

Developing a Virtual Assistant for Space Operations

Jeremy Ludwig, Bart Presnell, and Richard Stottler

Stottler Henke Associates, Inc., San Mateo, CA

ABSTRACT

Virtual assistants such as Alexa/Siri are now routinely used for everyday tasks. The objective of our work is to design and develop a virtual space assistant (VIRSA) to support situational awareness and command and control decision-making during space operations as part of the DARPA Hallmark program. VIRSA is a domain-specific assistant that leverages open-source artificial intelligence techniques to perform tasks that are time-sensitive, data-intensive, recurring, or otherwise challenging. The contributions of this paper focus on the development and refinement of VIRSA during Phase II of the Hallmark program along with challenges that inform future work.

1. INTRODUCTION

Space has become an increasingly congested and contested environment as nations and commercial entities have become more and more reliant on space-based capabilities. Enhanced space-domain awareness is necessary to manage these challenges, but the tools and processes available today will be insufficient for the future. For example, there is an abundance of information available to space operators, but it can be underutilized given the operational tempo and the complexity of decisions being made at vast distances and orbital speeds.

As part of the DARPA Hallmark program, our objective was to develop a virtual space assistant (VIRSA) to enhance decision making in space operations by (1) filtering and searching heterogeneous data and (2) providing assistance to operators for tasks that are time-sensitive, data-intensive, recurring, or otherwise challenging; and by (3) reducing operator workload with an intuitive user interface. Our requirements necessitated that the virtual assistant be tailored to real-world operational tasks and that it be complementary to existing capabilities as well as intuitive, unobtrusive, quick, reliable, and convenient.

2. BACKGROUND

One of the focus areas of the DARPA Hallmark program is the development and evaluation of testbeds and tools for space operations concepts. The Phase I participants included two testbeds, an ontology team, a number of tool teams, and two cognitive evaluation teams. Most of the Phase I participants continued to Phase II, with notable exceptions that included new tools to address capability gaps as well as reductions to a single cognitive evaluation team and a single user interface team. It should be noted that this paper presents a narrow slice of the overall program; there are other participants in the Hallmark program and additional focus areas that are outside the scope of this paper.

Two testbeds form the technological foundation. Each of the testbeds uses a distinct cloud-based technology stack to deploy tools, support inter-tool communication, and provide access to shared data. Each testbed is also responsible for running scenarios, where realistic data is used to simulate future space situational awareness and command and control challenges.

Building on the testbeds, a variety of teams developed tools (including VIRSA) that visualize and analyze data to support decision-making during space operations. The tools work together in a complementary fashion to provide the full range of functionality needed on a space operations floor. All of the tools are deployed on both testbeds, requiring a flexible implementation to support the distinct technology stacks, communication methods, and data access methods. Additionally, all tools interact with the operator through the user interface provided by the SOLAR application [1].

A series of evaluation events allowed space operations personnel to test the Hallmark environment, using the testbeds, tools, and data to make real-time decisions during situational awareness and command and control scenarios. The evaluation events were significant undertakings, where scenarios were presented in the context of a

simulated operations floor and staffed with personnel who have real-world experience. The cognitive evaluation team was responsible for carrying out these weeklong evaluation events. This included providing training, supporting scenarios, measuring performance, compiling feedback and guidance for future work, and then sharing the overall results with the greater Hallmark team.

3. RELATED WORK

Virtual assistant technology for everyday life is nearly ubiquitous in modern smartphones and computers. Two of the most well-known assistants are Apple's Siri and Amazon's Alexa, both of which answer questions or perform actions as directed by the user via voice commands and audio responses. For example, a user can ask for 'brunch places nearby.' Google search is another example of a commonly used assistant, leveraging typed commands and responses displayed on a computer screen. In this case, a search for 'brunch' yields a summary of nearby restaurants, a collection of images, and a list of questions other people have asked about brunch along with their answers. Domain-specific virtual assistants generally perform a constrained set of complex tasks that call for a deep understanding of a particular domain. When they work correctly, virtual assistants help users get their tasks done without having to dig through a variety of websites and applications or navigate a cumbersome user interface. This is exactly the kind of utility that VIRSA should provide; however, VIRSA answers questions about resident space objects rather than brunch.

From a technical perspective, VIRSA leverages existing open-source projects to achieve this functionality. Two specific examples are Apache Lucene [2] and Mycroft Padatious [3]. Lucene is a well-established and supported search engine, used in VIRSA to index and search heterogeneous data streams. Padatious is a neural network based intent parser, used in VIRSA to parse natural language input into known, structured, intent. For example, if the user types 'show me xyz,' the inferred intent is that VIRSA should search for xyz and then display the results in the appropriate user interface component. Padatious is developed as part of the Mycroft AI open source virtual assistant [4].

The contributions of this paper focus on the development and refinement of VIRSA during Phase II of the Hallmark program along with challenges that inform future work. This paper builds directly on the work carried out for VIRSA in Phase I of the Hallmark program [5].

4. PHASE I: SUMMARY

One issue with developing a virtual assistant for space operations is that the functionality requested by operators is limited only by the imagination. It is also difficult to know ahead of time what tasks are going to be time-sensitive, data-intensive, recurring, or otherwise challenging in the Hallmark environment until after the tools and scenarios have been developed. Two features became the focus of the Phase I development and demonstration efforts [5] as depicted in Fig. 1: RSO Summary and Keyword Search. **RSO Summary** is an at-a-glance summary of key information about a resident space object. Operators used this information to quickly assess a situation, solve problems in the face of unexpected events, dynamically re-plan, and anticipate implications of potential courses of action. **Keyword Search** addressed the need to perform efficient searches across heterogeneous data sources.

A series of evaluation events held every three months during the Hallmark program allowed space operations personnel to test the Hallmark environment, using the testbeds, tools, and data to make real-time decisions during situational awareness and command and control scenarios. The evaluation events were significant undertakings, where scenarios were presented in the context of a simulated operations floor and staffed with personnel who have real-world experience. Their expertise was matched against increasingly complex scenarios and an ever more capable Hallmark environment. The feedback for VIRSA from the Phase I events confirmed the utility of the RSO Summary and Keyword Search results—along with a list of requests for Phase II.




Fig. 1. Design mockup of RSO Summary (top) and Keyword Search (bottom) results.

5. PHASE II: METHODS AND RESULTS

The VIRSA Phase II implementation included adding three new features as suggested by the Phase I results, in addition to the continued refinement of RSO Summary and Keyword Search. Each of these features is described below, followed by the Phase II results.

The first new feature is **Intent Recognition**. When a user types text into the search bar, VIRSA should attempt to answer the question or perform the task for a narrow, pre-defined set of intents. If that fails, VIRSA falls back on performing search. For example, the user types “Show me the USAF missile defense sats in GEO” in the search bar. VIRSA then parses the free text input and recognizes intent (*filterRSOs*) and arguments (*searchText*=‘USAF missile defense sats in GEO’), resulting in a call to the method *filterRSOs(searchText)*. This method searches for RSOs and then displays the results in the common operational picture. Additionally, VIRSA should perform intent recognition for other tools and forward the parsed message (as opposed to implementing a method to fulfill the intent). In this case, a tool would describe i) the intent they are looking for and ii) how to find it. VIRSA will then send a message to the tool when that intent occurs. For example, if a user types, “what is the pedigree of alertXYZ,” VIRSA could parse the input to identify the most likely intent (*pedigreeRequest*, *alertArgument*=‘alertXYZ’) and then forward this message to another tool to fulfill. During Phase II, the intents were primarily used to direct searches to particular data sources to reduce false positives—e.g., direct the search to look at only RSOs, alerts, user manuals, and background information. VIRSA both fulfilled intents and forwarded intents to other services. Collectively, the RSO summary, keyword search, and intent recognition features are deployed as the VIRSA Search Service.

The second new feature is a **Topic Service**, which tracks popular topics from sources such as search and chat and displays them as autocomplete options when operators type in the search box. The idea is to provide visibility and ease of use for trending topics. Fig. 2 illustrates how recent/popular searches and topics are rendered in the SOLAR user interface. By clicking in the search bar, the operator can quickly see what other operators have been searching for and chatting about. Only the top items in each category are displayed. Typing into the search bar will filter the trending topics, focusing on the items of greater interest to the operator.

The third new feature is **Search Templates**, which preemptively suggests intent-based searches to the operator as additional autocomplete options. In Fig. 2 this is seen as the *information on <query>* item. If the user selects this line, then the search bar is partially filled out with ‘information on,’ where the user provides the rest of the statement. In keeping with the Hallmark microservices architecture, the Topic (and Search Templates) Service is deployed separately from the Search Service.




Fig. 2. Popular searches and chat topics, along with a search template in italics, in the SOLAR user interface.

Taken together, these three features move the functionality of VIRSA closer to that of familiar general-purpose search engines. Qualitatively, the new features were well received and operators pointed out that the pre-populated/trending autocomplete lists (e.g., Fig. 2) were helpful, robust, and seemed to work well. The final discussion on VIRSA indicated that operators would like even better matching—something “more Google-ish.” We took this comment as positive evidence that VIRSA is on the right track and as a warning that operators have high expectations based on tools they are already familiar with. Quantitatively, we saw significant and sustained use of both the Search and Topic Services throughout the evaluation events.

6. CONCLUSION

VIRSA illustrates concrete features that a virtual assistant would use to improve decision-making on the space operations floor. Within the context of the Hallmark environment, we have identified the highest impact initial features, evaluated these features in a series of weeklong evaluation events, identified additional assistant functionality for Phase II, and then developed and evaluated this functionality while continuing to refine the Phase I implementation. Based on the realism and complexity of the evaluation events, we believe that the implemented features and lessons learned will be useful in other space operations environments as well.

However, we do see several challenges in transitioning an assistant like VIRSA into operational use. First, there are obvious data classification issues involved in indexing and searching real data. This significant problem is being addressed through other programs; VIRSA would benefit from their results. Second, VIRSA would benefit from real-world data that identifies the pain points of actual operators and the types of tasks that would be useful to them. In order to make VIRSA “more Google-ish” as requested in the evaluation events, we need to better understand what the operators want to do. Third, it would be easier to apply VIRSA to an existing set of tools and processes rather than one that is constantly changing. It is difficult to identify tasks that are time-sensitive, data-intensive, recurring, or otherwise challenging when everything is in flux. Finally, VIRSA is envisioned as the natural language processing “glue” between the user input and the available tools combined with the available data. Fulfilling the tasks needed by the operators requires services that can perform the underlying tasks together with the data to support the services. All four of these challenges can be addressed—but they do highlight that making a domain-specific virtual assistant is a significant undertaking.

7. ACKNOWLEDGEMENTS

This material is based upon work supported by the United States Air Force Research Laboratory and DARPA under Contract No. FA8750-17-C-0167. The views, opinions, and/or findings contained in this article/presentation are those of the author/presenter and should not be interpreted as representing the official views or policies, either expressed or implied, of AFRL or DARPA. DISTRIBUTION STATEMENT A (Approved for Public Release, Distribution Unlimited).

8. REFERENCES

- [1] “Using Augmented Reality for Space Control,” *Charles River Analytics*, Aug. 26, 2019. <https://www.cra.com/company/news/using-augmented-reality-space-controlbr-charles-river-analytics-develops-ar-app-help>.
- [2] “Apache Lucene - Welcome to Apache Lucene.” <https://lucene.apache.org/>.
- [3] “MycroftAI/padatiuous,” Jun. 10, 2020. <https://github.com/MycroftAI/padatiuous>.
- [4] D. Schweppe, “Mycroft – Open Source Voice Assistant,” *Mycroft*. <https://mycroft.ai/>.
- [5] J. Ludwig, D. Cline, R. Stottler, T. Spradley, L. G. Militello, and J. Diulio, “A Virtual Assistant for Space Situational Awareness,” presented at the Advanced Maui Optical and Space Surveillance Technologies Conference (AMOS), Maui, Hawaii, Sep. 2019.