

CAPABILITY CONSIDERATIONS FOR ENHANCING SAFETY ON LONG DURATION CREWED MISSIONS: INSIGHTS FROM A TECHNICAL INTERCHANGE MEETING ON AUTONOMOUS CREW OPERATIONS

Shu-Chieh Wu^(1,2), Alonso H. Vera⁽²⁾

⁽¹⁾San Jose State University, San Jose, California, USA, Email: shu-chieh.wu@nasa.gov
⁽²⁾NASA Ames Research Center, Moffett Field, California, USA, Email: alonso.vera@nasa.gov

ABSTRACT (15 LINES OR 150 WORDS)

As future flight crews on long duration deep space missions are expected to operate more autonomously, considerations must be given to onboard capabilities and human-computer teaming that will fortify the safety net traditionally provided by the Mission Control Center. In August 2018, the Human Factors and Behavioral Performance Element of NASA's Human Research Program convened a Technical Interchange Meeting (TIM) on Autonomous Crew Operations at NASA Ames Research Center to address how intelligent technologies can be utilized to augment crew capabilities to support real-time anomaly response. In this paper, we highlight three topic areas discussed at the TIM that have direct implications for future crew anomaly response capabilities: smart structures, cognitive assistants, and manpower.

1. INTRODUCTION

Among the many challenges posed by long duration deep space exploration, communication delays, in particular, can cause considerable disruption to the current operation of crewed missions. Space missions historically have relied on the Mission Control Center (MCC) to direct every aspect of the operation in near real-time, from activity planning and procedure execution to anomaly response and troubleshooting [1]. The ability for the MCC to control the mission from the ground will be impacted or made impossible by one-way light time delays—for example, as much as 22 minutes when Mars is at its maximum distance from Earth. Historically, we have seen that unanticipated anomalies can defy even the best thought-out fault detection and resolution systems. As unanticipated anomalies will invariably arise in complex engineered systems, a lack of real-time communication will significantly weaken the support MCC represents: a safety net for the flight crew through its diverse areas of expertise and deep resources, especially during roughly the first hour following an event. In preparing for crewed space missions that go beyond low-Earth orbit to the Moon and Mars, considerations must be given to the nature, design, and implementation of the types of capabilities needed onboard the space vehicles/habitats, and the resulting concepts operations, to fortify the traditionally ground-based safety net weakened by communication delays.

In August 2018, the Human Factors and Behavioral Performance Element of NASA's Human Research Program convened a Technical Interchange Meeting (TIM) on Autonomous Crew Operations at NASA Ames Research Center. The goal of the meeting was to gather input from industry, academia, and branches of the Department of Defense (DoD) to address how intelligent technologies can be applied to augment crew anomaly response. In this paper, we highlight three topic areas discussed at the TIM that have direct implications for future exploration missions: smart structures, cognitive assistants, and manpower. We begin with an overview of anomaly response processes.

2. ANOMALY RESPONSE PROCESSES

Anomaly response refers to activities that operators undertake in response to a system fault, an off-nominal behaviour, or a cascading set of system disturbances (Watts-Englert, Woods, and Patterson, see [2]). They commence following the detection and recognition of an anomaly to fulfill broadly one of two functions: 1) troubleshooting (diagnostic search) for the underlying cause and 2) contingency management. Troubleshooting, characterized by an interaction of prediction and observation, is accomplished by solving three subproblems: generating hypotheses by reasoning from a symptom to a set of causes; testing each hypothesis to see which one(s) can account for all available observations; and discriminating those hypotheses that survive testing [3]. Contingency management concerns what to do next to manage the situation even when the underlying cause may not have been identified. Its activities include risk assessment, plan selection, plan modification, contingency evaluation, and safing/protecting the system. According to Watts-Englert and colleagues, the processes of troubleshooting and contingency management do not unfold in a linear sequence but often proceed in parallel and feed into each other.

The concepts above aptly describe the MCC's anomaly response process, as exemplified by its handling of a cooling system failure on the International Space Station (ISS) [1, 4]. On December 11, 2013, the flight control team in Houston detected an alarm and quickly determined that the external cooling system Loop A had shut down (*system disturbance*), resulting in losing half of the external station cooling capacity. It appeared that

the fault detection software automatically turned off power to the Loop A pump module that circulates the ammonia through the radiator because the ammonia was getting too cold (*symptom*). The team isolated the problem to the Flow Control Valve (FCV) that controls the flow of cold ammonia from the radiator entering the primary system (possible *cause*). To troubleshoot, the team first tried to restart the pump module and command the FCV movement using various methods to no avail (*testing hypotheses/contingency evaluation*), while at the same time shifting heat loads to Loop B (the remaining cooling system) and powering down equipment to reduce the overall amount of heat generated (*safing*). The anomaly ultimately took 14 days of 24/7 MCC support to resolve, including 2 Extravehicular Activities (EVAs) lasting 12+ hours in total, to replace the pump module (*plan selection*).

Several aspects of the response to anomalies like the cooling Loop A failure could potentially be facilitated by intelligent technologies. One relates to the monitoring and detection of anomalies. In current operation of the ISS, the ground handles most alarms unless communication is disrupted due to scheduled or unanticipated events. The Caution and Warning system (C&Ws) on the ISS issues four classes of alarms; among them, class 1 (emergencies) and class 2 (warnings) require immediate action by the crew and/or ground to avoid injury or death of the crew or damage to the vehicle. There are approximately 80 different types of emergencies and 800 different types of warnings [1]—all those that could be anticipated in advance. Adding unanticipated anomalies, the volume of work and the speed required to address it could overwhelm a small flight crew operating without ground support if unassisted by on-board technologies.

Another concern is the range of expertise and the amount of resources nominally required to handle anomalies. Flight control operations have evolved over time but the basic organizational structure remains. For ISS operations, there are 18 primary flight control positions/consols in the Flight Control Room (or Front Room, the room typically seen in media coverage) [1, see Table 3 on p.xxv], of which 12 are assisted by one or more additional operators in the Multipurpose Support Room (or Backroom). Six of the positions/consols manage core systems (power, computer control, communication, attitude control, thermal control, and life support) related to the safety of the vehicle and survival of the crew. Resolution of major anomalies often requires tapping into the Mission Evaluation Room (MER) for in-depth engineering analysis support. MER personnel retain and manage design specifications, manufacturing documentation, and general system knowledge and provide subject matter support on how various systems or components function or respond [1]. It will be a

challenge for a small crew of 4 or 6 people to cover such a range and depth of expertise.

A final aspect relates to the level of manpower required to respond to anomalies quickly, which like the organization structure also remains relatively constant. Nominal operation of the ISS is handled by about 60 flight controllers (48 in the Front Room, Backroom, and MER in Houston, 12 in the Payload Operations and Integration Center in Huntsville). Anomalies that require formation of a dedicated team (such as in the cooling Loop A failure example) could involve up to 150 personnel [Bobby Fard, personal communication, March 8, 2019] working 24/7 for days or weeks.

Due to the complexity of vehicles and the criticality of problems, anomaly response in space missions requires a large amount of distributed expertise, resources, and manpower be brought together and dispensed simultaneously and quickly. Intelligent technologies have the potential in supporting several aspects. They could provide contextualized information behind cautions and warnings and give the anomaly response process a head start. They could also help amass the wide range and depth of specialized expertise as well as investigative resources and bring them to bear more quickly. In the next section, we discuss two technology areas that could potentially provide those capabilities.

3. INTELLIGENT TECHNOLOGIES

3.1. Smart Structures

The essence behind smart structure technologies is to turn sensed data into information and use it to guide decisions and actions, much like what is needed in fault detection. Dr. Mario Berges, Professor of Civil and Environmental Engineering at Carnegie Mellon University discussed current and next generation technologies behind smart building structures and the challenges involved in advancing from sensed buildings to autonomous buildings. According to Berges, Internet-of-Things is beginning to enable much more automation in buildings, though not autonomy, because the latter remains difficult to set up. The difficulty also lies partly with the limitations of current data-driven solutions, specifically machine learning, in extracting useful information from data.

To illustrate, Berges cited two case studies; both concerned inferring the sensed stimuli with respect to what type the sensors were and what they measured. The first one was on Building Automation Systems (BAS), which can help building managers and owners reduce energy consumption. In an ideal framework, a self-managing BAS can be deployed to any building to automatically manage the information processing. That flexibility is enabled by an information mediator layer that handles the integration of heterogeneous information sources and information sharing among three self

managing functions - self-recognition (of own components and their configurations so that the needed information can be automatically retrieved), self-monitoring (of the working status of the components), and self configuration (of the information base based on the outputs generated by the other functions). However, because there is little standardization on the format of device metadata (i.e., information that helps contextualizes measurements or control signals sent from/to a device, such as the location within a building, the physical phenomenon being sensed, etc), such a framework must contend with unstructured and inconsistent labels from heterogeneous systems.

The second case study concerned designing non-intrusive load monitoring (NILM) for residential buildings. The objective of NILM is to provide appliance-level energy metering using data from only a whole-house meter [5]. There are two general approaches, event-based and event-less. Event-based approaches rely on detecting events (i.e., abrupt changes in power consumption) then classifying them based on appliance signatures, whose definition would require pre-identified labels generated for local features of events. Event-less approaches rely on inferences generated by factorial hidden Markov Models made computationally tractable by first constraining the state space using domain knowledge.

Both case studies, Berges argues, illustrate the importance of domain knowledge. Even though data abound in the physical world, it is information derived from this resource that generates value [6]. And the latter process requires significant domain expertise.

3.2. Cognitive Assistants

What does it take to augment human capabilities? Experts from IBM and NASA Langley provided an in-depth look at the technology, design, and deployment behind cognitive assistant systems based on IBM Watson cognitive computing technology.

Dr. Bill Murdock, Researcher and Computer Scientist at IBM Watson Research Center laid the foundation on how cognitive assistants support a user's information needs. He contends that information needs constitute a positively skewed distribution with a "tall head" and a "long tail". Tall head represents common questions. Because the questions are foreseeable, it is possible to optimize for each information need, provide highly curated responses, and perform with extreme accuracy. Long tail represents rare events/faults. Because they are unforeseeable, it is only possible to optimize for all of such instances together. Consequently, retrieved answers can only be moderately accurate (but often accurate enough), though what is lacking in accuracy may be compensated by providing more answers to a query. Tall head information is amenable to being implemented in conversational systems by listing and enumerating all instances that will lead to a particular piece of

information. Long tail information is more suited to be implemented in discovery systems, providing broad coverage of potential answers.

Dr. Jon Holbrook, Cognitive Scientist at NASA Langley Research Center and Dr. Graham Katz, Senior Managing Consultant at IBM put the discovery systems that Murdock discussed into an operational context. They described the development and demonstration of a Pilot Expert Advisory System based on Watson Discovery Advisor (WDA) technology, an application of the long-tail Discovery type of system [7]. The Pilot Expert Advisory System was billed as a human-autonomy teaming system that monitors and assesses in real-time states of the human, vehicle, and automation systems and links them with external sources of information to provide flight crew with relevant information in anomalous situations. It was designed to be able to answer questions posed by pilots in natural language and find answers in text sources. In building the corpus of expert knowledge that consists both general and domain specific aviation information, unstructured text from FAA publications (regulatory documents and airman's information manuals), relevant incident knowledge from the Aviation Safety Reporting System (ASRS), aircraft-type specific knowledge, as well as NASA select documents were ingested into the WDA system. Subject Matter Experts (SMEs) were consulted to construct a list of domain-specific terminology for natural language processing and to provide correct answers to domain specific-questions for training machine learning models. Tested against a use-case based on a real incident, the demo system was able to generate hypotheses about possible systems related to a particular fault message and on factors prone to cause that fault, with the correct answers listed at the top of candidate hypotheses. However, Katz acknowledged a couple of issues that helped put the initial success in perspective. First, technical specifications and formal engineering terminology did not always match up to the colloquial descriptions that flight crew used. Second, it was difficult for the SMEs to think of questions that they do not usually ask; that is, difficult to think beyond "tall head" questions.

Dr. Jeff Kephart, Distinguished Research Staff Member at IBM Watson Research Center introduced the concept of embodied AI. Rather than a simple Q & A system, embodied AI allows a cognitive assistant to have a brain, sensors (eyes, ears), effectors (hands, feet), and even emotional intelligence. It is effectively a software agent that co-inhabits a physical space with people and uses its understanding of what is happening in that space to act as a valuable collaborator on cognitive tasks. Kephart showcased several embodied AI prototypes and research projects. He began the presentation with a hypothetical Mars crew scenario in which an embodied AI agent senses an astronaut's behavior (looking worryingly at a

gauge) and offers assistance. The exchange is carried out in natural dialogs and requires the agent to be able to sense the immediate physical space (spatial intelligence) and perform a variety of processes according to context (human behavior analysis, emotion analysis, planning, simulation, reasoning, explaining, diagnosis, preference elicitation). He then showed several more embodied AI prototypes in the areas of exoplanet exploration, mergers and acquisition, oil and gas field development. The compellingness of the demos notwithstanding, Kephart acknowledged there remain many embodied AI research challenges: sensing and interpreting the user's environment (multimodal adaptive sensor fusion and rich transcription), interacting with the user (spatial AI and contextual interaction and models of self, world, and people), collaboratively executing high-level cognitive functions (e.g., planning, decision-making), building the software/hardware architecture (spanning Edge and Cloud), and measuring and improving the effectiveness of human-agent interactions.

3.3. Limitations

Both smart structures and cognitive assistants exhibit similar limitations in what machine/deep learning can accomplish. In the context of smart structures, deep learning systems that take in building energy and circuit load health data cannot answer new questions, only the question(s) they were trained on (as neural nets). The interpretation of answers provided by these systems remains reliant on human domain expertise. Furthermore, it remains the case that most building and circuit representations are top-down and therefore poor at supporting bottom-up questions (e.g., what other outlets are on the same circuit as this one?). Similar limitations are also found in cognitive assistants like IBM Watson, which can be trained to assist with diagnosis by providing answers to common questions but will falter at addressing unanticipated, rare events. Both topic areas acknowledge the unsurpassed role humans (specifically, using domain expertise and creative problem solving) play in bridging the gap of machine intelligence.

4. MANPOWER

Discussions of technologies often focus on what capabilities they provide and rarely on what is required to harness the capabilities, yet it is the latter that determines the ultimate success (or failure). Case in point, autonomous crew operations will undoubtedly require a slew of technologies to enable capabilities new both to the vehicle/habitat and the flight crew, particularly for troubleshooting during emergencies. How to determine whether the crew of four will be able to use them effectively at times of need? The issue of manpower is a novel one to space operations that have traditionally relied on (and benefited from) access to near limitless real-time ground support but a central and crucial one to the Navy. In her presentation, Dr. Nita

Shattuck, Professor at the Naval Postgraduate School helped lend support to the issue of manpower by describing a case study based on the Littoral Combat Ships (LCS).

The Littoral Combat Ship (LCS) is a relatively small and agile Navy surface ship specifically designed to operate in the littoral (near shore) area not accessible to Navy cruisers and destroyers. The LCS is a focused-mission ship, equipped to perform one primary mission at any given time; primary missions include antisubmarine warfare (ASW), mine countermeasures (MCM) and surface warfare (SUW) against small boats (including so-called "swarm boats"). It achieves its versatility through modular "plug and fight" mission packages, including unmanned vehicles (UVs); the ship's mission orientation is changed by swapping out its mission package [8].

The LCS is developed by two industry teams and therefore comes in two different designs. The Freedom class design, developed by Lockheed, is based on a steel semi-planing monohull with an aluminum superstructure, while the Independence class design, developed by General Dynamics/Austal, is based on an all-aluminum trimaran hull. The two designs also use different built-in combat systems (i.e., different collections of built-in sensors, computers, software, and tactical displays).

In 2001, the Navy began an effort referred to as the optimal manning initiative to reduce crew sizes aboard various legacy surface and amphibious ships [9]. The LCS employs automation to achieve a reduced-sized crew. The aim was to achieve a core crew size of 40 sailors. With the additional sailors as needed to operate the ship's aircraft and mission packages, a total crew of about 88 sailors would be needed, compared to more than 200 for the Navy's legacy frigates and about 300 (or more) for the Navy's current cruisers and destroyers.

Unfortunately, both LCS developments have been plagued with design and operational issues. During sea trials, Freedom class ships suffered repeated engine failures and Independence-class hulls exhibited massive corrosion and transmission failures, necessitating design modifications for both classes. Several crew errors during operations have resulted in significant repairs. These problems caused the Navy to conduct an engineering stand down of all LCSs in September 2016 to assess and mitigate systemic deficits [10]. A Government Accountability Office investigation was also conducted [11]. Both found that crew training was insufficient, and the Navy ordered that every sailor be retrained. It was also found that the core crew of 40 sailors and officers were too few to safely operate the ship without overworking personnel. Eventually, the complement was increased to 70 in 2016 [9]. Moreover, because ship operation proved so demanding, six LCS -- three of each

type -- are now dedicated solely to training new crews and another four to testing.

Considering the troubled operation history of the LCS, the objective of Shattuck's case study was to investigate what the right number and correct composition of crew is for the workload required. Conventional manpower analysis captures routine duties and events; level of manning is typically determined using the average. Critical phenomena are infrequent but carry dire consequences. How does a system manned according to the average respond to transient phenomena? To answer that question, Shattuck developed three workload models of the LCS crew based on the IMPRINT Pro Forces Module. The basic underlying concept is that crewmembers spend all of their time in some sort of "planned" activities/events, i.e., the ones that typically occur in in the ship's daily schedule. The planned activities are periodically interrupted by unforeseen events and emergencies (i.e., unplanned events). The three models had increasing levels of operational realism and complexity. The first, baseline model consisted planned activities and some regularly occurring unplanned events. The second model incorporated some irregularly occurring unplanned events. The third model further incorporated "black swans" -- very rare events that involved all crew, 12-24 hours in duration (triangular distribution). Shattuck found that even under the baseline model, watchstanders worked on average 2.6 hr/day more than the Navy Availability Factors (NAF) daily duty hour provision. Under the second model, engine, gas turbine system techs, and electrician's mates had the highest average daily workload. Under the third model, Shattuck found significant sleep loss and excessive sustained wakefulness; about 30 crew members did not sleep for over 40 hours. Moreover, crew responded mainly to the major events and only critical watches could be maintained.

Even though many problems of the LCS can be attributed to human-systems integration (HSI) related issues -- modernized interface found unusable by the operators, limited design review by HSI professionals, systems overdesigned for its purpose, incomplete training, and consequential operator fatigue and exhaustion over operation, there are manpower specific issues as well. For them, Shattuck highlighted two recommendations from US Navy's Strategic Readiness Review released in December 2017 [12]. One is to establish a process to measure the true workload of ships' crews, both periodically and after upgrades and modernizations, to determine if manpower models adequately predict personnel requirements at sea and in port. The other is to adjust ship manning levels to allow for adequate crew rest, performance of extraneous and collateral duties, and training that occurs while onboard ship, and to include some excess capacity.

5. CAPABILITY CONSIDERATIONS

What capabilities need to be onboard and how will they team with the crew to maintain the level of safety currently provided by the MCC through anomaly response support? NASA Procedure Requirements (NPR) 8705.2C on Human-Rating Requirements for Space Systems, the agency's current policy directive for carefully managed missions where safety risks are evaluated and determined to be acceptable for human spaceflight, dictate the following requirement regarding anomaly resolution:

"The space system shall provide the capability to utilize health and status data (including system performance data) of critical systems and subsystems to facilitate anomaly resolution during and after the mission" [13, Section 3.2.10].

It should be noted that the NPR defines the space system to include both the crewed space system and all space-based and ground-based systems that functionally interact with the crewed space system during the mission [13, Section 3.1.3]. In other words, it assumes that in anomaly resolution safety is achieved by capabilities present in all parts of the space system combined. It follows that more (if not all) of the same capabilities should be allocated to the crewed space system in future deep space operations where the assumed functional interaction will be absent in the first hour following an event. Here we propose three potential concepts of operations (ConOp) for the crew-ground-vehicle collaborative anomaly response in order of the amount of onboard capabilities required, and discuss what functions intelligent technologies could support.

We propose, at a minimum, the vehicle should provide enough capabilities to support the flight crew in safing the vehicle and themselves when major unanticipated anomalies occur. For example, in the cooling Loop A failure case, in addition to having access to system health and status data, the crew should have tools and methods to evaluate how the failure will impact overall station cooling and to determine what avenues are available to preserve it. Here, discovery systems could assist the crew by pooling information on the cooling subsystem design and vehicle heat load management through mining non-textual data (e.g., engineering schematic diagrams) for knowledge remains a challenge.

With more onboard capabilities, the crew could perform preliminary troubleshooting after safing. The focus is for the crew to troubleshoot anomalies for the purpose of collecting information to be later sent to the ground for further investigation, asynchronous in time. Here, smart structure technologies could be applied to provide better resolution on system health and status.

At the highest level, it is possible to envision a crewed space system with sufficient capabilities for the flight

crew to resolve anticipated and unanticipated anomalies on their own. A combination of smart structure technologies and “tall head” systems could be used to automatically handle anticipated cautions and warnings.

How will the crew be incorporated as part of onboard capabilities amid technologies? The lesson of the LCS highlights the issues of HSI and workload. When the total manpower is a crew of four, the same issues are amplified and new issues arise in different areas. In selection: what should the composition of the crew be in terms of expertise? In operation: what role does each one play in anomaly response? how to flexibly adjust the team (and teamwork) if one (or more) crew cannot perform at full capacity? How can trust be built between crew and technology?

6. FINAL THOUGHTS

Even though sending humans into space requires nothing short of engineering marvels, intelligent technologies that are ubiquitous in our digital lives are still a relative new comer in space operations, currently adopted in only a handful (but growing) applications [14]. While being full of potential, considerations must be given to carefully assess what their costs and benefits are (for an example of trade analysis, see [15]) as well as how best to integrate them (ideally, through an iterative HSI process, as described in [16]).

7. ACKNOWLEDGEMENT

The Autonomous Crew Operations Technical Interchange Meeting was hosted by the Human Factors and Behavioral Performance Element of NASA’s Human Research Program and sponsored in part by the Human Systems Integration Division of NASA Ames Research Center. Portions of this paper were adapted from a draft NASA Technical Memorandum by the same authors, entitled “Supporting Crew Autonomy in Deep Space Exploration: Preliminary Onboard Capability Requirements and Proposed Research Questions – Technical Report of the Autonomous Crew Operations Technical Interchange Meeting”. We thank Mary Connors and Tina Panontin for comments and suggestions.

8. REFERENCES

1. Dempsey, R. C., Contella, D. E., Korth, D. H., Lammers, M. L., McMillan, C. R., Nelson, E., ... Van Cise, E. A. (2018). *The International Space Station: Operating an Outpost in the New Frontier*. Retrieved from <https://www.nasa.gov/feature/new-nasa-e-book-offers-inside-look-at-space-station-flight-controllers>
2. Sgobba, T. (2018). Chapter 13 - Human factors in mission control centers. In B. Kanki, J.-F. Clervoy, & G. M. Sandal (Eds.), *Space Safety*

and Human Performance (pp. 569–620). Butterworth-Heinemann. doi:10.1016/B978-0-08-101869-9.00013-3

3. Davis, R., & Hamscher, W. (1988). Chapter 8 - Model-based Reasoning: Troubleshooting. In H. E. Shrobe & the American Association for Artificial Intelligence (Eds.), *Exploring Artificial Intelligence* (pp. 297–346). Morgan Kaufmann. doi:10.1016/B978-0-934613-67-5.50012-5
4. Anthony, S. (2013). International Space Station’s cooling system fails, but the crew is safe for now - ExtremeTech. *ExtremeTech*. Retrieved from <https://www.extremetech.com/extreme/172694-international-space-stations-cooling-system-fails-but-the-crew-is-safe-for-now>
5. Bergés, M., Goldman, E., Matthews, H. S., & Soibelman, L. (2010). Enhancing Electricity Audits in Residential Buildings with Nonintrusive Load Monitoring. *Journal of Industrial Ecology*, 14(5), 844–858. doi:10.1111/j.1530-9290.2010.00280.x
6. Bergés, M., Lange, H., & Gao, J. (2018). Data-Driven Operation of Building Systems: Present Challenges and Future Prospects. In I. F. C. Smith & B. Domer (Eds.), *Advanced Computing Strategies for Engineering. EG-ICE 2018. Lecture Notes in Computer Science, vol 10864* (pp. 23–52). Presented at the EG-ICE2018 25th International Workshop on Intelligent Computing in Engineering, Cham, Switzerland: Springer International Publishing.
7. Katz, G., Ding, C., & Doyle, A. (2018). *NASA Pilot-Engaged Expert Response Using IBM Watson Technology: Prototype Evaluation of Knowledge Retrieval System* (Technical Report No. NASA/CR-2018-220097). Retrieved from <https://ntrs.nasa.gov/search.jsp?R=20180007515>
8. O’Rourke, R. (2016). *Navy Virginia (SSN 774) Class Attack Submarine Procurement: Background and Issues for Congress* (p. 46).
9. United States Government Accountability Office. (2017). *Navy Force Structure: Actions Needed to Ensure Proper Size and Composition of Ship Crews* (Report to Congressional Committee No. GAO-17-413).
10. LaGrone, S. (2016, September 5). Navy Ordered Littoral Combat Ship Engineering Stand Down. *USNI News*. Retrieved from <https://news.usni.org/2016/09/05/navy->

ordered-littoral-combat-ship-engineering-stand-retraining-ongoing

11. United States Government Accountability Office. (2014). *Littoral Combat Ship: Deployment of USS Freedom Revealed Risks in Implementing Operational Concepts and Uncertain Costs* (No. GAO-14-447). Retrieved from <https://www.gao.gov/assets/670/664672.pdf>
12. US Navy. (2017). *Strategic Readiness Review*. Retrieved from <https://s3.amazonaws.com/CHINFO/SRR+Final+12112017.pdf>
13. NASA. (2017). Human-Rating Requirements for Space Systems (NPR 8705.2C). Retrieved from https://nodis3.gsfc.nasa.gov/displayDir.cfm?Internal_ID=N_PR_8705_002C_&page_name=main
14. Chien, S., Doyle, R., Davies, A. G., Jonsson, A., & Lorenz, R. (2006). The Future of AI in Space. *IEEE Intelligent Systems*, 21(4), 64–69. doi:10.1109/MIS.2006.79
15. Korsmeyer, D., Clancy, D., Crawford, J., & Drummond, M. (2005). Crew Centered Operations: What HAL 9000 Should Have Been. In *1st Space Exploration Conference: Continuing the Voyage of Discovery*. American Institute of Aeronautics and Astronautics. doi:10.2514/6.2005-2546
16. Pew, R. W. (2008). Some New Perspectives for Introducing Human-Systems Integration into the System Development Process. *Journal of Cognitive Engineering and Decision Making*, 2(3), 165–180. doi:10.1518/155534308X377063