

A “Stories-as-a-Service” Architecture for Enriching Training Simulations

Benjamin Bell, Ph.D., Debbie Brown, Ewald Enzinger, Ph.D.

Eduworks Corporation

Corvallis, OR

Benjamin.Bell@Eduworks.com, Debbie.Brown@Eduworks.com, Ewald.Enzinger@eduworks.com

ABSTRACT

The widespread use of simulation has transformed the training landscape with realistic and immersive depictions of a broad range of task environments. Training outcomes have been further enhanced through blending simulation with intelligent tutoring systems. A powerful and widely practiced learning resource, though, is under-represented in training-based simulation: first-person narrative accounts of experienced practitioners who encountered whatever specific challenges a learner is facing in a simulated exercise. Such accounts can be vivid, memorable, and credible. Making these stories available to learning environments when relevant and needed, though, is a non-trivial enterprise that requires infrastructure for collecting, semantically tagging, and distributing story content.

This paper introduces Augmenting Next-Generation Learning Environments with Stories (ANGLES). We present an architecture for content collection, tagging and distribution. ANGLES enables content to be collected from practitioners recording personal experiences from commodity devices (tablets, smartphones) through a mobile app that contributors use to capture, describe and upload their stories. We also discuss the algorithms that automatically transcribe and semantically tag the collected content. Through standards-driven APIs, ANGLES provides on-demand content to learning environments as well as providing a direct browsing interface for interacting with a collection. This Stories-as-a-Service capability enables a broader community of contributors than is typical of contemporary video collections; employs knowledge-driven speech-to-text coupled with semantic analysis that automatically tags content; and provides on-demand content to learning environments through API calls. This paper demonstrates to the modeling and simulation community how emerging technologies can combine to make the power of stories accessible across learning environments.

ABOUT THE AUTHORS

Benjamin Bell, Ph.D. is the president of Eduworks Government Solutions and a recognized expert in the application of artificial intelligence to decision support, simulation, training, and human-machine interaction. He has practiced in the field of artificial intelligence for over twenty years, leading funded research and development in applied settings, largely for military applications. He holds a Ph.D. from Northwestern University and is a graduate of the University of Pennsylvania.

Debbie Brown is a software engineer and senior learning technologist at Eduworks Corporation focusing on web application implementations that use AI and machine learning to semi-automate user-centered workflows and enable advanced adaptive eLearning tools and experiences. She has been conducting eLearning R&D projects for 20 years, and operating as a software engineer in government, academic, and workforce settings for 30 years. She holds an MS in Instructional Systems and a BS in Computer Engineering from Mississippi State University.

Ewald Enzinger, Ph.D. is a Research Engineer at Eduworks Corporation. He served as the technical lead for the ANGLES work reported in this paper. He is an expert in speech and natural language processing and currently serves as PI on a project funded by the Defense Advanced Research Projects Agency (DARPA). He has previously held research appointments at the Acoustics Research Institute of the Austrian Academy of Sciences. He holds a Ph.D. in Electrical Engineering from the University of New South Wales.

A “Stories-as-a-Service” Architecture for Enriching Training Simulations

Benjamin Bell, Ph.D., Debbie Brown, Ewald Enzinger, Ph.D.

Eduworks Corporation

Corvallis, OR

Benjamin.Bell@Eduworks.com, Debbie.Brown@Eduworks.com, Ewald.Enzinger@eduworks.com

OVERVIEW

There is an abundant supply of expert practitioners’ stories, a rich and renewable source of knowledge. The training and education community though has utilized this resource in only very limited ways – largely through personal interaction (military schoolhouses, apprenticeships, lectures) or static compendia (publications, recorded interviews). Though expert practitioners can and do enrich training and education of novices with whom they have direct contact, technologies exist or are emerging that could work in concert to unlock the power of stories through scalable architectures and multimedia processing.

Augmenting Next-Generation Learning Environments with Stories (ANGLES) addresses the persisting under-utilization of a plentiful resource: expert stories for training and education. The system includes a story collection approach that utilizes mobile devices and applies contributor-defined and automatically-derived indices. A cornerstone of our work is a novel Stories-as-a-Service architecture that provides client learning applications with a service that retrieves stories in response to requests published by the client. ANGLES enhances distributed learning by enriching training and education technologies with timely, relevant stories. This enhancement, moreover, can be easily accessible to training system developers and made largely transparent by the ANGLES application programming interface (API).

ANGLES includes a streamlined story capture capability that employs mobile device media capture, supports user-selected indexing, and applies automated approaches to generating meaningful indices. ANGLES also includes a web-accessible story server organized thematically or through keyword signatures that retrieves relevant stories and corresponding follow-up questions. And the ANGLES API exposes this capability to current and future learning environments.

CURRENT PRACTICES FOR EMPLOYING VIDEO STORY CONTENT

Collecting Content

Most instances of organizations utilizing video stories use content comprised of professionally-produced, bespoke video collections. Video stories for educational or informational purposes are usually captured using high-definition cameras under controlled conditions with professional technicians to achieve a high-quality appearance (for instance, each U.S. military branch maintains websites with videos collected in this manner). ANGLES adopts an approach to content premised on a crowd-sourced collection model where video is recorded using lower-fidelity equipment such as mobile devices. Our vision of cultivating a broad pool of story tellers (within any given community of practice) is better served by a low-barrier, anytime/anywhere video collection model.

Metadata Tagging

While there have been and continue to be initiatives to distribute video narratives online, the content generally speaking is organized by surface features (speaker, affiliation, topic, date, location). In the few instances where stories have been indexed using principles from various research disciplines, these efforts have been one-off demonstrations or commercial products with bespoke, purpose-built content. Finally, these efforts have been exclusively stand-alone, browseable collections, not created as services to support other instructional applications. In contrast, ANGLES helps advance distributed, anytime/anywhere learning by providing client applications with an enhanced service to make training and education more relevant, credible and vivid.

An example of how video story collections for informational or educational purposes organize content based on surface-level features is the simple tagging seen in Youtube channels. Youtube is the largest collector of online video but content is not tagged at levels characterizing educational utility. Some commercial platforms like Panopto apply automated speech recognition to enable searching of video content. Panopto and its competitors can thus support text-based searching of video but do not employ a semantic framework for characterizing the content. ANGLES's approach to tagging content has been most influenced by research that explored organizing expert vignettes according to a fixed set of indices, e.g., context, causes, results, concepts, examples, opportunities and warnings (Schank, Ferguson, Birnbaum, Barger, & Greising, 1991; Slator & Riesbeck, 1991; Ferguson, et al., 1992; Slator, 1994). This approach offers a rich set of semantic relations. However, using this approach in practice is labor-intensive.

ANGLES adopts a dual approach to tagging content metadata: (1) employ a more streamlined indexing scheme. A streamlined variant of previous indexing research was demonstrated in work focused on aviation training (Bell, Gold & Kaplan, 1998). The indexing was built around a questions-asked/questions-answered schema (Bernstein & Osgood, 1995), supported by a simplified web-enabled database management tool; (2) blend user-provided tags and automatically-generated tags. The ANGLES collection tool allows story-tellers to self-tag their content using the semantic schema created during the project; and through application of text analysis, will in future work be augmented with machine-generated tagging.

Content Integration

Web portals and bespoke video story libraries generally present stories in a question-asking or browsing interface that allows users to navigate the collection by topic, theme, or question. Youtube offers a flexible API that channel curators have used for such topical and thematic organization. The ANGLES stand-alone story browser that allows direct navigation through the video collection is only one aspect of the ANGLES project; the Stories-as-a-Service feature is a new capability, distinctive to approaches that employ a story collection within the confines of a singular application or web portal. The ANGLES story service is a web service that responds to requests from any learning environment using the ANGLES API.

RELATING STORIES TO QUESTIONS

To encode relationships between story objects and objects containing the questions that it answers and raises (Figure 1), ANGLES employs a third object class, relations, in order to provide for a flexible and extensible data structure for describing a link between a question and a story.

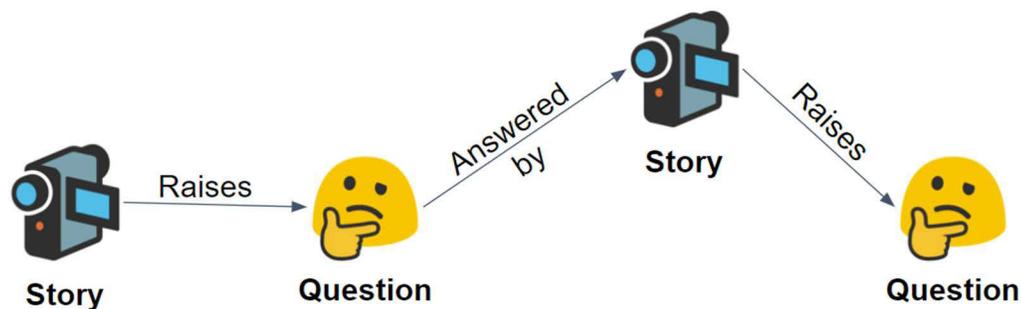


Figure 1. Relationships between Stories and Questions

There is a many-to-many topology between stories and questions, with each link characterized by a category (e.g., context, indicators, response measures, alternative courses of action, lessons learned). As illustrated in Figure 2, the relation object allows for flexibly representing categories to provide additional context about the relationship between a story and a question.

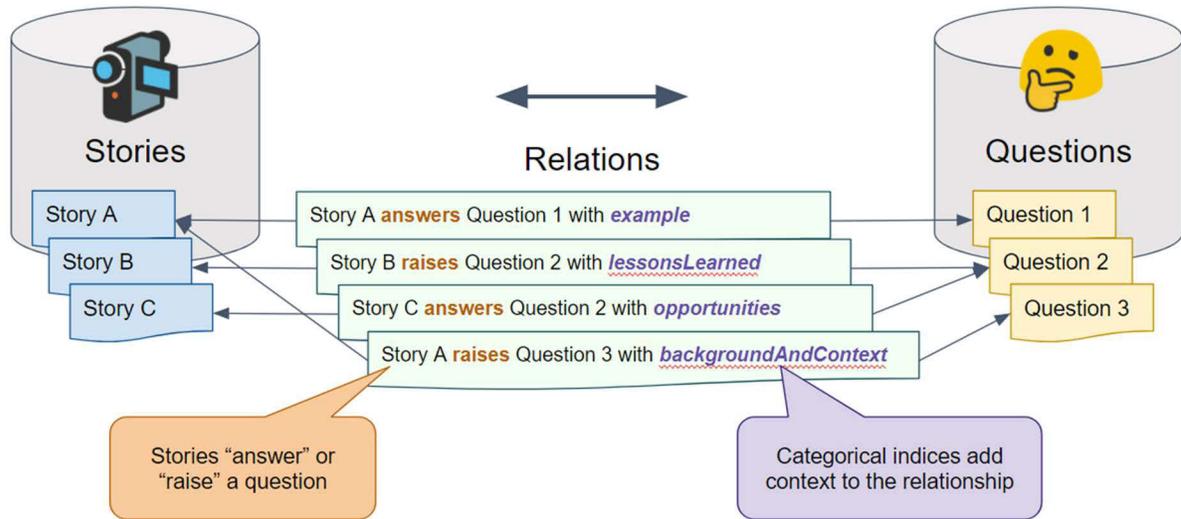


Figure 2: Categorical metadata tags based on story and question relationships

REPRESENTING STORY CONTENT AND ATTRIBUTES

To service API requests sent to the ANGLES server, stories must be tagged using markers that facilitate matching story content with questions sent by learning applications. The data structures representing stories, questions and relations are listed below:

Story:

- name
- description
- keywords
- startTime, endTime, and duration
- embedUrl
- storyTargetAudience
- storyTargetExpertise
- storyCategory
- cyberJobTitle
- educationalAlignment
- interactivityType
- learningResourceType

Question:

- text
- keywords
- storyTargetAudience
- storyTargetExpertise
- storyCategory

Relation

- source
- relationType
- target
- storyCategory

Each of these data structures maps to a schema.org representation, and is represented at Open Linked Data JSON objects in the ANGLES repository.

STORY COLLECTION

A tenet of ANGLES is that storytellers should be able to contribute content anytime, anywhere, using commodity smartphone or tablets. The ANGLES mobile app enables a user to capture live video or select a pre-recorded video, upload the video (after on-device compression), and tag video stories with keywords and metadata after the video is uploaded. Speech to text (STT) is part of the ANGLES server API and creates a transcript of the uploaded video. A pronunciation dictionary and language model used in the STT component can be automatically customized to a corpus of content. Our prototype employs the cybersecurity and the language model enables recognition of domain-specific vocabulary, such as “malware” and “honeypot”, as shown schematically in Figure 3.

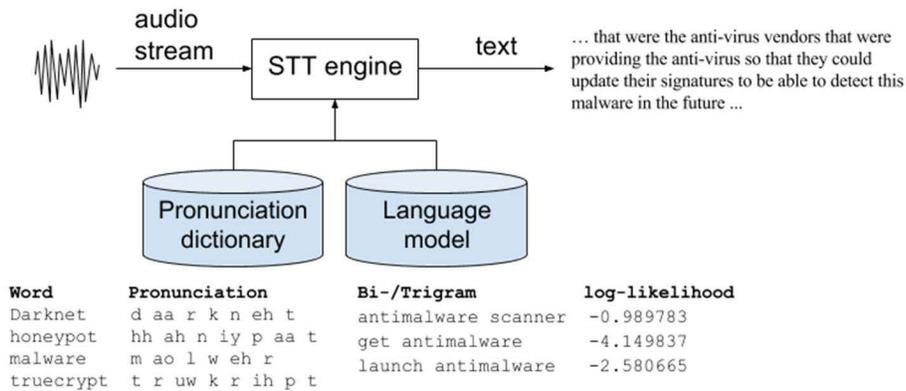


Figure 3. Speech to Text processing in ANGLES

We used Eduworks’ keyword extraction algorithms to automatically obtain a list of keywords that are suggested to the user for tagging, as shown schematically in Figure 4.

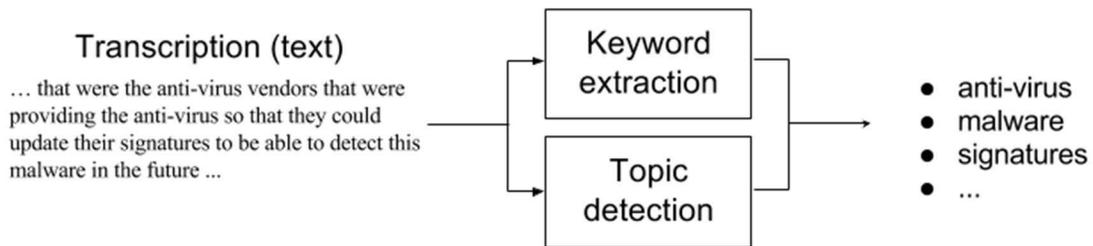


Figure 4. Keyword extraction processing in ANGLES

The screen displays from the resulting story collection mobile app are shown in Figure 5.

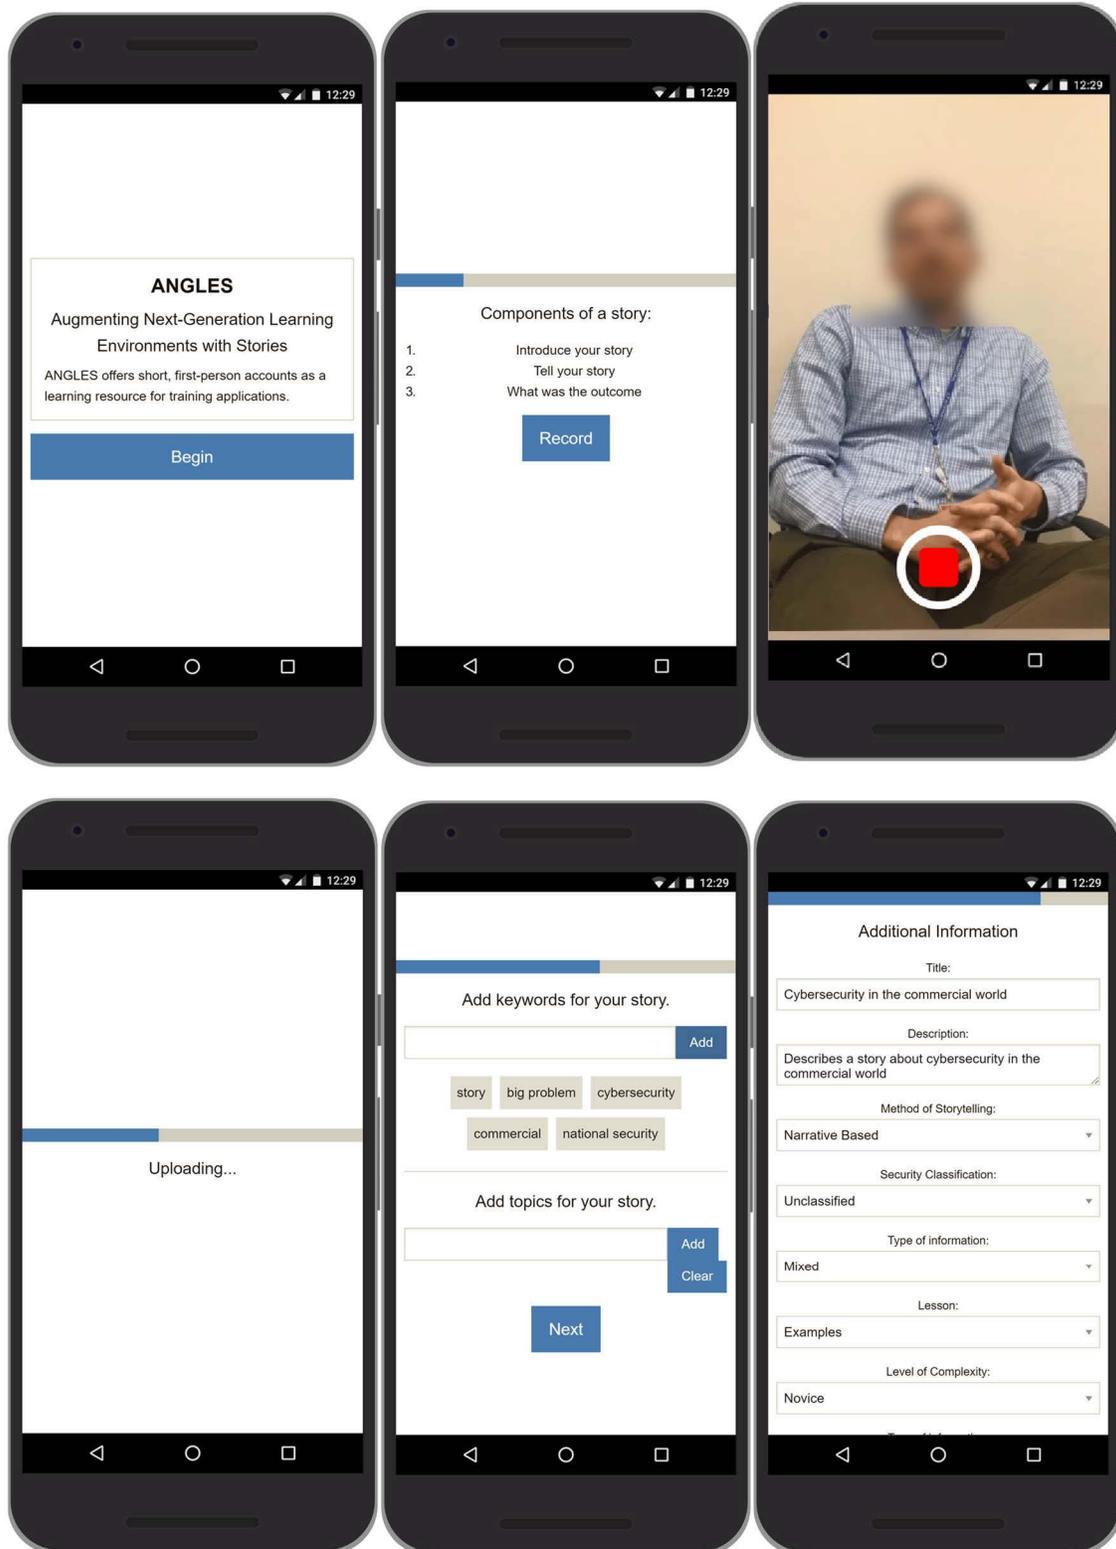


Figure 5. Example screen displays from the ANGLES story collection mobile app.

STORY BROWSER

ANGLES' story browser enables a user to ask questions, search stories by their category, traverse stories that are linked via categorical relations, and view video stories (Figure 6).

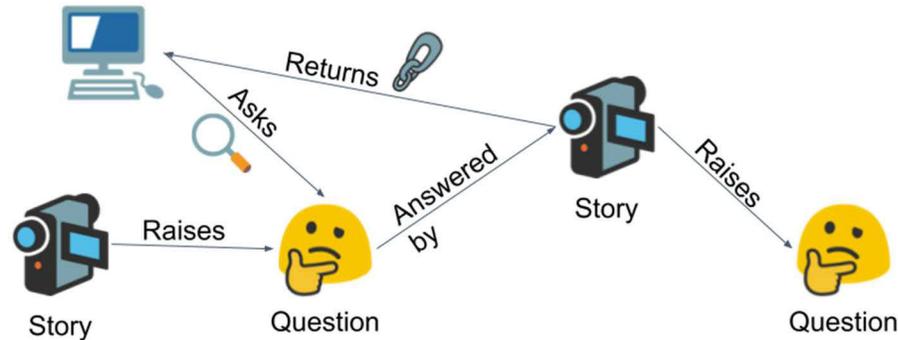


Figure 6. Schematic depiction of story browsing in ANGLES.

We created a web application enabling the user to retrieve video stories by entering a question or keywords in a question-driven interface or by exploring stories in a topic-driven interface. The interface presents screens displaying video stories and additional questions raised by these stories. Sample screens from the story browser appear in Figure 7.

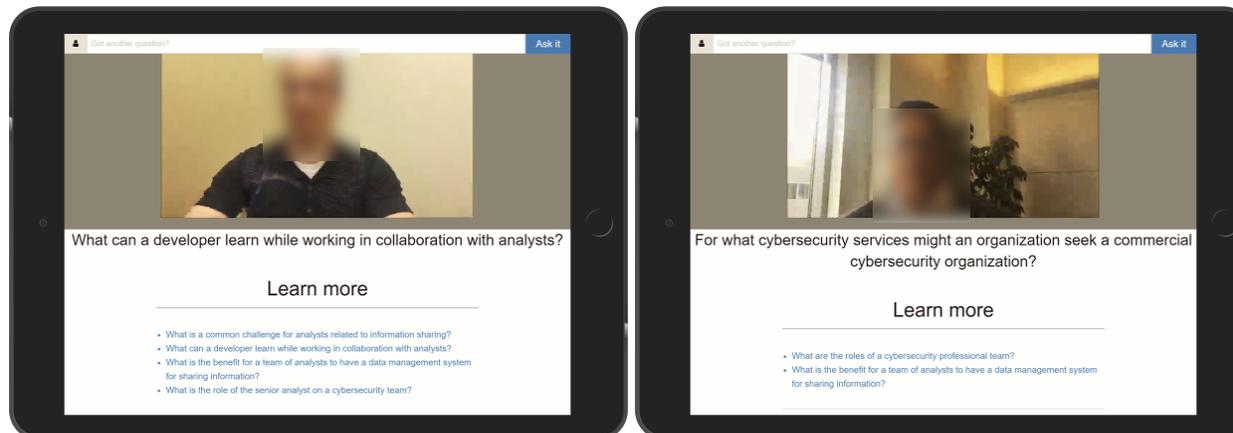


Figure 7. Examples of story browsing in ANGLES.

INTEGRATION

The ANGLES API enables learning environments to forward a user's question to ANGLES, which responds with an appropriate story that encapsulates an answer to the user's question. The question could be one asked directly by the user or inferred by the learning environment. The API also exposes the server-side algorithms used to tag meta-data for the uploaded videos. The data representations are based on Open Linked Data formats and reuse and extend common schemas. The API was designed as a set of web service calls based on Representational State Transfer (REST) to provide a streamlined and robust way for interaction and integration with other learning environments, as shown schematically in Figure 8. This enables ANGLES to be readily combined with software in a Service Oriented Architecture (SOA).

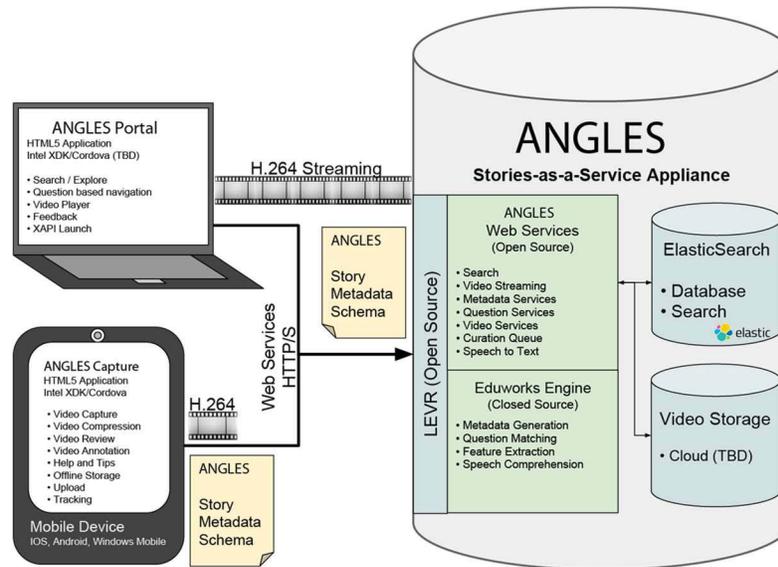


Figure 8. ANGLES service oriented architecture diagram

ENHANCED SEMANTIC TAGGING AND ALIGNMENT

A capability we have designed, but not yet implemented, is improved automated tagging of story content using semantic analysis techniques. Tagging of video stories with metadata and alignment with existing internal and external frameworks (e.g. competency frameworks, lists of questions) currently depends on manual tagging by users. Automating this process allows for faster and more comprehensive tagging, as well as accelerating retroactive alignment with new frameworks or questions. Our design employs probabilistic alignment of automatically STT-transcribed text from videos to external frameworks using a combination of ontological expansion, topic detection, and deep and shallow learning techniques to create maps of high and low-level concepts from videos and match these to external frameworks. This approach goes well beyond keyword matching used by many retrieval applications, but recognizes underlying concepts without strong dependency on matching specific words or syntax. Two use-cases for aligning video transcriptions to competencies and to questions stored in the system are shown in Figure 9.

CONCLUSION

Informal learning, through sharing of first-person experiences, can play an important role in sharing insights among peers or in expert-novice interactions. This paper demonstrates how emerging technologies can combine to make the power of stories accessible in modern, mobile, just-in-time learning contexts. ANGLES' baseline capabilities demonstrated the use of existing open source software for video capture, cross-platform deployment, video compression, audio stream extraction, video streaming, speech to text, automated keyword generation, and other methods of automating categorization and metadata generation. The broader vision of ANGLES is to allow a diverse community of story-tellers to share their experiences using smartphones or tablets, and to automatically tag and organize that content to provide a Stories-as-a-Service on-demand resource that enriches learning environments with vivid, relevant first-person accounts.

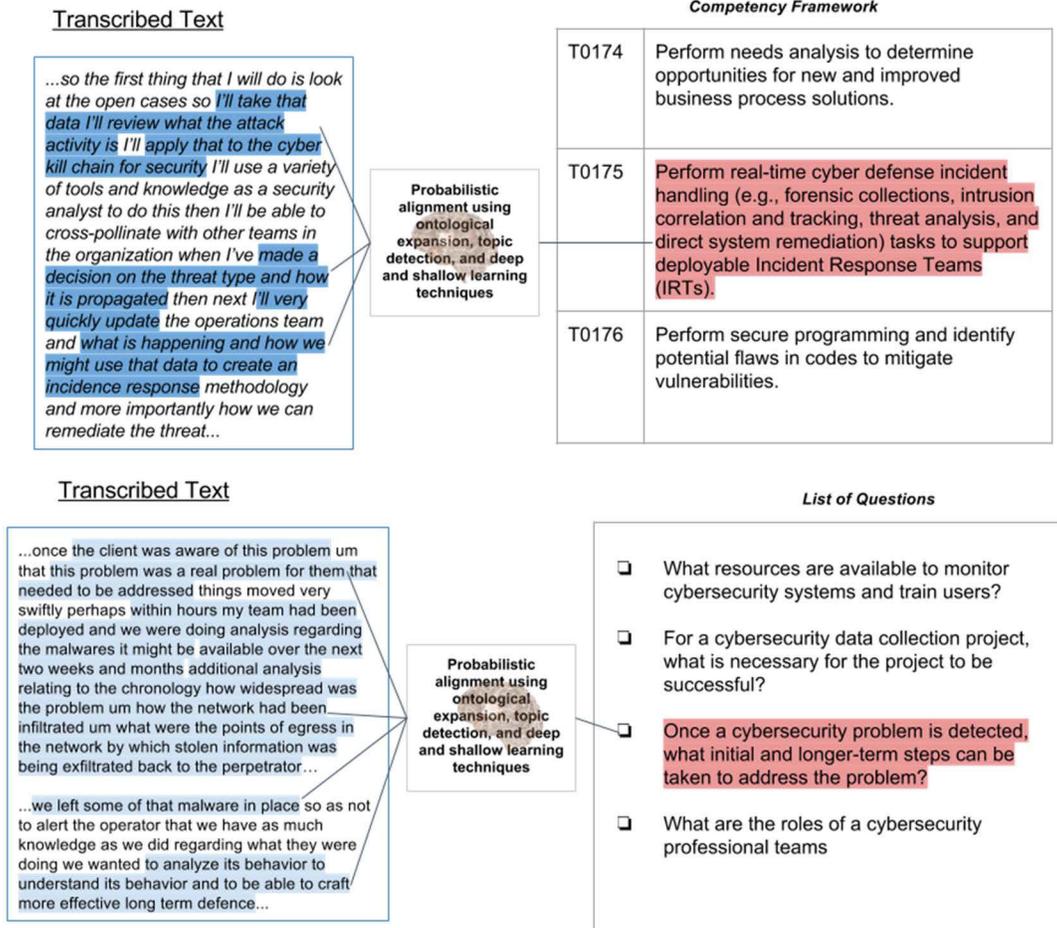


Figure 9. Probabilistic alignment of automatically STT-transcribed text from videos to external competency frameworks (top) and to questions (bottom).

REFERENCES

- Bell, B.L., Gold, S., and D. Kaplan (1998). "Hangar Flying as Aviation Training: Capturing expertise via online video libraries." Presented at the *International Conference on Human-Computer Interaction in Aeronautics*, Montreal, Canada, 1998.
- Bernstein, C. L. and Osgood, R. E. (1995). Support System Indexing for Learn-by-Doing Educational Environments, *AAAI Fall Symposium on AI Applications in Knowledge Navigation & Retrieval*, Cambridge MA, 1995.
- Ferguson, W. and Bareiss, R. and Birnbaum, L. and Osgood, R. (1992). ASK Systems: An Approach to the Realization of Story-Based Teachers, *Int'l J. of the Learning Sciences*, 1(4).
- Schank, R., Ferguson, W., Birnbaum, L., Barger, J., and Greising, M. (1991). ASK-TOM: experimental interface for video case libraries. In *Proceedings of the Thirteenth Conference of the Cognitive Science society*. Northvale, NJ: Erlbaum.
- Slator, B.M. (1994). Tools for Managing Stories and Questions. In *Proc of Sixth International Conference on Tools with AI (TAI-94)*. New Orleans, November 6-9, pp. 237-238.
- Slator, R. & Riesbeck, C. (1991). Taxops: a case-based advisor. *International Journal of Expert Systems*, 4(2), 117-140.