

CHAPTER 3 – Exploring Assessment Mechanisms in the Total Learning Architecture (TLA)

Gregory Goodwin¹, J. T. Folsom-Kovarik², Andy Johnson³, Sae Schatz³, and Robert Sottolare¹

US Army Research Laboratory¹, Soartech², Advanced Distributive Learning Initiative³

Introduction

The focus of this chapter is on the challenges and potential solutions to conducting real-time and long-term assessments of performance, learning, and domain competency in the Total Learning Architecture (TLA). TLA, a distributed learning environment, is being developed by the US Office of the Secretary of Defense to support capabilities for instruction anytime and anywhere. TLA is an evolving set of standardized specifications that enable responsible sharing of essential learning data between applications using common interfaces and data models. The applications that could be part of the TLA environment range from simple desktop applications to immersive simulations to mobile apps, and would serve as either service providers or consumers. Expected services include applications like intelligent tutoring systems (ITSs; e.g., AutoTutor, Cognitive Tutor, or Generalized Intelligent Framework for Tutoring [GIFT]-based tutor), which provide information to other services and consume information from other services.

The TLA is expected to provide services including experience tracking, competency assessment, learner modeling, and content brokering. All of these fundamentally involve learner assessments. Experience tracking (via the experience application programming interface [xAPI]) provides a standard for encoding and storing data about learners' interactions with learning experiences and applications, providing fine-grained evidence that can make assessment precise and timely. TLA will also establish a common way for systems to reference and represent competencies and competency relationships, supporting assessment sharing. Learner models will contain data about assessed mastery of competencies as well as traits, preferences, individual differences, and demographic data. Learner models that are broadly accessible to learning applications will support up-to-date and accurate competency assessment. Content brokering (i.e., recommending future experiences and training) also depends on learner assessments. Content brokering will support just-in-time learning and sequencing of learning events. Competency models will enable content to be tailored to the individual learner's needs.

Training applications like GIFT operating in the TLA environment will both consume learner data available through TLA and provide learner data as they complete training. It will be challenging to insure that all training applications will be able to both obtain necessary learner data from TLA as well as insure that they all output learner measures that can be used by other applications within TLA. This chapter explores some of the challenges of integrating a training application like GIFT into the TLA. This includes discussion of the discovery and development of methods to assess competency based on xAPI statements and recommendations for augmenting xAPI statements to facilitate interoperability among training applications and the TLA through methods such as semantic analysis.

The Evolution of the Current Defense Training Architecture

The rapid growth of the World Wide Web in the early 90's opened new opportunities for delivering computer-based instruction (CBI). As this medium was increasingly used in both higher education and business, learning management systems (LMSs) were created to facilitate the delivery of CBI (Kamel, 2008). By the year 2000, there were over 100 LMSs on the market though the vast majority of the market was dominated by a handful of systems (Falvo & Johnson, 2007).

There is no industry standard for what should or should not be in a LMS; however, there are some core functions that most of them possess. These include the ability to schedule courses and enroll students, support collaborative learning, store and deliver content, support learner evaluation/certifications, manage student records, and support career planning (Hunke & Johnson, 2006; Kamel, 2008).

In 1999, the Army published its requirements for the Army Learning Management System (ALMS) and began production in 2004. By the mid to late 2000's, all services and the Department of Defense (DOD) were looking to adopt LMSs for the delivery and management of online and blended courses (Hunke & Johnson, 2006; Kamel, 2008; Shanley, et al., 2012; Graul, 2012).

Though LMSs are widely used by government, academia, and industry today, they have their limitations. LMSs were developed to support enclosed learning systems on which the business of education depends. Specifically, LMSs control access to, schedule, and deliver curricula and by extension degrees and certifications earned.

Increasingly, it is recognized that training and education take place continuously in peer-to-peer interactions, during the execution of one's job, and in any number of self-development activities. Through the internet, individuals can access how-to videos, blogs, forums, webcasts, etc., to get answers to questions or training on just about any topic. The current generation of LMSs are not equipped to monitor or manage any of these kinds of learning interactions. There is a need for LMSs to evolve.

A learning architecture helps to define the basic functionality of the next-generation integrated learning environment (Kamel, 2008). The DOD Advanced Distributed Laboratory (ADL) has described just such an architecture. It is known as the Training and Learning Architecture (TLA).

TLA Overview

The TLA is a vision of the next-generation integrated learning environment. As mentioned previously, the capabilities of current LMSs define and constrain e-learning and blended learning environments today. LMSs provide standardization and structure to the way learners interact with content (Kamel, 2008). Unfortunately, these attributes work against a learning environment that seeks to deliver training in a way that uniquely adapts to each individual, an environment envisioned in the TLA.

The TLA includes five basic functions: experience tracking, learner profiles, content brokering, and competency networks (Johnson, 2013). To facilitate innovation and make it easier and less costly to maintain and improve, the TLA will be based on non-proprietary, open-source approach to delivering services.

The first function of the TLA, experience tracking, is all about being able to monitor learner activities wherever they may occur. In most LMSs today, measures of learners consist of summary scores, course completion, hours of instruction, and the like. These are essentially the same measures that have been recorded since formal education began. With the massive growth of online content, learners often spend a considerable amount of time outside of the LMS before, during, and after a course to prepare for the course, improve their understanding, refresh their training, seek tutoring or support on challenging topics, and even to share their own knowledge with others. By monitoring these kinds of behaviors, quite a few interesting outcomes are suddenly possible. For example, it would be possible to determine what course content is particularly challenging for students. It would also be possible to then automatically tailor training for those individuals. One could also determine what material is retained well and what seems to be most easily forgotten.

The ADL has created a specification for experience tracking called the xAPI, which is a non-proprietary specification for tracking and storing experiences across learning platforms (e.g., simulators, virtual worlds, web content, mobile devices, games, and observer-based measures). The learner's activity stream (a series of xAPI statements) is stored in a JavaScript Object Notation (JSON) database called a learning record store (LRS). Each statement includes the actor, verb, object, and optional information about results, context, etc. The power of xAPI is that it can record human performance at both the micro and macro levels. For example, a single keypress or behavior can be an xAPI statement, but it can also be a statement that the individual completed a course, received a certification, etc.

The second capability of the TLA is learner profiles. A learner profile is a map of the learner's background, experience, knowledge, and traits. It goes beyond a simple student record held in current LMSs, which primarily include courses completed, grade point average, etc. The learner profile in the TLA would include any information about the learner that may impact how and what training should be provided to the student. For example, cognitive abilities like intelligence, reading speed, and reading level. It also includes prior experience and knowledge in different domains like math or human psychology, small unit tactics, or how to operate a specific system. This profile is clearly dependent on the measurement of learners and so a LRS is a key enabler of the development of learner profiles. A learner profile is more than just a collection of measures. For example, over time, learners forget and lose skills when they don't have an opportunity to practice. A learner profile needs to take into account skill and knowledge retention over time if it is to be accurate.

The third capability of the TLA is competency networks. Each branch of the military has defined competencies that they need in their respective workforces. Some competencies may be common across all service members while others may be very specific to a single specialty. Competencies are defined by organizations as are the training and measures needed to develop and maintain those competencies. Competencies are complex and are usually developed over years through a combination of formal education, mentoring, job assignments, peer interactions, and more. A competency network is a way of representing those types of experiences and how they feed into the development of various competencies. By mapping the learner profile to the competency network, it will be possible to determine the competencies of the individual. The TLA will provide each service with a way to represent the training and measurement of their respective competencies.

The final TLA capability is content brokering. Content brokering has to do with tailoring the delivery of content to individual learners. Content brokering is dependent on knowing what the learner has done (via experience tracking), what the learner currently knows (via the learner profile), and what the learner needs to be provided with (via the competency network). Content brokering uses content registries and repositories to make training recommendations. Because so much content is now available outside of the walls of most LMSs, tools that enable machines to find and understand learning content automatically will play a big role in content brokering. These tools will do more than simply perform a semantic analysis, they will consider who uses them, how they are rated, and which ones seem to provide the biggest benefit to learners.

Current Learner Assessments in GIFT

Although GIFT was developed prior to the TLA, the assessment mechanisms within GIFT are compatible with various aspects of the TLA. While the TLA may be concerned with experience tracking across a range of formal and informal learning activities, GIFT is also using evidence-based assessment techniques to determine the current and projected performance of each learner and team based on their historical performance in a domain and based on behavioral markers and physiological measures, which indicate cognitive, affective, and physical states of learners.

In many ways, the assessments and services proposed for the TLA mirror those in GIFT. The key difference being that where the TLA is concerned with the progression of an individual at the curricular and career levels, GIFT is concerned with the progression of a learner at the course or lesson level. To do this, GIFT uses its own versions of experience tracking, learner profiles, content brokering, and competency networks to know where the learner's knowledge and skill are at each point in the course (learner module), where the learner needs to go (domain module), and what methods and content (pedagogical and domain modules) to employ to help the learner get there (see table 1). Figure 1 illustrates how GIFT uses these assessments and processes to tutor an individual using what is known as the learner effect model (Sottolare, Brawner, Goldberg & Holden, 2012; Sottolare, Goldberg, Brawner & Holden, 2012; Sottolare, Brawner, Sinatra & Johnston, 2017).

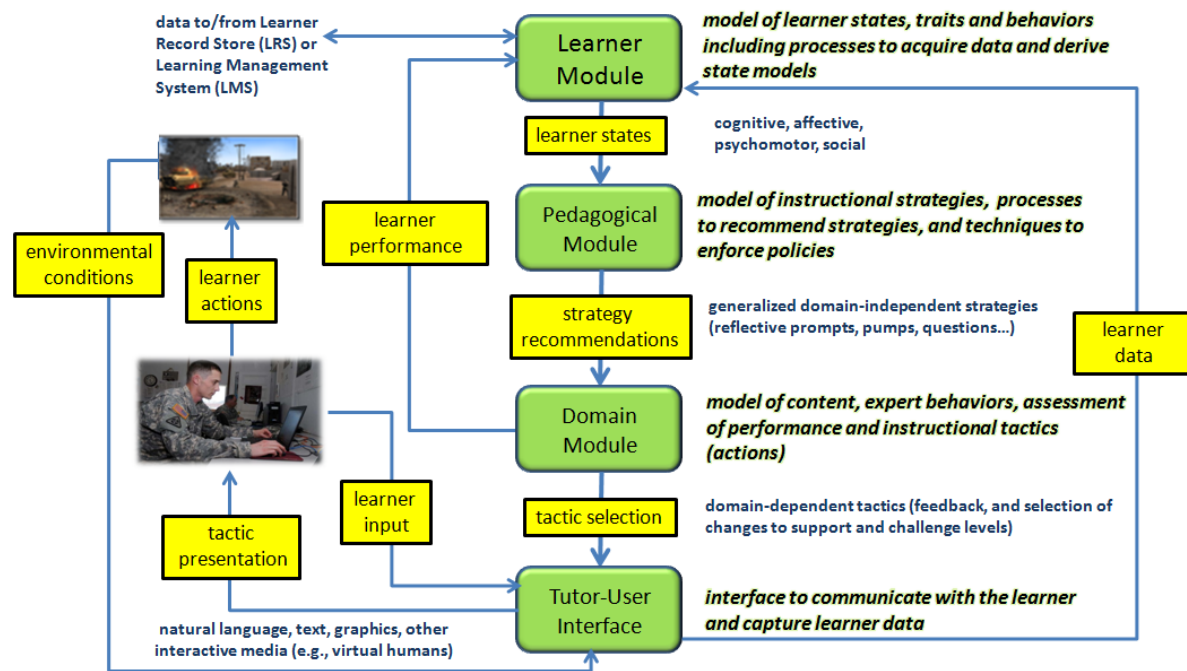


Figure 1. Instructional processes and architecture of GIFT.

GIFT is a framework that modularizes the common components of ITSs. These components include a learner module, an instructional or tutor module, a domain module, and a tutor-user interface. One of the main motivations for creating this framework was to lower the cost and labor needed to create ITSs by facilitating reuse of components and by simplifying the authoring process (Sottolare, Goldberg, Brawner & Holden, 2012).

The learner module represents the current state of the learner and also tries to predict future states; however, the learner module is not typically conducting assessments in GIFT. Assessments are done by the domain and sensor modules which pass their assessments to the learner module. The learner module uses those assessments as well as demographic and historical data about the learner to provide a classification of the learner's cognitive, affective, psychomotor, and competency states.

Right out of the box, GIFT can collect student measures in several ways. First, GIFT uses surveys for soliciting responses from learners. Survey questions can be used to collect demographic information or administer standard psychometric questionnaires like the NASA Task Load Index. Survey questions can also be used to test comprehension or knowledge either pre, post, or during training. These assessments of

comprehension are done in the domain module and then passed to the learner module for management/maintenance of learner state representation (Sottolare, Brawner, Goldberg & Holden, 2012; Sottolare, Goldberg, Brawner & Holden, 2012).

GIFT can also log all student keyboard and mouse actions that occur when learners interact directly with the tutor-user interface. When GIFT passes the learner to another training application, like a simulator, GIFT can collect learner interactions with those applications via an application programming interface (API). In fact, GIFT comes with ready-made gateway modules for some popular applications like Microsoft PowerPoint and VBS2. Once again, when these interactions are used to assess learner's understanding of concepts, knowledge, or skill by the domain module, those assessments are passed on to the learner module to update and maintain a representation of the learner state.

Finally, GIFT provides a standardized means for collecting data from a variety of commercial sensors that record physiological (e.g., electroencephalogram, electrocardiogram, electromyogram) and behavioral (e.g., eye-tracking, Xbox Kinect, force transducers, accelerometers) data. Sensor data is filtered, segmented, and/or extracted by GIFT's sensor module to provide a basic assessment of the raw sensor data. This basic assessment can then be used by the domain module to assess the learner's state, and this state assessment is then passed to the learner module. Raw sensor data are also logged for post-training analysis.

A simple framework for understanding how different types of assessments are used by GIFT is depicted in Figure 2. This framework divides assessments across two dimensions. First, assessments are divided into pre-training and in-training categories and second, they are divided into domain dependent or independent categories (Goodwin, et al., 2015).




Figure 2. Conceptual framework for assessments.

As can be seen in this figure, assessments that are completed prior to training primarily influence larger macroadaptive strategies that GIFT would employ. For example, if training on the operation of a tactical radio was being delivered to a medic, examples and exercises would be relevant to the medic mission as opposed to the infantry mission and so forth. Assessments done in training would primarily influence micro-adaptive tactics such as whether to increase the difficulty of the instruction, whether to provide hints or feedback, etc.

Because GIFT and the TLA both engage in learner assessments, content brokering, learner modeling, and experience tracking, it is necessary to determine how these two systems would work together in a coordinated fashion. These topics are discussed in the following sections.

Integrating GIFT with the TLA: Experience Tracking

The xAPI enables a revolutionary new capability for *experience tracking* that forms the core of data sharing and interoperation in the TLA. Experience tracking refers to encoding and storing in a sharable manner some fine-grained information about actions a learner takes or events that impact a learner during a simulation, across a college course, or simply while consuming a video or text. In turn, understanding what individual learners experience from moment to moment can provide the grist to choose content that relates to the experienced context, promptly identify and respond to key moments, or proactively create conditions for learning. As such, the experience tracking approach colors the TLA perspective on competency management, learner modeling, and content brokering.

General challenges associated with experience tracking include scope, tractability, and privacy. Scope refers to the questions surrounding what types of data should be included in experience tracking, what level or granularity, what frequency of update, and so on. While recording more data certainly creates more raw material to work with, indiscriminate capture increases the difficulties related to tractability and privacy. Tractability refers to issues introduced with storage size, computation speed, network latency, human attention if required, and any other resource that can be strained by large amounts of data. Experience tracking can be made more tractable with selectivity about what experiences are useful to store or using batching methods. Privacy concerns relate to the storage of experiences that reasonably impact learning but might not be desirable to share publicly, such as arrival at a real-world location or failure on a high-stakes test. Methods being explored to protect privacy include identifying opportunities to store categories of information rather than unneeded specifics and introducing access control that limits who can read recorded experiences.

Specific to adaptive training systems, there are exciting instructional design challenges surrounding how adaptation that responds to a different training system can improve learning. What is acceptable or helpful in terms of adapting learning to respond to a past experience in another system, or even shifting control over the learning experience between systems? At the technical level, an important challenge is the need for one adaptive system to understand the experiences recorded by another training system. Maximally effective adaptation requires knowing the learner's past experiences, which are stored in recorded xAPI messages. However, it can be easy to share this valuable knowledge about experiences but hard to understand what someone else has shared. Some reasons for this are varying interpretation of the standard xAPI vocabulary and the ease of creating new vocabulary to encompass every situation. The semantics of a message might also vary subtly with context, such as the sequence of past experiences that led up to it or even which version of the software wrote it. As xAPI finds widespread acceptance, how to use each part of the vocabulary is being defined and agreed upon in communities of practice. However, new users and new use cases will continue to introduce variations.

GIFT is capable of writing and reading xAPI messages. Examples include messages for GIFT course completion as well as for responses to tests or surveys. The specific xAPI statements generated by GIFT are currently determined by course authors. Because there are very few consumers of xAPI data streams, course authors have little guidance as to what kinds of experience statements should be produced. For the TLA to make use of the experience reporting capability of GIFT or other adaptive training systems, it will be necessary for the TLA to be able to make requests to those systems. Additionally, it will be helpful for the ADL to provide some guidance on what kinds of experience streams it needs. For example, ADL is working to determine the balance between recording domain-specific data versus measures that would have more

general, domain-independent uses. Another challenge is developing standardized ways of reporting measures within specific training domains. There are always a variety of ways to report skill levels, knowledge, and experiences relevant to specific domains. To make it easier to compare and aggregate such measures across delivery platforms, standards are needed. Currently, the ADL is using communities of practice to develop standards for reporting activities, context, and processes so that a common vocabulary can be used by all practitioners.

Integrating GIFT with the TLA: Competency Management

The TLA uses common information about competencies that can help to coordinate teaching and training across systems. Competencies in the TLA refer broadly to skills, knowledge, abilities, and other targets of learning. TLA competency information includes human-readable descriptions, machine-readable definitions, and relationships between competencies such as levels and prerequisites. Clearly defining competencies helps make learner assessments portable across training systems.

Typical challenges surrounding competency management include agreeing on the meaning of a competency and who is allowed to define the meaning. Many different interpretations and weightings of component factors are valid for different uses. For example, a competency reflecting knowledge of a cyber intrusion tactic for the general population might require knowing that it exists or knowing rules to avoid the attack. For more specialized learners, related but different competencies might need to be defined that require understanding of underlying principles or even ability to carry out the tactic. In addition, competencies may be associated with different learner performance across populations, which complicates leveling and norming. As such, competency definitions may be created that are so broad they become vacuous or so specific they are not amenable to reuse across the TLA. Human-readable and machine-readable competency definitions must be carefully aligned. Also, careful judgment may be needed in maintaining the competencies because any change to a definition might disrupt another system that uses the same competency, while excessive versioning or division of the competencies might break links between systems that should survive that change.

Within adaptive training systems that use managed competencies, different structures may be needed to make the competencies applicable to different models of learning. Systems may require different relationships to be defined or interpret relationships differently (e.g., creating correlations between skills as historical data or for the purpose of machine learning and inference). They may need specialized data about relating competencies to specific content (e.g., which skills contribute to good performance in a simulator) or about content characteristics (e.g., what modes of presentation are appropriate). So, a key challenge is selecting the data to store for each competency and then collecting that data. The large number of competencies and all their details are unlikely to be encoded completely or correctly, or they may become out of date in a process called *concept drift*. Because of these challenges, it is likely that adaptive training systems require a way to impute or learn and refine competencies.

In considering the prospect of GIFT making use of competencies that are defined and managed in a TLA-enabled location outside of GIFT, policy and technical challenges arise. Are GIFT competencies defined by a single central body with authority or are unit-specific differences expected? Are competencies defined by one military group allowed to be reused, reinterpreted, or even changed by outsiders? Finally, it is known that the existence and definitions of some competencies may not be published or shared, for example, when they are classified. The TLA supports such an environment by allowing multiple components to manage competencies separately, so that for example an unclassified component might manage most competencies while a component on a secured system could manage classified competencies. However, such a separation will introduce technical challenge surrounding replication and coordination of data for processes that need secure access.

Integrating GIFT with the TLA: Learner Profile

Learner models contain data about assessed proficiency or mastery of competencies as well as traits, demographic data, preferences, learning goals, and transient states of the learner.

Broad challenges of learner models include those that are well discussed in the literature as well as interesting new challenges introduced by the TLA focus on fine-grained experience tracking. First, there is a need to create unified measures to express each learner state or trait, and translate between measures. Since it is unlikely that every different training system will use the same scales or metrics to describe learner characteristics, methods must be created to let different measure interact. Does a four out of five on one scale equal an eight out of ten on another scale? Where does “meets expectations” fall on either scale? Standardization in the TLA aims to let systems answer some such questions automatically. Next, making sense of learner experience records requires rolling up raw data from one or many experiences into actionable information. This requires understanding different models of change, including differing stages to express how learning takes place and differing models of skill or knowledge decay. Finally, variations in processing across training systems should reasonably lead to varying levels of trust in learner model contents that are shared from other training systems. It should be possible to evaluate information in a learner model based on its recency, authority, and so on. Therefore, systems participating in the TLA require architectural capabilities to tie each learner competency estimate back to the components and the evidence that produced it.

An important challenge specific to adaptive training systems is the *cold start* problem, where computer systems must take time to identify learner traits necessary for adaptation. The TLA is specifically designed to mitigate cold start difficulties by sharing the needed information with participating systems. A second challenge is the often negative perception of assessment associated with high-stakes testing. However, the value of the TLA does not require high-stakes tests or formal assessments. The TLA may support more formative assessment that acts to address some concerns about time spent on testing. Finally, there exist a challenge surrounding learner models that may contain personally identifying information or otherwise sensitive information such as records of failures on important training events. As described previously, personal control over private information or intelligent filtering such as sharing the lack of a success to date, rather than a definitive failure, could provide approaches to address such challenges without reducing the important value of a shared learner model.

In the GIFT framework, modeling learner characteristics outside of the established GIFT model suggests a possible opportunity for automated intake of new learner model definitions. However, the learner characteristics in GIFT are by design very general to all training domains. As a result, it may be that design changes will be needed before GIFT can incorporate new shared characteristics that let GIFT take advantage of the learning context other systems know. For example, it may be desirable to introduce a layer of interpretation that can translate learner model characteristics for GIFT instead of, or as an intermediate step toward, adding the characteristics as first-class members in the GIFT learner module.

Integrating GIFT with the TLA: Content Brokering

Content brokering refers to recommending future experiences and training based on learner goals, characteristics, and assessments.

Similar to competency management, managing characteristics of content is difficult to do at scale. Metadata and paradata describing each piece of content needs to be authored and stored. As we have argued, such automated collection and maintenance of such data is likely to play an important role in TLA-enabled train-

ing systems. Automation will help ensure that data is accurate and up to date, which is of increased importance when different systems need to coordinate to understand and recommend or broker learner experiences.

An interesting research question surrounds the assembly of a unified learning experience out of the atoms provided by different adaptive systems. Content brokering is likely to take advantage of, and be challenged by, new modes of learning such as second-screening or switching between systems when an experience is incomplete, as opposed to linear completion of a single recommended learning pathway. Unified language and surface presentation of instructional content is likely to require precise content descriptions that are not currently available. Some even argue that unified user interfaces, iconography, or fonts might be needed to avoid distraction and extraneous cognitive load caused by switching between training systems.

The granularity of content to be brokered is an interesting question. While humans are able to make “close enough” fits between learner needs and content, or quickly identify atoms of content to a very fine-grained level, the same can be difficult for automated systems. For example a human may find it appropriate to direct a learner, “read the first two pages, the last three are not related to our work now.” Adaptive training systems need a way to link deep into content and markup the different competencies related to fractional portions of content that is experienced. This is an interesting challenge because content portions that are viewed or need to be viewed may be continuous and have graduated effects rather than discrete. As an example, watching the first 60 seconds of an instructional video may have different effects from watching the first 90 seconds. The video may or may not be possible to break into discrete segments that have beginning and end points a machine can identify. In a training simulation, it may be desirable for a learner to experience one particular path or a few out of thousands of possible paths. Content brokering needs a method to understand these impacts and select, suggest, or influence certain paths at a fine grain that expert human instructors can achieve.

GIFT may provide a central method for the TLA to start and control content in adaptive systems. However, this GIFT approach to content brokering currently requires a GIFT-specific interoperability program that costs developer time to create. The GIFT framework also assumes that content it brokers within its framework will give GIFT certain types of control over the learning experience, such as allowing GIFT to terminate running content. This kind of centralized control makes sense in GIFT and is used in useful ways such as switching to another activity without leaving too many windows open. However, not every system that participates in the TLA will find it possible to use this structure. A minimal set of direct interoperability controls should be identified and possibly modularized into a reusable content brokering component that could reduce the developer effort to participate in GIFT content brokering.

Conclusions and Recommendations

The framework being proposed here for the integration of GIFT (or any adaptive training system) and the TLA sees the TLA as handling management of training and education at the course and above levels while a system like GIFT handles management of training at the course and below. For example, the TLA might determine that a learner needs some block of training in GIFT to develop or maintain a particular competency and so it would in essence bring the learner to the GIFT classroom for training. As part of this handoff, the TLA would provide GIFT with a record of relevant prior experiences and training of this learner. At the conclusion of training, GIFT would return the learner to the TLA and would update the learner profile accordingly.

As discussed previously, there are some challenges that need to be addressed in order for this framework to become reality as shown in Figure 3. First, for the TLA to deliver the learner to the GIFT classroom, the

TLA needs to be able to know about the training available in GIFT and how it maps into various competencies. To the degree that there are not clearly defined methods for identifying competencies and defining competency networks, it is not clear what the link between a GIFT course and a competency network would be. Even assuming such competency networks exist, there is still the challenge of determining how to evaluate the content of a GIFT course so that it could be mapped to specific nodes on the network.

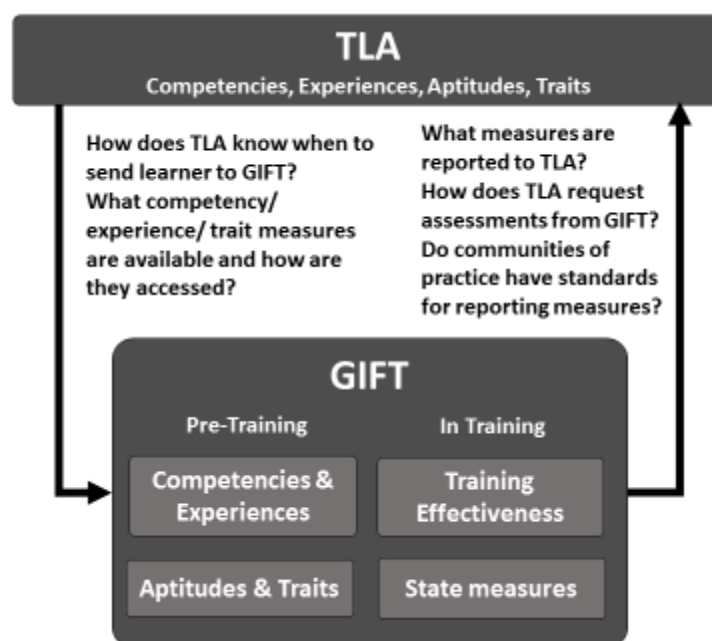


Figure 3. Challenges for integrating GIFT and TLA.

In terms of experience tracking, GIFT can generate xAPI statements that can be fed into an LRS. So, on the surface, it would appear that GIFT is already largely TLA-compliant. The biggest challenges that exist with regard to experience tracking are not specific to GIFT, but are general challenges of experience tracking. Specifically, standards need to be developed for reporting of activities within specific training domains to facilitate comparisons of measures across applications. Other issues involve determining the granularity of learner activities that should be reported and the identification of critical contextual information that is needed for proper interpretation of activities. As communities of interest develop conventions and standards for experience tracking, GIFT will need to comply with those standards, but that shouldn't be a difficult to do.

In conclusion, it is clear that GIFT and other adaptive training systems must be part of an integrated learning environment like that envisioned in the TLA to realize their full potential. In this environment, GIFT will know about learners as soon as they enter the GIFT classroom, maximizing GIFT's training efficiency by eliminating the need for it to spend time interrogating the learner before delivering its training. This environment will also enable GIFT to adapt training across training venues or modalities. For example, suppose a course includes a sequence of classroom training followed by simulation training followed by live training. If learner's activity streams are being recorded in the classroom, GIFT would be able to adapt the simulator session based on their classroom performance. Similarly, if performance is tracked during live training, GIFT could both use that data to evaluate the effectiveness of the training it delivered in the simulator and even to recommend additional remedial simulator training when needed.

References

- Falvo, D.A. & Johnson, B.F., (2007). The use of learning management systems in the United States. *TechTrends*, 51, 40–45.
- Goodwin, G., Johnston, J., Sottolare, R., Brawner, K., Sinatra, A., Grasser, A. (2015). Individual Learner and Team Modeling for Adaptive Training and Education in Support of the US Army Learning Model: Research Outline (Special Report 0336). US Army Research Laboratory: Aberdeen Proving Ground, MD.
- Graul, M. (2012). Framework for adaptive learning Content management and delivery (FALCON – Report AFRL-RH-WP-TR-2012-0192). Air Force Research Laboratory, Wright Patterson Air Force Base: Dayton, OH.
- Hunke W. & Johnson, R. (2006). Learning Management Systems White Paper (Report TAT-09429). Information Assurance Technology Analysis Center: Herndon, VA.
- Johnson, A. (2013). The Training and Learning Architecture: Meeting the Needs of the Next Generation of SCORM. Slides from a Webinar. Downloaded from: <http://adlnet.gov/wp-content/uploads/2013/02/TLAWebinarFeb2013-1.pdf> on 4 Jan, 2017.
- Kamel, M.N. (2008). Learning Management Systems: Practical Considerations for the Selection and Implementation of an E-learning Platform for the Navy (Report NPS-GSBPP-08-011. Naval Postgraduate School, Graduate School of Business and Public Policy. Monterey, CA.
- Shanley, M.G., Crowley, J.C., Lewis, M.W., Straus, S.G., Leuschner, K.J. & Coombs, J. (2012). Making improvements to the Army distributed learning program (Report MG1016). Arroyo Center, Santa Monica, CA: Rand Corporation.
- Sottolare, R., Brawner, K., Goldberg, B. & Holden, H. (2012). The Generalized Intelligent Framework for Tutoring (GIFT). US Army Research Laboratory.
- Sottolare, R., Brawner, K., Sinatra, A. & Johnston, J. (2017). An Updated Concept for a Generalized Intelligent Framework for Tutoring (GIFT). Orlando, FL: *US Army Research Laboratory*. May 2017. DOI: 10.13140/RG.2.2.12941.54244.
- Sottolare, R., Goldberg, B., Brawner, K. & Holden, H. (2012). A modular framework to support the authoring and assessment of adaptive computer-based tutoring systems (CBTS). In Proceedings of the *Interservice/Industry Training Simulation & Education Conference*, Orlando, FL, December 2012.