Interoperability of Learning Systems

Sharing User Model Data between Diverse Applications via a Secure Digital Learning Cloud

Summary

Intelligent Tutoring Systems (ITS) developers and other learning systems professionals devote a great deal of time and effort to evaluating the cognitive abilities and personality attributes of learners who access their systems. The information gained about a particular learner is often stored in a user model that can be referenced later to help adapt the learning system to the specific needs and abilities of each learner.

While this adaptive process can have a major impact on the effectiveness of individual learning systems, the information gained about individual learners is not normally accessible by any other learning systems. Consequently, adaptability of any given system is often limited by practical constraints on how much assessment can be completed in a single application. This paper explains how MARi™, which launched on July 4 in private beta, allows virtually any software application to easily share observed user attributes with other systems to help create a better, more personalized learning environment for each learner, while protecting contributing organizations' investment and the intellectual property of the data they share with others.

Introduction

The MARi™ Digital Learning Cloud exists for the specific purpose of helping the learning systems community share data between diverse applications in a manner that is both easy to use and flexible enough to meet a wide variety of software adaptability needs. Adaptability of learning software is often a key factor in the overall effectiveness of the software in regard to helping learners understand the material being taught in a reasonable amount of time. Many learning applications include introductory modules that perform some type of assessment about each learner before proceeding with the main material to be learned. The information gained in these introductory modules can then be used to tailor the presentation, contextual cues, and difficulty level to help ensure that each learner receives the proper material to help the learner grasp the material most effectively.

Unfortunately, it is often not practical to provide extensive introductory assessments that include valuable, but tangential, assessments of variables like cultural background, effective reading level, or personality traits. However, there are existing assessments that do a great job of evaluating those very attributes in the marketplace today. So why not use a common, secure learning cloud to keep track of those attributes as they are observed in various assessment tools and then share those observations with other learning systems that can use them to improve each learner's overall educational experience.

With the help of MARi™, an Intelligent Tutoring System developer can create a biology tutor that is sensitive to cultural differences or different reading comprehension levels without having to be an expert in the evaluation of those attributes. The ITS developer can simply plug into MARi™ and retrieve the results of prior assessments for each learner when the learner launches the ITS. Now, the ITS developer can be more certain that the learner's ability to understand, or misunderstand, presented material on life sciences is actually based on the learner's grasp of the underlying principles that govern how living things function, and not due to language comprehension or cultural differences.

Finally, once the biology ITS finishes teaching the learner, the ITS can generally draw some conclusions about each learner's proficiency in biology. Those conclusions and direct observations can be stored in MARi™ to help some other learning application more effectively adapt to the learner. The cycle of retrieving personalized adaptation data, adapting to each learner, and reporting observations continues over time with each consuming learning system also being a contributor of assessments and observations so that the learners' digital learning clouds continue to be enhanced and become more and more accurate.

How Does MARi Work?

MARi™ is composed of a small set of API entry points that allow registered applications to read and write data from digital learning clouds. Each application selects two sets of attributes that define which attributes it will read and which attributes it will write into each learner's digital learning cloud. These attributes are selected from within the MARi™ Open Personal Attribute System (OPAS), which maintains a complete list of all personal attributes that are available for use.

OPAS also allows researchers and developers to add new attributes to MARi™ through a curated process that ensures that all new attribute requests are indeed useful and are not already represented by an existing attribute entry. There is no set standard or restriction on attributes, but a set of guidelines and procedures have been established to keep the OPAS effective.

Once the read and write contracts for attribute sets have been established for a particular application, the MARi™ API can be used to exchange data with MARi™. All MARi™ API entry points are RESTful API calls that utilize HTTP GET and POST protocols. So the API can be used with virtually any programming language and any environment that has access to the Internet. Data sent and received by MARi™ is packaged as standard JSON objects. However, to ensure that each learner's data and privacy is protected, participating applications are required to use OAuth authentication to access MARi™ data.

The diagram on the next page shows a conceptual model of the way MARi™ works. In the center of the diagram, the MARi™ data is represented as a series of data layers with each layer containing data from a different source application or a predictive model. The top layer represents a view of the shared data that can be accessed by multiple applications. There can in fact be many different shared views of the data, based on permission settings from both the contributing sources of the data and the individual learner that the data describes. The actual decisions about who to share the data with are dependent on the specific purpose of the captured data. Most of the time, the individual learner is in complete control over who sees the shared data, but under certain circumstances the contributing sources of the data can specify that their data can be shared with a specific set of applications that cumulatively constitute a suite of interacting applications.

On the left side of the diagram, different learning systems (either from the same developer or from different developers) begin the process by contacting the OAuth server to acquire an access token. That token is required on all subsequent attempts to access the learner's MARi™ data. Then, each learning system can access the shared data that is available to all the applications in the suite, or they can access private data that is only available to their specific application. This tiered approach to data visibility allows learning systems to store as much learner specific data as they wish, while choosing to only share specific pieces with others.

In addition to gaining access to data from other applications within their own suite of applications, each learning system has the ability to access data that has been provided by other assessment tools that are not specifically part of their suite of applications. Over the course of their lifetime, a learner may participate in many assessments of their cognitive skills and personality traits. They may also provide valuable information about different aspects of their lives via self-reporting applications; and of course, they may play serious games that capture interesting observations about them. All of this data is available to learning system developers (assuming the individual learner agrees to allow it) to help create adaptive systems.

The overarching goal of MARi™ is to provide a secure data management facility that encourages developers to share useful data between applications while providing the compartmentalization of sensitive data to ensure that the source of each piece of data can control the access to the data it provides.

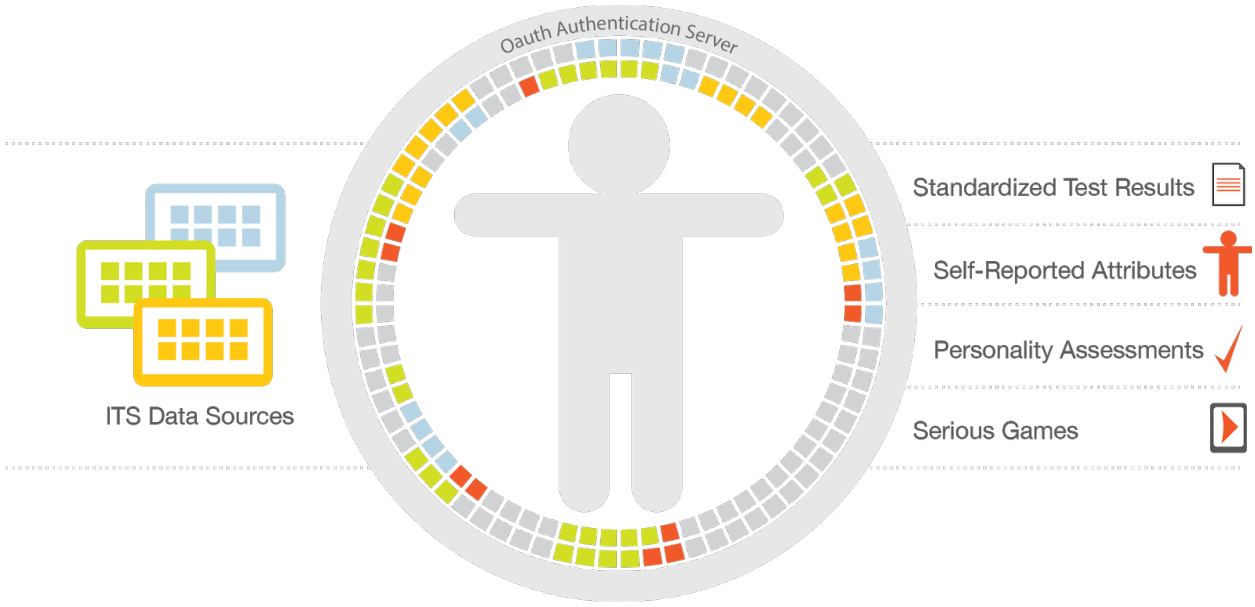

MARi Conceptual Model

Integrating MARi with ASSISTments

The MARi™ platform provides support for both question-level performance tracking and cumulative student capability tracking by presenting multiple data views of each student's personal learning cloud. As students use ASSISTments to complete their assignments, the results of each question can be logged into the student's MARi account for later analysis. Once the question-level data has been logged, it remains available to ASSISTments, or any other application designated by the ASSISTments team, for use in research, product enhancements, or real-time adaptation of learning software to improve each student's learning experiences.

In addition to question-level logging, ASSISTments can also record observations about the student's overall understanding of the subject material by considering the results of multiple questions, the number of hints required, and time to complete question sets. This overall capability information is what MARi calls "personal attributes" because they describe some piece of information about the student that might be of interest to ASSISTments or other learning applications in the future. All attributes are time stamped and fitted with a tailored expected memory decay function to ensure that consuming applications are aware of how fresh the data is. ASSISTment ARRS-based reassessments are an ideal method for periodically checking the retention of student knowledge and refreshing that knowledge at appropriate times.

Furthermore, statistical models and machine learning techniques can be applied to both the question-level data and other personal attributes to predict latent knowledge that the student is likely to possess, based on the capabilities already observed by any MARi-enabled application. Since MARi™ is an open architecture, statistical models used to predict student knowledge may be created by any registered researcher. The results of the models are stored alongside all the directly observed knowledge attributes, but the underlying logic of the model is protected and only available for review by the organization that developed the model. The results of each model are fully attributed to the source of the model, and tagged to indicate that the value of the attribute has been predicted rather than directly observed. The modeled knowledge attributes can then be shared, by the developing researchers, with any other registered application or research group that might find the predicted attributes useful in enhancing the user's learning experience.

While MARi™ strives to provide students with as much visibility into their own data as possible, MARi™ also provides mechanisms to ensure that developers and researchers who generate attribute data can always access the data they generate and share that data with other entities. MARi™ also respects the intellectual property concerns of organizations that generate student data and therefore offers mechanisms to limit or eliminate student visibility into captured data when appropriate. Sponsoring organizations like school districts, research teams, and corporate development teams can be assured that the data they capture will always be available to them and their partners, along with a wide array of other personal knowledge attributes made available from the MARi™ community at large.

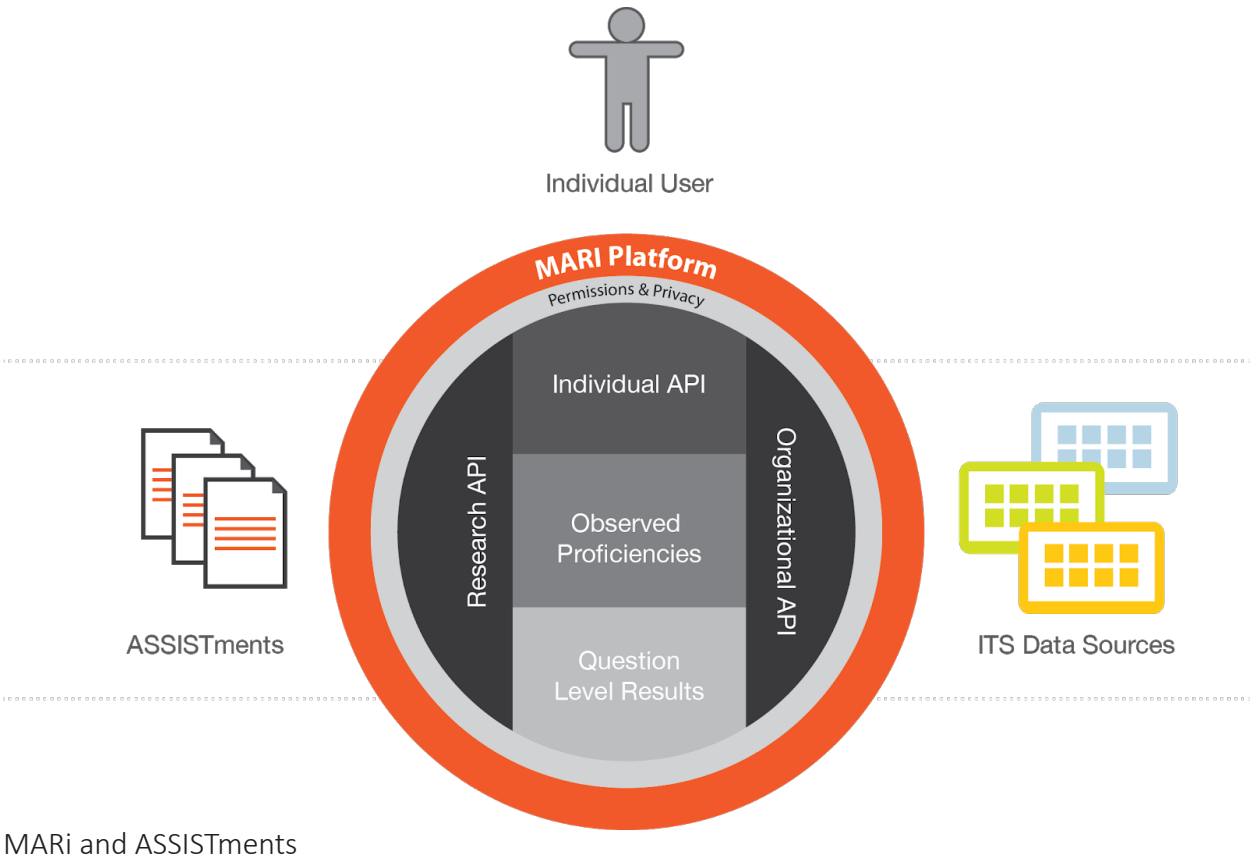

MARi and ASSISTments