Invited talk. Proceedings of the 27th Annual Conference of the Japanese Society for Artificial Intelligence, 2013.
https://doi.org/10.11517/pjsai.JSAI2013.0_2A2

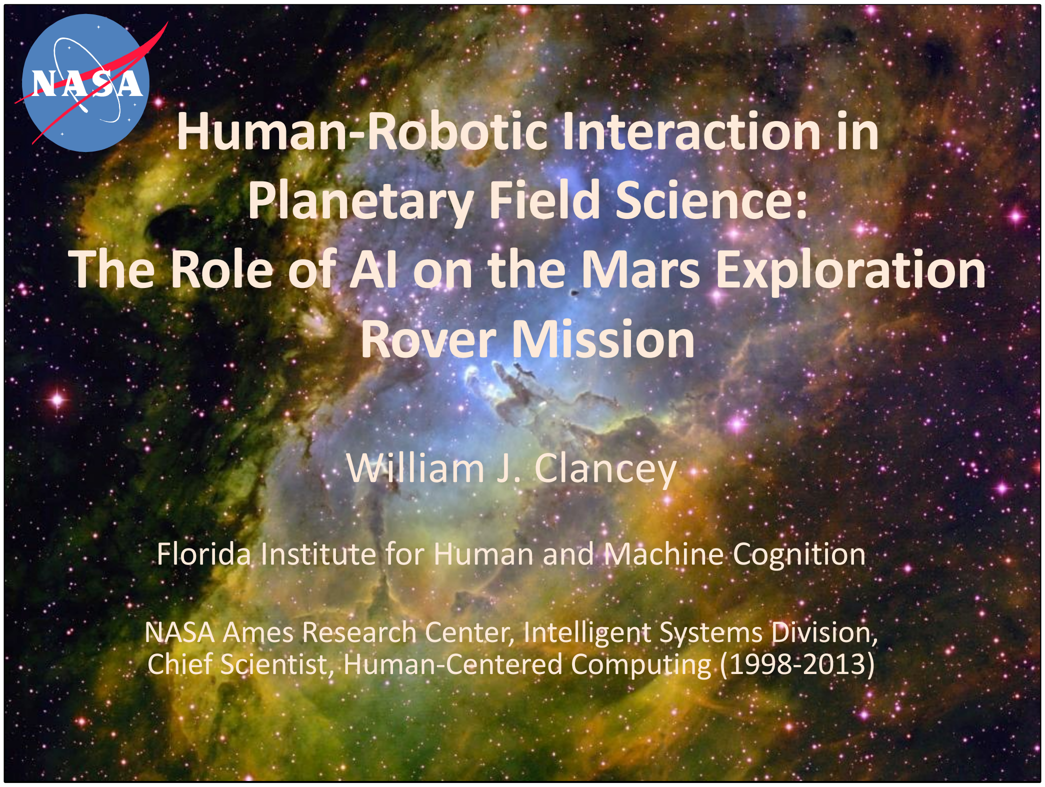

My presentation is about the Mars Exploration Rovers (also called the MERs), twin mobile robotic laboratories that we sent to Mars in 2003, one of which is still operating today.

Using AI planning, process control, and robotic navigation methods MER enables people to travel and study the surface of another planet. We call that study "field science."

In my talk I emphasize how the idea of human-robotic interaction is transformed in MER. The great distance between Earth and Mars causes a time delay in sending and receiving commands and data; so we do not interact with the rovers in real time. We send them programs and they carry out the precise movements, sampling, and analyses we have specified.

In particular, I'm going to tell you how a systems thinking approach made the Mars Exploration Rover missions so successful. We learned from those missions that rather than replacing people, robotic systems can help people form a stronger team and work together.

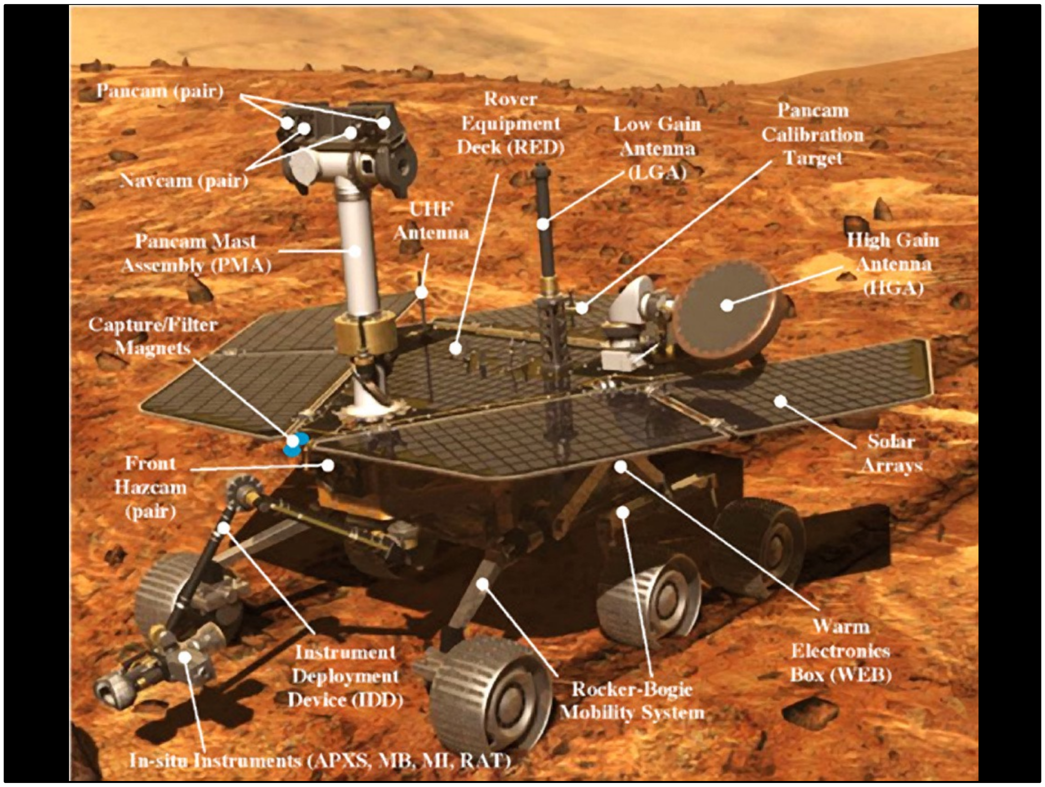

The two MERs—Spirit and Opportunity—arrived on Mars in January 2004. The MER is a mobile robotic laboratory, which we program from Earth to take photographs, scrape rocks, and analyze chemical composition. Of course the MERs can also move, and through virtual reality tools we have experienced the first overland expedition on another planet. The rovers have served as our physical proxies on Mars. The scientists and engineers back on Earth are collectively the mind of the rover, controlling the exploration and interpreting the results to decide what to do next.

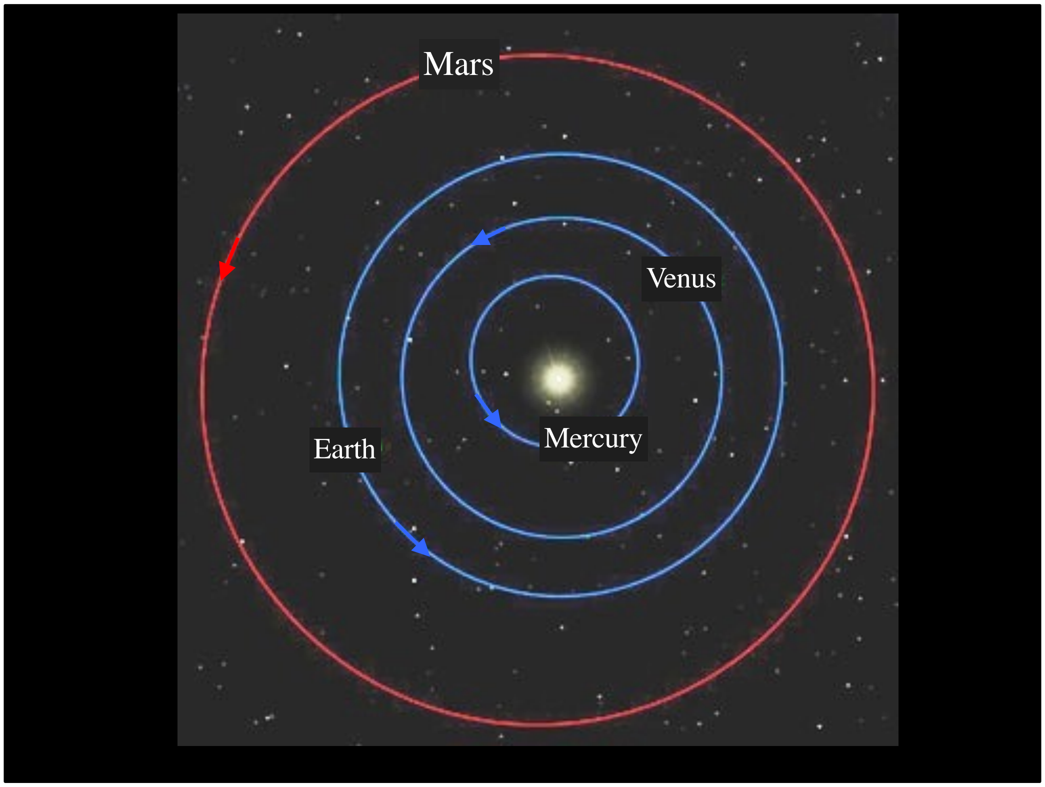

In the scale of the Universe, Mars is right next door—about 9 months travel using conventional chemical rockets.

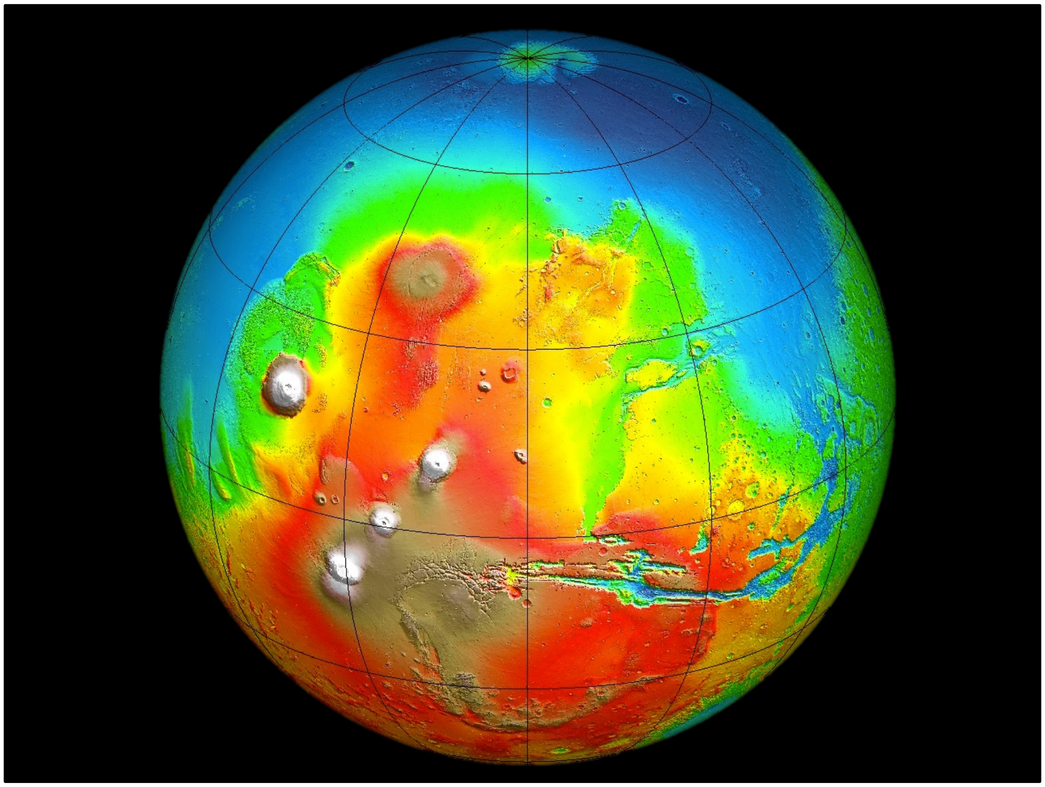

One striking elevation map created from orbit shows the lower areas colored blue, and suggests that large parts of the northern hemisphere might have been covered in seas 3 or 4 billion years ago; and there is evidence for ancient shorelines.

So what happened? Mars is in the "habitable zone" of the sun. Did life form on Mars? Why was its atmosphere lost? Are microorganisms living today below the surface? If life formed there, did it form separately from Earth—or are we related?

These are big questions and this is why many of us are excited about Mars.

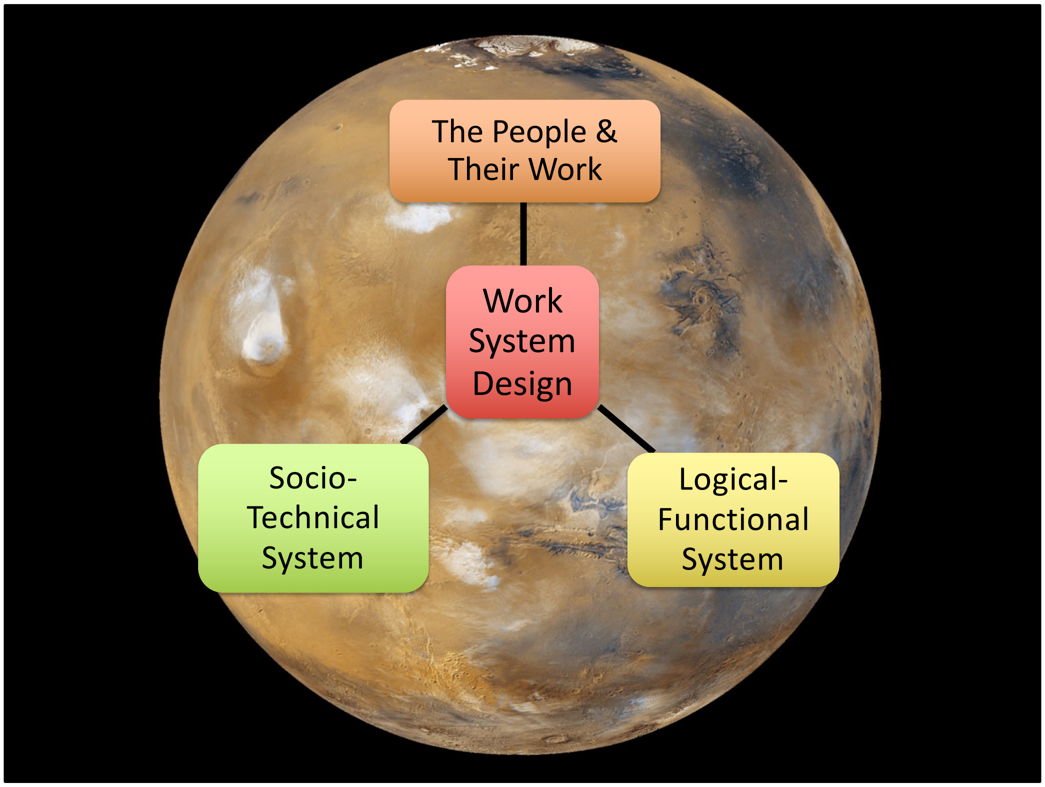

How do you design a mission consisting of people, technology, and processes so that scientists are able to work on Mars without leaving the Earth?

A very useful perspective is to view the combination of people, technology, and work processes as a *system* that is designed and operated holistically, balancing what the people want to do and how they would like to work with the technical methods available.

To provide a system that enables the scientists to carry on their work on Mars we need to combine the logical-functional perspective of engineering with a socio-technical perspective that enables us to understand and facilitate how people and technology work together.

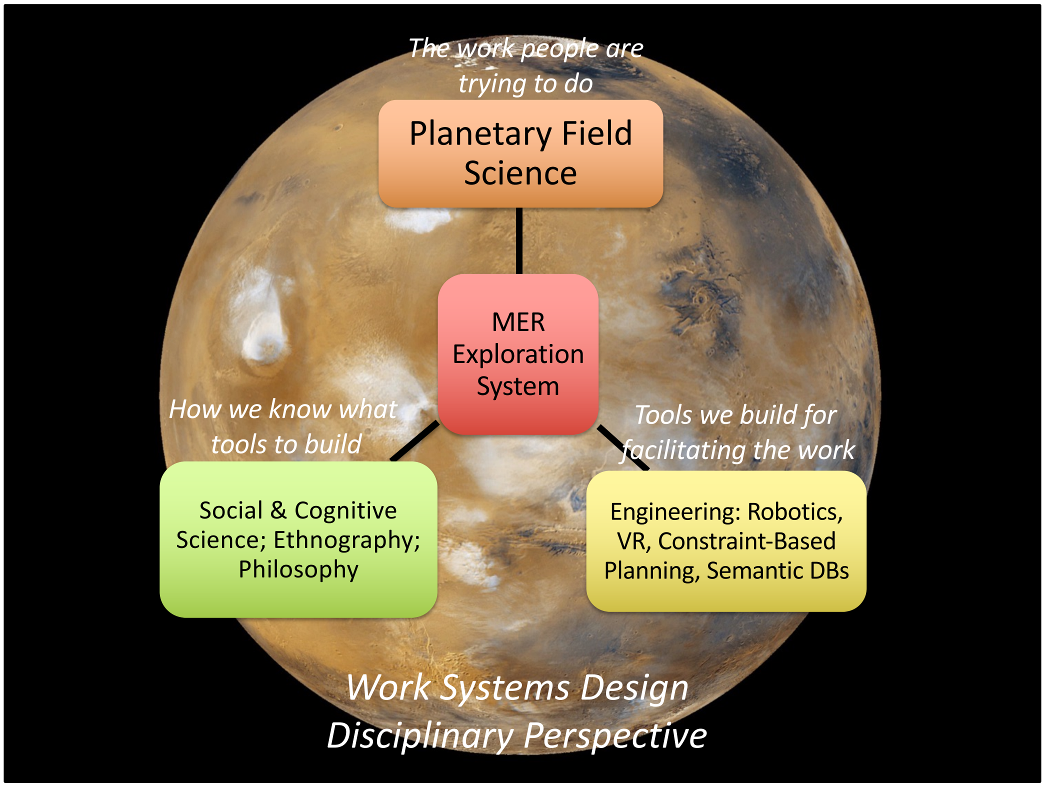

To make that a bit more explicit: The work the people are trying to do is planetary field science—it includes geology, chemistry, biology, and meteorology, with blended specialties such as astrobiology and geobiochemestry.

The engineering work combines a wide variety of methods for developing and controlling scientific instruments, communication, power & navigation systems, and computer software. It includes the technology we send to Mars and have in satellites orbiting Mars and Earth. It also includes the computer programs that scientists and engineers use for visualizing what is happening on Mars, planning future automated laboratory actions, and accessing and analyzing the data returned.

One of the most important ideas I want to emphasize is that the engineering expertise for building such tools must be coupled with expertise for knowing what tools to build—or all of our futuristic ideas and expensive efforts will be wasted. Knowing what tools to build requires complementary expertise from the cognitive and social sciences, including ethnographic methods for observing and analyzing people at work, and even philosophy for being able to properly frame complex questions relating human intelligence and our motivations to automated systems.

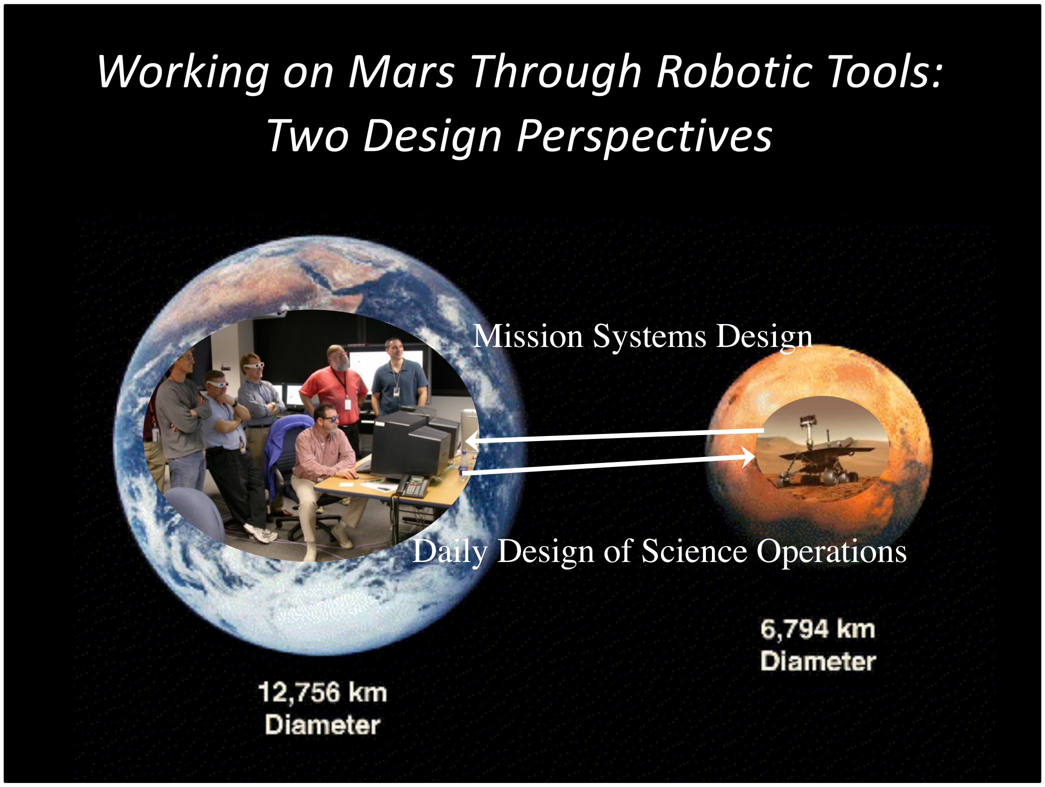

Mars is about half the size of Earth. Its distance from the Earth ranges from 56 million km to 400 million km. Communications therefore take roughly 5 to 20 minutes each way. A Mars year has roughly twice as many days as an Earth year. A Mars day is called a sol and is about 24 hours and 40 minutes. So the planet rotates slightly slower than the Earth and the seasons are roughly twice as long.

Designing the Mars Exploration Rover mission involved two levels of design problems: First is the problem of designing the overall Mission System– including the rovers, ground systems software, organizations and roles, processes and schedules. Second is the *daily problem* of designing a program of operations for each rover.

The overall design of the rover, including the choice of the instruments and how they work together strongly affects what operations you can do every day – and that affects how the scientists and instrument teams work together. How different disciplines work together affects how we study the planet and what we can learn about it.

So how do we design tools to help scientists study Mars?   On the one hand engineers invent new instruments that can operate in Mars conditions. On the other hand, we work with scientists to see how they actually explore a landscape.

My own work at NASA began in 1998 by studying how scientists explored Haughton Crater in the Canadian Arctic. In what are called "analog missions," computer scientists and space scientists have traveled to Haughton to learn and experiment with tools they might need on Mars. This includes tools to record data, navigate, and schedule their time and inventory supplies.

As an example of such tools, my team at NASA Ames developed a voice commanding system called Mobile Agents that could serve as a kind of astronaut's assistant on Mars. We were using some of the same technology that later became Apple's SIRI, but we went far beyond that in using computer agents to integrate GPS devices, biosensors, cameras, robots, email, and so on.

In my initial study of how the geologists study a landscape, I documented how they collected and organized sample, to be analyzed with instruments in laboratories back on Earth…

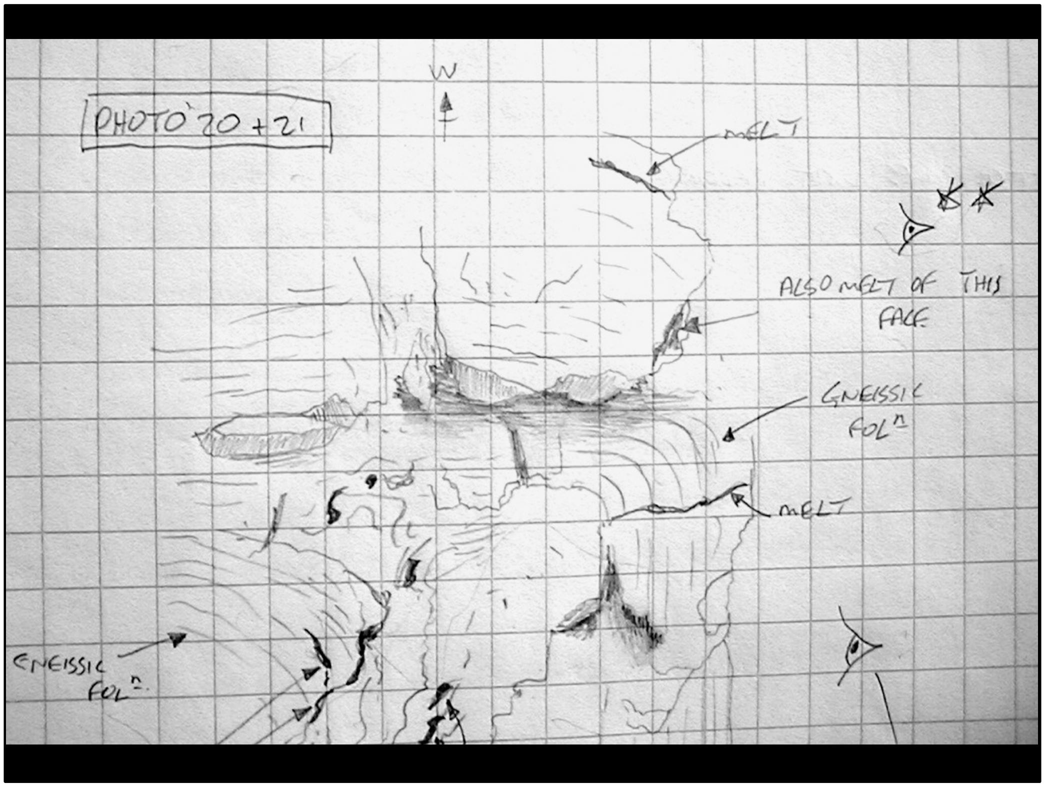

I studied what they diagrammed and described in their notebooks, and how this related to their published work.

I observed especially how they tended to work alone or in small groups.

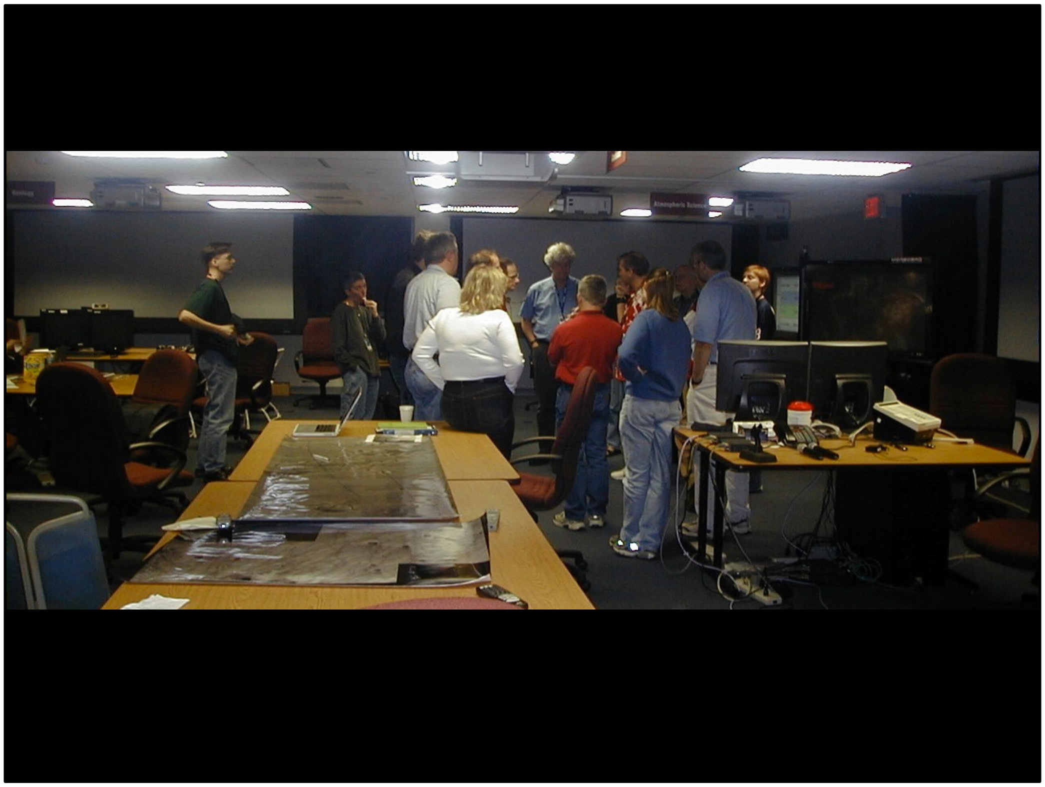

In 2004 I was fortunate to observe the Mars Exploration Rover scientists working in Pasadena, California, I was immediately taken by the contrast from what I observed in Earth field science: For MER the scientists are indoors, in a dark room, part of a large team, doing everything by consensus.

People from different disciplines are required to work together—geologists who in the Arctic would race to the nearest outcrop on a hill to survey the landscape, were working with mineralogists who wanted the rover to stop and take a new sample every few meters, and among them were laboratory scientists who had never done field work before.

So working remotely through a rover creates a new way of doing field science.

This new practice changes the scientists and leads them to relate to their tools, the rovers, in unexpected ways.

How could the scientists work together under these conditions? How could they accept the anonymity of the mission team where their names would never be associated in the public press with any of the decisions of what rock to analyze, how long to stay, what the data means?

How could people used to seeing, touching, and roving at will study a landscape through a programmed laboratory?

How was it possible at all to do field science on another planet while remaining here on Earth—let alone keep the team engaged for over nine years?

The key is that, although the scientists can’t directly see and control what is happening, the design of the rover instruments and computer software makes it possible for them to be virtually present on Mars. Through the combination of stereo and spectral images, and being able to move about, scuff the soil, and scrape rocks, they experience being there.

For example, Steve Squyres, the Principal Investigator of the mission, described their landing – “We realized we had landed in a crater, probably Eagle Crater, and that’s where we were. And then we noticed, 800 meters away—maybe we can make it—there’s Endurance Crater… Wouldn’t it be great to actually get there!” Squyres’s descriptions are all first person – “We had landed… We noticed… Maybe we can make it…” In this imaginative projection the scientists become the rover. And this experience of being on Mars is essential to the success of the mission—it enables them to actually do field science.

To know what rocks and soils are nearby, and what they can reach, or how long it might take to get somewhere, they use a combination of 3-D images, computer graphics, and simulations, often overlaying them.

These visualizations and associated tools provided by the Science Activity Planner allow the scientists to point to places, give them names, and control precisely where new photographs are taken and where the instruments are placed.

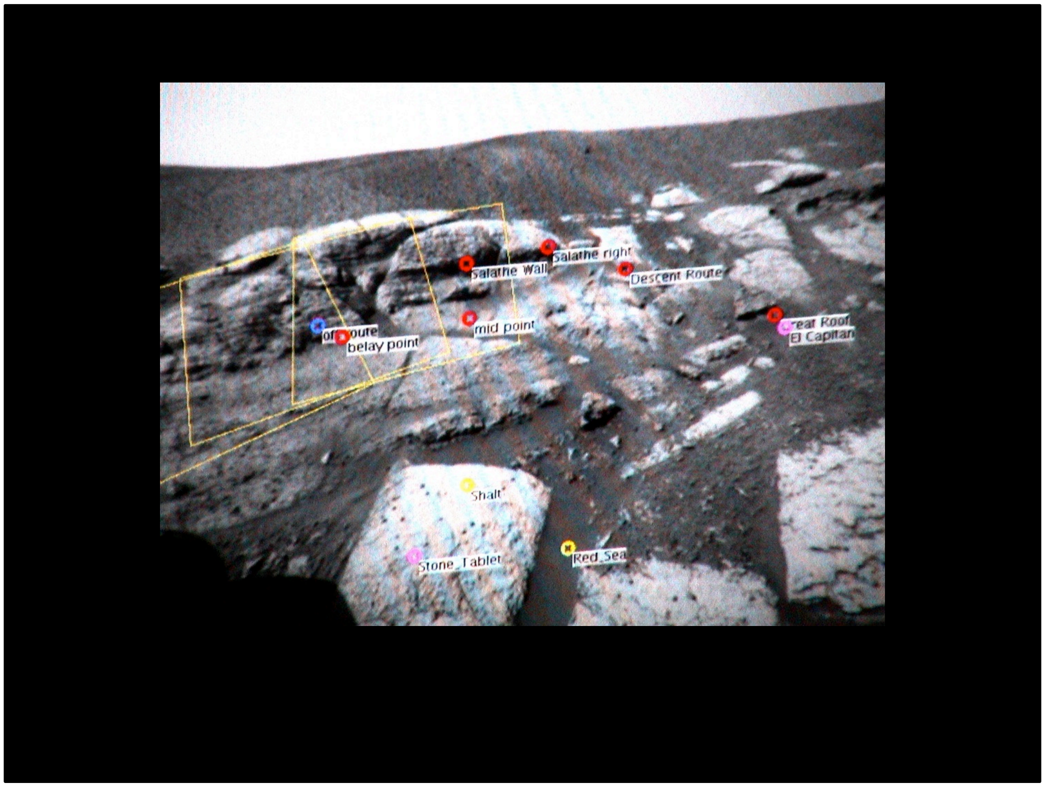

So for example, they can draw a yellow bounding box to specify where a camera should zoom in for a more detailed image.

Each photograph can be used like a map of an area on Mars because its location relative to the rover is precisely registered in the planning program. As we move in from panoramas used for navigating to images of outcrops and rocks, we can see and mark up details.

Even small rocks and patches of soil might be named and become targets for an atomic spectral analysis or a microphotograph.

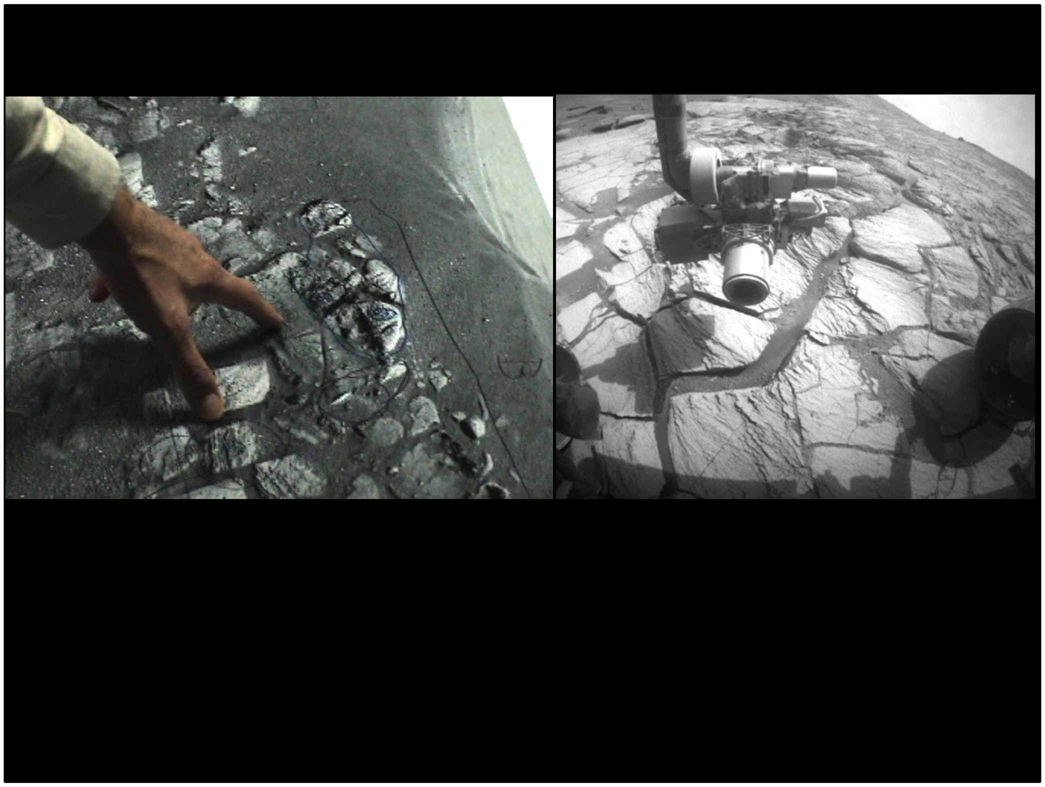

Combining these planning tools with their imagination, the scientists can work as if they were on Mars.

Jim Rice, a geologist on the mission said: “I put myself out there in the scene, the rover, with two boots on the ground, trying to figure out where to go and what to do, and how to make THAT what we’re observing with the instruments. Day in day out, it was always the perspective of being on the surface and trying to draw on your own field experience in places that might be similar.

David Des Marais, an astrobiologist, described it this way: “The first few months of the mission, they had these huge charts on the wall, engineering drawings of the rover, with all of the dimensions. We’d have some geometric question, “Well can we see this; can we reach this? Is this rock going to be in shade or is it in the sun?” We’d go stand and stare at those charts. And over time we stopped doing it so much because we began to gain a sense of the body. That’s projecting yourself into the rover. It’s just an amazing capability of the human mind—that you can sort of retool yourself.

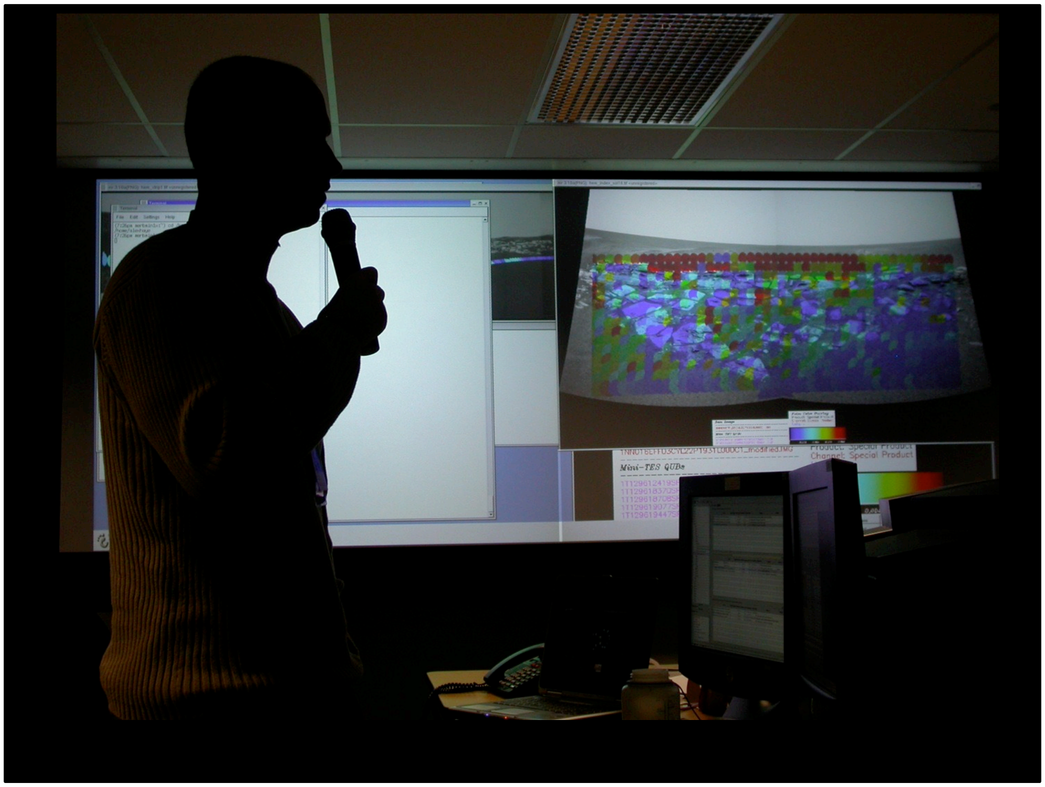

Acting through the robots they control, the scientists look around, manipulate materials, and move over the landscape.

They might pretend to be the rover, crouching down and gesturing with an arm to better imagine what's reachable.

And through the eyeglasses of special cameras, they can directly see iron minerals in the rocks…

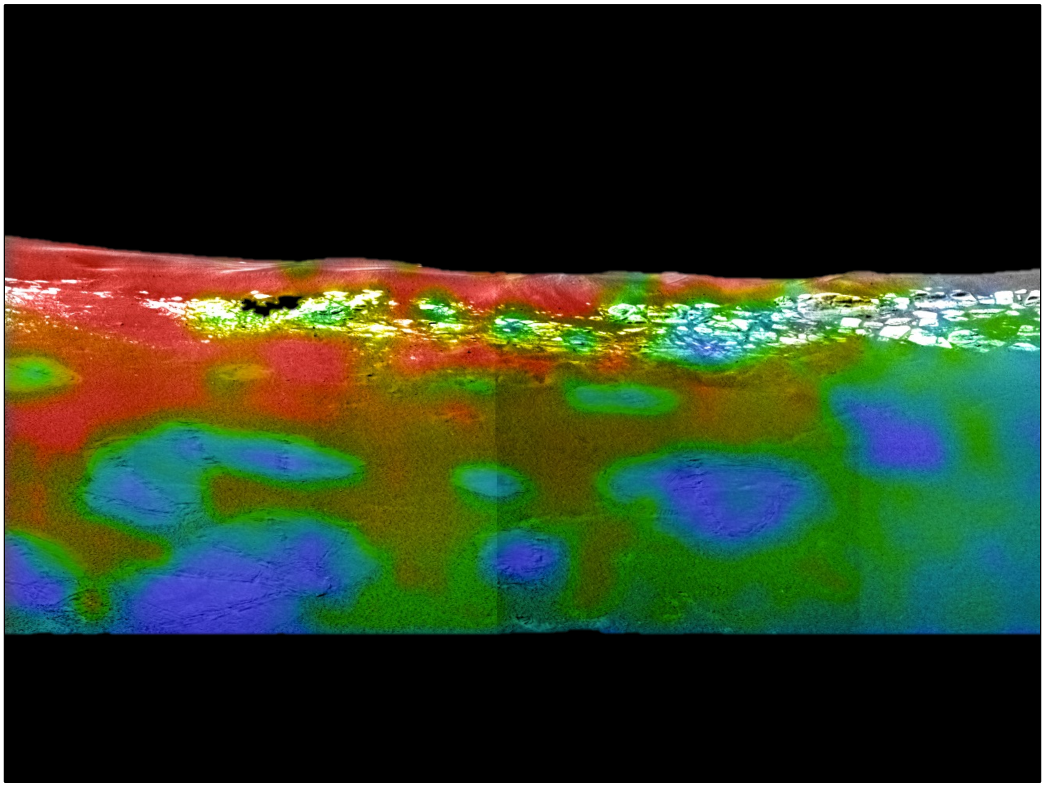

…they are transformed in a way to a kind of cyborg on Mars.

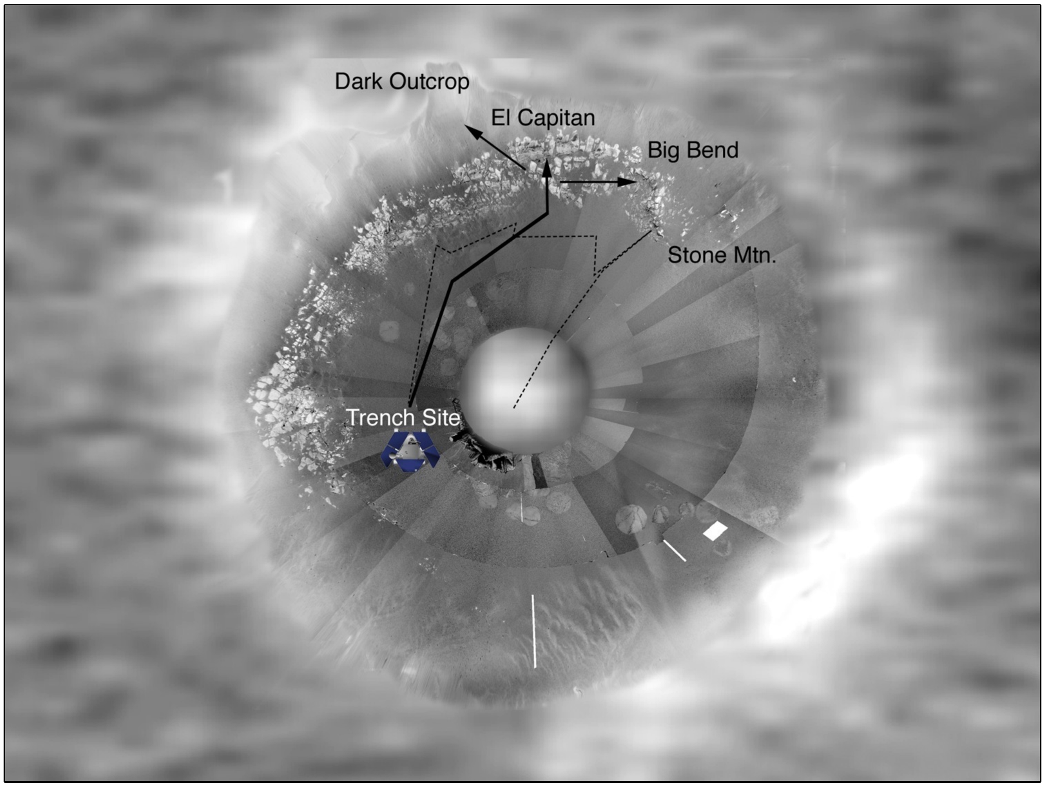

Computer graphics are composed with images taken from orbit to orient the planning and programming of routes and targets. This view of Eagle crater from 2004 combines a low resolution orbital image of the crater with images of the outcrop along the inside rim taken by MER.

Such third-person views provide an important way of locating the rover, and then by projection, locating yourself in the work on Mars.

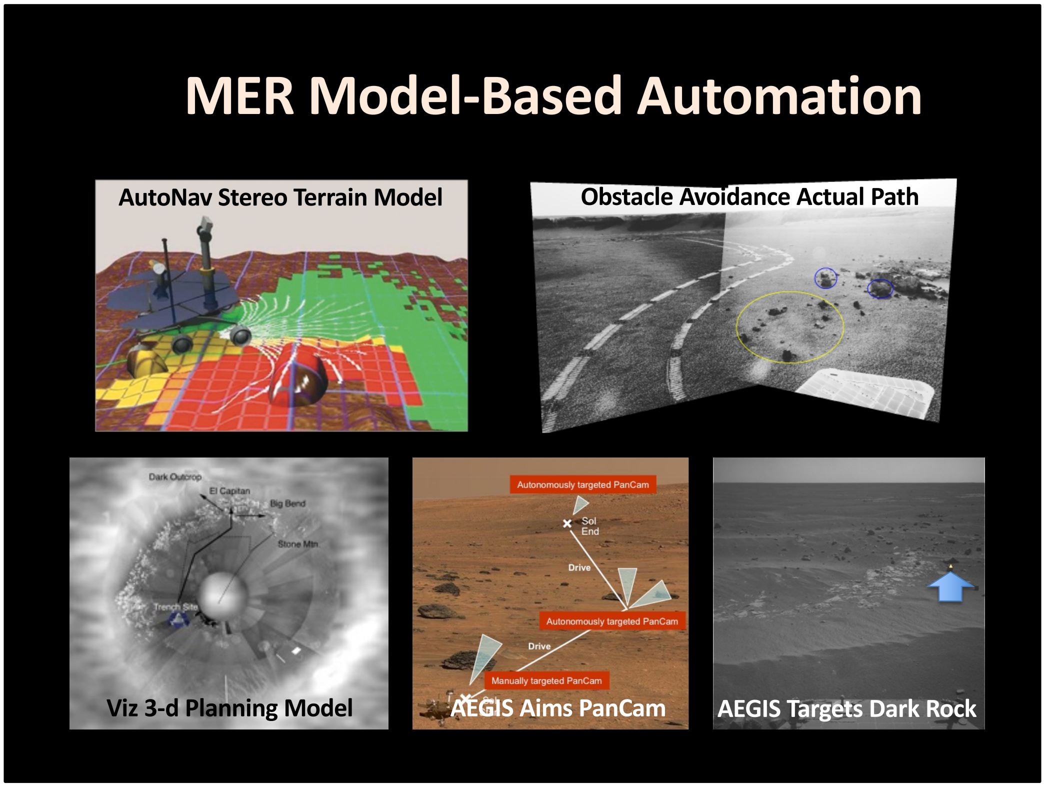

So you can see that several AI programs were used during the MER mission. I have already mentioned the Science Activity Planner (SAP), the computer program used by scientists and engineers for creating commands referring to features and targets. The combination of SAP using VIZ for graphic overlays was essential for providing a sense of virtual presence on Mars.

(note contnues on next page)

Returning to the main question—how is the MER actually programmed to carry out the operations the scientists want? After they specify operations, there's usually more requests than can be fulfilled in a Mars day, so they must work with engineers to develop a prioritized plan that respects the time, power, and memory constraints of the rover.

If you visited the science and engineering coordination meeting during the prime mission, the first 90 days of landing on Mars in 2004, you would see the scientists up front, on the bridge of a ship as it were, with huge displays of the martian surface that lay before them.

Behind the scientists, as if below decks, were the engineers, arrayed behind big square monitors showing the ship's power, memory, and an evolving timeline plan for tomorrow's work.

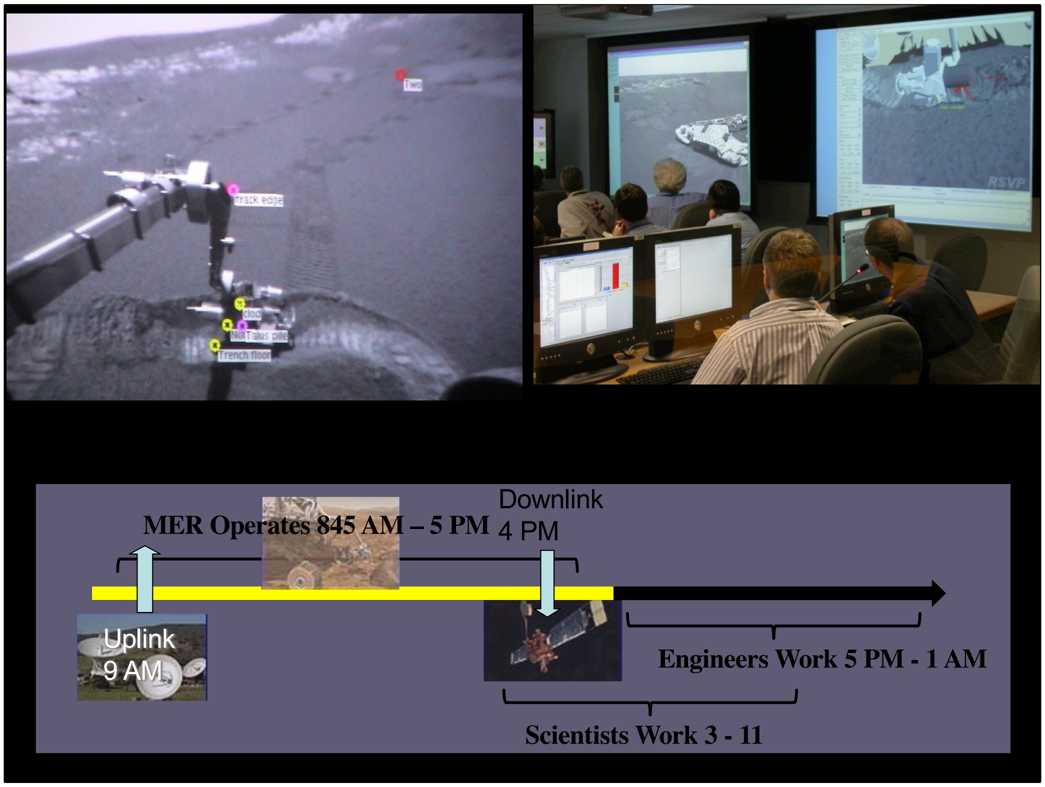

This coordination meeting would occur about 6 PM local Mars time every day. The scientists arrived at work about 3 hours earlier, mid-afternoon rover time, ready to receive photographs and instrument analyses, the result of the rover's work during the day.

The MERs are solar powered, so they work on roughly a 9 - 5 schedule, local Mars time. Every morning each rover would receive a new program for the day's work. So between dinner time on Mars and sunrise the scientists and engineers must finalize a plan—followed by a second shift of engineers who refine and test the program before it is sent to Mars.

If that weren't enough, keep in mind we were simultaneously operating two rovers on Mars for over five years. Spirit and Opportunity were in fact two missions, operating in parallel—they had their own meeting rooms for science operations at JPL, their own engineering programming teams, 6 PM coordination meetings, and their own cache of free ice cream.

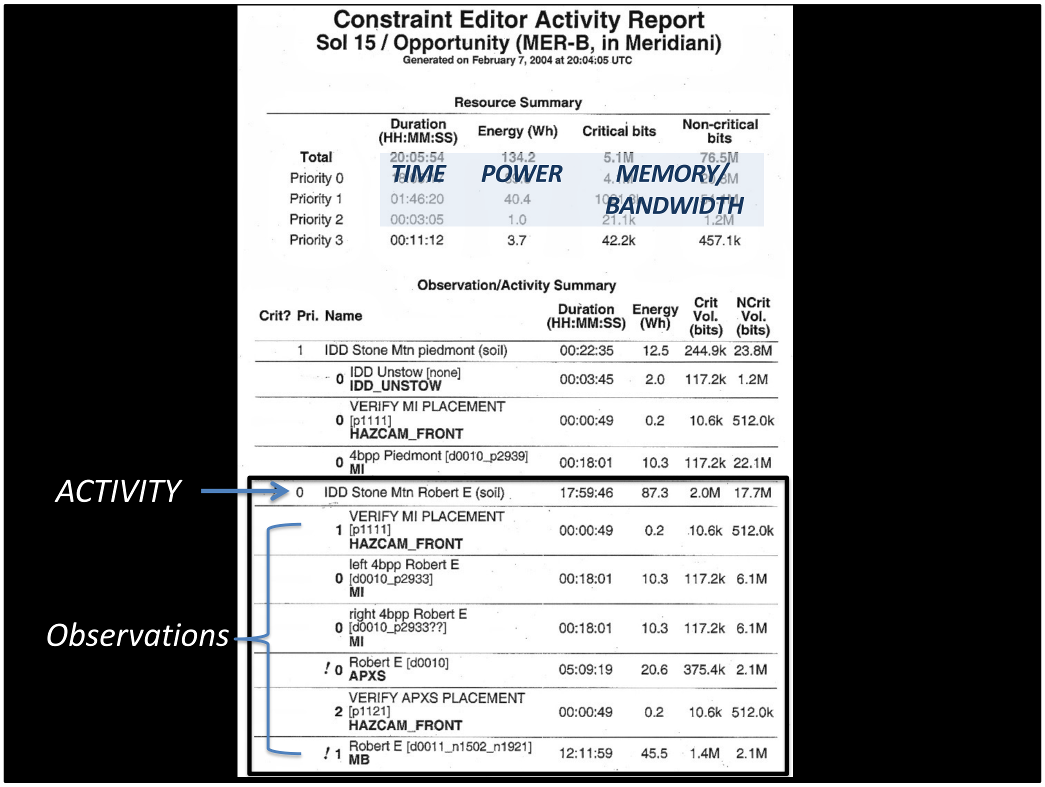

Constraint Editor Activity Report
Sol 15 / Opportunity (MER-B, in Meridiani)
Generated on February 7, 2004 at 20:04:05 UTC

Resource Summary

	Duration (HH:MM:SS)	Energy (Wh)	Critical bits	Non-critical bits
Total	20:05:54	134.2	5.1M	76.5M
Priority 0	[illegible]	[illegible]	4.[illegible]	[illegible]M
Priority 1	01:46:20	40.4	[illegible]	[illegible]
Priority 2	00:03:05	1.0	21.1k	1.2M
Priority 3	00:11:12	3.7	42.2k	457.1k

Observation/Activity Summary

Crit?	Pri.	Name	Duration (HH:MM:SS)	Energy (Wh)	Crit Vol. (bits)	NCrit Vol. (bits)
	1	IDD Stone Mtn piedmont (soil)	00:22:35	12.5	244.9k	23.8M
	0	IDD Unstow [none] **IDD_UNSTOW**	00:03:45	2.0	117.2k	1.2M
	0	VERIFY MI PLACEMENT [p1111] **HAZCAM_FRONT**	00:00:49	0.2	10.6k	512.0k
	0	4bpp Piedmont [d0010_p2939] **MI**	00:18:01	10.3	117.2k	22.1M
	0	IDD Stone Mtn Robert E (soil)	17:59:46	87.3	2.0M	17.7M
	1	VERIFY MI PLACEMENT [p1111] **HAZCAM_FRONT**	00:00:49	0.2	10.6k	512.0k
	0	left 4bpp Robert E [d0010_p2933] **MI**	00:18:01	10.3	117.2k	6.1M
	0	right 4bpp Robert E [d0010_p2933??] **MI**	00:18:01	10.3	117.2k	6.1M
!	0	Robert E [d0010] **APXS**	05:09:19	20.6	375.4k	2.1M
	2	VERIFY APXS PLACEMENT [p1121] **HAZCAM_FRONT**	00:00:49	0.2	10.6k	512.0k
!	1	Robert E [d0011_n1502_n1921] **MB**	12:11:59	45.5	1.4M	2.1M

The result of the 6 PM meeting between the scientists and engineers is a formal plan called an “activity report.” It specifies the sequence of operations and the time, power, memory/bandwidth for every step called an “observation.” Here is where the AI research on planning and scheduling was applied for MER. Practically speaking it was more important than AutoNav because without a constraint-based planner we would not have been able to reprogram the rover overnight.

A key operating constraint is that the two rover teams shared a single mission control center—where the engineers attend to monitors like those you might find in Houston's mission control, their computers connected to banks of satellites and earth stations that communicate with the rovers.

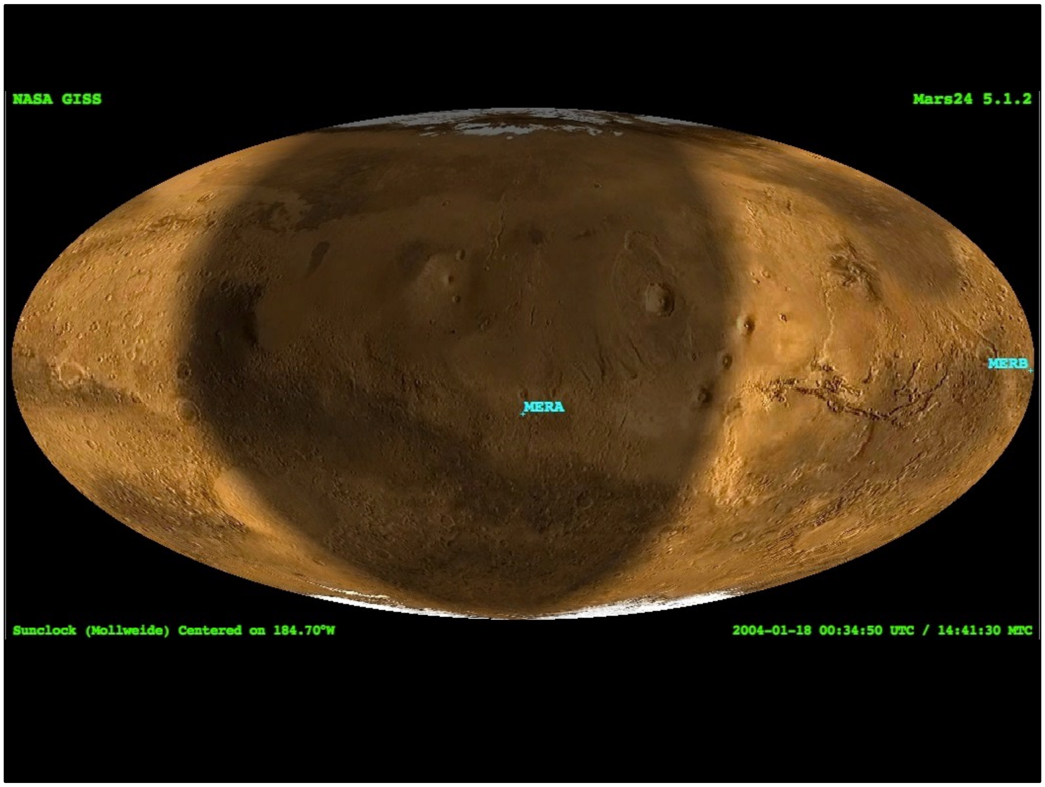

This common engineering activity required the two missions to be coordinated in a special way. If you look at the map of Mars with the landing sites of the two rovers at Meridiani Planum and Gusev crater, you'll notice that we landed the rovers about 180 degrees apart near the equator.

Most people will realize that the sun angle will be most favorable for solar power near the equator, but few realize how the geography of the mission relates to the problem of commanding the rovers every day.

Placing Spirit and Opportunity on opposite sides of Mars allows a single command center and management organization operating around the clock to focus on one rover at a time—controlling them is separated by a half day on Mars. This illustrates why understanding and designing the mission comprehensively as a total system is so important—the choice of landing sites affects the scheduling of facilities and operations.

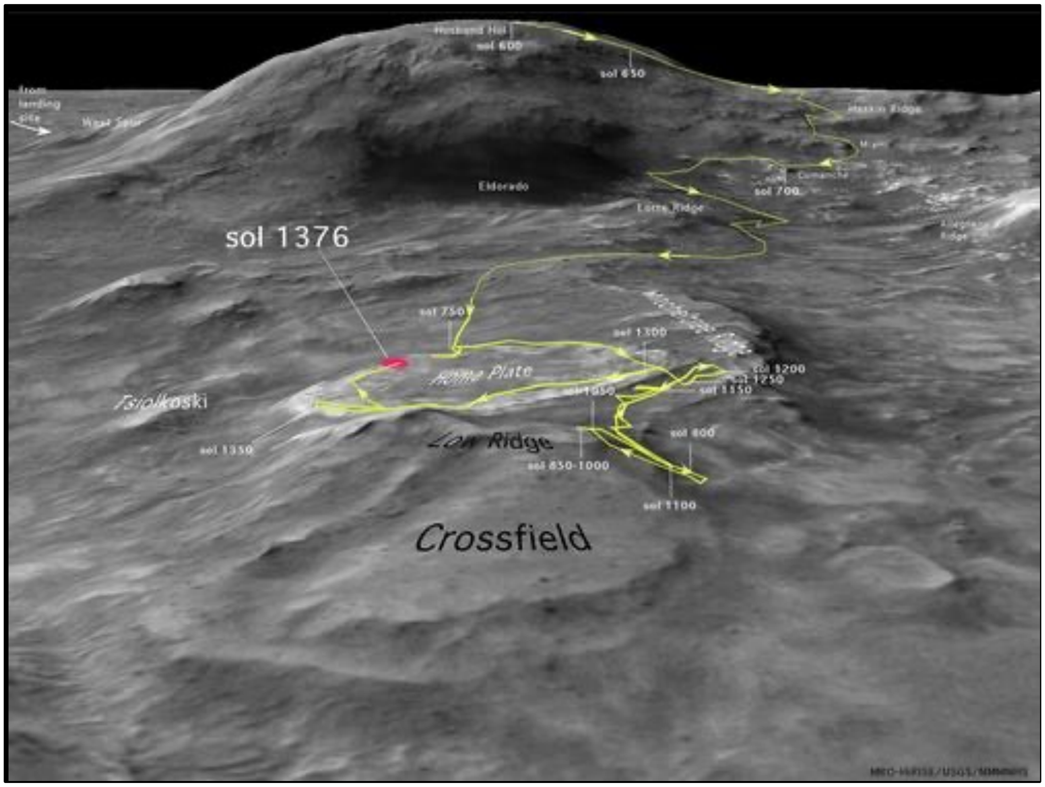

Let’s shift now to reviewing how the two expeditions on Mars with Spirit and Opportunity developed since the landing in 2004.

Here’s a map based on a photograph taken from Mars orbit. It shows how Spirit traveled from its landing site in Gusev crater about 100 meters up and over the Columbia Hills and down onto a kind of plateau called Home Plate. This map is a good example of how orbital photographs are important for doing field science on another planet. It shows how we use orbiting satellites and surface rovers together to study Mars.

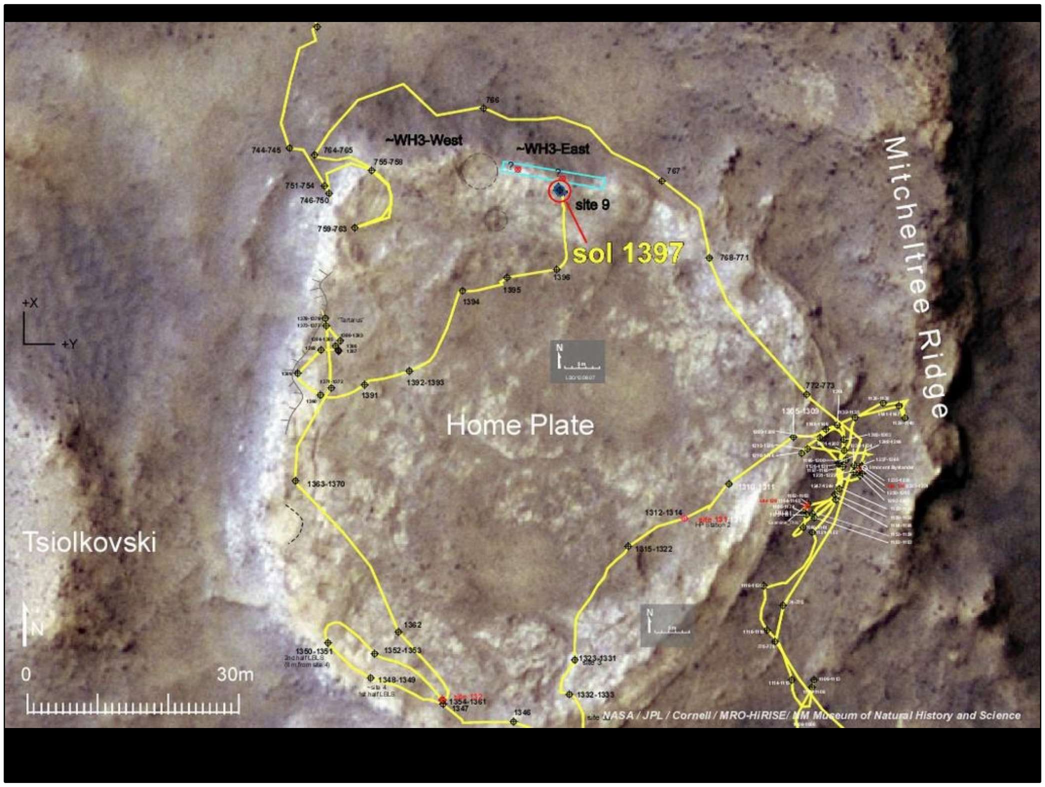

Home Plate is about 100 meters across and might be a remnant of large hot springs like those we find in Yellowstone Park in the United States. We spent about five years exploring the Columbia Hills and Home Plate with Spirit, way beyond the planned 90 martian days. This was possible because dust devils cleaned the solar panels, so the batteries could be recharged and the rover could survive the winters.

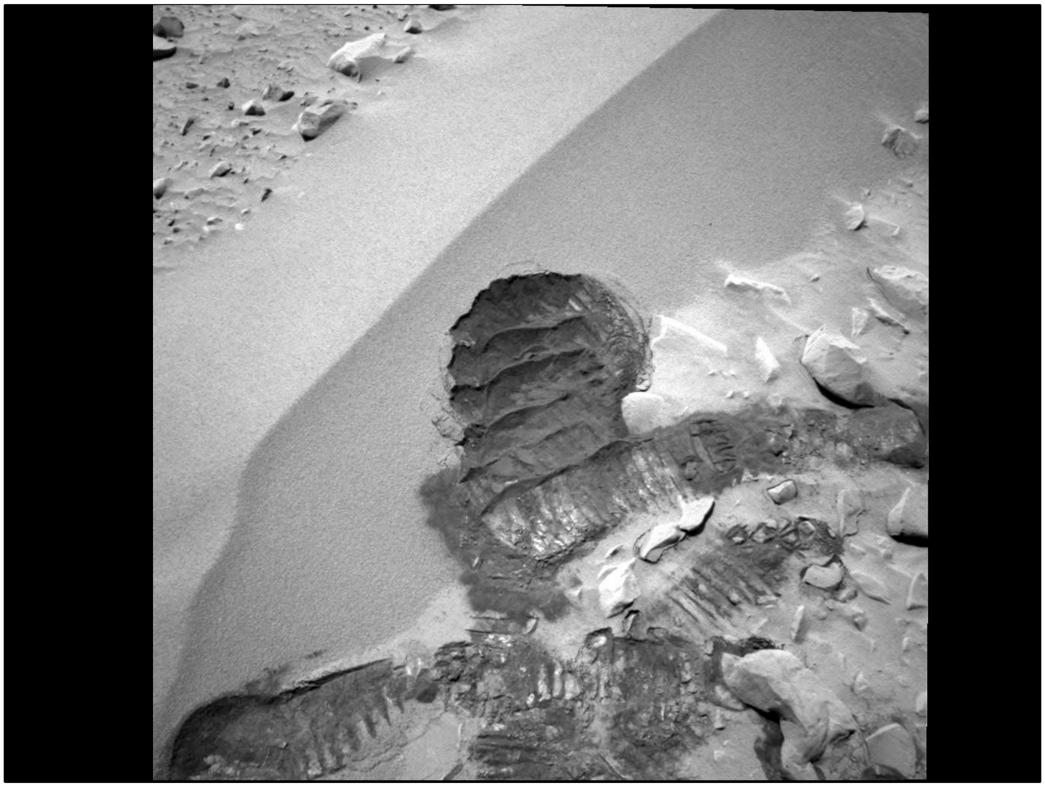

This is one of my favorite photos. It shows how we can use the rover to scuff the sand and get an up-close view, so it feels like we're looking at our own boot prints on Mars.

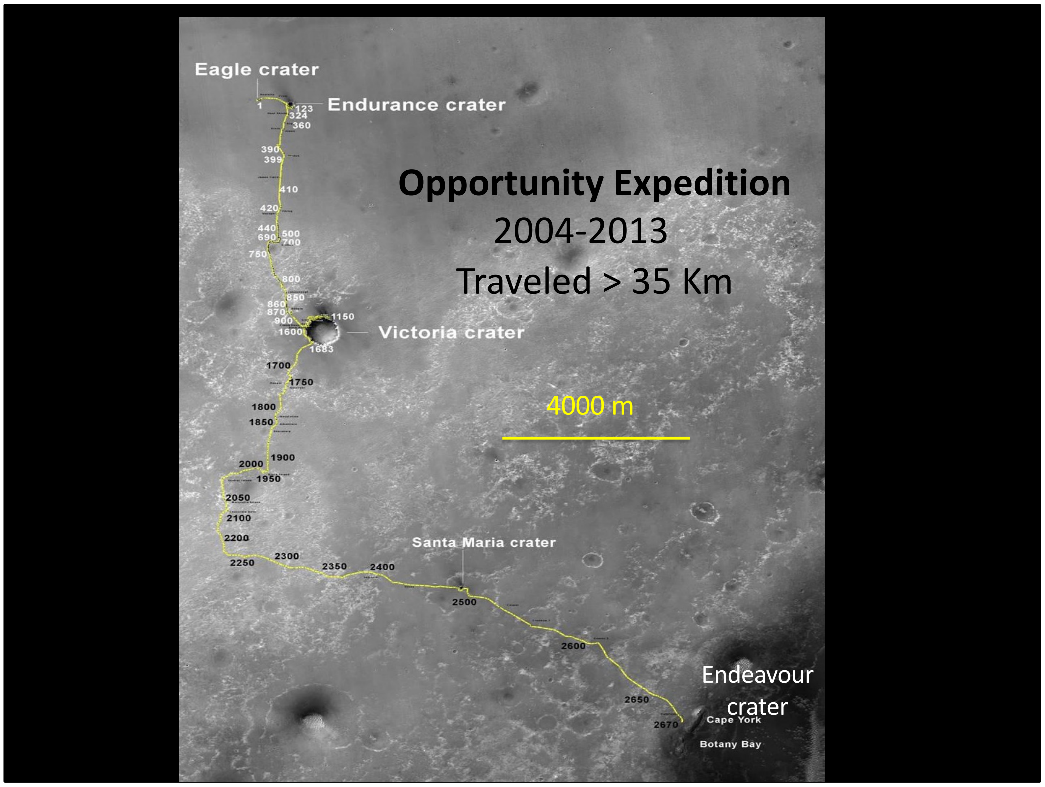

Spirit eventually became stuck in the sand in 2009 and we lost contact with it in 2010.

The expedition by the science team using the Opportunity rover has been even more spectacular. We have traveled over 35 km and are still using it to explore Mars. We have studied several very different craters and surface topographies during these nine years.

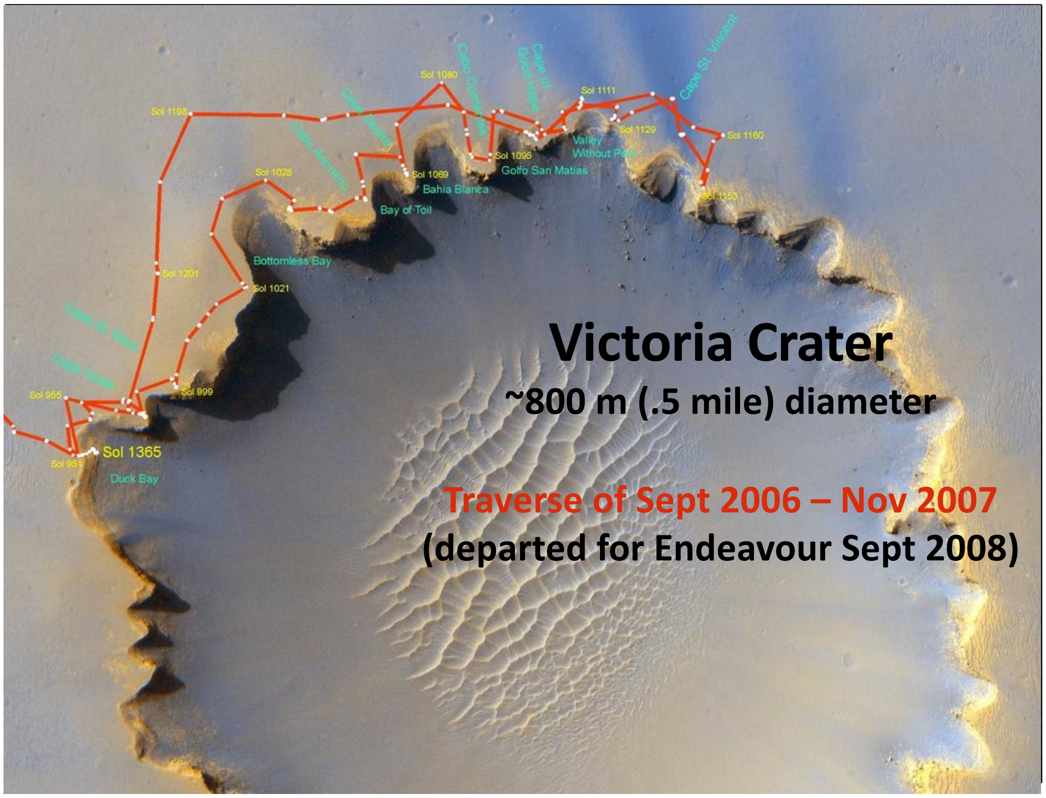

Here you see how we traveled around Victoria Crater.

After two years, we proceeded to Endeavour crater, where we arrived in 2011, three years after leaving Victoria!

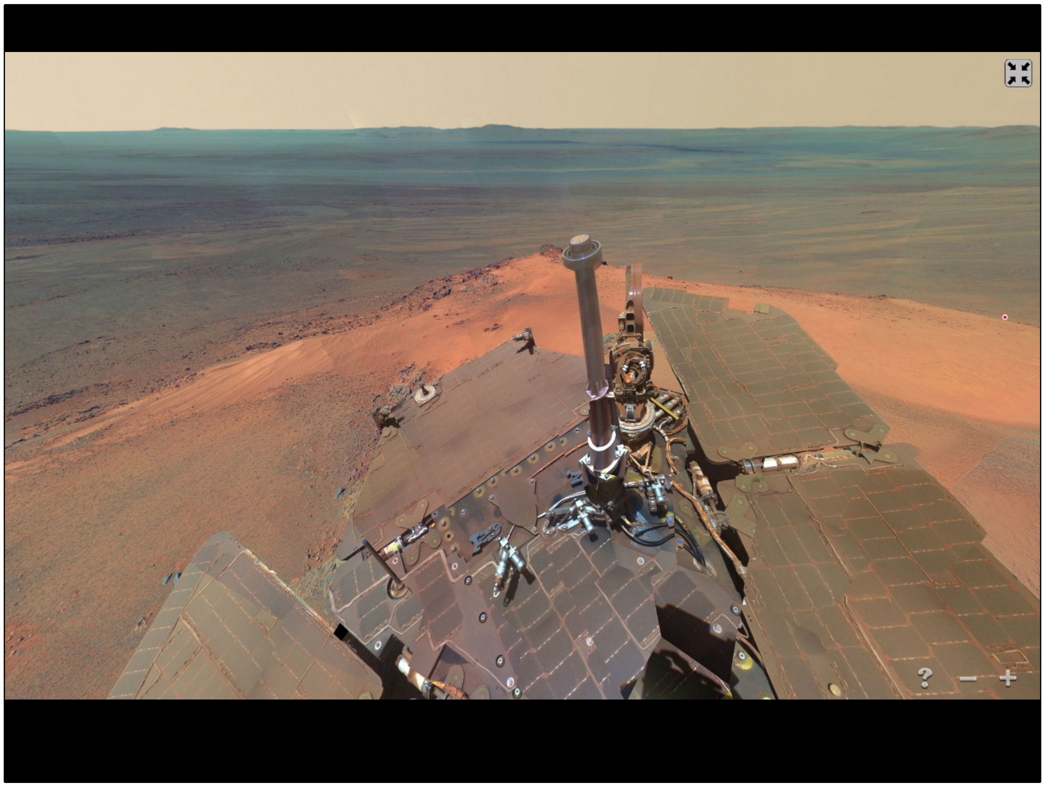

Here is Opportunity's view from the edge of Endeavour whose rim you can see in the distance. Endeavour crater is about 22 km across. (Screen shot from http://panoramas.dk/mars/greeley-haven.html)

This image is presented in exaggerated color to make some differences between materials easier to see.

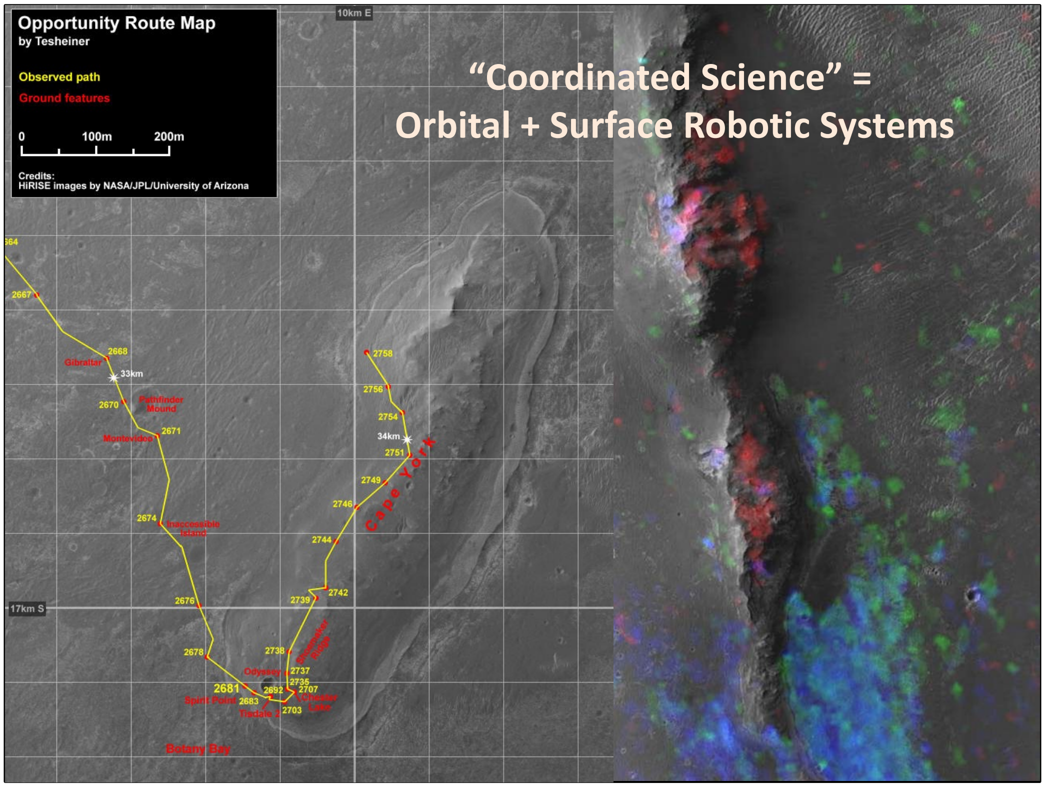

The study of Endeavour is an excellent example of how we use orbital photographs to efficiently target Opportunity’s path. From orbit (image on the right) we detected an area whose spectral signature resembled clays we find on Earth. Clays are formed by gradual chemical weathering of rocks in the presence of water. So we drove Opportunity directly over to where we saw the clays from orbit and are continuing to study this area today.

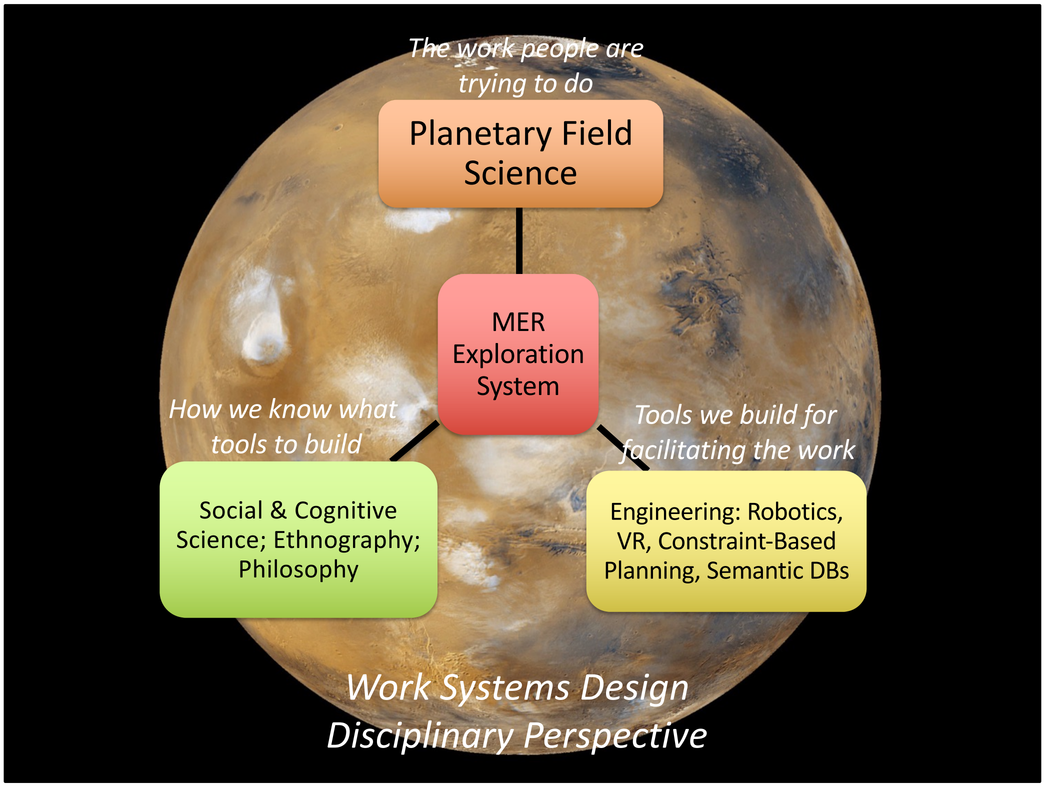

To summarize what I am saying about human-robotic interaction, let's return to the diagram I showed earlier. If we are thinking as systems engineers like Squyres, we realize we are developing a work system, not just an isolated robot or computer tool. So I refer not to the rovers alone, but to the "MER Exploration System" – which relates people, technology, and work processes.

The work the people are trying to do is planetary field science; everything starts by understanding what they want to do and how to help them work together.

The engineering work combines a wide variety of methods for developing and controlling scientific instruments, communication, power & navigation systems, and computer software. The planning and scheduling tools are based on detailed models of the subsystems and how they interact. These systems used model-based programming methods derived from AI research.

I emphasized that the engineering expertise for building such tools must be coupled with expertise for knowing what tools to build. This requires complementary expertise from the cognitive and social sciences to properly frame complex questions relating human capabilities and preferences to automated systems.

The Pivotal Design Principle: "One Instrument, One Team"

"You've got these sensors and each of them provides complementary bits of knowledge, so that the totality is more than the sum of the individual parts. *You're going to use the payload to fullest advantage, if people look at it as being entirely at their disposal….*Everything works together."

Steve Squyres, *MER PI*

Steve Squyres, the PI for MER, understood that systems thinking was necessary for a successful mission. He referred to it as "Coherent systems engineering—instruments relate to one another, so people play together. Everything is at your disposal. We don't have pancam guys arguing with mini-TES guys, but rather geologists arguing with chemists about exploration…" He found that his role as "science systems engineer" involved helping the engineers understand the scientific needs:

"What would typically happen is that the engineers would each be given part of the problem to solve, and they would go off and solve their part of the problem. But they didn't maybe have the broad overview of what we were trying to do scientifically. They had a specific job, a set of requirements to meet: to build a bracket, build a widget, build a wheel, build a computer, build a rover that would meet those requirements. But, they might not have enough insight into the science to realize that, even though this ONE PART of the rover meets the requirement, if you could make it just five percent better or ten percent better it would have a huge influence on the quality of the science."

"When it came to the *design* of the payload, rather than have a bunch of PIs each go off, and they each design their instrument and then they propose it, and you've got to figure out some way to make it work together, I was able, at the *outset* to design the microscopic imager, the RAT, the Mössbauer, the APXS, so they had fields of view that fit together nicely so that everything worked when you went to look at a given spot on the Martian surface. Just design the whole system so the pieces fit together."

MER: Serendipitous Benefits of Daily Commanding

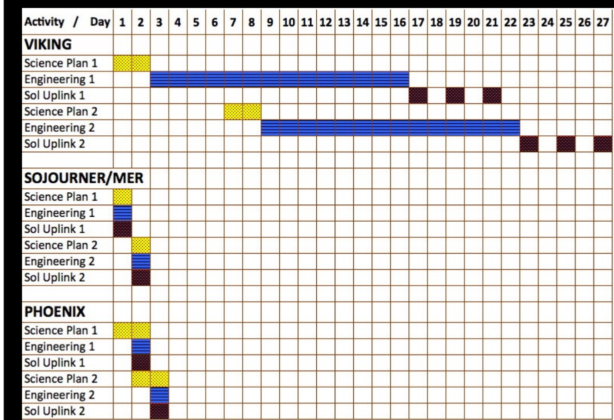

Viking veterans were skeptical; mission design presentations did not mention cognitive or social benefits

- Intensive daily activity, with sufficient feedback to assess the previous day's results
- Significantly increased the scientists' creative engagement of "being on Mars" (virtual presence)
- Enhanced conflict resolution, collaboration, and productivity, and hence quality of the field work

A successful mission depends not just on the rover and the instruments and how these are automated as a robotic laboratory on Mars, but also the *tools for programming* the rover, the *roles* of different people, their shared *facilities*, and especially the *schedules* by which everything comes together.

A key insight from MER is that commanding the rover every day is important for doing the scientific work.

At first, the veterans from the 1970s Viking mission were skeptical about the MER work system design—programming tools were much more primitive in the 1970s. It required more than two weeks to convert an initial plan into a tested program that could be uploaded to the Viking lander on Mars. But we showed during the Sojourner/Pathfinder mission in 1997 that daily commanding was possible.

MER is much more complicated than Sojourner with several interacting instruments, an arm, and lots of timing, power, and bandwidth management required. In the end daily commanding was surprisingly important for keeping the scientists engaged over a long time and working together as a team.

By the way, you can see that for Phoenix, which was stationary—it was called a "lander" not a mobile laboratory—we started planning a day earlier because the area we might study was predictable—it was the same every day.

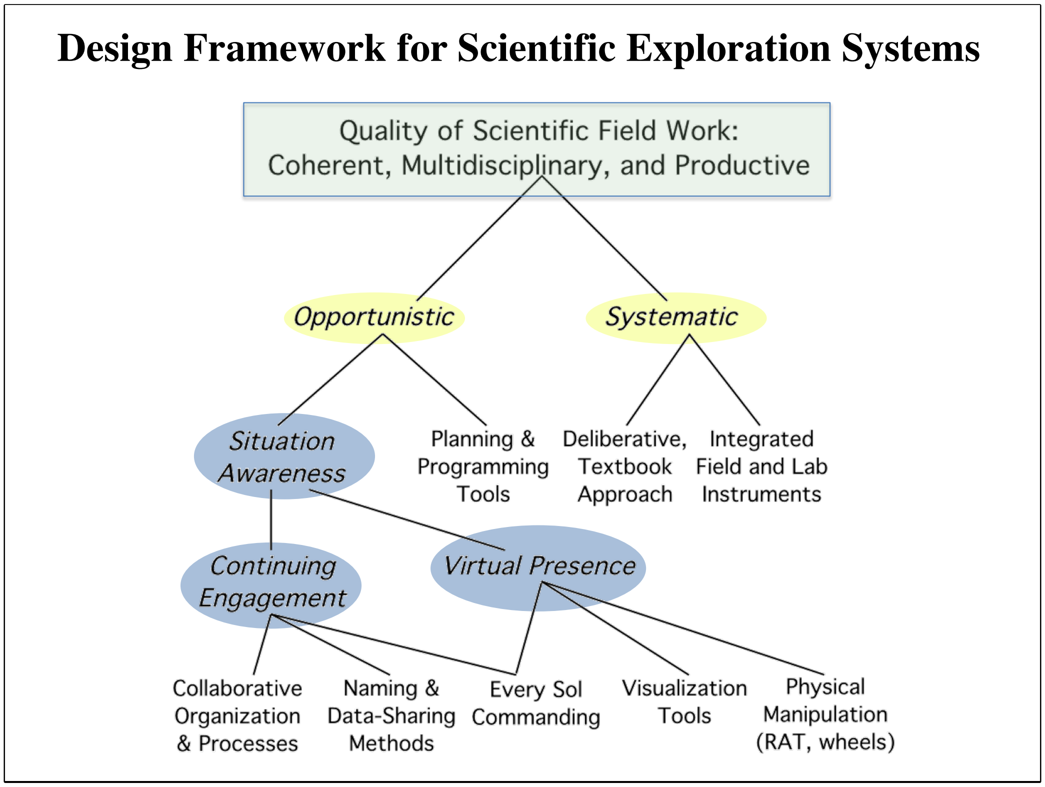

It was by contrast to Phoenix that I realized how opportunism and systematicity are different strategies for conducting field work. Design of Phoenix operations was exclusively systematic; MER was designed to allow opportunism, primarily because it's a rover—you have no idea how far you will go and what you might encounter. Phoenix had a fixed reachable workspace and bounded lifetime of five months (because its batteries were depleted and it froze during the martian arctic winter).

I think this diagram should make clear that engineering is itself a complex multidisciplinary enterprise—yes, many engineers are specialists working on power, navigation, instrument automation, and so on. And everything must be integrated and designed so it interoperates within the constraints of space and weight, time, energy, memory, bandwidth, and so on. But above all, the system must help the scientists to study Mars.

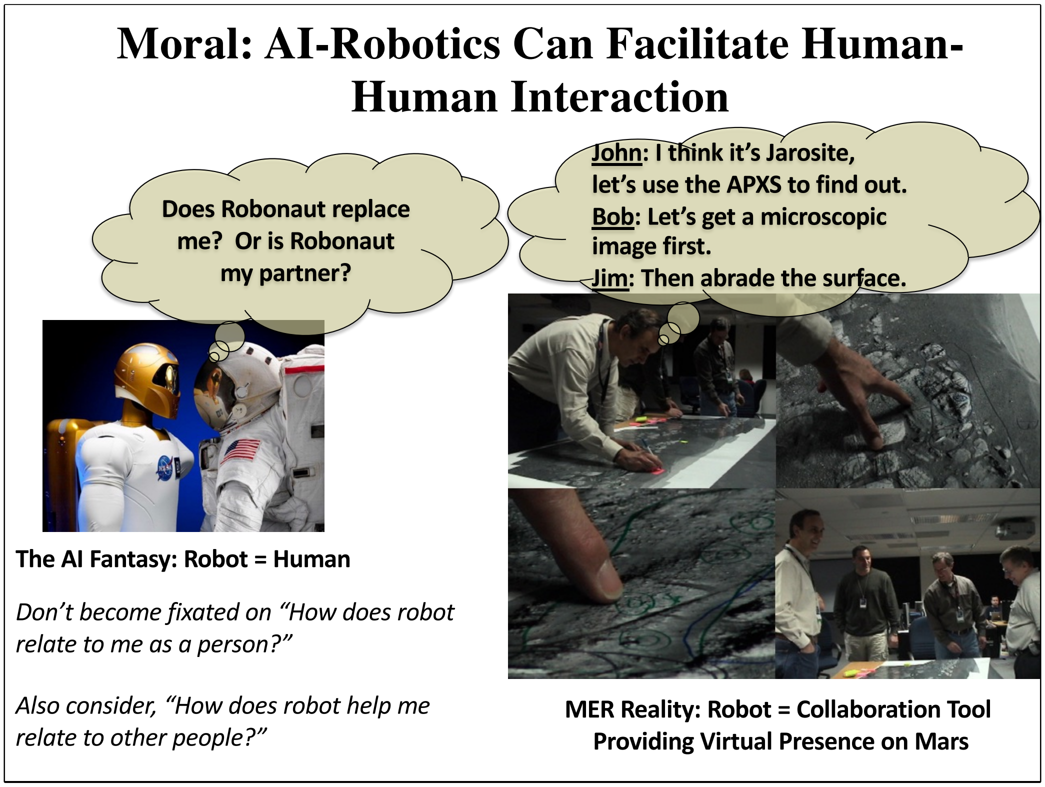

Many people have talked about space exploration in terms of people partnering with robots or robots replacing people. In MER we see that equating robots with people is too simplified—MER is our physical surrogate on Mars, but cognitively and socially, MER's design as a work system enables scientists from different disciplines to work together and to work with engineers. MER is a collaboration tool—it helps people work together, it doesn't replace them. We're learning the same lesson from social media and networking how computers can bring together the thoughts and work of many people in a collective enterprise.

Rather than replacing the scientists, the rover's automation makes them agents on another planet; and by the synergistic design of the instruments and computer programs, the laboratory becomes a tool that promotes multidisciplinary collaboration. The rover's instruments work together and the Science Activity Planner program models HOW the instruments operate and interact. This enables different scientific disciplines to use the laboratory together, as if it were a single instrument.

So when we think about Human-Robotic Interaction, I suggest that we try to focus on what people are doing, focus on their practices, how they want to interact with the world and each other in time and space. Automation is great when (like AutoNav) it complements what people are able or want to do. Rather than trying to replace people entirely, automate what you can and focus on how you can help people communicate and work better together.

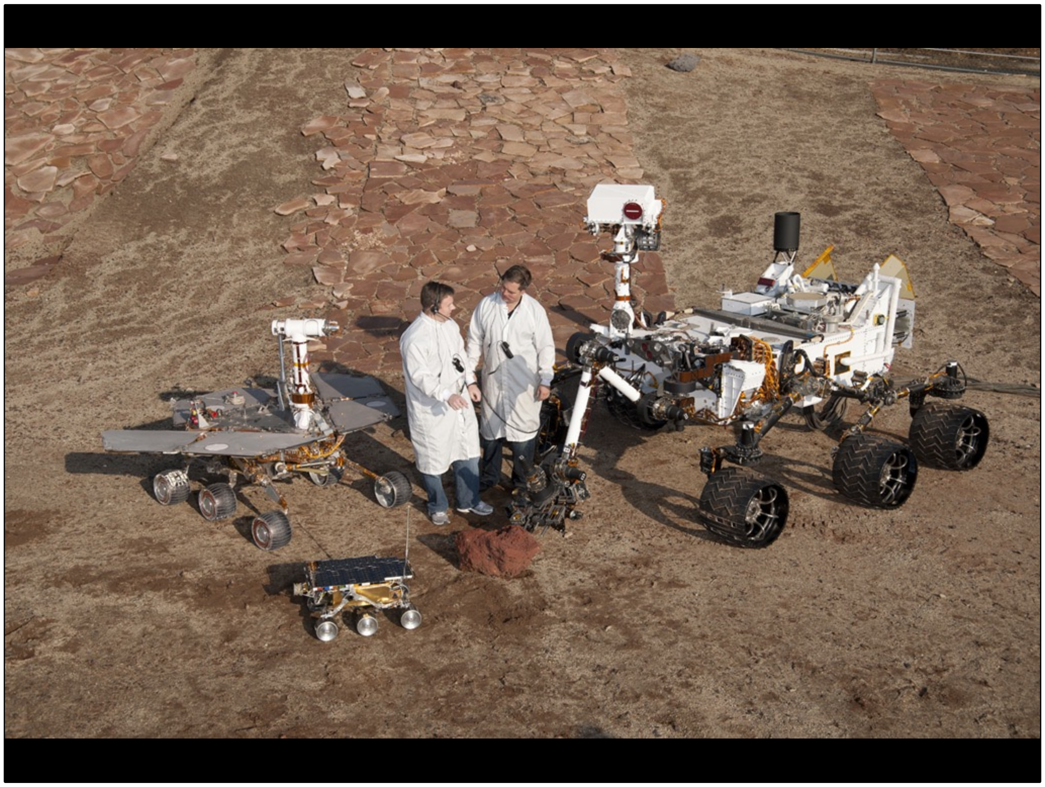

I could not end my talk without mentioning the most important and sophisticated mission we have ever sent to the surface of another planet. As most of you know we have another rover on Mars today called the Mars Science Laboratory.

Here you see MER on the left and MSL on the right, nicknamed Curiosity. Curiosity is about five times heavier than MER. It is powered by a radioisotope generator that converts heat to electricity, so it can operate for several years and even at night. MSL carries much more sophisticated instruments for sampling and atomic analysis.

In the bottom left you see tiny Sojourner the first martian rover, which operated in 1997 as part of the Mars Pathfinder mission in which methods and tools for reprogramming the rovers on a daily basis were first tested.

=============

Sojourner and its flight spare, named Marie Curie, are 2 feet (65 centimeters) long. The Mars Exploration Rover Project's rover, including the "Surface System Test Bed" rover in this photo, are 5.2 feet (1.6 meters) long. The Mars Science Laboratory Project's Curiosity rover and "Vehicle System Test Bed" rover, on the right, are 10 feet (3 meters) long.

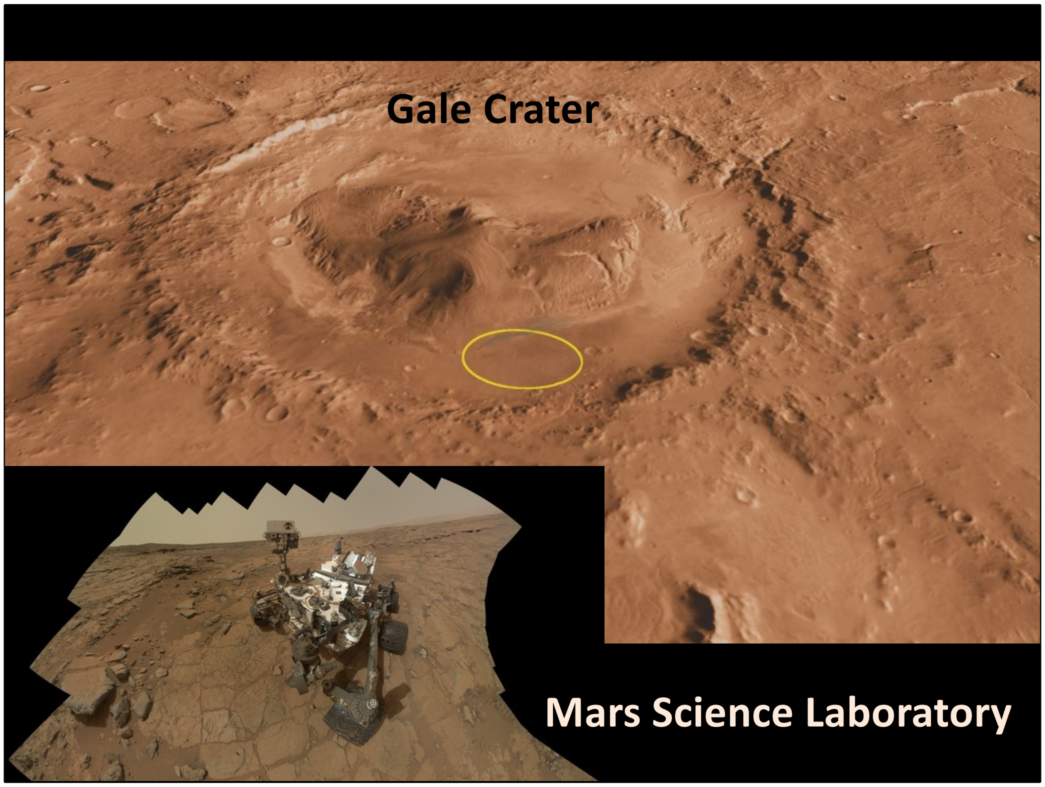

Curiosity landed in Gale Crater on August 5, 2012. Using the drill and the chemical laboratory onboard, we have determined that water was at this location for a long time. And unlike where Opportunity landed at Meridiani, it was not acidic or salty, but more conducive to life. In other words, with Opportunity we showed that water was on Mars' surface. With Curiosity we have learned that the water at Gale Crater provided a habitable environment.

In coming years, we will climb and explore the central mound of Gale Crater. It's called Mount Sharp and towers 5 kilometers over the landing site and has been built up for perhaps 2 billion years. We will study the layers to learn how the deposits were formed and how the climate and landscape changed over time. Working through the robotic Mars Science Laboratory, we are on an expedition, doing field science on another planet.

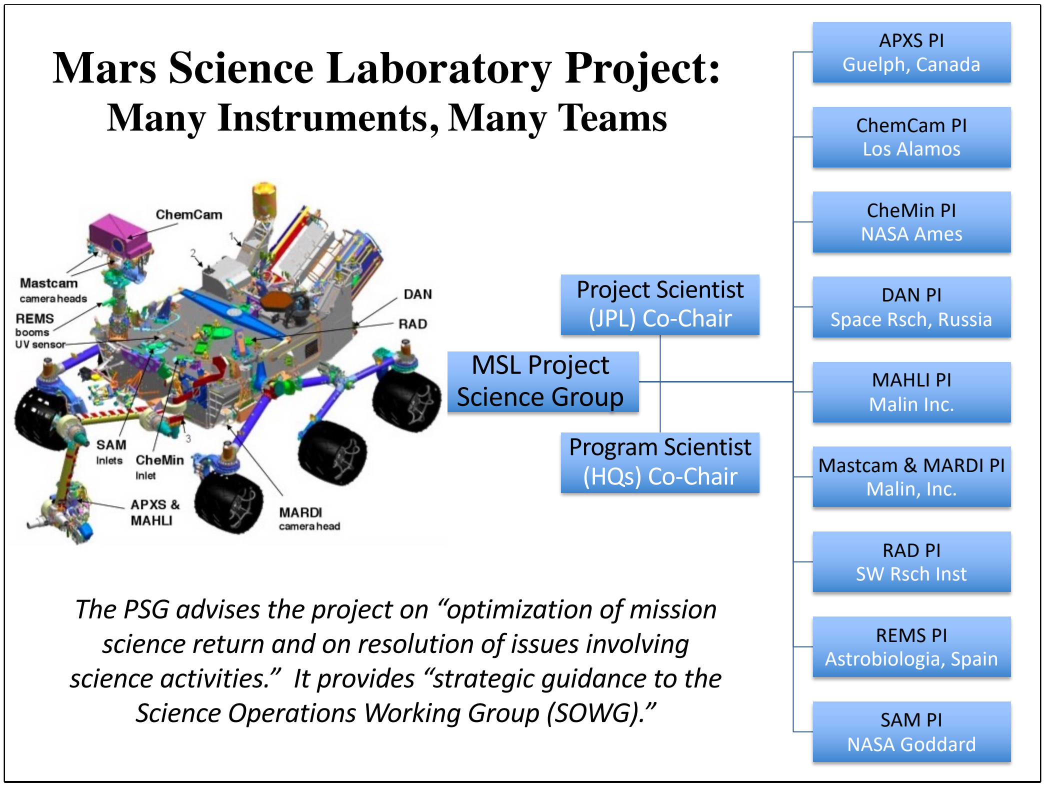

The design of the MSL project is very different from MER. An entire layer of organization has been added between the engineers and the scientists. Rather than just one Principal Investigator, like Steve Squyres on MER, there are nine PIs, representing nine instrument teams. The PIs along with the project scientist at JPL and a program scientist at NASA headquarters constitute a Project Science Group called the PSG. "The PSG advises the project…." (see slide).

It is not clear why this extra layer of management is required. MER has six instruments but does not require six PIs. One difference is that MER was a university-led mission, so the work system was designed by Squyres, a Cornell professor.

MSL has been brought entirely inside JPL like Cassini and most of the previous planetary missions. This is a very different work system. It is not necessarily better or worse. We might never find out because JPL would not allow an ethnographic study to be carried out as Squyres allowed during MER.

But I do know from talking to one of the scientists that this more bureaucratic organization has isolated some of the scientists from the decision making, which is hierarchical, rather than the consensus approach used for MER. Instead of "One instrument, One Team," MSL is a rover with many instruments and many teams. It seems likely that this work system design makes MSL less effective for promoting collaboration.

If you would like to learn more about the ideas I presented today, you can read my book, *Working On Mars*, which was published as a trade book last September by MIT Press.

The book includes 50 illustrations and 25 color plates, including many of my own photographs. I wrote this specifically for students interested in science and engineering careers, as well as professionals interested in how robotic technology has revolutionized the work we can do on another planet. I hope you will enjoy reading it.

So with this image of the shadow of the Mars Exploration Rover during a sunset on Mars, I will be happy to take any questions.

(END)