness. With this approach, other research programs relating to ADL will be able to take advantage of the architecture we create to enable their own functionality and eventually to reach transition.

In the initial FLUENT research effort, prototype algorithms and software will be produced and shared. The prototype will demonstrate a proof of the concept and will be tested and evaluated without running human subjects, instead using simulated data. In subsequent iterations, the algorithms and software will be made successively more robust, efficient, and capable. The result of the full effort will be validated by instructional experts, evaluated with a human-participants study, and transitioned to potential customers and users. Software will be provided as open source. FLUENT will provide a capable technical system contributing to TLA, as well as a basis for ongoing research in meta-adaptation.

This work will improve the TLA learning experience for learners, instructors, content creators, human performance researchers, and other personnel who use the system to understand how learners are performing and engaging with content.

ACKNOWLEDGEMENT

This material is based upon work supported by the Advanced Distributed Learning (ADL) Initiative under Contract No. W911QY-16-C-0019. Any opinions, findings and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the Advanced Distributed Learning (ADL) Initiative.

REFERENCES

- Anderson, J., & Lebiere, C. (1998). *The atomic components of thought*. Mahwah, NJ: Lawrence Erlbaum.
- DeJong, G., & Mooney, R. (1986). Explanation-based learning: An alternative view. *Machine Learning*, 1(2), 145-176.
- Folsom-Kovarik, J.T., Wray, R.E., & Hamel, L. (2013). Adaptive assessment in an instructor-mediated system. Paper presented at the *16th International Conference on Artificial Intelligence in Education (AIED)*, Memphis, TN.
- Jones, R. M. & Fleischman, E. S. (2001). Cascade explains and informs the utility of fading examples to problems. *Proceedings of the Twenty-Third Annual Conference of the Cognitive Science Society*, 459-464. Hillsdale, NJ: Lawrence Erlbaum.
- Jones, R. M., & Langley, P. (2005). A constrained architecture for learning and problem solving. *Computational Intelligence*, 21, 480–502.
- Kolodner, J.L. (1993). *Case-based learning*: Kluwer Academic Publishers.
- Laird, J. (2012). *The Soar cognitive architecture*. MIT Press.
- Mitchell, T. M., Keller, R. M., & Kedar-Cabelli, S. T. (1986). Explanation-based generalization: A unifying view. *Machine Learning*, 1(1), 47-80.
- VanLehn, K., & Jones, R. M. (1993). Integration of analogical search control and explanation-based learning of correctness. In S. Minton (Ed.), *Machine learning methods for planning*. Los Altos, CA: Morgan Kaufmann.
- VanLehn, K., Jones, R. M., & Chi, M. T. H. (1992). A model of the self-explanation effect. *Journal of the Learning Sciences*, 2, 1–59.