

Report Documentation Page			Form Approved OMB No. 0704-0188		
Public reporting burden for the collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to a penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.					
1. REPORT DATE 29 OCT 2014		2. REPORT TYPE Summary		3. DATES COVERED 01 APR 2014 - 30 SEP 2014	
4. TITLE AND SUBTITLE Knowledge-Aided Interface for Big Data Streams			5a. CONTRACT NUMBER W911QX14C0022		
			5b. GRANT NUMBER		
			5c. PROGRAM ELEMENT NUMBER		
6. AUTHOR(S)			5d. PROJECT NUMBER 0010444320		
			5e. TASK NUMBER		
			5f. WORK UNIT NUMBER		
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Mod9 Technologies 2150 Shattuck Ave, PH Berkeley, CA 94704			8. PERFORMING ORGANIZATION REPORT NUMBER		
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) Sue Kase			10. SPONSOR/MONITOR'S ACRONYM(S) ARL-APG		
			11. SPONSOR/MONITOR'S REPORT NUMBER(S)		
12. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release, distribution unlimited					
13. SUPPLEMENTARY NOTES The original document contains color images.					
14. ABSTRACT					
15. SUBJECT TERMS					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT SAR	18. NUMBER OF PAGES 2	19a. NAME OF RESPONSIBLE PERSON
a. REPORT unclassified	b. ABSTRACT unclassified	c. THIS PAGE unclassified			

Summary Report
DoD SBIR: OSD13-LD2

Arlo Faria (arlo@mod9.com)
Mod9 Technologies

This SBIR Phase I project explored a knowledge-aided interface for Big Data streams, specifically considered in the context of multimedia signal processing for real-time alerting capabilities. The effort successfully demonstrated: a flexible methodology for audio and video signal collection on low-powered computing devices; a comparison of several stream data processing frameworks, resulting in the selection of the Amazon Kinesis service; state-of-the-art speech recognition and keyword search, as validated in the NIST OpenKWS evaluation and competitive with performers in the IARPA Babel program; design and initial prototypes for semantic knowledge integrated with context-sensitive visualization in a novel user interface.

The technology developed is capable of capturing live broadcast news sources such as streaming Internet video. This is implemented using desktop operating systems that can run on very low-powered computing hardware (typically less than 10W power consumption). By monitoring audio from the sound card or taking screenshots, the system conveniently allows collection from diverse input sources, such as: USB-connected television/radio tuners via USB-connected adapters, or analog audio via microphones or line-in jacks. The operator is able to configure and verify the process by simply seeing and hearing the proper output over the device's connected speakers and/or monitor. This "WYSIWYG" approach (or perhaps also "What You *Hear* Is What You Get") proved to be simpler and more flexible than alternatives such as writing customized server software modules for handling each kind of individual input source.

Real-time stream processing is a major recent advance in scalable Big Data systems research. In addition to notable open-source projects such as Apache Storm and SparkStreaming, we considered Amazon Kinesis, which was only recently launched as a hosted service in the Amazon Web Services ecosystem. Several factors ultimately led to our decision to focus exclusively on building solutions for Amazon Kinesis, including its ease of deployment and compatibility with potential future integration in secure private cloud infrastructures that may be viable for government customers. The Phase I effort resulted in a simple demonstration of signal collection and upload to Amazon Kinesis; the proposed Phase II effort will seek to demonstrate how the provisioned throughput capacity of such a system can adequately scale to hundreds of simultaneous streams.

Because multimedia signals are unstructured data, specialized processing is necessary to extract content such as words spoken in audio signals. Such technology is our team's unique specialty. Having spun off from prior research in the IARPA Babel program, our team demonstrated its ability to deliver rapidly deployable and effective spoken keyword search capability for multiple languages. Participating in the NIST OpenKWS evaluation, our system exceeded the program's base period target performance and was most notably distinguished by the fact that it was developed in just four days – rather than fully utilizing the evaluation's four-week schedule. We look forward to continuing development of such speech recognition systems, as well as complementary technologies such as speaker identification.

In addition to demonstrating scalable and effective signal collection and processing, our Phase I effort also focused on the design of a knowledge-aided interface that would facilitate human-machine interactions situated in a context such as an analyst working in a tactical operations center. We successfully demonstrated the use of FrameNet, a semantic knowledge database, for search query expansion and interactive refinement. Our Phase II proposal seeks to further develop such a user interface, with an integrated display of information such as location-aware social media filtering.

Below is a schematic depiction the developed technology:

