National Aeronautics and Space Administration

NASA KNOWLEDGE JOURNAL

Office of the Chief Knowledge Officer

"To capture and share what we know now to ensure mission success"

km.nasa.gov

WINTER 2015

NASA KNOWLEDGE **JOURNAL**

ISSUE NO. 1

STAFF

NASA CHIEF KNOWLEDGE OFFICER AND PUBLISHER
Dr. Edward Hoffman
ehoffman@nasa.gov

EDITOR-IN-CHIEF
Don Cohen
doncohen@rcn.com

MANAGING EDITOR
Mark Schwartz
mark.i.schwartz@nasa.gov

EDITOR
Susan Snyder
susan.snyder-1@nasa.gov

CONTRIBUTING EDITORS
Dan Daly
daniel.daly@inuteqllc.com

Ramien Pierre
ramien.pierre@inuteqllc.com

SENIOR KNOWLEDGE SHARING CONSULTANT
Jon Boyle
jonboyle1@verizon.net

KNOWLEDGE SHARING ANALYST
Ben Bruneau
ben.bruneau@inuteqllc.com

KNOWLEDGE PROGRAM MANAGER
Yvonne Massaquoi
yvonne.massaquoi@inuteqllc.com

BUSINESS MANAGER
Kimbley Lewis
kimbley.lewis@inuteqllc.com

DESIGN AND PRINTING
Hirshorn Zuckerman Design Group, Inc.
www.hzdg.com

TABLE OF CONTENTS

All images and photo credits are NASA unless otherwise stated

On the cover: In this April 25, 1990, photograph taken by the crew of the STS-31 space shuttle mission, the Hubble Space Telescope is suspended above shuttle Discovery's cargo bay some 332 nautical miles above Earth.

WELCOME FROM

NASA's Chief Knowledge Officer

Many people think of knowledge management as lessons learned databases and other online tools for locating expertise. Such online tools are indeed one valuable category of an organization's resources, practices, and processes for capturing and sharing knowledge, but they are only one of many. It is my hope that the articles in this new journal will give some sense of the variety and scope of knowledge services and activities at NASA.

Databases, repositories, archives, and other online sources preserve and make accessible expertise over time. At NASA, where almost everything is built on learning from missions, the value of most projects and programs can last for many years if they are advanced (and updated) by communicative teams committed to sharing what is being learned. What is at the heart of knowledge services is gaining know-how and know-what from practitioners reflecting and sharing stories and insights at the core of their knowledge.

Too often, though, critical knowledge is omitted or cannot be fully captured in documents, and it is difficult to understand knowledge out of context. We have all experienced moments when finding what you are looking for may seem difficult or impossible.

This first issue of the *NASA Knowledge Journal* focuses on real work being done at NASA to overcome these challenges. The first article describes the successes of our video channel *JPL Tube*, with a new feature providing keyword searching. A related article follows the innovations around search and accessibility, illustrating the importance of search methodologies and data-driven visualization.

Much expertise, especially the subtleties of how to carry out complex work, can only be effectively shared person-to-person, by working together or talking together, despite all our best technological advances. Goddard Space Flight Center's Pause and Learn sessions are an important example of the power of face-to-face knowledge sharing. Similar stories shared here include those on physical space for sharing, such as Spaceport Innovators at Kennedy Space Center, and Goddard Space Flight Center's process of learning from the cancellation of GEMS.

The motivation for people to capture, transfer, and apply knowledge is crucial to any knowledge sharing enterprise. The effectiveness of intrinsic vs. extrinsic rewards for sharing knowledge remains central to achieving our goals. It is part of what makes a healthy knowledge culture, which includes sharing stated goals, recognition of contribution, mutual trust, and tolerance for and recognition of the value of mistakes. This journal addresses those challenges. NASA remains at the leading edge of computation, as a concluding article shares what IBM's Watson, a supercomputer, is realizing at Langley Research Center.

What is captured in this journal is representative of what NASA is doing well now and has done in the recent past. For more on knowledge services at NASA, please visit our website **km.nasa.gov** and follow **NASAKnowledge** on Facebook and Twitter.

Warmly,

Ed Hoffman
NASA Chief Knowledge Officer

JPL Tube Pioneers Searchable Video Capture

BY DAVID OBERHETTINGER

The NASA/Caltech Jet Propulsion Laboratory (JPL) is the lead NASA field center for the robotic exploration of the solar system. JPL recognizes the importance of capturing and retaining key knowledge related to the design and operation of spacecraft systems. Once an aircraft manufacturer loses a key system engineer, there are likely competitors from whom to obtain a replacement. But once JPL fails to adequately document knowledge on an esoteric topic such as Mars Entry, Descent, and Landing (EDL), there exists no similar organization that can supply such expertise.

The Curiosity rover was subjected to stringent sterilization procedures.

"VIDEO IS THE NEW DOCUMENT"

Due perhaps to the pervasiveness of e-mail and presentation tools, NASA engineers today may be less disciplined in preparing formal technical memoranda than their predecessors 20 years ago.[1] Five years after a decision is made, when the mission launch date arrives, no one may be able to answer the question, "So why did we set the value of that software alarm to 'Off' for launch?" Video recording offers significant advantages over documents in capturing technical decision making and training: It requires no scribe, and it more accurately captures the nuances and emotional content of interpersonal communications. And with high-definition cameras available in cellphones, video recording today requires little specialized equipment or expertise. Arranging for a professional JPL videographer to film a one-hour meeting in a conference room may cost $1,000, in contrast to quite adequate video and audio quality from a cellphone propped up on the conference table. Once recorded and archived, however, video documentation is not very accessible. It is usually impractical to find a key piece of information from within hundreds or thousands of hours of video content.

JPL TUBE

Hence, JPL teamed with Microsoft Research and NASA IT Labs to develop *JPL Tube*, an online tool with a format and operability similar to YouTube®. Any U.S.–based person within the JPL firewall can upload or view videos. As a video is uploaded, *JPL Tube* automatically generates a scrolling transcript that keeps pace with the frames in the film. This transcript aids in user comprehension of the technical material. Also, this produces significant labor savings in achieving compliance with the Rehabilitation Act of 1973, which requires Federal agencies to make their electronic and information technology accessible to people with disabilities. But more significant, associating a textual transcript with a video makes the video full-text searchable. A search of the site by a user produces one or more video snippets that contain the search string. Clicking on a selected snippet keys the user to a specific video, and to a specific point within the video where the topic is discussed.

"This transcript aids in user comprehension of the technical material… associating a textual transcript with a video makes the video full-text searchable. A search of the site by a user produces one or more video snippets that contain the search string."

JPL teamed with Microsoft Research and NASA IT Labs to develop JPL Tube*, an online tool with a format and operability similar to YouTube®.*

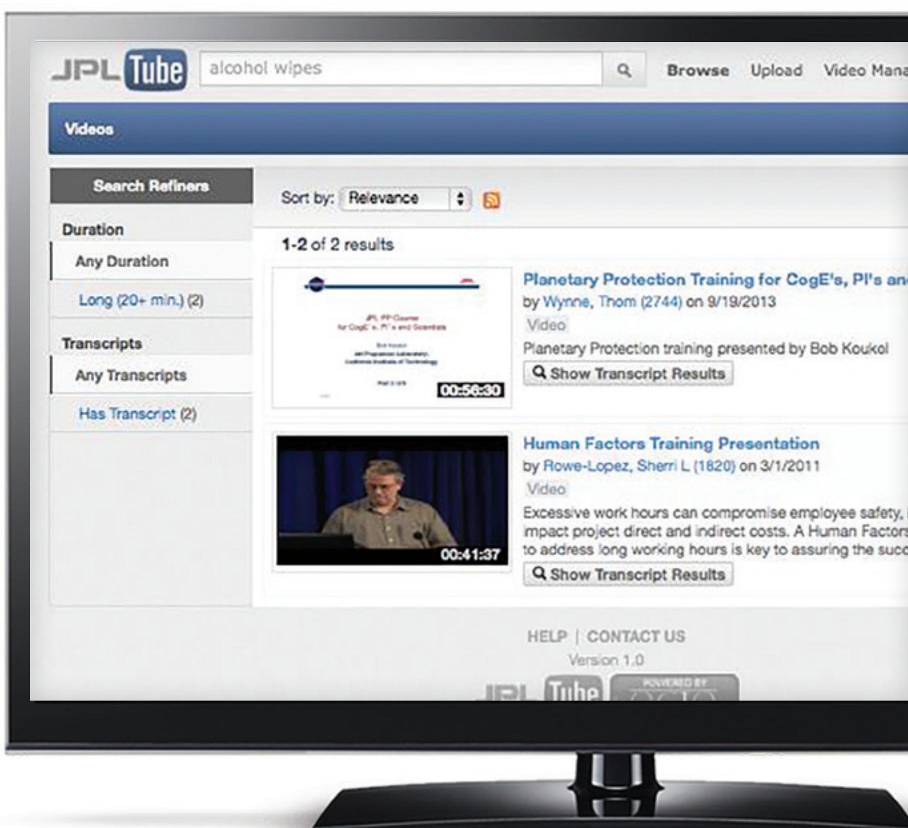

[1] G. Robinson, NASA Office of the Chief Engineer, Memo No. 6208 ("Documentation"), September 16, 2005.

For example, NASA policy and good practice requires JPL to sterilize spacecraft surfaces to prevent contamination of alien environments by Earth organisms. Hence, a *JPL Tube* user might ask, "How effective are alcohol wipes in sterilizing spacecraft?" Robert Koukol taught an eight-hour seminar at JPL in 2013 titled "Planetary Protection Training for Cognizant Engineers, Principal Investigators, and Scientists." The seminar was recorded and posted to *JPL Tube*.

If a user were to enter the search string "alcohol wipes" into the *JPL Tube* search window, the search results would include two video snippets. One of the two is a snippet from approximately two hours into Mr. Koukol's presentation.

The transcribed video snippet includes the following remarks by Mr. Koukol:

"Alcohol wipes! Everybody says, 'Oh, so the alcohol kills the organism?' Ah, if you want to save bacterial endospores, if you want to save the spores, a really good way to do it is to store them in alcohol. So it really doesn't kill off the spores at all. Now, it will kill off vegetative cells—the kind of stuff we have on our skin that we are worried about. Or if you know one of the big things now is that everyone has anti-microbial soaps and anti-microbial sprays, and everything else. You go to the grocery store and they have this little thing and you pull the wipe down and wipe the handle on the shopping cart…"

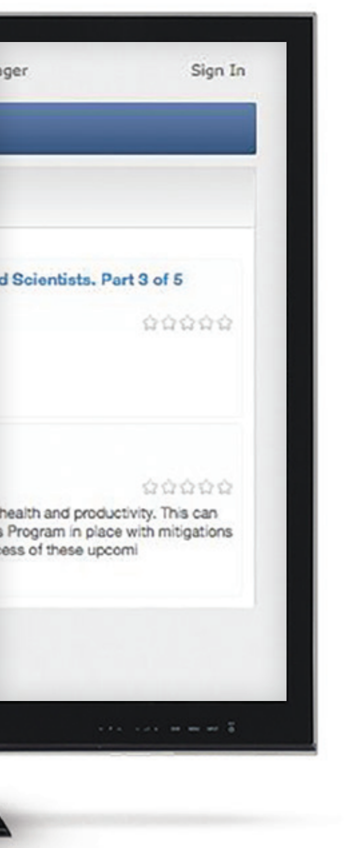

MICROSOFT'S MAVIS

For the transcription, *JPL Tube* employs the Microsoft Audio Video Indexing Service (MAVIS®), which uses Deep Neural Net (DNN) technology and Probabilistic Word-Lattice Indexing to convert digital audio into text. Artificial neural networks employ statistical algorithms that mimic biological nervous systems in the way that they learn by recognizing patterns within a large amount of input data. In contrast, more conventional speech recognition technology requires multiple passes over an audio file, providing incremental improvements in the transcription. Microsoft claims that the DNN approach offers greater speed and accuracy.

JPL is the first subscriber to employ MAVIS in a complete system that allows untrained users to add videos to the collection. Uploading a video and adding related metadata takes a few minutes. MAVIS is a product of Microsoft Research; under the current MAVIS license, JPL is charged about $10 per hour of video. *JPL Tube* came online in 2012, and JPL personnel have presently uploaded over 1,200 videos. Over the quarter ending September 30, 2015, an average of 324 videos per day were viewed by JPL users of the system. Because *JPL Tube* access is limited to U.S. persons at JPL, JPL does not require International Traffic in Arms Regulations (ITAR) clearance before videos can be posted. This has eliminated a potentially time-consuming and expensive constraint to video capture.

JPL recently incorporated a number of *JPL Tube* enhancements, including:

- An increase in the maximum permitted upload size from 4GB to 8GB. Previously, a multi-hour video had to be uploaded as discrete segments.
- A capability for the contributor to return to the uploaded video and edit the transcript.
- A multi-threading capability for *JPL Tube*, in which multiple videos can be uploaded and processed simultaneously.

NASA TUBE

NASA asked JPL to develop a *NASA Tube* with similar features to the JPL product. A pilot version of *NASA Tube* went live in September 2015 with 141 videos pre-loaded. Due to some NASA security compliance issues, the pilot version does not at this time allow NASA personnel to upload additional videos. But they can review the other tool capabilities at http://nasatubedev.nasa.gov.

JPL and NASA missions are high-risk endeavors because the spacecraft are typically one-of-a-kind, high-unit-value engineering designs intended for operation in an extremely hostile environment. Lost knowledge presents not only the additional cost of restoring it but also a higher risk because the legacy knowledge had been verified through system test and spaceflight. *JPL Tube* has the capability to transform the retention of technical know-how by capturing key knowledge at a very reasonable cost without the need for a human scribe. The utility and benefits from *JPL Tube* processing, storage, and search have been demonstrated for more than two years, and JPL is eager to extend this capability Agency-wide.

David Oberhettinger is the Chief Knowledge Officer at NASA Jet Propulsion Laboratory.

Better Access to Critical Expertise

BY DAVID MEZA

One question I repeatedly receive from employees at NASA is, "Why can't I ever find what I need?" People are accustomed to instant access to information from the World Wide Web through the multitude of search engines available. They cannot understand why it is so difficult to replicate this ease of access at work. The short answer: numbers. Google, for example, indexes billions of URLs and documents and accepts 5 billion queries per day. Those huge numbers make it possible for them to continually refine their search capabilities. At Johnson Space Center (JSC), our search infrastructure indexes less than 1/10 of our known data and accepts a few thousand queries per month. Even if we could afford to index all the millions of documents JSC has, I am not convinced our search results would greatly improve, because the query sample set would remain small. Given this problem, the NASA-JSC Knowledge Management office has embarked on a project to improve the accessibility and visibility of critical data at NASA. The project is broken down into three phases:

1. Understand the employee base.
2. Research and evaluate search methodologies.
3. Provide data-driven visualization.

The need to improve search at NASA is clear and critical. A recent report by IHS, a company providing business information and analytics, underscores my greatest concerns for the NASA workforce. It reveals that an engineer's average time spent searching for information has increased 13 percent since 2002. Additionally, 30 percent of total R&D spent is wasted on duplication of research and work previously done, and 54 percent of decisions are made with incomplete, inconsistent, and inadequate information. I fear the problem is only growing as the quantity and complexity of information continue to increase. According to Gartner, the world's leading information technology research and advisory company, the amount of available data will increase 800 percent by 2018 and 90 percent of it will be unstructured. Our project seeks a NASA-specific solution to the problem.

UNDERSTANDING THE NEEDS

To understand the employee base, we conducted interviews to understand their needs and frustrations. We learned that employees with different common job profiles search differently. For example, a scientist or researcher typically explores data looking for common themes and connections, while an engineer or project manager is more likely to be looking for specific answers to pressing questions. A semantic approach—one that considers context and related terms to determine meaning—is more useful for the former, while a computational or cognitive approach is more useful for the latter. Our team then set out to develop use cases for the identified job profiles to determine what they searched for, how they wanted results presented, and where the content they searched for was located. That process led us to identify both the data critical to our users and the authoritative sources for the data.

METHODOLOGIES

Our benchmarking revealed that, on average, technical professionals must consult 13 unique data sources to get the information they need to make informed and consistent decisions. Understanding the different search methodologies has been crucial to making improvements. We worked with a number of vendors, developers, and search gurus to evaluate the capabilities of semantic, faceted, cognitive, and computational search with the use of natural language processing and advance text analytics,

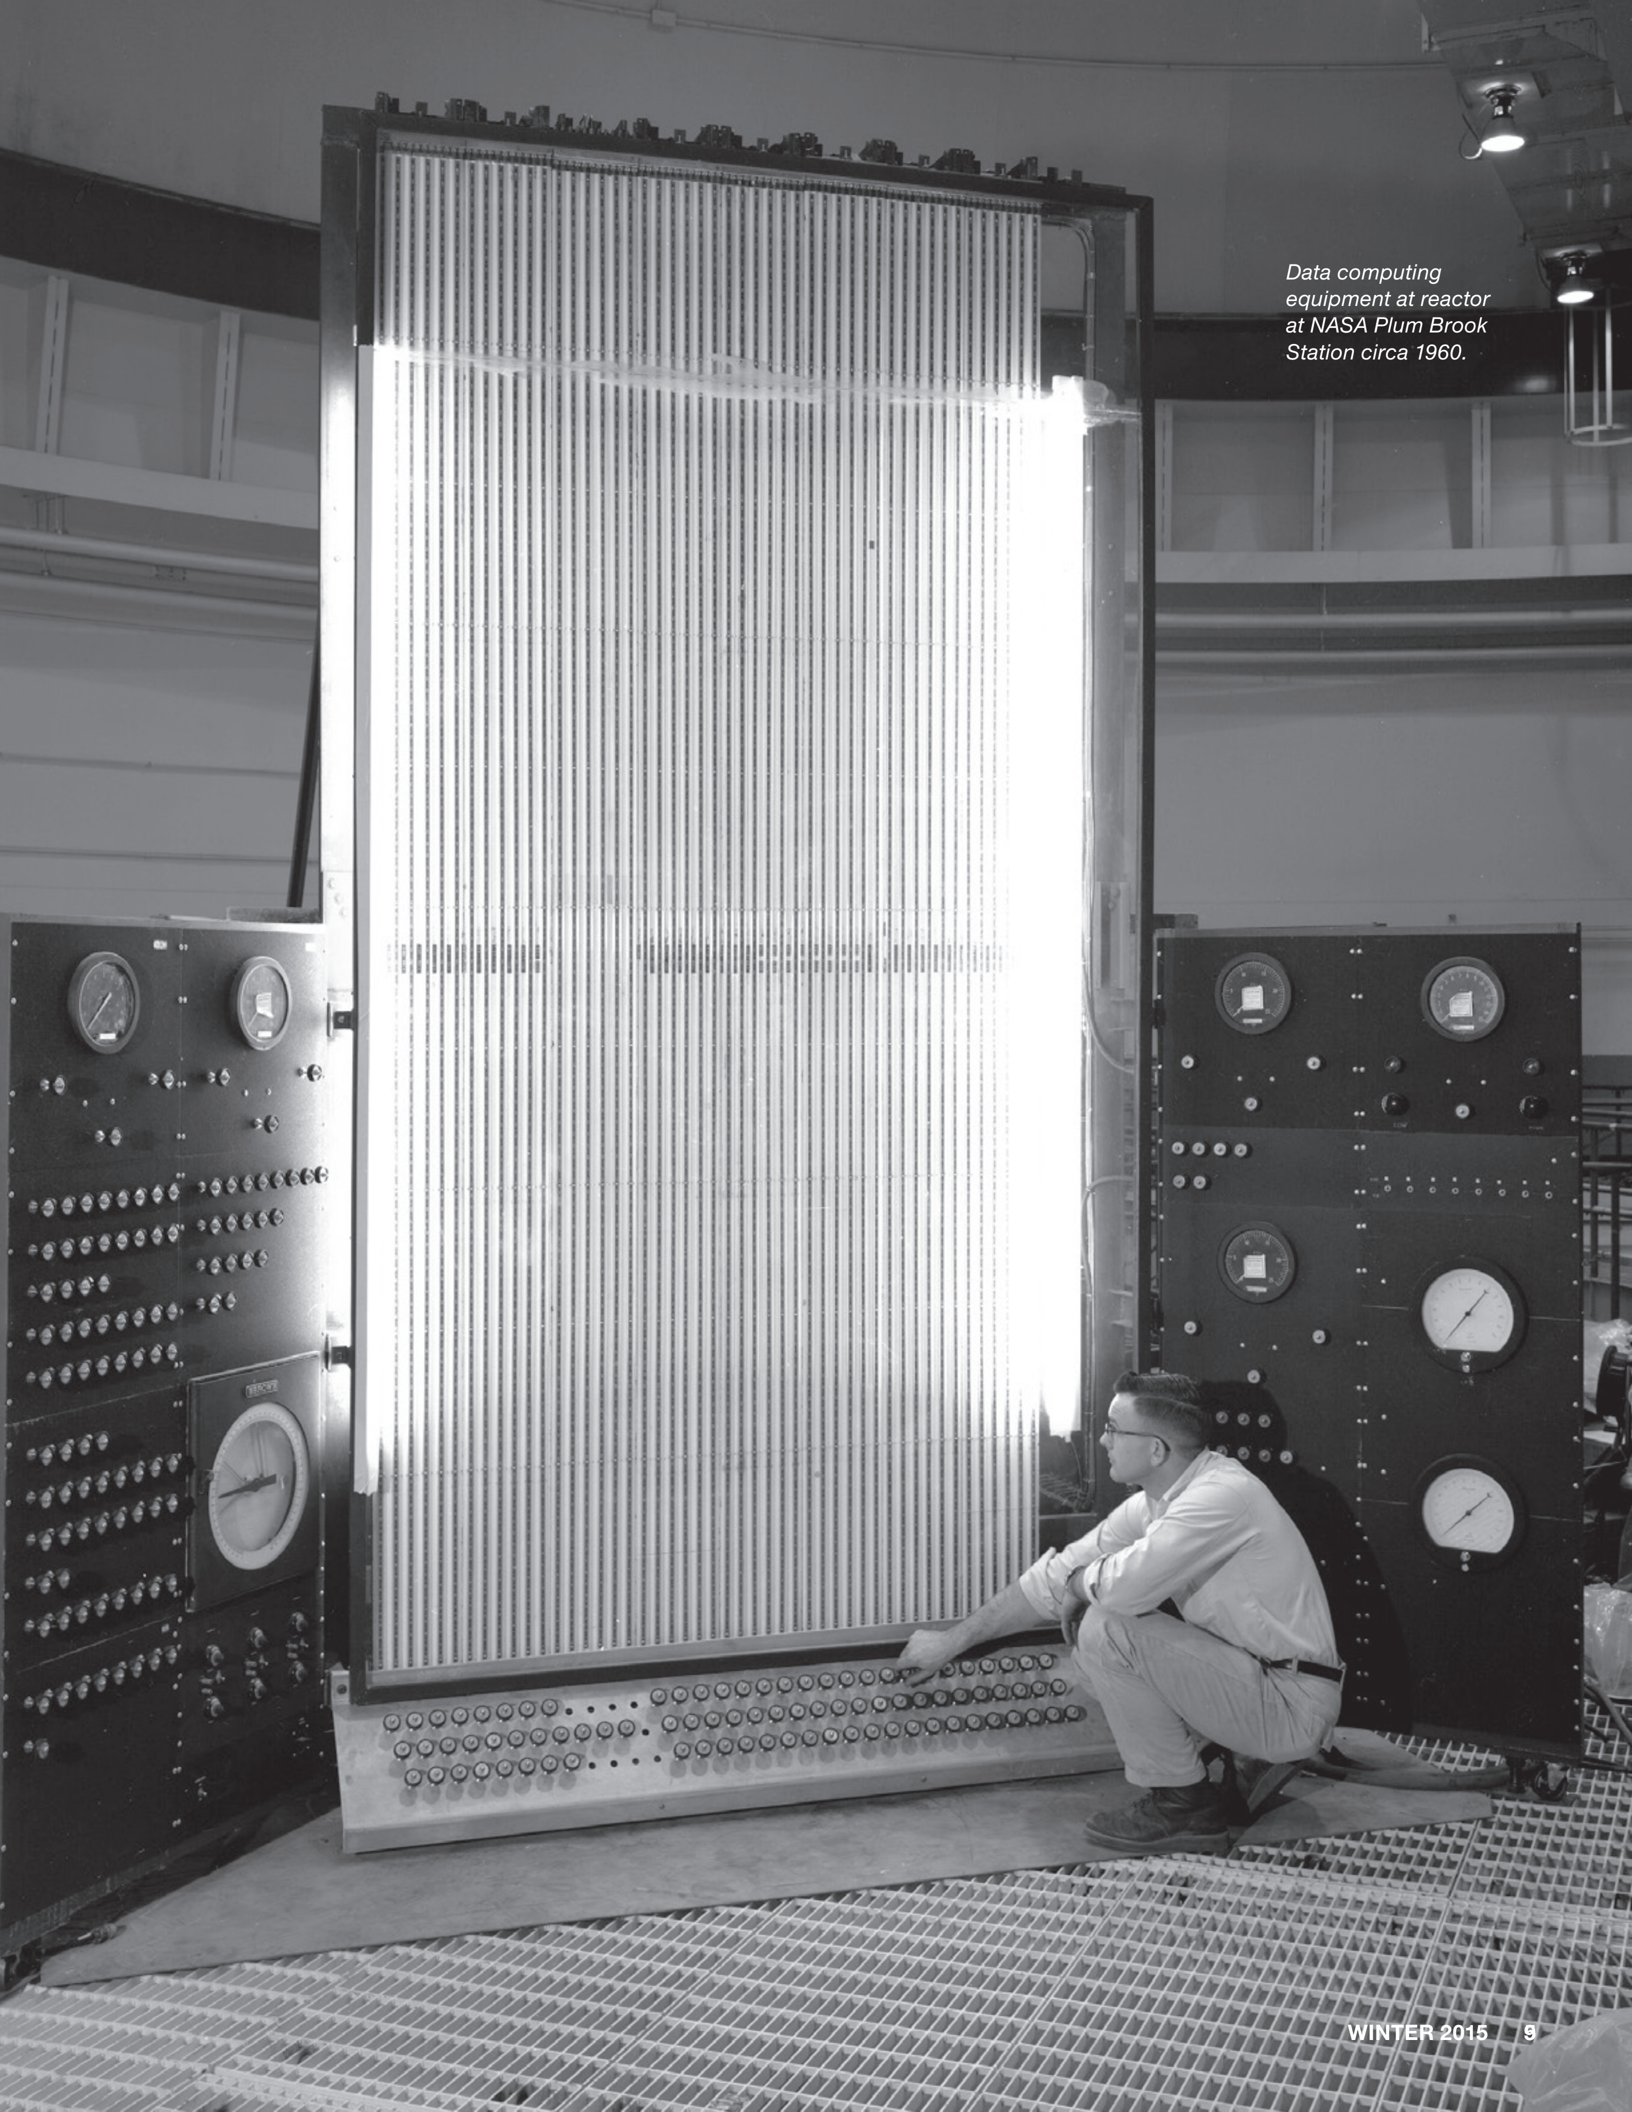

Data computing equipment at reactor at NASA Plum Brook Station circa 1960.

comparing their search results to our current keyword search engine. The team began to realize there is no single solution for search that will meet all the needs at JSC. All of those approaches have their pros and cons; which of them are best in a given situation comes down to the type of query and the data being searched.

We arrived at some important conclusions:

- A master data management plan is essential. A common process for metadata, format, and structure has to be implemented to improve the results of the search query.
- With the quantity of data growing exponentially, focusing on identifying our critical data is imperative.
- Once identified, standards should be in place for how data is created, where it is stored, and who has access to it.
- The use of analytics to explore, explain, and exhibit the data is essential.
- Information should be presented in the manner most useful to the user.

VISUALIZATION

The last phase of our project, in process at the time of this writing, focuses on how the data influences presentation, or, as I call it, *data-driven visualization*. This is not a new term, but one that took on more meaning for me as the project progressed. Users want more than the common hierarchal list based on keyword relevancy; they also want to visualize connections, themes, and concepts. The advent of social networks, YouTube, and ubiquitous access has transformed the way we find information and the way we learn. A 5-minute video replaces an instruction manual and #stayinformed replaces the storytelling and news. Fortunately, our technology development has kept pace with users' requirement to visualize data differently. Network analysis, graph databases, HTML5, and text analytics are but a few of the capabilities available to transform the hierarchal list into a conceptual framework that allows our users to start with a natural language query, visualize a network of concepts and themes, and traverse the connections quickly to find their desired answers much more quickly than in the past.

We have applied this approach to a collection of over 2,000 lessons learned, filterable only by date

SAMPLE TOPIC MODEL VISUALIZATION

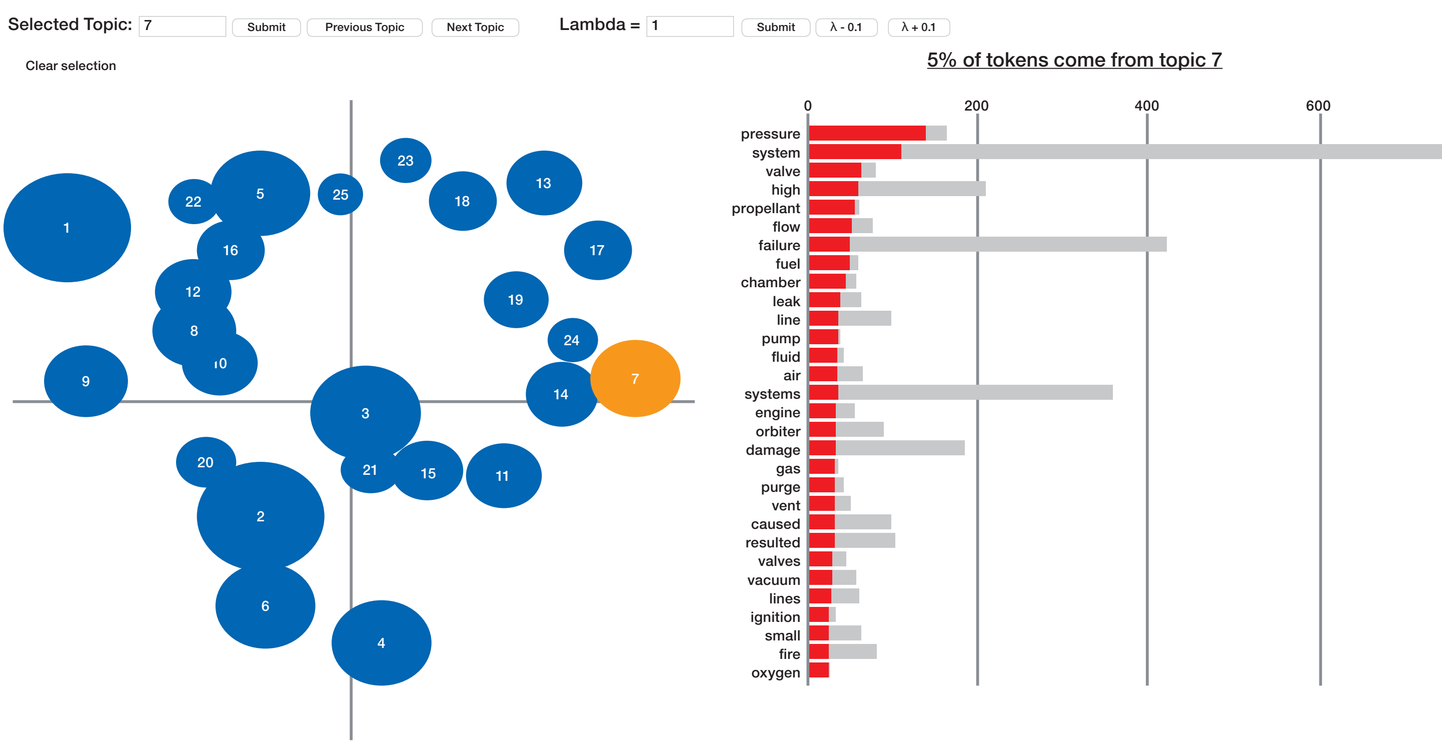

NASA'S KNOWLEDGE MANAGEMENT GRAPHIC MODEL

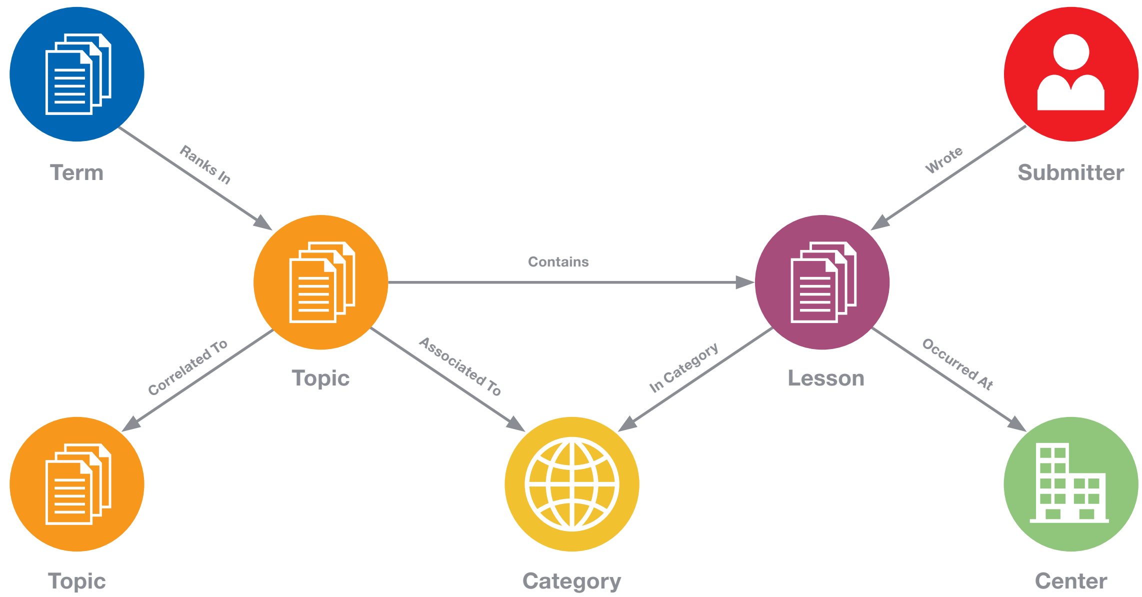

and center. A common complaint we receive from users is how time consuming and difficult it is to search through and collect information from the lessons learned. For example, a keyword search on "pressure valve failure" returns a list of over a hundred documents. A majority of them have only one of the keywords, but a user would have to read through the document to determine if it was relevant to them.

Utilizing a topic-modeling algorithm know as Latent Dirichlet Allocation, we analyzed the documents. Topic modeling provides an algorithmic solution to managing, organizing, and annotating large archival text. The annotations aid in information retrieval, classification, and corpus exploration, which, in the context of topic modeling, uses text analytics to discover concepts' associations, correlations, etc., drawing out latent themes based on the co-location of terms within the documents.

Topic models provide a simple way to analyze large volumes of unlabeled text. A "topic" consists of a cluster of words that frequently occur together. Using contextual clues, topic models can connect words with similar meanings and distinguish between uses of words with multiple meanings. A topic, which is a probability distribution over the words in the document, is assigned to each document. Documents may be assigned multiple topics, but one topic will have a greater probability. The topic is identified by its associated terms and their probability. Now, when I look at the topic containing "pressure, valve, and failure," I know there is a higher likelihood that the documents associated with this topic pertain to pressure valve failure and not just one of the words. Additionally, topic modeling in this example drops the number of documents to review to fewer than 20.

Another benefit to this approach is the ability to correlate the topics with each other and with metadata, such as Project Phase, Center, or Safety Issues. By correlating topics to each other, users can quickly expand their investigation to other topics that may have useful ancillary information. We accomplish this by applying data-driven visualization techniques. In our preliminary testing, users have been able to jump easily to a topic of interest, cutting their search time dramatically. We are still doing testing to quantify the time reduction.

More work still needs to be done, but the team is excitedly looking forward to the future. In the next phase, we will be conducting pilot tests of various visualization platforms to determine capabilities, feasibility, and scalability of the systems. Once these tests are completed, we should have a better idea on the requirements necessary to scale the systems up for use across NASA.

David Meza is the Chief Knowledge Architect at NASA Johnson Space Flight Center.

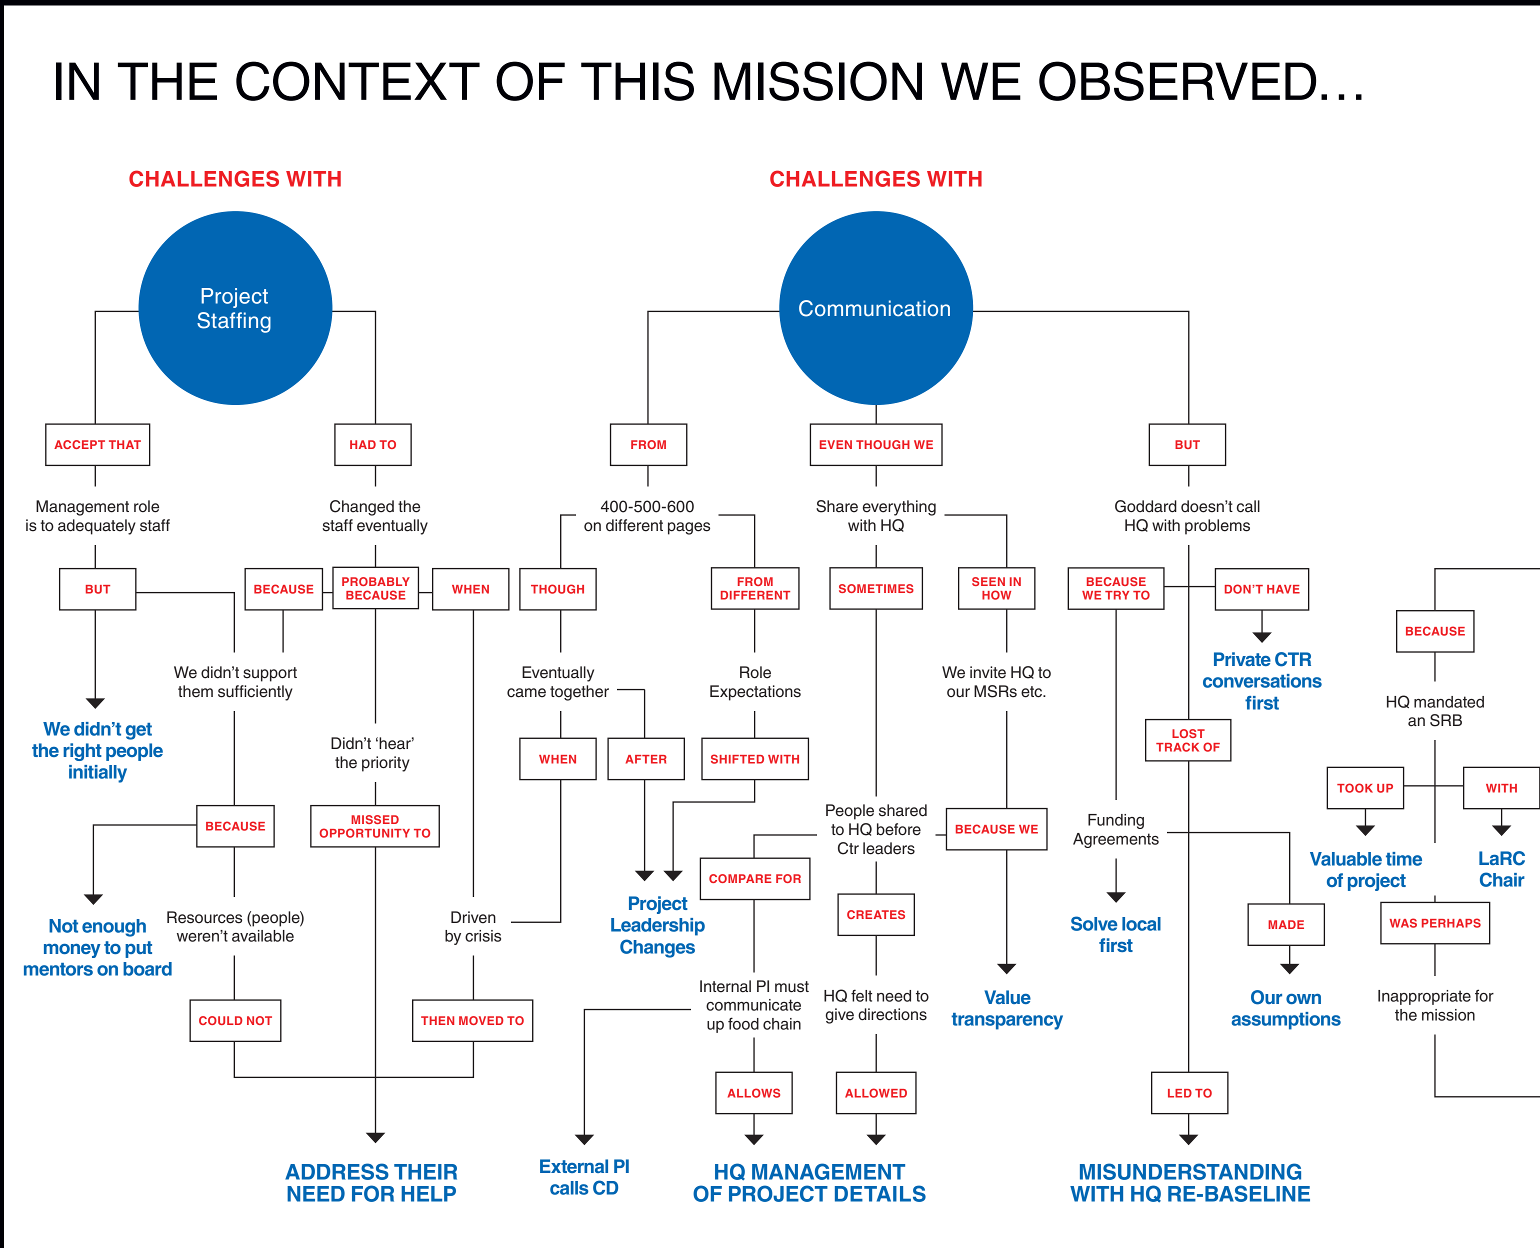

Learning Lessons from GEMS

BY ED ROGERS

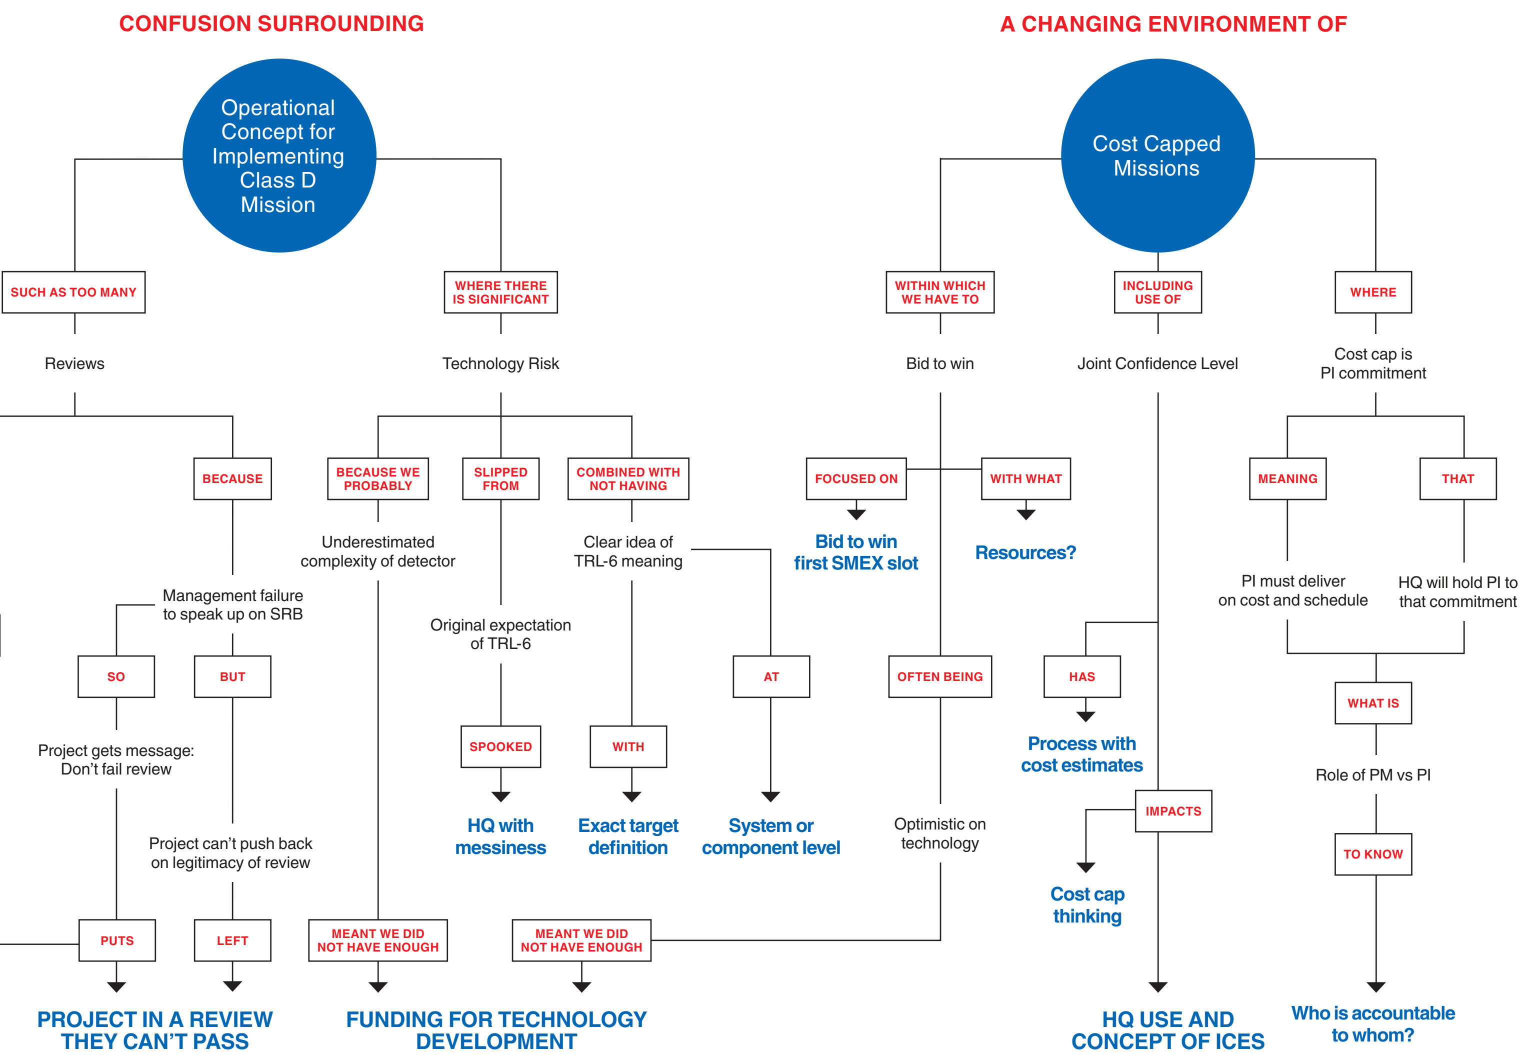

The Gravity and Extreme Magnetism SMEX (GEMS) mission failed to pass confirmation in May 2012. Among other things, the mission would have helped determine how spinning black holes affect space-time. The mission had been competitively selected for Phase A Study efforts in May 2008 and then selected for mission award in June 2009. Along the way, the project got behind and overran its budget. HQ eventually came to believe that it could not be completed on time and within cost, and GEMS was cancelled. The story of the mission itself is captured in a case study written at Goddard in the aftermath of processing the lessons learned. What I would like to consider here is *how* we went about getting lessons from this experience and making sure those lessons were socialized and integrated as widely as possible. I hope that story may help others learn how to get lessons from difficult experiences and reapply them.

The "failure to confirm" decision created a lot of angst at Goddard. Specifically, we struggled to understand how the Center could expend all that effort to win a mission competitively and then fail to execute well enough to keep it going. The Center

Director wanted answers and lessons so that something like that would not happen again. He directed me to oversee the process to "make sure we get the real story and understand it." Initially, the failure led to a lot of finger pointing: within Goddard, between HQ and Goddard, and even at external players. We needed a better approach that would get beyond the blame game and facilitate genuine reflection and learning. The approach I took had four distinct phases:

1. Hold Pause and Learn (PaL) sessions to air the issues within the project team.
2. Generate concept maps to capture the context of the story.
3. Share the maps to generate conversations at both the Center and NASA HQ.
4. Develop the story into a case study that could be taught to future groups and used in NASA training.

The PaL with the GEMS project team and the concept maps that emerged from it revealed that there were unresolved issues with how Center management and NASA HQ dealt with the project. This finding suggested we needed to hold a PaL with Goddard senior leaders to elicit their views. Center director Chris Scolese agreed to use an hour of one of his staff meetings to discuss GEMS. That PaL allowed the Center leadership to openly discuss ways in which they had not supported the GEMS team or been as responsive as they needed to be to keep the mission viable.

Fortunately, I had been sent on a detail to Science Mission Directorate (SMD) at NASA HQ around this time and had good access to the senior leadership within SMD. So we were also able to hold a PaL with leaders at HQ. These three PaLs produced nine concept maps, three from each PaL to capture answers to three basic questions:

1. What happened? What did we observe going on?
2. What decisions did we make? That is, what was our role?
3. What should we do differently? What were the lessons we learned?

The maps in each set were titled Observation, Decisions Made, and Lessons Learned.

The map on the previous pages gives some sense of the patterns of cause and effect these tools reveal. This particular map details the effects of center management decisions regarding staffing, communication, and the operational and budget challenges of a Class D mission.

Each group reviewed its own maps but also studied the maps from the other groups. This process helped socialize the different views and make clear what different people saw as significant parts of the story.

I then held a joint session with the project team and Center leadership to gather around the various concept maps and discuss what we should learn from them. These conversations were both therapeutic and instructive. The project team learned that the Center leaders in fact took responsibility for their role in the failure and wanted to learn how to improve. The project team learned things about the way they operated that contributed to the premature end of the GEMS mission. They both learned the point of view of NASA HQ, which was very different from the initial blame talk going on at the Center. The conversation took on the much more productive tone of trying to understand where others were coming from and what we could learn from the experience. Both the project and Goddard leadership wanted to involve NASA HQ as well, so a further joint meeting on GEMS Lessons Learned was held at NASA HQ.

The joint GEMS Lessons Learned Session was held at HQ on November 5, 2013. The room was packed with Goddard Center leadership, HQ SMD leaders, the science PIs, GEMS project management, and project team personnel. I had gone through all the maps and consolidated the comments into six categories of what I judged to be the major common issues around which lessons needed to be learned. I checked these with key leaders of the project, Goddard and HQ prior to the meeting. I then presented these six broad lessons for discussion to the joint group at HQ. I posted the large concept maps around the room as "evidence," though everyone in the room had already seen them.

THE SIX LESSONS WERE:

LESSON 1: Recalibrate the Small Explorers Mission (SMEX) Model because it is difficult to use a SMEX, with its tight cost caps, to mentor inexperienced leads. To execute a SMEX now requires very experienced managers who are typically engaged on larger missions. Re-examine the funding profile for practicality rather than just forcing the available funding into a profile that may or may not be executable by the mission.

The map on the previous pages gives some sense of the patterns of cause and effect these tools reveal.

LESSON 2: Keep the larger science community aware of status and the context of changes. Keep them involved in descopes/replans and recognize that the story of compelling science must be retold frequently or it may lose its shine. Keep the Program office in the loop to HQ and Center management and keep vendors and outside stakeholders up-to-date on mission status and descopes.

LESSON 3: Clearly define cost cap expectations. Have a kickoff meeting on Independent Cost Estimates between stakeholders (including ICE developer) to establish a joint understanding and clear process. Clearly define the nature of review meetings regarding cost. Clearly specify when the cost cap is being revisited.

LESSON 4: Pay closer attention to staffing, especially during the early phase of mission when it is critical to get ramped up and not fall behind schedule or run over cost. Don't wait until a crisis emerges to address staff needs. Manage key staff changes, being careful to avoid mass (all-at-once) changes. Take time to establish expectations and understandings.

LESSON 5: Clarify in-house PI mission management. Specify the PI role in communication with HQ and how it relates to the communication roles of project manager and Program Office. Conduct practical PI training at the Center, setting expectations and important role functions.

LESSON 6: Establish a clear process and expectations for Joint Confidence Level and specify how they will figure into HQ Gateway Decisions. Distinguish between the JCL and ICE role for the mission.

As I got up to present the summary charts of the top six lessons we had all agreed on (from the map discussions), I laid down one ground rule: Anyone in the room could suggest actions, but only in regard to their own work. In other words, there would be no sharing of gripes about other organizations couched as lessons learned. All the leaders stepped forward and owned the lessons from GEMS. Each of the key players took their own action items and lessons from the general six presented. Later, I completed the case study and released it for use at Goddard. We use it internally in a leadership training program for people all across the Center. The original project manager comes and shares personal insights to the story. It is now available on the public site of case studies from the GSFC OCKO.

So what happened? Did we learn the lessons? Well, GEMS became a common topic of discussion within SMD and Goddard and shorthand for kinds of mistakes we did not want to repeat. HQ owned up to their actions that had contributed to the outcome, as did the Center management. With the burden of "it's all our fault" removed from the heads of the project team, they also were able to acknowledge missteps and misunderstandings. Once blame and its negative consequences are taken out of the equation, people become able to examine and admit their mistakes. Even the contractor weighed in with their own lessons learned on what they saw and should have done to help mitigate the situation.

In the end, it's hard to calculate how many people were affected by the lessons of GEMS. What was critical was that many people saw their leaders actively learning lessons right along with them. The blame culture disappeared, accountability increased and the likelihood that a GEMS will be repeated is much lower on all fronts. We still have the case study to remind future project personnel of the lessons so they won't wear off with time.

Often in these kinds of situations, people are quick to jump to conclusions with the partial information they have. People are also prone to use existing models or allow biases to influence their thinking. Prior experience can actually cloud our thinking, especially when it comes to learning complex lessons. We are creatures of simplification. A simple story (even if untrue or partially true) will remain in circulation long after the facts have been revealed simply because it's simple and fits some comfortable belief. An important lesson from the GEMS story is that it takes effort to get meaningful lessons—lessons that can actually change future behavior.

Ed Rogers is the Chief Knowledge Officer at NASA Goddard Space Flight Center.

Spaceport Innovators Keep on Innovating

BY DAVID J. MIRANDA AND DON COHEN

Since 2011, a group of Kennedy Space Center (KSC) employees calling themselves the Spaceport Innovators have been organizing and attending talks on a wide range of subjects. They have also put together an annual event called the Innovation Expo at KSC that features speakers and displays highlighting work at the center as well as new ideas that—they hope—will help spark innovation. Over time, membership in the group has grown to almost 400 people, with nearly every KSC organization on center represented.

The fundamental thrust behind Spaceport Innovators' activities is that new thinking springs from combinations of existing ideas that have not been brought together before. So the group's emphasis has been on exposing members to projects and ideas from many areas, especially from fields that they may think are unconnected with their own. From the beginning, the organization has sponsored lunchtime talks by speakers both from KSC and outside the center. Typically, 20 to 30 individuals participate in each meeting.

The biggest challenge has been to attract busy people to talks that they think are not relevant to their work. That stumbling block makes clear that there is still much work to be done to foster innovative thinking, since it is precisely by connecting those supposedly unconnected ideas that innovations happen.

To promote exposure to new ideas, the group has added brief TED Talk–like videos on a variety of subjects to the beginning of meetings. A recent meeting about its new *iGuide* to innovation, for instance, started with a video on jazz. The fact that jazz musicians invent music as they play by listening and responding to each other (rather than just going off in their individual directions) suggests a model for the mixture of freedom and discipline that characterizes innovative collaboration.

EXPOS

To encourage knowledge-sharing on a wider scale, the group has organized annual center-wide innovation expos, which began in September 2012. Laboratories and other facilities at the center offer tours and put together exhibits that highlight their work. In many cases, people learn for the first time what their colleagues in other departments are doing. Organizers hope that those encounters will encourage ongoing conversation between groups and, ultimately, innovative collaborations.

The first one-day event in 2012 also featured a range of speakers on innovation, some from within NASA and some from other organizations including the U.S. Navy, Publix supermarkets, the Walt Disney Company, and Universal Orlando.

In subsequent years, the expo has become a multi-day event. The 2013 expo lasted four days and drew nearly 1,000 KSC employees. It featured as speakers and exhibitors KSC partner organizations including United Launch Alliance, Boeing, Space Coast Energy Consortium, and the University of Central Florida.

INNOVATION SPACES

One of the core principles of innovation is that it happens most often in groups. The idea of the solitary genius creating the next new thing in his private workshop is mainly a myth; collaborations that blend and test ideas from various sources are the driving engines of successful innovation. So it is important to have physical spaces where people can work together and equally important that those spaces support innovative processes. Spaceport

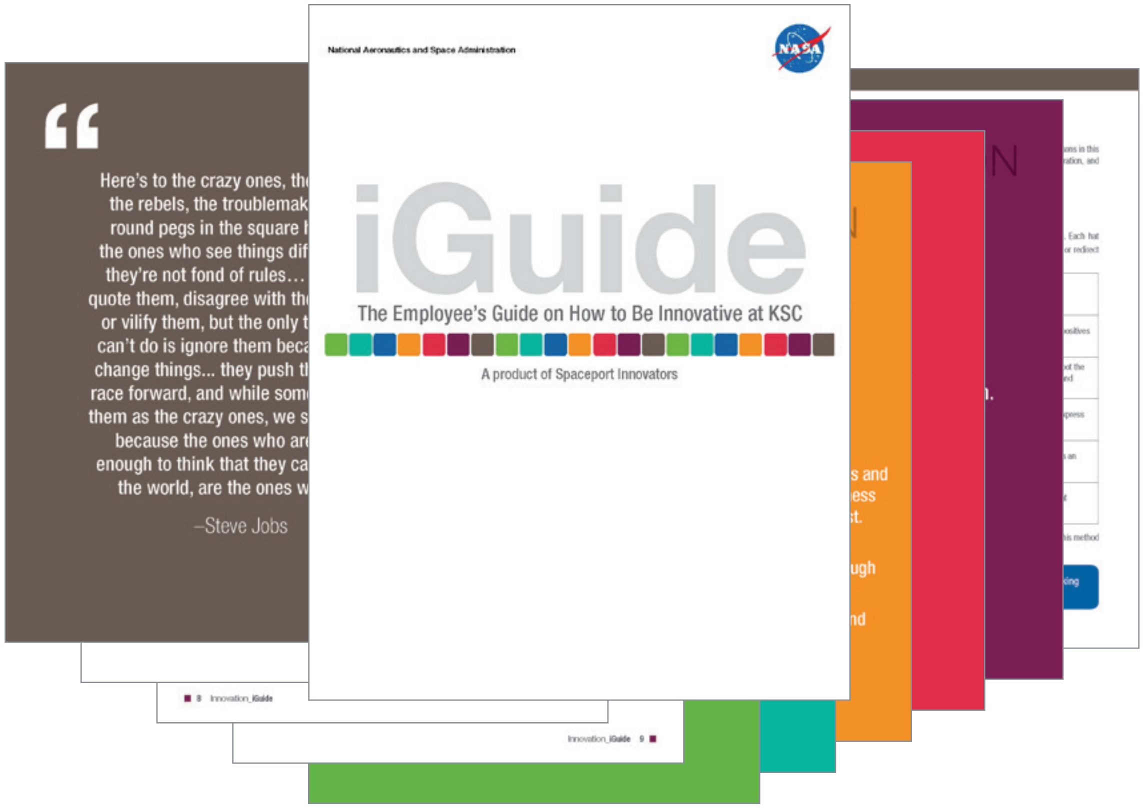

Innovators have been helping to develop new spaces at KSC especially designed for innovation.

One such space was developed for KSC Swamp Works in 2011. It features open work spaces and white boards everywhere to encourage collaboration and conversation. Spaceport Innovators are working on two more spaces, one on the north side and one on the south side of the space center, to make these innovative areas easily accessible to any groups that want to use them.

THE *iGUIDE* AND BEYOND

At the end of 2014, Spaceport Innovators released the *iGuide*, described as an employee's guide on how to be innovative at KSC. It describes processes for coming up with useful new ideas and developing them into innovations. (As the guide says, an idea alone is not an innovation; the idea must be turned into a new process or product that has an impact on people.) The *iGuide* focuses on innovation as a collaborative process and outcome. In addition to offering resources and techniques to encourage innovation, it addresses some of the factors that stand in the way, including fear of failure or ridicule and negativity that suppresses or outright rejects new ideas before they have been tested.

Perhaps as important as the advice the guide offers to innovators is the fact that it exists as a tangible proof of the importance of innovation at KSC. The opening section of the engaging publication reinforces this point with quotations from more than a dozen senior leaders that encourage innovation and emphasize its importance to the future of Kennedy.

Freely available as a PDF file, the *iGuide* has generated positive response at KSC and other centers that are using the guide or planning to develop their own versions.

Spaceport Innovators has been and remains a completely volunteer effort, with no formal charter. Organizing members sometimes struggle to find time away from their other work to devote to it, but they value the freedom that unofficial status gives them to launch new projects without having to go through lengthy approval processes. Their autonomy has made it possible to work on innovative spaces at KSC and to publish the *iGuide*; the group is likely to undertake other pioneering approaches to innovation in the future.

David J. Miranda is the Project Lead for the IDEAS project at NASA Kennedy Space Center.

Don Cohen is Editor-in-Chief of NASA Knowledge Journal.

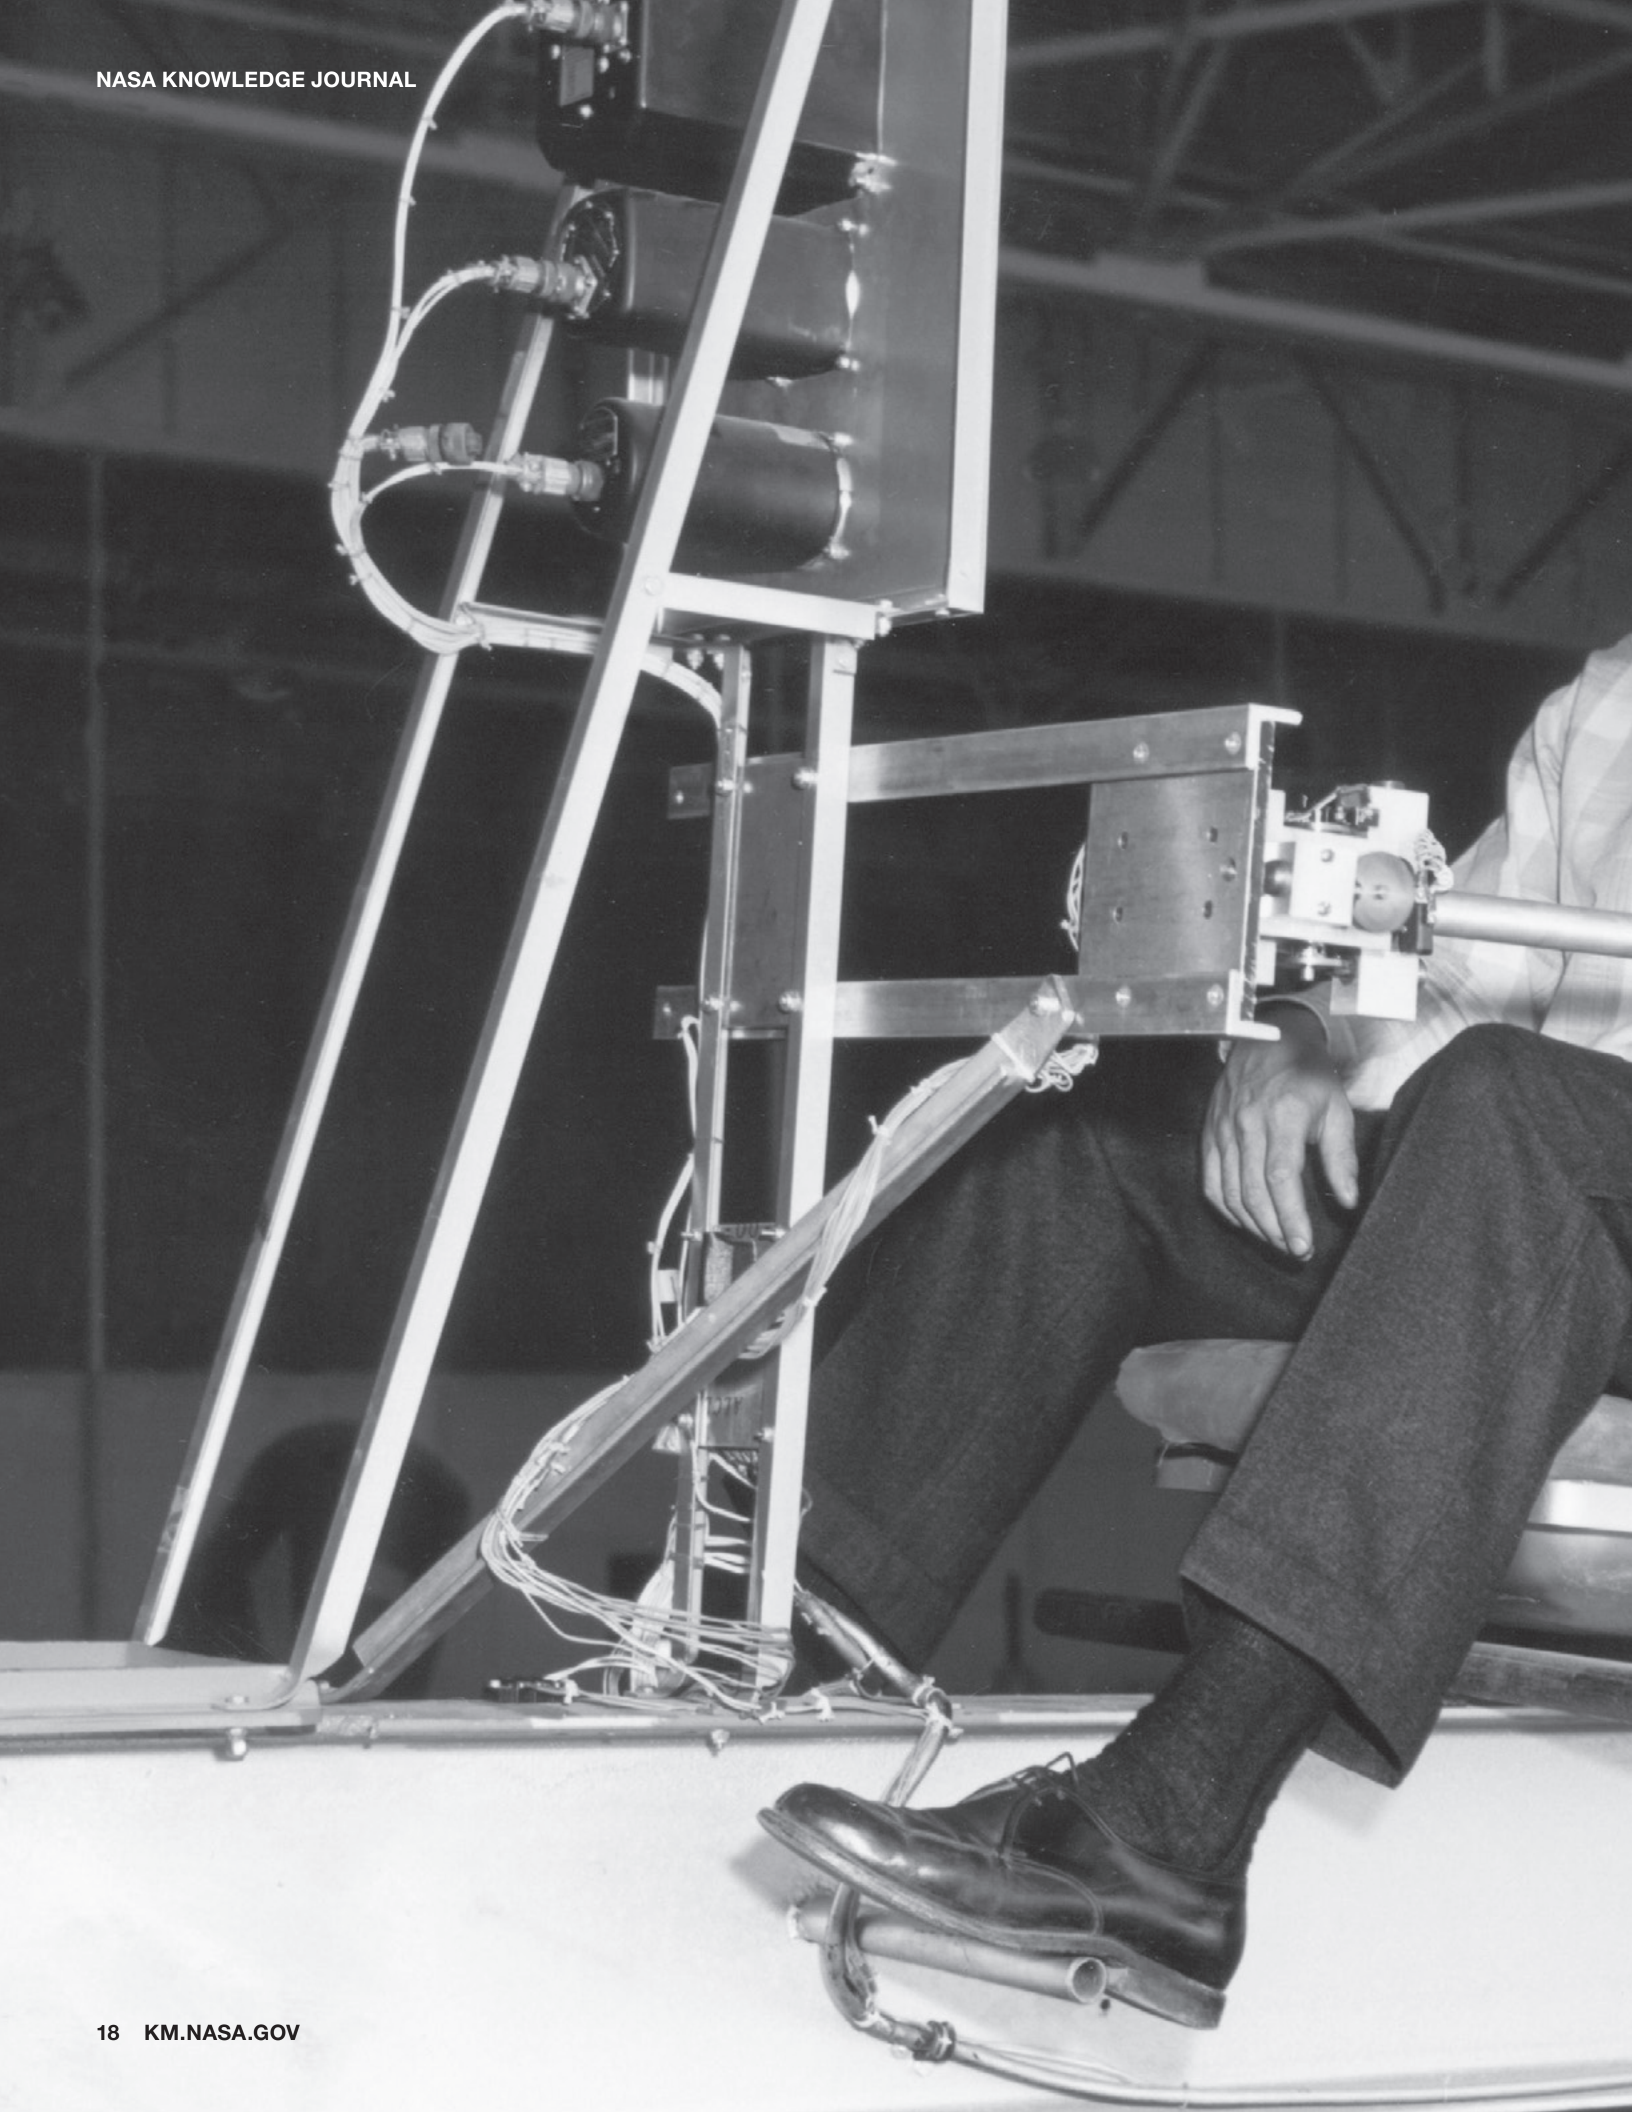

What Motivates Knowledge Sharing at NASA?

BY MICHAEL BELL

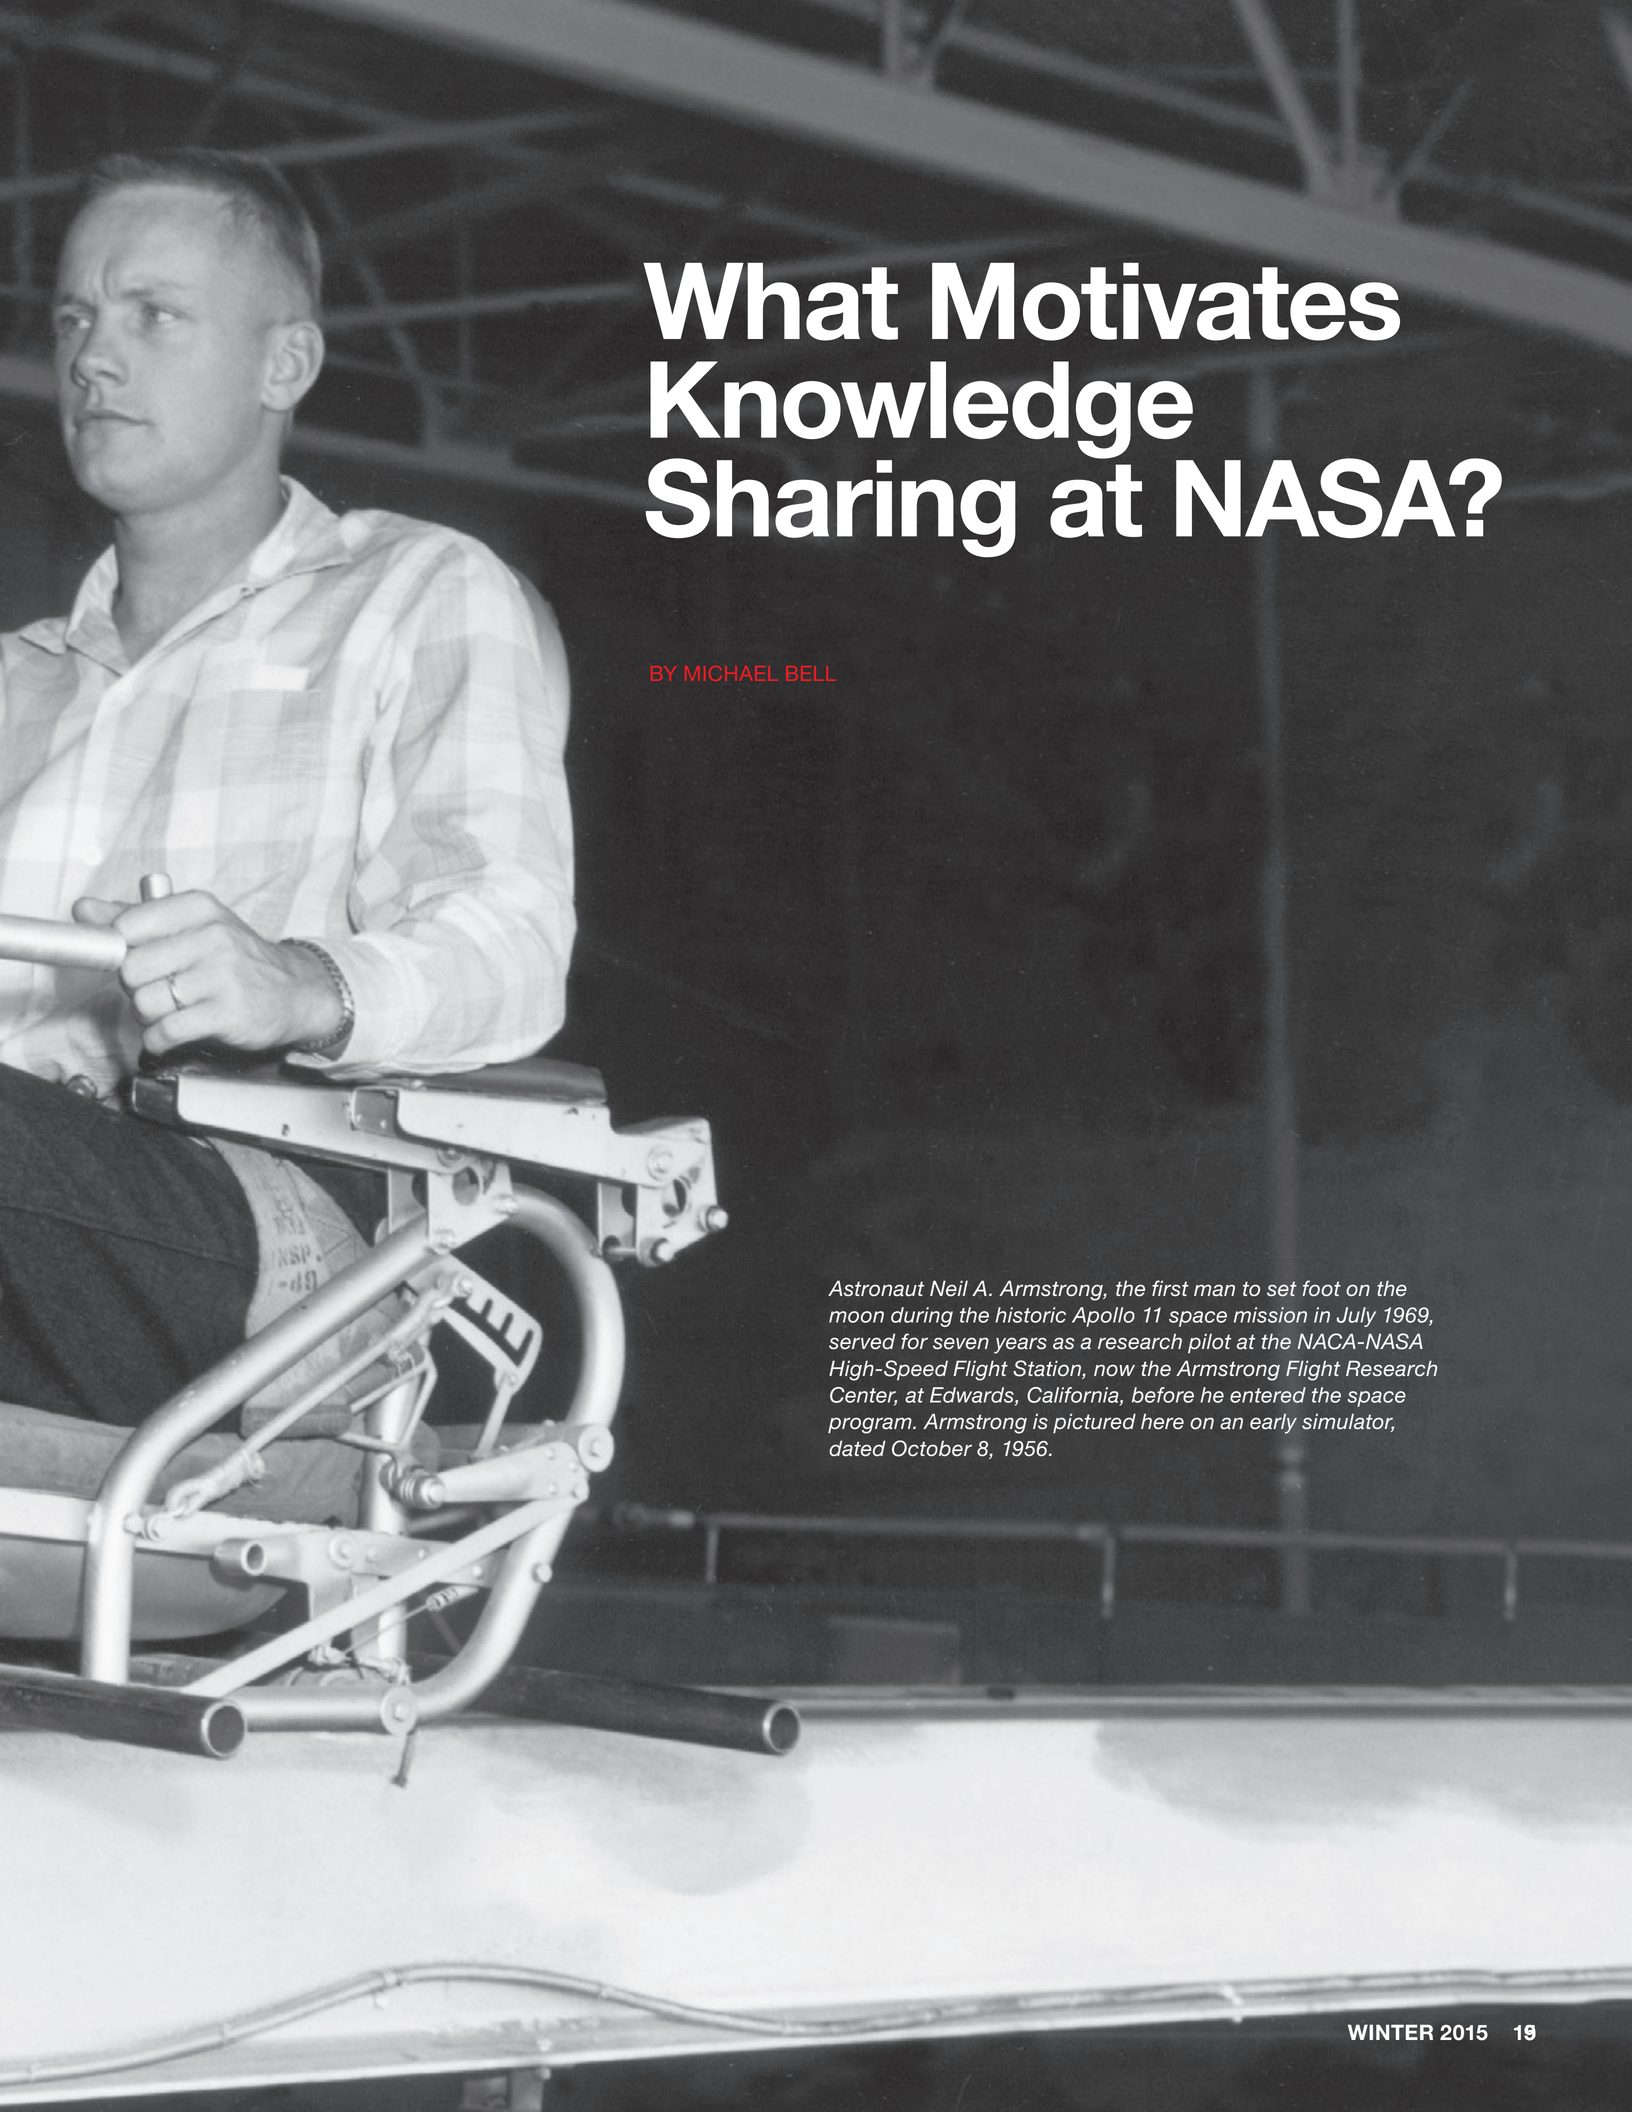

Astronaut Neil A. Armstrong, the first man to set foot on the moon during the historic Apollo 11 space mission in July 1969, served for seven years as a research pilot at the NACA-NASA High-Speed Flight Station, now the Armstrong Flight Research Center, at Edwards, California, before he entered the space program. Armstrong is pictured here on an early simulator, dated October 8, 1956.

From the early days of knowledge management in the 1990s, practitioners have looked for effective ways to motivate knowledge-sharing in their organizations. They wondered what kinds of incentives would encourage busy professionals to take time to document their expertise for lessons learned databases and other repositories or to work directly with colleagues who could benefit from what they know. Or—to mention another possible impediment to sharing—what would persuade them to share with others the exclusive knowledge that gave them status and made them valuable to the organization.

One kind of strategy involves giving tangible rewards for sharing knowledge: for instance, cash in exchange for lessons learned, gift certificates, or participation in lotteries that offer a variety of prizes. Some of these efforts are versions of what is sometimes called "gamification"—creating a contest or game that potentially rewards desired behaviors or competition in order to engage participants in a process. In Sweden, for instance, drivers who observe posted speed limits are entered in a lottery for cash prizes drawn from a pool fed by speeding fines. In theory at least, the power of gamification should come from both the potential for rewards and the fact that many people like games; they choose to play them because they are fun.

But experience has shown that using games and offering material rewards to encourage knowledge sharing in organizations can have serious downsides. Paying people a bit of cash for sharing their expertise or entering them in a lottery for gift certificates can seem to trivialize the essential activity of collaboration and even be seen as an insult to professionals who take pride in their work. And games with prizes sometimes tempt people to try to "game" the system. Offering cash for lesson submissions often means that quantity goes up and quality goes down. And one NASA center that tried giving cash for contributions to the lessons learned repository found many attempts to cheat the system.

So it is not surprising that some studies have found that tangible rewards actually demotivate people, discouraging the behaviors they are intended to promote.

Experience has shown that the right intrinsic motivations are often more effective than prizes. Erika Blaney's article, "Five Intrinsic Motivators and How They Impact Employee Engagement," describes those motivations this way:

1. Autonomy ("I control")
2. Mastery ("I improve")
3. Progress ("I achieve")
4. Purpose ("I make a difference")
5. Social interaction ("I connect with others")

The Q&A sites stackexchange.com and stackoverflow.com are examples of vibrant knowledge sharing systems. The people who share knowledge there are passionate about sharing their expertise and get intrinsic rewards for helping others and being recognized as helpful. Sharing travel knowledge and reviews on the Trip Advisor website is also voluntary. My colleague who travels frequently feels good about helping her fellow travelers find clean, safe hotels. Her reviews are rated as very useful and accessed often! This is an example of feeling good about making a difference.

Recently the Morpheus project manager came to Kennedy Space Center (KSC) from Johnson Space Center to share lessons learned. Morpheus is developing a prototype lander that can land and take off vertically. The PM shared some technical lessons but really became passionate when he shared the lessons learned about how he communicated the project's risk profile to stakeholders and senior management. He described how he was successfully able to "fail forward" because of the rapport he established with his stakeholders even within a NASA culture that is risk averse and under a 24-hour-a-day media microscope. My impression was that he was spreading the good news with the hope of making the agency better—a goal that the audience shared.

I believe that a common sense of "purpose" and a desire for "social interaction" (numbers four and five on Blaney's list) are especially important driving forces for the sharing that happens by way of the NASA Engineering Network Communities of Practice and on Yammer.

Knowledge sharing at NASA benefits from the fact that employees have an important and beneficial shared goal: the advancement of science and space exploration. This sense of a valuable and even noble common purpose builds trust—your colleagues want the same thing you do—and encourages the open exchange of

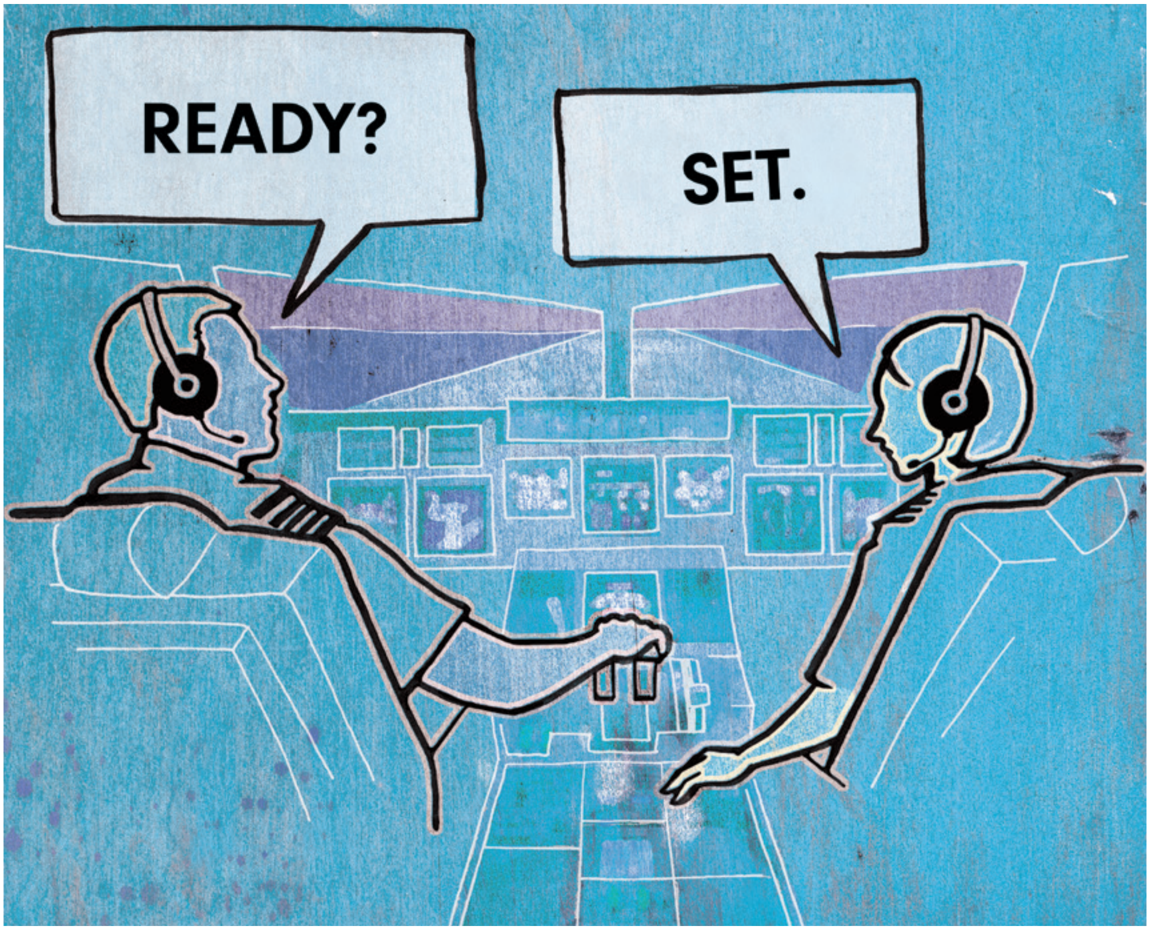

expertise needed to achieve the agency's goals. The fact that quite a few NASA retirees—"graybeards," as they are sometimes called (reflecting a time when the vast majority of NASA engineers were male)—continue to offer advice to current project teams is one example of this spirit of shared enterprise that continues even after people stop drawing a salary.

Because NASA's long and complex projects generally require sizeable teams of engineers, scientists, and managers, open collaboration is an essential ingredient of success. There is no way anyone can go it alone at NASA and hope to achieve anything of value. Shared accomplishments or no accomplishments are the only choices.

Here at KSC, we see both some of the impediments to knowledge sharing and examples of sharing driven by intrinsic motivation. Lack of time and sometimes the failure to realize the potential value of learning from a particular event can inhibit knowledge exchange. And it is possible we do not yet put enough emphasis on sharing knowledge with others and seeking the best knowledge from others. An employee once told me, "I don't recall anyone being recognized for looking at a lessons learned." Also, the fact that we have so many systems and ways to access knowledge may actually work against the desired result. Faced with that multiplicity of sources and the complexities of finding exactly what they are looking for, people tend to limit their search to their small networks of trusted colleagues.

But we also see people helping others learn what they know through postings in the KSC Forum and employees taking time to interact outside of their workgroup during the "Innovation Expo," and many KSC mentoring success stories through NASA Connect. We hope that sharing stories of those valuable interactions will encourage more and more of those rich exchanges.

Michael Bell is the Chief Knowledge Officer at NASA Kennedy Space Center.

The Evolution of Pause and Learn at Goddard

BY BARBARA FILLIP

A Pause and Learn (PaL) session is a team conversation. On the surface, it is deceptively simple, just conversation among members of a team, typically conducted after a significant event or project milestone. But it is different in many ways from other conversations—from a staff meeting, for example. A PaL is facilitated by a knowledgeable outsider who helps the team identify and capture knowledge gained at key project stages. Essential elements of the conversation are documented in the form of Knowledge Maps (KMAPs), which can then be shared to benefit others. A Knowledge Map is a visual representation of a conversation, highlighting valuable insights and comments within the context of that conversation.

The PaL practice was introduced at Goddard in 2005-2006 by Ed Rogers, the Center's Chief Knowledge Officer. As of December 2014, more than 90 Pause and Learn sessions have been facilitated at Goddard, 24 of which took place in 2014 (a record year). Carrying out PaLs for almost 10 years has taught us a lot about what works and what doesn't. We have learned which core elements are essential and where useful changes are possible.

The fundamental elements that make a PaL effective are:

- A facilitated conversation
- Key questions
- Team ownership of the process
- Clear and immediate benefits.

Outcomes can be less than optimal when any of these core elements is disregarded. At the same time, it is important to keep learning from the process and be flexible. In a 2012 *ASK Magazine* article, I described an adaptation of the process to allow one team to learn directly from another by observing a PaL. Since then, we have tried other innovations, some with more success than others.

A FACILITATED CONVERSATION

A PaL session is best facilitated by someone outside the project or organization participating in the PaL, someone who doesn't have a personal stake in the conversation but whose primary goal is to act as a facilitator and guide to maximize the benefits participants get from the conversation. At the same time, the facilitator must be sufficiently knowledgeable about the general operations and processes of the organization to understand its issues and goals. This facilitator function was initially provided by the Chief Knowledge Officer and now continues within the Flight Projects Directorate by a Knowledge Management lead.

This facilitated approach works well, with the corollary that all participants must understand that they are coming to the conversation as equals. The organizational hierarchy is temporarily set aside and everyone in the room has an equal opportunity to speak up. To ensure an open and honest conversation, expectations need to be set ahead of time with the team leaders. In practice, this means that it is best for the team leaders to set the tone by sending the PaL invitation to the participants and explaining the intent, but allowing others to take the lead when the conversation starts.

Open conversation is less likely when management layers above that team are invited to the process. In some cases, it is also valuable for the team to meet and talk without their immediate leadership/management in the room.

Getting the right participants in the room is critical. When a mission is being built out-of-house and

Goddard's role is limited, it may be difficult or even not advisable to suggest a Goddard PaL. Since the primary objective is for Goddard to continuously learn from its missions, a focus on Goddard's lesson can be achieved by focusing on the Goddard team's lessons rather than those of the entire mission team.

We recently sat down with a mission manager to talk about what he learned as mission manager for a PI-led out-of-house mission. There is value in doing that even when a traditional PaL with all the key stakeholders is not possible, but the benefits of PaLs usually come from group discussion.

One PaL last year failed to achieve its intended objectives when key members of the team were not able to attend. We decided to gather insights from the rest of the team by conducting one-on-one interviews, but the process quickly fell apart. The failure demonstrated that a PaL conversation with all key team members in the room is different from and more than a collection of individual perspectives. It allows individual perspectives to be expressed, but also gives the team an opportunity to correct misunderstandings and come to a common understanding of each other's perspectives. What can be achieved in 90 minutes of conversation is not achieved by compiling independent insights from team members gathered in one-on-one interviews.

That experience reminded us to stick to the basics of the PaL and watch out for the risks of diverging from the original approach. While it's still a good idea to collect insights from a core member of the team who isn't able to attend the PaL, it is not advisable to try to build a PaL out of individual interviews. If it is not a team conversation, it is not a PaL.

STARTING WITH KEY QUESTIONS

Another fundamental element of a successful PaL is a focus on key questions. The facilitator meets with the team leader (ideally face-to-face) to explain what a PaL is, how it is conducted, its benefits and expected outputs, and to discuss who should be invited as well as the dynamics of the conversation. The facilitator also asks a few questions to get a sense of the top issues that should be addressed. While the key questions are always the same—"What happened? What went well? What didn't go well? What could we have done differently?"—there can be variations. For example, a project moving from Phase B to Phase C might want to focus on "What are we going to do differently moving forward?"—a question likely to provide immediate benefits.

In the past year, we have worked with one office to make the PaL a complement to their existing lessons learned process that had one individual write up lessons learned in a pre-formatted presentation, which was then given to management. The process was primarily meant to capture lessons at the level of the technical workforce and worked well to identify technical issues and strategies to remedy them. It did not sufficiently capture the project challenges that originated at higher levels in the interactions between key stakeholders, however. When a Pause and Learn is conducted, these higher-level stakeholder dynamics are discussed (and ideally all the stakeholders are represented in the room). Sometimes teams are reluctant to talk about issues that are externally imposed on them and, they think, not open to change. Why bother talking about it when we know there's nothing we can do about it? But it is precisely those issues that need to be brought to the attention of management through as many channels as possible. A Knowledge Map is a tool, and one way to convey important messages.

As the PaL process has become better known at Goddard, team leads have sometimes asked for a PaL when they needed a facilitator to help with a group discussion. When the key questions are significantly different from those of a traditional PaL, though, it's probably not a PaL and should be called something else.

A TEAM OR PROJECT-OWNED PROCESS

Initially, PaL sessions focused almost entirely on supporting the team's learning. The conversation maps or knowledge maps developed were owned by the team and were not disseminated outside the team, in part because the lessons might not be fully understood without their local context. The preferred method for sharing lessons across projects was a face-to-face workshop where participants could have open discussions about their experience.

This internal PaL format is still used occasionally, but most project-specific PaL sessions are now accompanied by conversation maps meant for dissemination across projects. The web of maps that has evolved over the past two years is accessible to everyone at Goddard and by extension everyone at NASA.

For the Flight Projects Directorate, the PaL sessions and associated maps are the preferred approach for identifying and sharing lessons learned. It has therefore become important to document the sessions

While the key questions are always the same—"What happened? What went well? What didn't go well? What could we have done differently?"—there can be variations.

and lessons in a way that makes them accessible to other projects. "Accessible" means available for review and constructed in such a way that they make sense to someone outside the project.

Making the maps accessible to other projects can have disadvantages. Project team members can become concerned about who is going to see the maps and how the information might be interpreted or—more to the point—misinterpreted. The map review and validation process is designed to address these concerns. Team leaders review the draft knowledge maps before they are integrated into the aggregated map system known as the KMAP web. The review is meant to identify key elements that may have been missed, correct errors, and adjust the wording to avoid misinterpretations. The key is to keep the maps as close to the essence of the conversation as possible.

Once they have been approved by the team's leadership, the maps go through a second review by management, giving them the opportunity to ask for clarifications and add comments. Management does not redact the maps but has the ability to provide additional guidance when a lesson or insight from a project could be misread or misinterpreted.

CLEAR AND IMMEDIATE BENEFITS

There were 24 Pause and Learn sessions in 2014. One office has fully institutionalized the process and conducts a PaL after each key event. One project conducted two PaLs in one year, both follow-ups to key reviews. This is a relatively small, very fast-paced project whose leadership has found value in the PaL conversation within the team and in the careful articulation of key insights on the knowledge maps, which they have posted on the wall in the project's central office hub.

In one case this past year, a different PaL format was deployed to gather lessons learned. This effort was carried out by an external consultant, with the support of the Office of the Chief Knowledge Officer, within the Management Operations Directorate. A much more structured approach was used and a more structured, linear report was generated, with specific recommendations for action. The process required two full days of meetings. While this format has value under certain conditions, it is not realistic to ask project teams to set aside two days for a Pause and Learn activity. A typical PaL session is scheduled for 90 minutes with the entire team. The PaL Map review/validation process requires another hour or so of the project leaders' time.

The return on investment of a traditional 90-minute PaL is quite substantial, in part because the team's investment of time—a scarce project resource—is reasonable.

Reflection time tends to be undervalued, especially where keeping to the schedule is critical to success—a fact of life for many projects. But just as project managers regularly point out that a key lesson learned is to know how to spend some reserves early to address issues that would cost much more to deal with later on, time invested in reflection early on and throughout the project life cycle can save time and money later. And of course identifying the most effective way to spend reserves requires some reflection.

In the end, a project increases its chances of success when art meets science, when people are able to work well together to address tough technical challenges. Working well together means communicating well. That is why the Pause and Learn is so important. It develops or reinforces a team's ability to engage in open, honest conversations; it maintains or opens up ways for team members to talk to each other and address issues that may not necessarily come up in traditional staff meetings or in one-on-one conversations.

Barbara Fillip is the Knowledge Management Lead at NASA Goddard Space Flight Center.

Continued Learning and Improvement

While keeping the fundamental PaL principles in mind, we continue to look for ways to improve and expand this valuable practice. Here are a few potential improvements we are working on or considering:

APPROACHING THE LARGE PROJECTS AT A LOWER LEVEL.

Large, complex projects cannot easily do PaLs at the mission level. There would be some value in conducting PaLs on a more regular basis with instrument teams or subsystems, especially for in-house instruments.

DEMONSTRATING THE VALUE OF AGGREGATING LESSONS AND MAKING THEM ACCESSIBLE TO OTHER PROJECTS.

The return on a 90-minute investment of time is clear. The ROI of time and effort spent making lessons widely accessible needs to be demonstrated. There is more work to be done with the project teams to demonstrate the value of looking through and discussing insights and lessons from other projects and more to be done to ensure that knowledge maps present lessons and insights in a usable way.

ANALYZING AND USING LESSONS BETTER.

There is a need for more in-depth analysis of aggregated lessons, more use of lessons and insights in workshops, and more follow-up by management, which could benefit from reviewing and discussing the Knowledge Maps on a regular basis.

CONDUCTING TIMELY WORKSHOPS BASED ON THE PROJECT LIFECYCLES.

For example, if multiple projects will be going through Integration and Testing in the next six months, a workshop bringing together projects about to enter that phase and projects that have recently completed that phase would enable an exchange of valuable knowledge when it is needed.

INCREASING INTEGRATION OF THE PAL PROCESS/LESSONS LEARNED PROCESS WITH OTHER KNOWLEDGE-MANAGEMENT-RELATED EFFORTS WITHIN THE FLIGHT PROJECTS DIRECTORATE.

Integrating the PaL process with professional development, best practices, and associated knowledge networks and other activities can provide new benefits. This has already started to happen and will likely increase.

K2020 at JSC: Facing the Knowledge Challenge

BY DON COHEN

Ed Hoffman, NASA Chief Knowledge Officer (CKO), introduced the Knowledge 2020 meeting at Johnson Space Center in late April 2015, saying he expected it to build on the accomplishments of the first K2020 meeting in October 2014 at Kennedy Space Center. There too, NASA center CKOs and knowledge points of contact, as well as other NASA knowledge workers and managers, had met with experts from within and outside the Agency to discuss effective approaches to acquiring, retaining, and sharing valuable organizational knowledge.

Hoffman described the request of the Aerospace Safety Advisory Panel (ASAP) for a NASA long-term knowledge strategy, over the next five years, as an

Photo Credit: NASA CKO

opportunity to clarify goals and what NASA needs to do to reach them. Since its February meeting, the ASAP has recommended and is seeking more information about actions on these knowledge activities at NASA:

- Benchmarking
- Incentives and rewards
- Integration of knowledge services and resources in search and flow of work
- Formalized practices
- Effective knowledge capture, visibility, integration of lessons learned
- Continued focus on critical knowledge

According to Hoffman, NASA has made good progress on capture and retention, but much work remains to be done to ensure that NASA has and can effectively apply the knowledge it will need in the future. In the course of the three-day meeting, participants agreed that the keys to successful knowledge work in organizations include identifying and focusing on critical knowledge, incorporating knowledge activities as much as possible into the flow of work, and developing an organizational culture that is open to learning and sharing.

In a conversation with the group by phone, NASA Chief Engineer Ralph Roe emphasized the goal of greater integration among knowledge services at the centers—the importance of sharing methods and activities and looking for opportunities for direct dialogue and storytelling about the lessons learned that need to be shared. The K2020 meeting is, of course, one of those opportunities. Knowledge capture and management are critical, Roe said: "Developing the tools and methods and best practices to do that is essential for NASA to be successful."

CRITICAL KNOWLEDGE

Many initiatives in the early days of knowledge management tried to make as much knowledge as possible available throughout organizations. Those efforts were usually wasteful and often ineffective. Over time, knowledge practitioners have come to see the value of more strategically targeted approaches: giving their attention and resources to the specific knowledge that can help the organization achieve its goals and to the particular people who need that knowledge.

Several speakers at this K2020 meeting talked about how important that focus is in their organizations. "Things in knowledge management are moving so fast it's dizzying," said Marty Lipa, Merck's Executive Director of Knowledge Management. He stressed the critical need to link knowledge to outcomes, to understand the organization's aims and to work with the expertise and on the processes relevant to achieving results. The explicit goal of Microsoft's knowledge activities—"instant relevant knowledge in context"—also puts the emphasis on the knowledge that matters, according to Jean-Claude Monney, Global Knowledge Management Lead for Microsoft Services. And Tim Bridges, Director of Knowledge Management for Boeing, talked about the importance and the challenge of "stratifying content"—that is, identifying the knowledge that has the most value for the company and most needs to be retained, shared, and reused.

John McQuary, formerly Vice President for Knowledge Management and Technology Strategies at Fluor Corporation and now an independent consultant, made a related point during a panel discussion with Hoffman and fellow consultant Dan Ranta. Knowledge practitioners need to "think like a business," he said, and link their own goals to the organization's strategic aims. Ranta said that people at ConocoPhilips, where he had worked for eight years, had to make the business case for the knowledge help they requested before they could receive it.

Clarity about goals also helps make measurements of knowledge work possible. As Monney said, you can't measure your success unless you have a clear idea of what you are trying to accomplish. "You can't measure knowledge," Monney added, "but you can measure the impact of knowledge."

MARKETS FOR KNOWLEDGE

We may think that good ideas should sell themselves, that they will be readily embraced by our organizations because of their self-evident value. In his video talk

April 2015, Dan Ranta (far left), former Director of Knowledge Sharing, ConocoPhillips; and Jean-Claude Monney (center), Global Knowledge Management Lead, Microsoft Services, mingle with participants at the welcome reception of Knowledge 2020 2.0.

on knowledge markets, Larry Prusak argued that knowledge and organizations don't work that way. Ideas need to be actively sold to "buyers"—the managers and executives who have the influence and budgets needed to put them into practice. Many advocates of knowledge services have often confessed to experiencing—at one time or another—reluctance of leadership to invest in knowledge services, especially during the initial setup of a knowledge program. "There needs to be a marketplace," stressed Prusak, "and a price."

The status of the sellers also matters. A successful seller typically has been with the organization for at least 10 years. A long tenure is required to understand the organization's needs thoroughly as well as its ways of working and talking about work—all factors in presenting the right idea in the right way. In a discussion after the talk, participants agreed that how an idea is presented—how it is "framed," to use Prusak's term—is an important contributor to its success or failure. At NASA, for instance, "knowledge services" has proved to be a much more acceptable term than "knowledge management." And Bridges mentioned the importance of showing Boeing management that knowledge management is a form of risk management—a central concern of the aerospace company.

Time is also needed to establish the extensive personal networks and the trust required to communicate an idea effectively and convince others that it is worthwhile. The successful sharing of new ideas—and of any complex knowledge—is profoundly influenced by these social factors. Knowledge exchange is as much a cultural phenomenon as a technical or procedural activity.

CULTURE

The importance of culture explains why communities of practice have become such an important part of knowledge work. Communities of workers who share a professional identity and have come to know and respect one another through joint work and mutual assistance are predisposed to offer knowledge to one another and accept knowledge offered both in person and by way of electronic repositories and work spaces. Employees who have no direct connection are less likely to do either. That is why Monney calls communities of practice "the heart and soul" of knowledge work at Microsoft, where 87 communities have 18,000 members and 1,000 officially recognized subject matter experts. Monney notes that the median time it takes to get a response to a question posted on a community site is only 28 minutes.

> "Knowledge exchange is as much a cultural phenomenon as a technical or procedural activity."

Bridges talked about the importance of what are called Communities of Excellence at Boeing. Getting value from Boeing's diverse knowledge through a collaborative environment enhances company success, as does the similar knowledge-management objective of ensuring long-term continuity of key knowledge within all functions and businesses. Like Microsoft, Boeing has officially designated experts available through the Boeing Knowledge Network (BKN) via inSite. But he also talked about the challenge of overcoming some employees' reluctance to share expertise. Although they explain their unwillingness as due to rules against sharing information, the content in question is often not restricted.

As Lipa and his colleague at Merck Samantha Bruno noted, the biggest challenges are often cultural, and knowledge work should start with people's behaviors: understanding why people act the way they do so that knowledge solutions can better support their needs. Monney emphasized the value of cultivating personal networks, building trust, and developing a culture of responsibility to share.

Don Cohen's discussion of organizational cultures that encourage admitting mistakes provided another look at what can be called "knowledge-friendly cultures." Errors are an important source of learning if they can be freely discussed, and accepting them is an essential part of innovation, since the search for something new invariably involved making mistakes along the way. In his description of the elements of effective teaching of engineering, Anthony Luscher, professor of Aerospace Engineering at The Ohio State University, identified learning from engineering mistakes as an essential part of that education. "I've never seen anything that works the first time," he said.

Many organizations claim to encourage openness to errors, but actions speak louder than words. Organizations need to show that people can admit mistakes without being punished. No-fault discussions of mishaps in after-action reviews and pause-and-learn sessions are also important, as are trusting relationships and shared goals.

CASE: OCO-2

Ralph Basilio's discussion of the Orbiting Carbon Observatory-2, for which he was project manager, told one story of responding positively to failure. The first OCO satellite crashed into the Indian Ocean because the launch vehicle's faring failed to separate and the satellite's extra weight made reaching orbit impossible.

After that failure, "no one walked away," according to Basilio, and the team was determined to reduce the risk of a similar problem as much as possible. He quoted George Bernard Shaw: "Success does not consist in never making mistakes but in never making the same one a second time."

The OCO-2 team was devoted to its work in part because of the importance of the mission. The spacecraft was designed to measure atmospheric carbon dioxide more accurately than any existing technology and would provide new information about the global carbon cycle that could improve our understanding of its role in climate change.

Given a tight schedule and limited budget, OCO-2's success depended on learning some key lessons. Chief among them was not making changes to the original design without compelling reasons. "Make it work" was the team's mantra, not "make it better." Inevitably, changes had to be made—for instance, to interface with a different launch vehicle (a Delta II rather than the Taurus-XL used to launch OCO). The necessary modifications made it even more important not to change other things.

Basilio also identified using one electronic repository with links to other sources as a help in completing the work quickly; "one-stop shopping" for technical information proved to be an important time-saver.

INTEGRATION AND OPENNESS

"Integration" was one of the themes of the ASAP report, which included a recommendation concerning "one-stop shopping" for lessons learned, and was cited as a goal by Roe. It was mentioned by other speakers as an important part of their knowledge work. Ranta said that the organizations he works with already have good knowledge work going on "in pockets," but those activities are not as widespread or connected as they need to be.

Bruno said that creating "one Merck," making the organization "boundary-less," was an important goal. Bridges also spoke about "one company," stating: "Knowledge knows no organizational boundaries." Establishing process standards and encouraging the flow of people throughout the organization by giving them assignments in different areas were among their approaches to fostering this integration.

Another kind of integration—integrating knowledge tools and activities into the flow of work—was identified by Lipa as an important goal.

Openness to the world of valuable knowledge outside the organization—a kind of integration with that wider world—is also important. Bruno mentioned the value of being "porous," and Monney described the challenge of connecting productively with Microsoft's 400 outside partners. Dan Ranta recommended "sister network sharing," promoting questions to other networks.

LOOKING TO THE FUTURE

Overall, the meeting highlighted significant accomplishments in knowledge work at NASA, began benchmarking relationships with other highly technical organizations, and engaged in exercises to define approaches to the challenges NASA faces. Practitioners have learned from experience and from each other, especially in recognizing the complex blend of culture, process, and technology that define genuine knowledge activities.

There is still much to do. For all their improvements, knowledge tools are still not sufficiently integrated into the flow of work. The sharing of knowledge and successful approaches to knowledge work across diverse organizations like NASA is happening but not as widely as it should. Creating the knowledge culture is an ongoing task. As John McQuary said, "This is a journey; you're never done."

Don Cohen is Editor-in-Chief of NASA Knowledge Journal.

Putting Cognitive Computing to Work at NASA Langley

BY MANJULA AMBUR AND DON COHEN

NASA's need for the capacity to accumulate and make sense of vast amounts of material is what makes Watson a potentially valuable tool for NASA. The amount of technical content relevant to NASA missions and research projects is enormous and growing rapidly. The good news is that the information needed to solve most technical problems is almost certainly out there somewhere; the bad news is that the sheer volume of information can make finding and digesting the relevant content harder and more time consuming. That is why a number of us at the NASA Langley Research Center have been working with IBM on a pilot program to see if IBM's Watson computer can help our researchers and engineers to analyze and digest the material they need as efficiently and effectively as possible and help them to develop innovative solutions.

Most people first became aware of IBM's Watson computer when it outplayed human experts on the *Jeopardy!*® TV quiz show in 2011. The computer was able to answer questions posed in ordinary language that often demanded a knowledge of context and an ability to unravel complexities and ambiguities.

Watson understands questions, produces possible answers, analyzes evidence, computes confidence in its results, and delivers answers with associated evidence and confidence levels. It does this with a vast database of structured and unstructured content. Watson is used in a variety of fields, including medicine. It assists physicians, for instance, suggesting diagnoses and treatments.

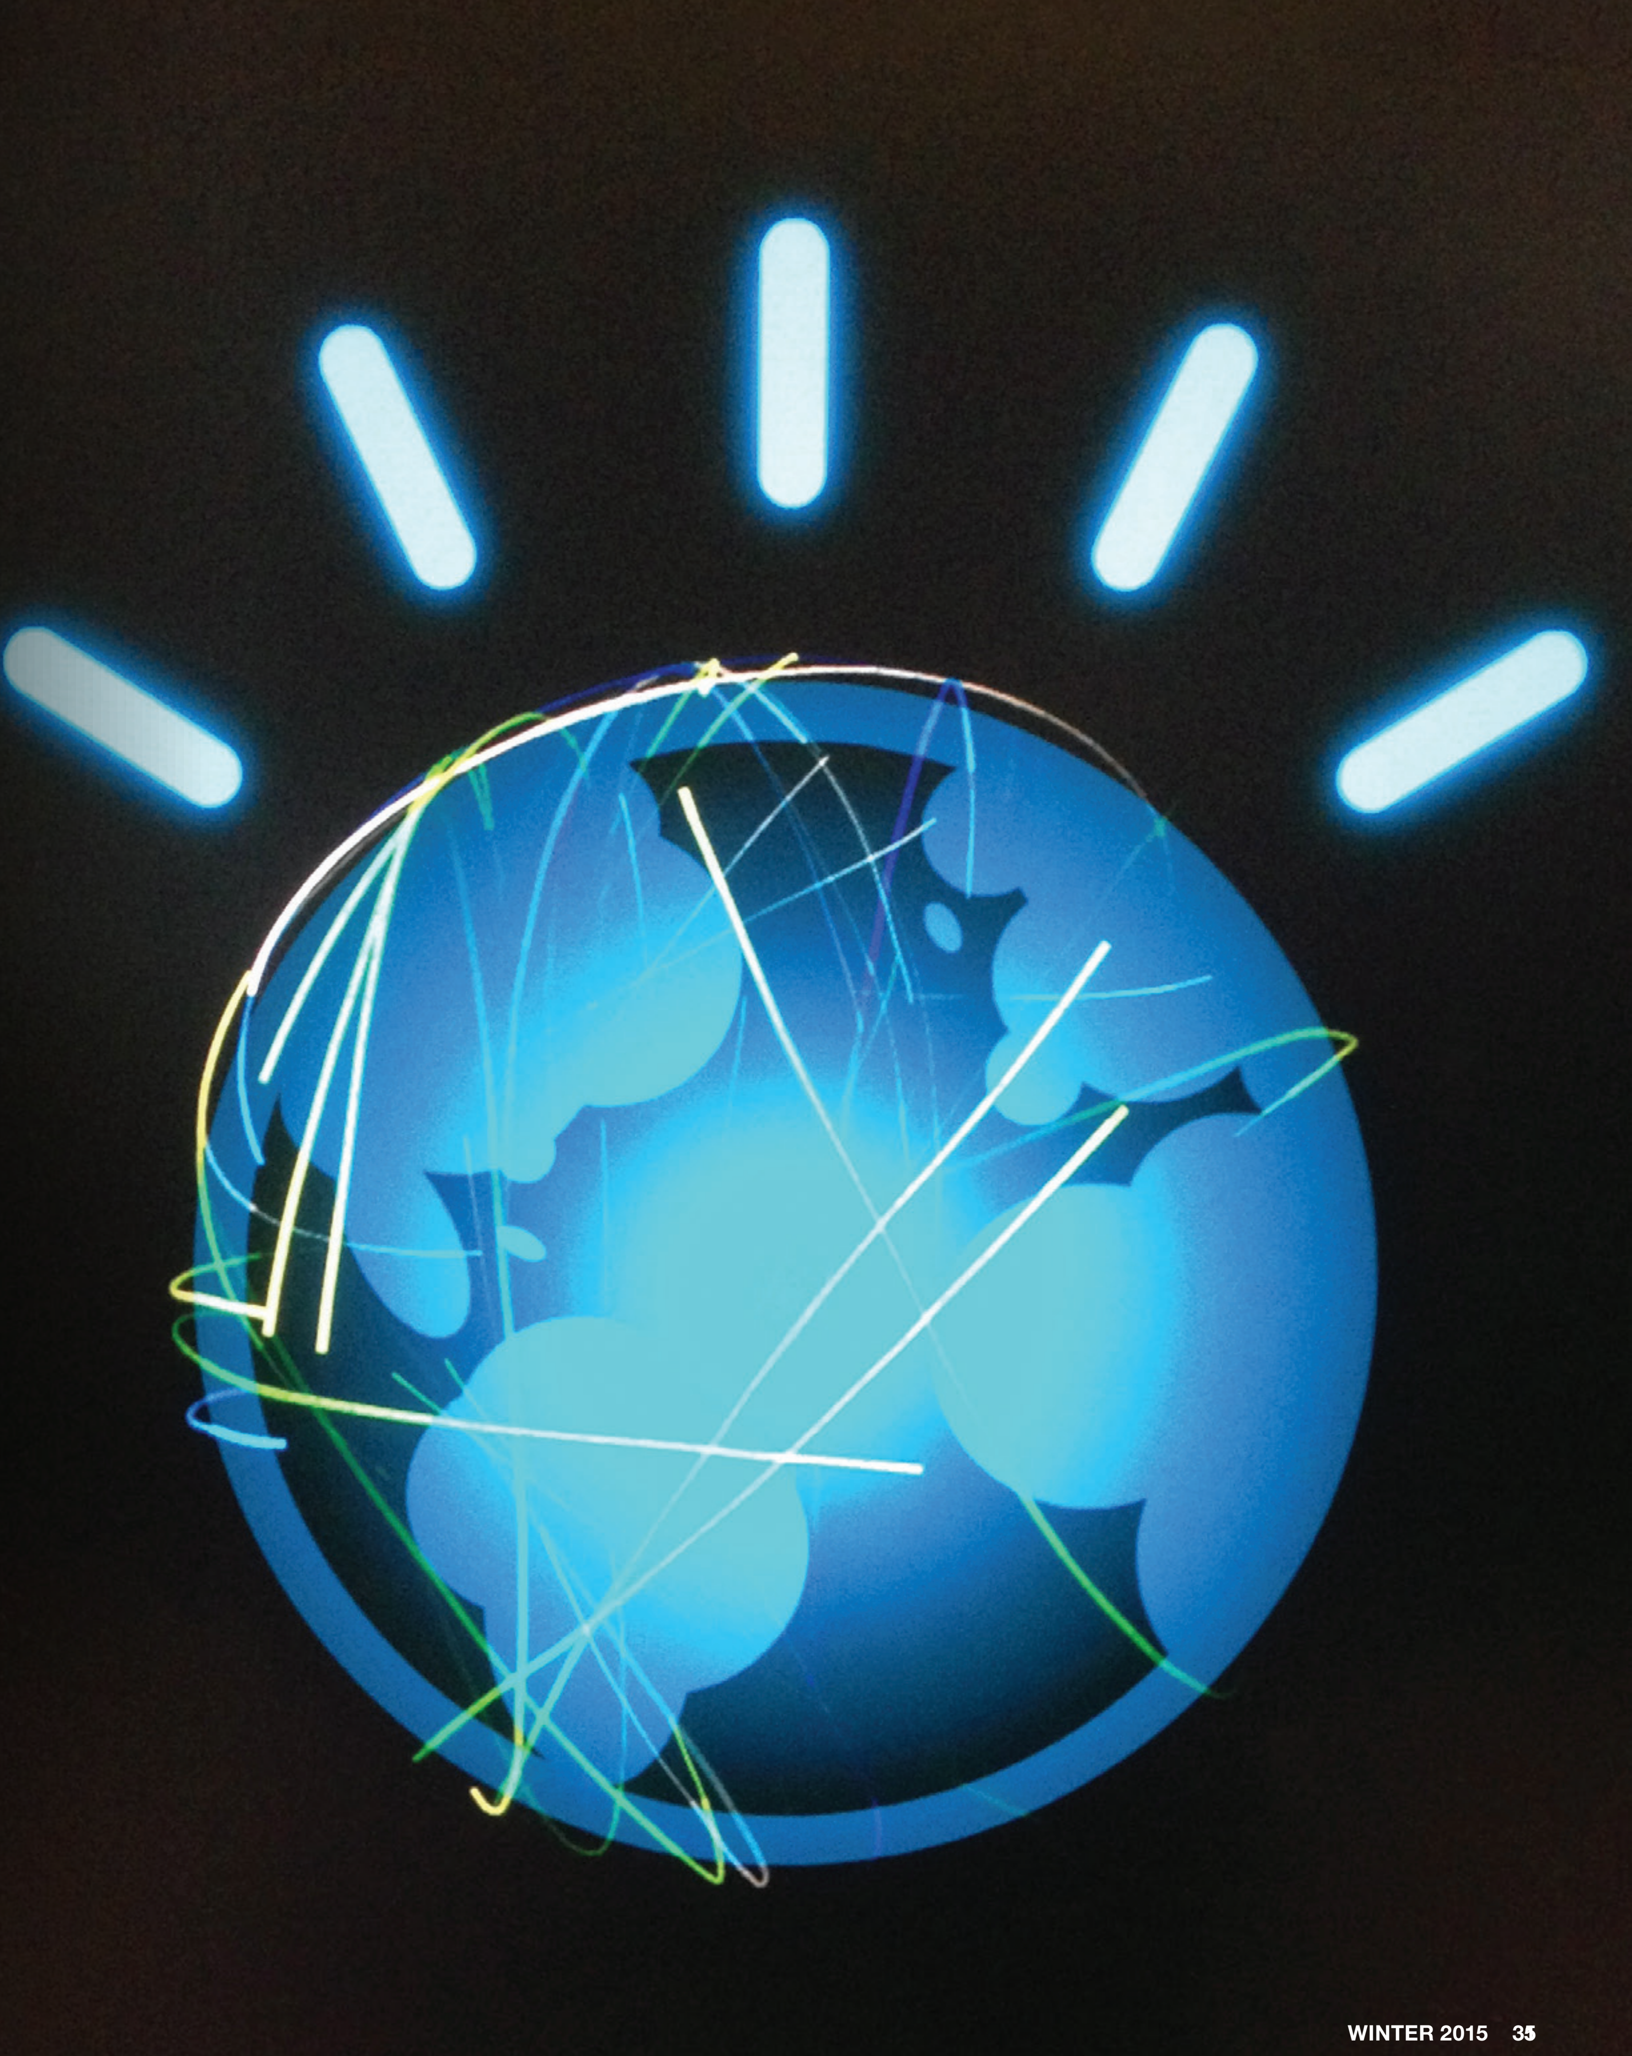

Like the quiz-show Watson, the healthcare Watson compares symptoms and other information with its vast stores of content that includes medical literature, clinical trial results, and patient records to arrive at its recommended treatment paths with traceable evidence that a physician can use to make the best decision possible. The computer is not "smarter" than the doctors. A human-machine collaboration and symbiosis combines the computer's mastery of vast quantities of relevant information—more than any human expert could absorb or remember—with the expertise, experience, and creative intelligence of the professionals who use it.

PILOT PROGRAM

We started to investigate the Watson technologies in 2011 and held a center-wide seminar on Watson that summer, including visiting Watson Labs in New York. In 2012, we chose to investigate the "content analytics" software part of the overall Watson suite for our proof of concept. We were comfortable choosing the content analytics element early, since we thought it would be the most helpful to address specific NASA information challenges for only a modest investment. Based on positive researchers' response to the proof of concept, we formulated the "Knowledge Assistant" in 2013—our Langley-specific version of Watson Content Analytics (WCA) to search and analyze unstructured information from multiple sources and to quickly understand and deliver relevant insight with customizable facets and concept extraction. This pilot started in 2014 in collaboration with IBM WCA experts and Langley researchers/subject-matter experts.

We chose two technical areas—carbon nanotubes and autonomous flight—to test the WCA pilot concept. Both of these fields are technology incubator areas with a high potential for Center research and so are worth the investment of time and money.

Currently, subject matter experts (SMEs) in these and other fields have to read the entire literature corpus to digest and identify important information essentially "by hand." Not only does the process take a lot of time, it is likely to miss valuable content and connections. We thought of the Watson knowledge assistant as a virtual helper to SMEs, one that could efficiently analyze the material and help them to readily find important connections between information as well as identifying experts in the field.

The first phase of our pilot, from January to August 2014, tested Watson Content Analytics' ability to make sense of large bodies of knowledge, provide brief summaries of the key points of documents, and help SMEs discover potentially valuable trends, relationships, and experts. Reviewing content relevant to nanotube technology (approximately 130,000 documents), for instance, the machine generated automated clusters of documents on topics like the strength of reinforced composites and identified experts based on their contributions to the literature. These results were validated by SMEs.

The NASA SMEs involved in the pilot see the value of this kind of analytics—the ability to review all relevant materials and uncover connections that a human searcher would likely miss. The Knowledge Assistant had the ability to analyze the content in a variety of ways and identify important authors and experts on a range of technical subjects.

Phase 2 of the pilot, which began in September 2014 and was completed in February 2015, has focused especially on making the discovered content accessible and useful to researchers and engineers. This phase, working with IBM experts, has led to important progress on developing a more intuitive user interface, one that includes visualizations that make important information and connections apparent.

Watson Content Analytics uses syntactical and statistical sorting algorithms for clustering articles that share key terms, parses by "concepts," and identifies top experts, possible collaborators, and linkages between concepts. For further knowledge extraction, these results have been augmented with semantic technologies such as taxonomies and custom concept rules, improvements that required collaboration with subject matter experts. Our pilot showed that close collaboration with content analytics experts and mission team members who need information is critical to success.

The Knowledge Assistant pilot with its improved user interface has been demonstrated widely to many SMEs and senior technical leaders; their feedback on its usefulness is very positive. One of the challenges we have to consider is the cost of purchasing the scholarly content the system needs access to for deep analytics. Buying all the scholarly content that might be relevant is prohibitively expensive. So our current plan is to offer WCA as a capability to analyze the content that is internal to NASA, open aerospace content including NASA reports, and the research collections of individual users.

IBM's Watson computer system competes against *Jeopardy!*'s two most successful and celebrated contestants—Ken Jennings and Brad Rutter.

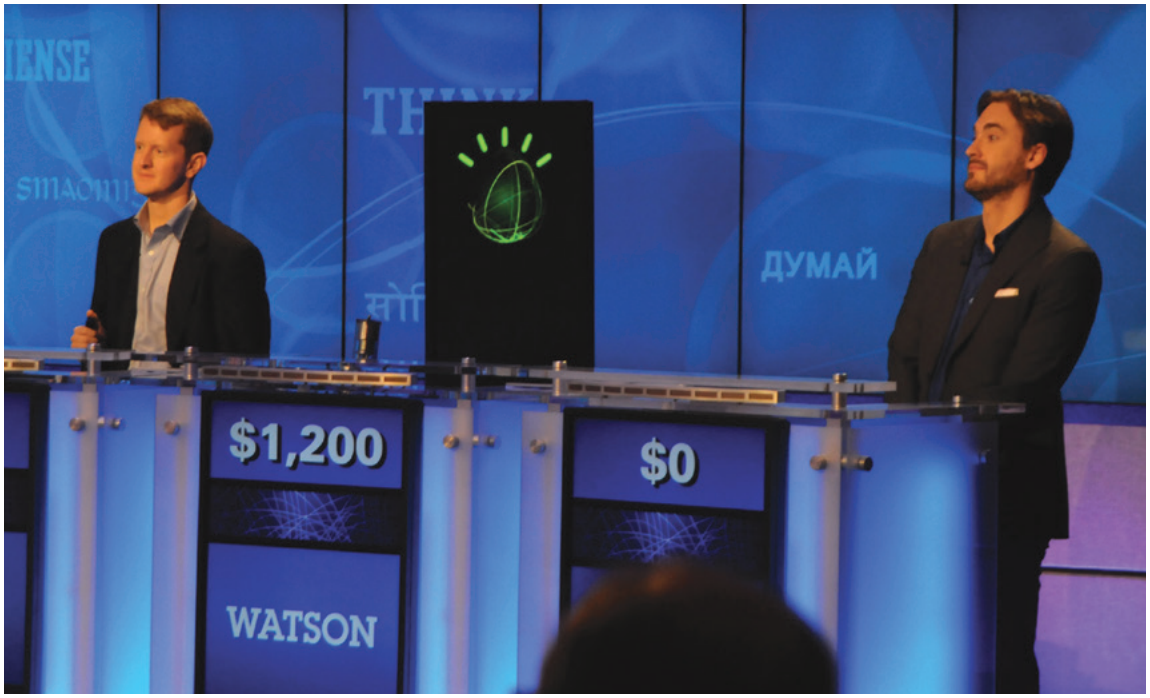

Photo Credit: IBM

THE FUTURE

The future potential of analytics technologies seems especially great. We are pursuing a full proof of concept to demonstrate the full functionality of these analytical capabilities that we believe will become an essential tool for researchers and engineers currently overwhelmed by the quantity and complexity of information in their fields.

We have just started to work with IBM Watson experts to develop an Aerospace Innovation Advisor Proof of Concept using IBM Watson Discovery Advisor technology, which is currently being used in medicine. Our aim is to demonstrate how natural language processing and machine learning technologies can be applied to aerospace research and development to accelerate the pace of discovery and innovation by analyzing and fully leveraging massive amounts of technical information. Watson Discovery Advisor technology is designed to boost analysis, provide valuable insights, and inspire research by finding connections, insights, and hidden relationship that human experts are unlikely to find. It can answer questions and even suggest questions that researchers have not thought to ask.

Looking further into the future, we can imagine a "Virtual/Intelligent Agent" that is a true collaborator with an expert in a human-machine partnership. Such a system might be able to read relevant scientific literature in a variety of foreign languages, understand mathematical equations and tables, and relate it to material in English. It might understand multimedia content: images, figures, formulae, and videos. And, most important, it might be able to provide direct answers or lists of possible answers to users' questions (rather than just lists of potentially useful documents). In those cases, as in its current use for medical diagnosis, the system would provide not just a set of possible answers but also information about the evidence it used to arrive at them, so the human experts will have the information they need to evaluate the conclusions. The system would be a true collaborator in our future research and engineering development efforts.

Note: Trade names and trademarks are used in this report for identification only. Their usage does not constitute an official endorsement, either expressed or implied, by the National Aeronautics and Space Administration.

Manjula Ambur is the Leader of Big Data Analytics and Chief Knowledge Officer at NASA Langley Research Center.

Don Cohen is Editor-in-Chief of NASA Knowledge Journal.

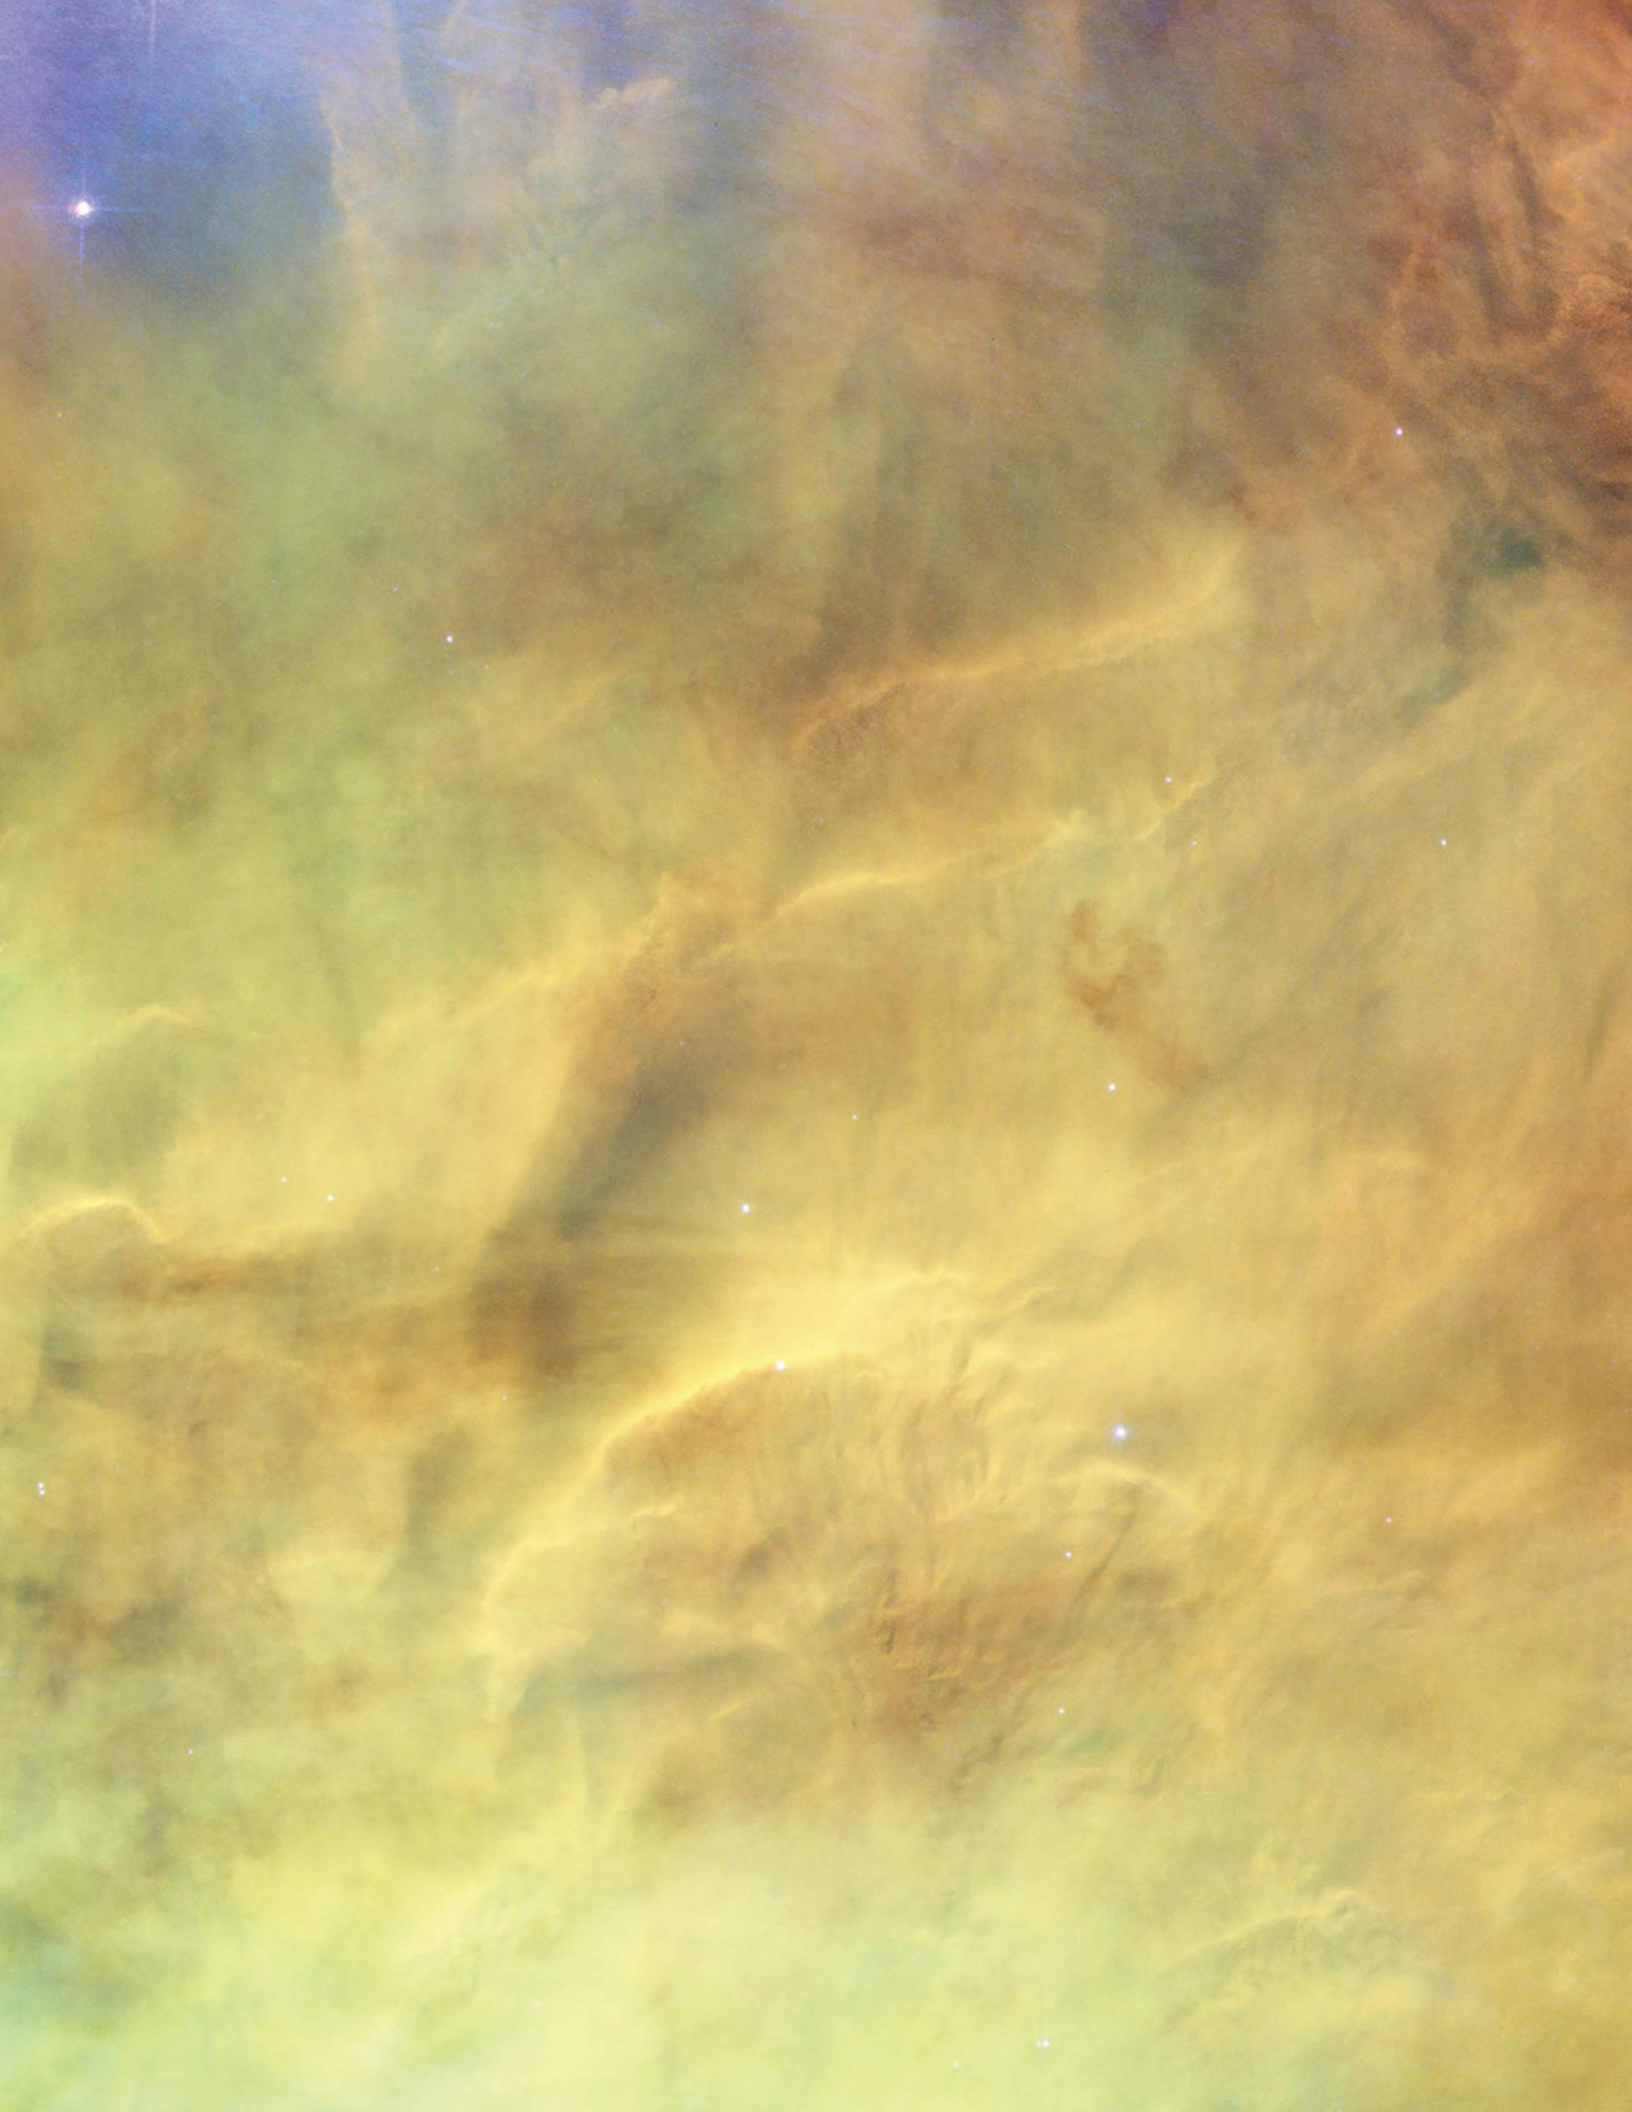

Like brush strokes on a canvas, ridges of color seem to flow across the Lagoon Nebula, a canvas nearly 3 light-years wide. The colors map emission from ionized gas in the nebula and were recorded by the Hubble Space Telescope's Advanced Camera for Surveys.

"Being here, living here, is something that I will probably spend the rest of my life striving to find just the right words to try to encompass and convey just a fraction of what makes our endeavors in space so special and essential," said Flight Engineer Peggy Whitson of Expedition 5, who lived six months in space on the station. Backdropped against the Caspian Sea, this full view of the international space station was photographed by a crewmember onboard the Space Shuttle Discovery after the undocking of the two spacecraft.

National Aeronautics and Space Administration
NASA Headquarters
300 E Street SW
Washington, DC 20546

www.nasa.gov

NP-2015-10-2100-HQ