MIT Open Access Articles

Visual Apollo: A Graphical Exploration of Computer-Human Relationships

Massachusetts Institute of Technology

Visual Apollo: A Graphical Exploration of Computer-Human Relationships

Yanni Alexander Loukissas, David Mindell

Introduction

As NASA struggles with an uncertain national policy environment, it is seeking new combinations of human and robotic modes of exploration. During the Apollo era, establishing human presence in space was an integral part of NASA's work.[1] More recently, the Mars Exploration Rover missions have demonstrated how rich and successful remote exploration of a planetary surface can be.[2] Space exploration is one of many examples of technical operations conducted in extreme environments that are raising new questions about the relative importance of human and remote presence. What does it mean to "be there?"

Researchers, designers, and operators of situated robotics and remote technologies are testing increasingly distributed configurations of human–machine teams in technical operations, from space exploration to surgery.[3] However, with new forms of automation come unexpected changes in the social organization of work. If we are to understand the implications of increased automation for control, responsibility, and safety, we need to look beyond the methodological boundaries of traditional computer–human interaction studies in both human factors and social studies of technology.

Human factors studies tend to focus on individual operators and quantitative representations.[4] They emphasize workload, interface, and situational awareness but frequently overlook the social organization of human–machine teams and the cultural production of operator roles. Yet these factors can have profound effects on the acceptance of new technologies, in both engineering decisions and national policy making. Although social studies of technology address these broader socio-cultural issues, they often do so in formats that privilege qualitative data, incremental analysis, and linear explanations at the expense of considering technical and temporal measures of events.[5] The study of distributed computer–human relationships requires new methods that are capable of picking up on multi-channel interactions. In our collaborative work, we are developing methods that bring together a combination of individual, social, quantitative, and qualitative data in rich, graphical, real-time representations.[6]

1   For more on the changing role of humans in spaceflight, see David Mindell, *Digital Apollo: Human and Machine in Space-flight* (Cambridge: MIT Press, 2008).

2   See William Clancey, *Working on Mars: Voyages of Scientific Discovery with the Mars Exploration Rovers* (Cambridge: MIT Press, 2012); William Clancey, "Becoming a Rover," in *Simulation and Its Discontents,* Sherry Turkle, ed. (Cambridge: MIT Press, 2008), 107–27; and Zara Mirmalek, "Solar Discrepancies: Mars Exploration and the Curious Problem of Interplanetary Time" (doctoral thesis, University of California, San Diego, 2008).

3   James Hollan, Edwin Hutchins, and David Kirsh, "Distributed Cognition: Toward a New Foundation for Human–Computer Interaction Research," *ACM Transactions on Computer–Human Interaction* 7, no. 2 (June 2000): 174–96.

4   Raja Parasuraman, "A Model for Types and Levels of Human Interaction with Automation," *IEEE Transactions on Systems, Man, and Cybernetics* 30, no. 3 (2000): 286–97; Thomas B. Sheridan, *Humans and Automation: System Design and Research Issues* (New York: Wiley-Interscience, 2002).

5   Bruno Latour, *Science in Action: How to Follow Scientists and Engineers Through Society* (Cambridge: Harvard University Press, 1988); Gary Downey and Joseph Dumit, *Cyborgs & Citadels: Anthropological Interventions in Emerging Sciences and Technologies* (Santa Fe, NM: School of American Research Press, 1997); Lucille Suchman, *Human–Machine Reconfigurations: Plans and Situated Actions* (New York: Cambridge University Press, 2007).

Figure 1
A screenshot of the Apollo 11 visualization
after it has played through. The original
visualization is an interactive, color
application with integrated audio. For a
video-snapshot of the application go to:
http://www.mitpressjournals.org/doi/
suppl/10.1162/DESI_a_00258
© 2014 Yanni Alexander Loukissas
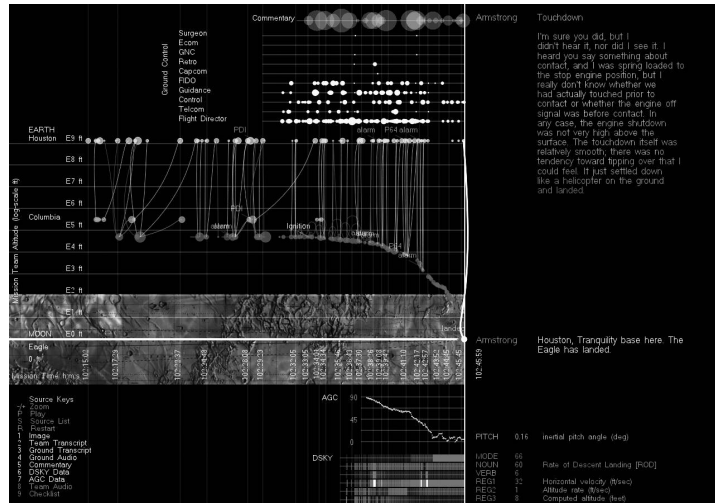
6   This paper elaborates on the contribu-
     tions and limitations of a stand-alone
     visualization of the first lunar landing,
     presented at CHI 2012 in the interactivity
     session. See Yanni Loukissas and
     David Mindell, "A Visual Display of
     Sociotechnical Data," in *Proceedings
     of the 2012 ACM Annual Conference
     Extended Abstracts on Human Factors in
     Computing Systems Extended Abstracts*
     (CHI 2012): 1103–6. Furthermore, it
     frames that project within a broader
     research strategy: to develop a visual
     language for bringing quantitative and
     qualitative data into view for interpretive
     studies of digital culture.
7   For an example of the use of animation
     in an airplane accident investigation, see
     Colgan Air Flight 3407, National Trans-
     portation Safety Board, Public Hearing,
     May 12–14, 2009, www.ntsb.gov/news/
     events/2009/buffalo_ny (accessed March
     4, 2012).
8   Data include material from Mindell,
     Digital Apollo, as well as from the
     following: NASA, Apollo 11 Descent and
     Ascent Monitoring Photomaps, NASA
     Manned Spacecraft Center, Houston, TX
     (1969a); NASA, Apollo 11 Technical
     Debrief, NASA Manned Spacecraft
     Center, Houston, TX (1969b); NASA,
     Apollo 11 On-Board Voice Transcription,
     NASA Manned Spacecraft Center,
     Houston, TX (1969c); NASA, Apollo 11
     Range Data, NASA Manned Spacecraft
     Center, Houston, TX (1969d); NASA,
     Apollo 11 Technical Air-to-Ground
     Voice Transcription, NASA Manned
     Spacecraft Center, Houston, TX (1969e);
     and Spacecraft Films, *Apollo 11: Men
     on the Moon,* Twentieth Century Fox
     Entertainment (2002).

In this paper, we explain our use of data visualization as an approach to the study of technologically mediated human roles and relationships. Although many quantitative researchers make use of visual methods, they often forgo the difficult integration of qualitative details and socio-cultural context in the process. In the domains of spaceflight and aviation in particular, visualizations are used in accident investigations; however, they often focus on the machinery rather than on the network of people involved.[7] Meanwhile, qualitative researchers have, by and large, avoided visualization. So, sensor data and numerical calculations are graphed and charted, while human communications and relationships remain unseen. Our approach is to develop a common format for looking at all these data to reveal enmeshed social and technical dynamics otherwise difficult to imagine or communicate.

To start with a long view of human–machine relationships and make use of available data, our first visualization is from the early history of spaceflight. Using the 1969 Apollo 11 lunar landing as an example leverages previous work done in *Digital Apollo: Human and Machine in Spaceflight,* as well as newly recovered downlink data, which reveals states of the Apollo guidance computer in sync with human interactions among the astronauts in the lunar module and the command module, as well as personnel in ground control.[8] In what follows, we describe the Apollo 11 visualization in detail and situate it within a history of influential precedents. We also highlight the patterns of computer–human interaction that our visualization reveals and account for those patterns it overlooks. The benefits of data visualizations are many: Sources can be extensive, presentations can be more accessible, and time can be introduced as an interactive variable. However, the format also imposes constraints. Creating a legible visualization of data requires leaving out much that could enrich our understanding of an event but that might not graphically fit. In

9   For historical collections of information visualizations, see Edward R. Tufte, *The Visual Display of Quantitative Information* (Cheshire: Graphics Press, 2001); Edward R. Tufte, *Envisioning Information* (Cheshire: Graphics Press, 1990); and Anthony Grafton and Daniel Rosenberg, *Cartographies of Time: A History of the Timeline* (Princeton: Princeton Architectural Press, 2010).

10 Noteworthy examples include Stuart K. Card, Jock Mackinlay, and Ben Shneiderman, *Readings in Information Visualization: Using Vision to Think* (San Francisco: Morgan Kaufmann, 1999); Robert Jacobson, ed., *Information Design* (Cambridge: MIT Press, 2000); and Ben Fry, *Visualizing Data: Exploring and Explaining Data with the Processing Environment* (Sebastopol, CA: O'Reilly Media, 2008).

11 For relevant technical precedents to our work, see Casey Reas and Ben Fry, *Processing: A Programming Handbook for Visual Designers and Artists* (Cambridge: MIT Press, 2007); Adam Fouse et al., "ChronoViz: A System for Supporting Navigation of Time-Coded Data," in *Proceedings of the 2011 Annual Conference Extended Abstracts on Human Factors in Computing Systems,* CHI EA '11 (New York: ACM, 2011), 299–304; Colin Ware, *Visual Thinking: For Design* (San Francisco: Morgan Kaufmann, 2008); and Stephen G. Eick, "Engineering Perceptually Effective Visualizations for Abstract Data," in *Scientific Visualization Overviews, Methodologies and Techniques, IEEE Computer Science* (1995): 191–210.

12 Prominent social studies of visualization in scientific and technical practice include: Elspeth Brown, "The Prosthetics of Managements: Motion Study, Photography, and the Industrialized Body in World War I America," in *Artificial Parts, Practical Lives,* Katherine Ott, David Serlin, and Mihn Stephen, eds. (New York: NYU Press, 2002); Peter Galison and Caroline A. Jones, *Picturing Science, Producing Art* (London: Routledge, 1998); and Bruno Latour, "Drawing Things Together," in *Representation in Scientific Practice,* Michael Lynch and Stephen Woolgar, eds. (Cambridge: MIT Press, 1990), 19–68.

the version of the Apollo 11 visualization presented here (see Figure 1), we chose to leave out physical interactions, historical and bibliographical information, and the virtual presence of the designers and developers of technologies involved. (Finding new ways to incorporate these data will be taken up in future visualizations.) Finally, in consideration of the benefits, absences, and limitations, we close with an explanation of necessary steps to make data visualization a viable, widespread format for the study of computer–human relationships across domains.

Approach

Data visualization—the design of digital techniques for imaging information—has only recently been formalized as an independent, interdisciplinary method of research. Today's visualizations build upon a long history of work in the development of scientific graphics and maps.[9] In the past decade, scholars have tried to establish a number of books as foundational texts for the digital turn in visualization.[10] Our approach to data visualization leverages recent technical advances, including better programming tools and a greater understanding of the nuances of time-coding, visual perception, and interpretation.[11]

Moreover, we maintain a critical perspective towards visualization and the data employed therein. We look to existing social studies of visualization to understand the opportunities and pitfalls in using technology to make the invisible world visible for science—both in studies of humans and of the material world.[12] Although visualization has been a subject of social inquiry for many years, it has often been overlooked as a possible tool. Some notable exceptions include network analysis, the use of geographic information systems, and, more recently, custom applications in the digital humanities.[13] However, none of these initiatives specifically frame distributed relationships in computer–human systems.

We are developing a series of operation-specific visualizations as a path toward establishing more general tools and techniques to help others see computer–human relationships in a range of technical domains. Our current focus is primarily on short operations (about ten minutes) to capture moment-by-moment interactions. Our approach represents computer–human relationships as a history of conversations. The voice channel on the Apollo 11 mission, connecting Neil Armstrong and Buzz Aldrin in the lunar module, Michael Collins in the command module, and Charlie Duke in ground control, clearly fits into this category. Less obviously, however, these crews are also communicating with the lunar module itself, with the people who designed its hardware, the programmers who wrote its software, the trainers, the mission planners, and numerous other participants whose judgments have contributed, explicitly or implicitly, to the total

human–machine system. We seek to highlight these hidden conversations, along with the more obvious ones. By clarifying the numerous human relationships embedded in any complex technical system, our approach seeks to enable researchers and designers to better understand how technical and social systems shape one another.

Our tools for data visualization are not merely a means of presentation; they are spaces for analysis and inquiry, objects to think with.[14] Indeed, our approach supports a designerly way of working and relating to information. The designer's approach is one of learning through making.[15] Hence, we are interested in making visualizations in an open-ended manner, rather than in following ready-made templates. Furthermore, in addition to developing our own visualizations, we are working toward an open-source visualization toolkit, one that can extend and improve on existing systems for building custom information visualizations.[16]

The Apollo 11 Landing Visualization

Mindell's book, *Digital Apollo,* narrates the history of the Apollo guidance and control system and its relationships to debates about pilots' roles in aviation in the twentieth century. The last half of the book focuses on the last ten minutes of a lunar landing—the most difficult and dangerous phase of the entire Apollo mission—from a ten-mile high orbit to a safe landing. Using a kind of technical ethnography, the book examines the design of the computers and software that executed the landings. It then analyzes data, transcripts, audio, video, and technical debriefs to create a thick, second-by-second description of the critical final phase of the lunar landing. The book describes not only how it was designed to work in principle, but also how it actually worked out on each of the six lunar landings, highlighting the anomalies, errors, surprises, and creative workarounds. The narrative includes technical debates over circuit design and software executives, traditional human-factors considerations like attention and workload, and also perspectives from the history and sociology of technology, such as tensions about the professional identities of the crew members, national political goals and Cold War agendas, the debates among designers of the system, and the knowledge distribution and power relationships between the crew members and ground controllers.

While *Digital Apollo* presents this narrative in a written form, here we seek to expand the understanding of one event during the Apollo missions by presenting it in a visual, interactive format. This format allows us to display the richness of available evidence over a fixed time period. Since the publication of *Digital Apollo,* the authors have obtained the rope listing (see

13 As an example of such custom applications, see Alice Thudt, Uta Hinrichs, and Sheelagh Carpendale, "The Bohemian Bookshelf: Supporting Serendipitous Book Discoveries Through Information Visualization," in *Proceedings of the 2012 ACM Annual Conference on Human Factors in Computing Systems,* CHI '12 (New York: ACM, 2012), 1461–70; "Mapping Controversies," www.demoscience.org/ (accessed February 4, 2013); and "HyperStudio—Digital Humanities at MIT," http://hyperstudio.mit.edu/ (accessed February 4, 2013).

14 For an in-depth discussion of the productive relationship between making and thinking, see Seymour Papert, *Mindstorms: Children, Computers, and Powerful Ideas* (New York: Basic Books, 1980).

15 The hands-on nature of design thinking is thoroughly explored in Donald A. Schön, *The Reflective Practitioner: How Professionals Think in Action* (New York: Basic Books, 1983).

16 Many others have built visualization toolkits, but for different ends. See Jean-Daniel Fekete, "The InfoVis Toolkit," in *Proceedings of the IEEE Symposium on Information Visualization,* INFOVIS '04 (Washington, DC: IEEE Computer Society, 2004), 167–74; and Jeffrey Heer, Stuart K. Card, and James A. Landay, "Prefuse: A Toolkit for Interactive Information Visualization," in *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems,* CHI '05 (New York: ACM, 2005), 421-30.
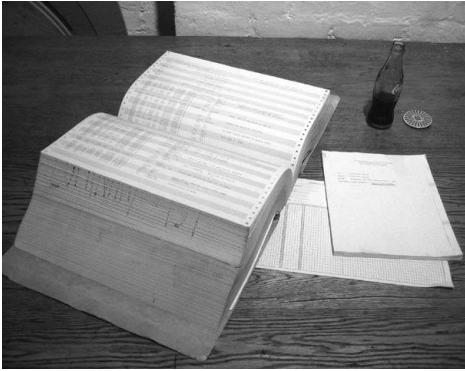
Figure 2 (left)
The rope listing, containing the landing code
for the Apollo 11 lunar module.

Figure 3 (right)
A sample page from the downlink data,
displaying synchronized values from the DSKY.

Figure 2), courtesy of Don Eyles, one of the key programmers of the embedded landing code in the lunar module, as well as the of downlink telemetry. The telemetry data came to us in the form of microfilm copies of old computer printouts (see Figure 3). These printouts contain details on the state of the computer systems every two seconds during the landing.

Individual pages of the downlink data were scanned as image files, and the results were put through an optical character recognition (OCR) program. However, even after enhancing the contrast and sharpness of the scans, the OCR was only about 75% reliable. Each digitized page of the listing had to be compared to the original document to catch remaining errors. The corrected listing data were finally formatted in a series of xml files. Much of the listing data remains unprocessed. Future work might be directed at other parts of the extensive document. We have processed and combined this new data with communication transcripts, post-flight debriefing interviews, and other supplementary sources to create an interactive visualization of the last ten minutes of the Apollo 11 landing.

The figures that follow are data streams excerpted from the whole visualization (shown in Figure 1). The x-axis represents time, starting on the left at mission time 102:15:20, just after the descent orbit insertion (DOI) burn. (Note that mission times are in "hours:minutes:seconds" since launch.) At this moment, the lunar module is in a relatively safe, ten-mile orbit above the surface of the moon.

The figure's vertical axis is divided into multiple sections. At the very top of the visualization, a commentary function (shown by itself in Figure 4) offers reflections on important moments during the landing, simultaneous to the events themselves. These reflections come from debriefing interviews with the
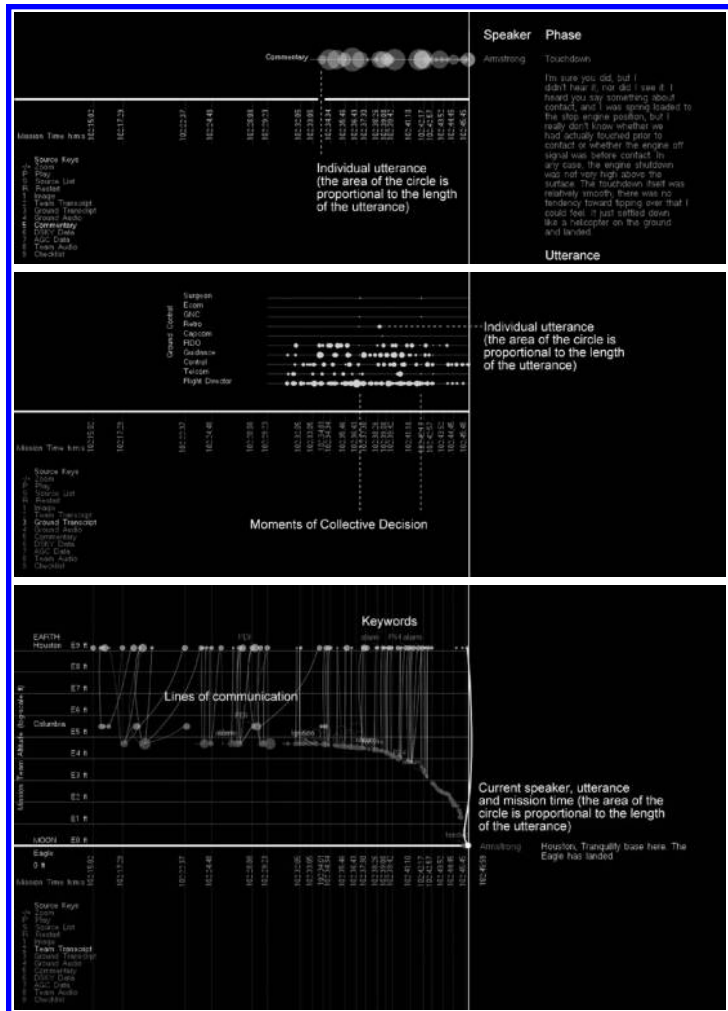
astronauts, as well as from analysis of the events in *Digital Apollo.* Together, these comments help to indicate where points of interest are and possible ways of interpreting the data (see Figure 4).

Just below the commentary, members of mission control in Houston are present, from the mission surgeon (top) to the flight director (bottom); these roles correspond to individuals seated at consoles in the main mission control center. (This section of the vertical axis is shown by itself in Figure 5.) Additional support teams, not represented, communicate with each of these team members from back rooms.

To visually accommodate actors at vastly different spatial distances, we have chosen the engineering convention of the logarithmic scale, which can depict a broad range of values while focusing resolution at the small end. The middle section (shown in Figure 6) is in exponential units of altitude in feet from the moon to the earth (approximately 10^9 feet away). Utterances from the landing CAPCOM (the astronaut on the ground who communicates with the crew in space), the command module pilot, the mission commander, and the lunar module pilot are plotted on this graph.
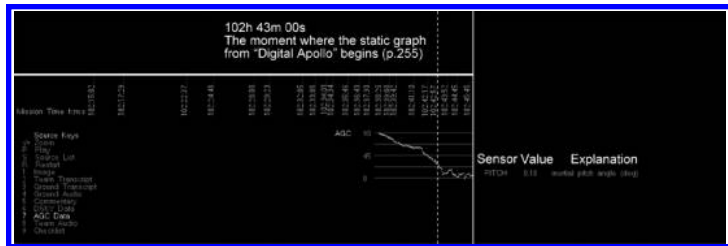
Communications from the command module, *Columbia* (in lunar orbit), are mapped between 10^5 feet and 10^6 feet. Dialog from the lunar module, Eagle, is visible starting at an altitude of 50,000 feet (between 10^4 feet and 10^5 feet on the left of the graph) and ending on the surface of the moon (0 feet on the right).

In Figure 1, this graph is partially overlaid on a composite image of the lunar surface.[17] The two graphics are rotated 90 degrees to one another, with the axis of rotation running through the landing point. The bottom section of the visualization contains data from the Apollo computer systems, the display/keyboard interface to the Apollo computer (DSKY), and the abort guidance computer (AGC). These computer systems put the team members in dialog with an extended network of scientists, designers, and programmers who are not otherwise present. The changing pitch of the vehicle (shown by itself in Figure 7) can be monitored just below the lunar surface landing site image. In sync with the human dialog, the DSKY displays modes and values that represent voices from the now-distant architects of the Apollo control systems (shown in Figure 8).

Each circle on the graph represents an utterance by one member of the team or ground control, with the size of the circle proportional to the length of the utterance. Lines connecting subsequent utterances represent inquiries and responses between team members. Small looped lines around the dialog from the lunar module represent communication within the vehicle between crew members. Specific events are labeled, such as computer program changes and the famous "1201" and "1202" program alarms, which caused some distraction to the crew at critical times.

The Apollo 11 landing visualization was implemented first in Processing, a Java development environment, and later using Open Frameworks, a graphical C++ library. It uses data from multiple files and file types: two separate communication transcripts, two audio files, an image of the lunar surface, and the printed records of two computer systems: the DSKY interface and the AGC. Our application reads these sources into memory as a series of events coordinated through common mission times. We have chosen an event structure to highlight the real-time, interactive
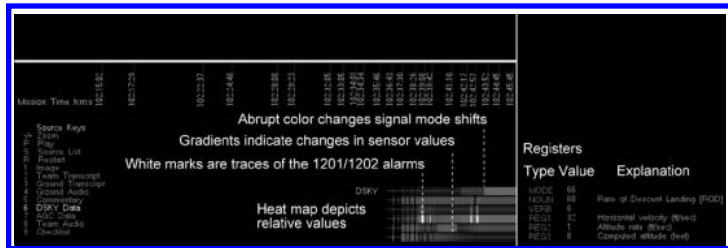
relationships among humans and machines. Our code uses this
structure to display data from each source in time and in relation-
ship to other sources. Events such as an utterance or pitch value
are mapped onto the display to convey their significance as part of
a sequence of operations.

Landing Patterns
The visualization shown in Figure 1 allows us to see aspects of the
Apollo 11 landing that would be difficult to make out otherwise.
As mentioned, the *Digital Apollo* narration of the same events
requires an entire chapter, without presenting much quantitative
data. The visualization displays macro patterns across qualitative
and quantitative data streams. For instance, communications are
visibly sparse for the first phase of the landing, owing to data
dropouts and problems aligning the lunar module's high-gain
antenna. In fact, the intermittent functioning of the communica-
tions channel added a significant source of workload and stress for
the crew. During the writing of *Digital Apollo,* these communica-
tions patterns became discernable only after long hours studying
and writing about the transcripts, whereas in visual form they are
more immediately evident. To some degree, the visualization sum-
marizes the work of the scholarship, which involved bringing into
coordination disparate data sets, including trajectories, communi-
cations, checklists, computer data, and debriefs. The visual presen-
tation also allows for easier tracking of how the workload and
authority were negotiated during the critical final phases of land-
ing, and how that workload was offloaded from the lunar module
to Houston in response to the program alarms.

Our analysis of the landing focuses on three important
phases, all of which are captured in the visualization. The first
phase begins on the far left of the visual, at about 102:33:05, when
powered descent initiation (PDI) is displayed. This event begins
the formal landing sequence, when the computer fires the lunar
module's rocket engine along its vector of travel, slowing it down.
The decrease in velocity reduces the lunar module's trajectory,
causing it to fall out of orbit. Once the PDI burn begins, the vehicle
can either land safely, crash into the moon, or execute a dangerous
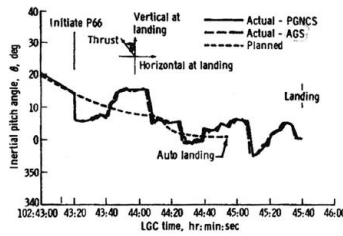and uncertain abort within the next ten minutes. The time from

Figure 9
Pitch data from Apollo 11. This graph corresponds with the data in the visualization, starting at mission time 102:43:00.

PDI to landing is the longest, most difficult, and most critical segment of the mission. On a scale of 1 to 10, Armstrong described it as a difficulty level of 13.

A second phase begins soon after PDI, when the computer begins generating a series of 1201/1202 "program alarms," indicating the computer is having a problem. The first of these occurs at 102:38:30, in the middle of the visual. Communications become intense as the crew decides whether to execute an emergency abort and as ground controllers try to diagnose the problem. Cognitive workload shifts from the astronauts and the onboard computer to the ground controllers ("we're go on that alarm"), who quickly and accurately assess that the computer is overloading in response to a data display process but is only dropping non-essential tasks. In reaction, ground control takes over the display monitoring task (Duke: "Eagle, we'll monitor your Delta H"), thus freeing up critical resources, such as processor cycles and human attention.

Finally, the last phase begins at 102:41:35, when the visualized transcript and the DSKY data both read P64. The computer is switching into a new program (#64), which begins to pitch the vehicle toward vertical so the crew members can look forward out their windows and visually identify their landing site. The visualization displays how the pitch data begin to change, in coordination with the program change and the crew members' verbal recognition of the event.

After the vehicle pitches over at P64, the astronauts' attention focuses on their interaction with the terrain and the computer. Communications with the ground are minimized, culminating in Armstrong decision to turn off automated targeting at P66 and land in semi-automated "attitude hold" mode. In Figure 9, the pitch graph from a contemporary post-flight analysis traces Armstrong's final efforts to navigate the rocky surface of the moon. The descent engine of the lunar module was mounted on a gimbal on the bottom of the craft. Thus, the pitch (rotation around a horizontal axis) controlled the direction of thrust, allowing Armstrong to change the lateral acceleration of the lander by rotating its orientation. A historical account of these dynamics is given in *Digital Apollo,* along with the static graph of the pitch.[18]

Our newly recovered pitch data agree with the graph in *Digital Apollo*—an important indication of the accuracy of the downlink data. However, when presented in the visualization, in sync with the dialog and other indicators, the rotation of the vehicle under Armstrong's command comes to life. As the craft approaches the surface of the moon, Armstrong first pitches back to slow the forward velocity and then forward to make it to the other side of a large crater (not clearly visible in the included moon image). In debriefings after the mission, he noted, "I had tipped it over like a helicopter." The abrupt fluctuations in the pitch data reflect accounts of these last few moments by Armstrong himself,

18   Mindell, *Digital Apollo,* 254.

who described his performance as "a little spastic in final approach," and by observers like Normal Mailer, who described the lunar module as "skittering like a water bug debating which pad it will light on."[19] But none of these accounts conveys the pace and rhythm of the event quite like our visualization, which brings the pitch data into coordination with the transcript. In fact, the pitch contrasts interestingly with the dialog, which conveys a more subdued account of the events. Indeed, most of Armstrong's indecision was expressed through his handling of the lander and not verbally. Only after Armstrong utters his famous "Houston, Tranquility Base here, the Eagle has landed," do expressions of joy and relief flood the transcripts. This contrast between data streams reveals that multiple channels are necessary to give a full account of such events.

Our visual presentation of the data surrounding the first lunar landing creates a rich space for further inquiry, supporting continued conversation and analysis of these three moments and others. Furthermore, the continuous negotiation among humans and computers displayed in this visualization is representative of interactions in any domain of work in which distributed teams search for a safe trajectory in uncertain social, technical, and environmental contexts.

Opportunities and Limitations
Our visualization of the Apollo 11 lunar landing represents a new kind of text for the study of computer–human relationships. Here, we highlight some of the representational characteristics that our approach brings to the analysis of such events. Three characteristics in particular—diverse data sets, a graphical format, and temporal interaction—are the most salient among these characteristics, and they have important implications for how we see real-time interactions. In this section, we explain each of these merits briefly but also discuss what our approach leaves out.

First, our visualizations integrate data sets from many sources in an orderly, comprehensible way. The Apollo 11 visualization depicts simultaneous communications, sensor values, and states of the guidance computer over the course of several minutes. It allows vastly different sources of data to be automatically read, synched, and searched in ways that would be extremely laborious if done manually. For example, every use of the term "alarm" is highlighted, regardless of the data set in which it shows up, whether a communication transcript or a computer register. The wealth of data brought together in the Apollo 11 visualization gives viewers a sense of immersion in the event.

Second, our work transforms the broad range of multi-format data into graphical patterns that can be analyzed visually. This representation involves the flattening of some data—for instance, turning an individual utterance into a circle of proportional size.

19   Norman Mailer, *Of a Fire on the Moon* (Boston: Little Brown & Co., 1970), 377.

Although this flattening produces significant losses, it also creates new opportunities—in this case enabling audiences to see all communications simultaneously. Our graphical format also preserves many of the spatial relationships of the original event, scaled appropriately. However, non-spatial conditions in the original event can just as easily be expressed spatially. The flattening of utterances into geometry is one example. Another example is the way the pace of descent toward the surface of the moon is depicted along the x-axis of the visualization. Here, time is expressed in terms of space; the further to the right an event appears, the later it occurred. The advantage of such transformations is that they do not need to be explained at length. Once the axes are labeled, the graphical relationship between time and space becomes obvious.

Third, our representation of data sets is temporal. The Apollo 11 visualization presents data in real time. The tempo of the interactions during the last few moments of the landing was difficult to communicate in a written format. Incorporating time as a dimension of our representation allows viewers to directly experience the rhythm of computer–human exchanges. Moreover, the interactive features of the visualization allow time to be manipulated as a variable. Viewers can move though the data at their own pace by adjusting the timeline or, in one version of the application, by zooming in to look at a detail. They can also choose which data sets to see in combination using built-in toggles. This self-directed interface to the data can be overwhelming for some, but it can also be empowering if viewers are guided adequately or know what they are looking for.

Even in this rich format, only a portion of the potentially useful data is included. Some of the data we would like to use are simply not available in this historical case. However, we also are still learning to find a balance between presenting all the relevant data and making a focused representation that does not overtax the attention of viewers. We believe that the current visualization is a good balance of information and simplicity. However, we can also examine some of the data that are left aside: physiological, embedded, historical, and narrative.

Physiological data, including gestures, attention, orientation, and even heart rate, convey considerable information about operator identities, social relationships, and even organizational cultures. In the example of the Apollo 11 visualization, we have limited data on the micro-behaviors of Aldrin and Armstrong during the descent. More information about their physical interactions could inform us about non-verbal communications and distractions that made a difference during the intense period of the landing.

In addition, we believe that the developers of participating technologies have an embedded presence during such events. Indeed, these remote contributors can and should be seen in

visualizations. In this case, we have the landing code in the form of the rope listing, but it has yet to be integrated.[20] Human relationships with "automation" are fundamentally human relationships with other humans. We do not see "automation" as an exogenous variable to social interaction, or as "tools" to be optimized from a human-centered point of view. From our perspective, the acts of making things (and systems) are similar to the acts of making human relationships. In the words of sociologist Richard Sennett, "both the difficulties and the possibilities of making things well apply to making human relationships," and "the capacities our bodies have to shape physical things are the same capacities we draw on in social relations."[21] Hence, both the human–human and computer–human interactions can and should be studied with the same methods. For instance, developers' judgments about the aptitudes and limitations of Apollo astronauts—their perceptual skills, their cognitive capacities, and even their ranges of motion—enabled and constrained every computer–human interaction during the first lunar landing.

Moreover, histories such as biographies of humans, development paths of machines, and even the evolution of organizations and political conditions are extremely relevant. These data are available. Much of it is in *Digital Apollo,* but we are still developing methods to make this data graphical in a meaningful way. Such data can be added in small text notes, but at the risk of overloading the current application. Another version of the visualization, not depicted in this article, includes a logarithmic timeline, so that cultural, political, and biographical data, as well as mission events such as training schedules, can be included on long timescales, while still allowing later focus on second-by-second operations.

Finally, the Apollo 11 visualization contains little in the way of a traditional historical narrative. Of course, the design of the visualization itself presents data in a curated way. However, the visualization presents no overt or linear story to support a singular audience interpretation. We believe that visualizations for more general audiences could benefit from directed storylines. Many audiences are not interested in sifting through the data and simply want to get the message. Visualizations need not be only research tools for browsing data; they also can help broader audiences understand complex events.

In future work, we will experiment with alternative means for data visualization that make more room, both visually and conceptually, for physiological, embedded, historical, and narrative elements. For now, we close by reflecting on what our example visualization contributes to current ways of seeing computer–human relationships in technical operations and the unanswered questions that can motivate our next steps.

20 Apollo computer programs were casually referred to as "ropes" because they were eventually hardwired in woven copper for flight. The rope listing is a printed text version of the landing program. See Don Eyles, ''Tales from the Lunar Module Guidance Computer.'' Paper presented at 27th Annual American Astronautical Society Guidance and Control Conference (Breckenridge, CO, February 2004).

21 Richard Sennett, *The Craftsman* (New Haven: Yale University Press, 2009), 290.

Conclusion

The example of the Apollo 11 lunar landing illustrates some of the benefits of our use of data visualization to study computer–human interactions. Our approach incorporates inclusive data sets, a graphical and accessible format, and a temporal and thus immersive experience of the data. Data visualization provides an opportunity to represent complex social and technical data in a format that is accessible to broad audiences. However, the format also presents new constraints. Data must graphically fit into a visualization. In the example provided, physiological, embedded, and historical data have not been fully integrated. Furthermore, although the format presents an opportunity to work outside the narrative structure of writing, some kind of sequential explanation would be useful—perhaps invaluable—for a general audience. With these opportunities and limitations in mind, we are continuing to refine our approach and are moving toward more general systems for visually representing data from computer–human interactions. We are now developing several classes of visualization in parallel, each of which is tuned to the expectations and ends of particular audiences.

To begin with, we must recognize the limits on the complexity of data visualizations, which are related to the choice of audience: Social scientists are willing to put up with a lot of complexity and ambiguity; technologists are more interested in data directed at decision making; general audiences want to get the broad message. In response, we are learning about each these audiences and their needs, from data collection to final presentation.

At any level of engagement, we see visualization as a component of a broader toolkit for explaining computer–human relationships. Even now, we are experimenting with ways in which it can be used to complement other means of data collection (e.g., expert interviews) in aviation, surgery, and undersea archeology. Indeed, our tools are not meant to replace or replicate the full work of anthropologists, sociologists, and historians. Rather, our goal is to help a broad range of audiences develop richer perspectives on human operators and invite deeper joint qualitative and quantitative studies.

In summary, this work responds to a pervasive problem: We cannot easily comprehend the full spectrum of human participation in real-time, technical activities unfolding in complex environments. If researchers, designers, and operators of automated systems, as well as the public, are to understand the human implications of new technologies, they need more inclusive and accessible ways of accounting for the distributed and varied positions of skilled humans in technical operations. New research opportunities in data visualization have relevance for any endeavor where human operators face social and technical challenges in the face of remote presence, simulation, automation, and related technologies.