Connecting Human-Robot Interaction and Data Visualization

Daniel Szafir*
ATLAS Institute and Department of Computer Science
University of Colorado Boulder
Boulder, Colorado, USA
daniel.szafir@colorado.edu

Danielle Albers Szafir*
ATLAS Institute and Department of Computer Science
University of Colorado Boulder
Boulder, Colorado, USA
danielle.szafir@colorado.edu

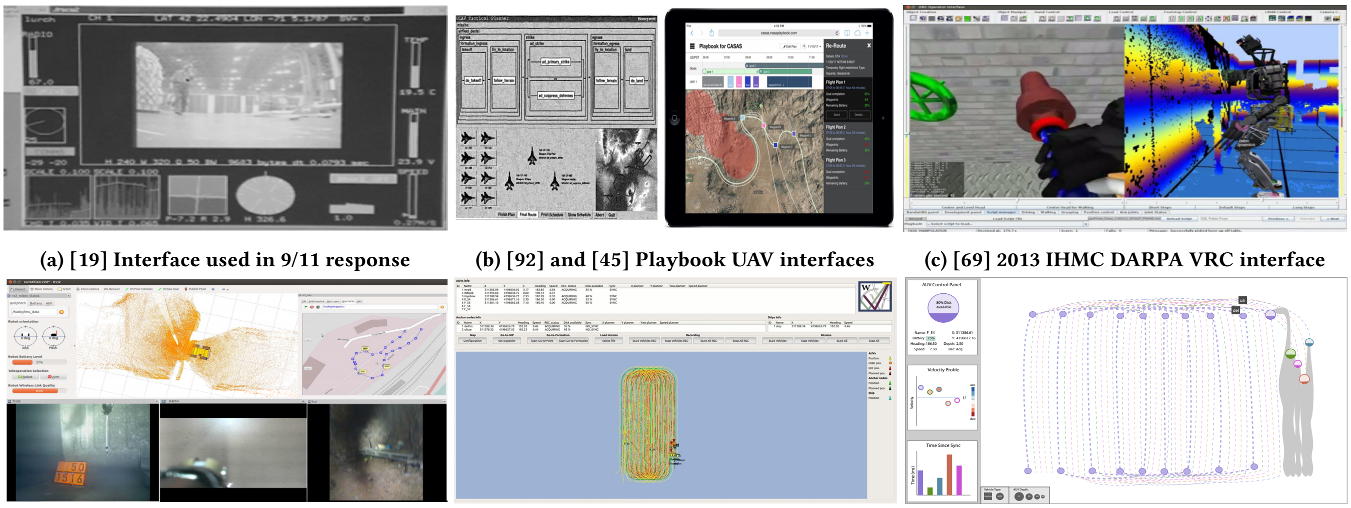

(a) [19] Interface used in 9/11 response **(b) [92] and [45] Playbook UAV interfaces** **(c) [69] 2013 IHMC DARPA VRC interface**

(d) [49] ICARUS C2I interface **(e) [136] WiMUST marine robot interface** **(f) Notional redesign of the WiMUST Interface**

Figure 1: Examples of archetypal robot interfaces are shown above. This paper highlights opportunities for improving robot interface design by integrating knowledge of data visualization and recognizing the importance of data analysis tasks in HRI.

ABSTRACT

Human-robot interaction (HRI) research frequently explores how to design interfaces that enable humans to effectively teleoperate and supervise robots. One of the principle goals of such systems is to support data collection, analysis, and human decision-making, which requires representing robot data in ways that support fast and accurate analyses by humans. However, the interfaces for these systems do not always use best-practice principles for effectively visualizing data. We present a new framework to scaffold reasoning about robot interface design that emphasizes the need to consider data visualization for supporting analysis and decision-making processes, detail several data visualization best-practices relevant to HRI, identify a set of core data tasks that commonly occur in HRI, and highlight several promising opportunities for further synergistic activities at the intersection of these two research areas.

*Both authors contributed equally to this research.

HRI '21, March 8–11, 2021, Boulder, CO, USA

ACM ISBN 978-1-4503-8289-2/21/03…$15.00
https://doi.org/10.1145/3434073.3444683

CCS CONCEPTS

• **Computer systems organization** → **Robotics**; **External interfaces for robotics**; • **Human-centered computing** → **Visualization**; **Graphical user interfaces**.

KEYWORDS

Robot Interface Design; Data Visualization; Human-Robot Interaction (HRI); Information Visualization; VIS; InfoVis

ACM Reference Format:
Daniel Szafir and Danielle Albers Szafir. 2021. Connecting Human-Robot Interaction and Data Visualization. In *Proceedings of the 2021 ACM/IEEE International Conference on Human-Robot Interaction (HRI '21), March 8–11, 2021, Boulder, CO, USA.* ACM, New York, NY, USA, 12 pages. https://doi.org/10.1145/3434073.3444683

1 INTRODUCTION

Robots are increasingly helping humans explore environments, collect data, and manipulate the physical world. Historically, robot deployments have required significant human oversight and direct intervention, leading to human-robot interfaces that focus primarily on supporting human supervision and teleoperation of robot activities. However, advances in sensing, actuation, and autonomy are rapidly expanding robot capabilities and creating new opportunities for scientists, engineers, and analysts to collaboratively engage with robots to accomplish domain-related tasks, such as mapping archaeological sites [27] or inspecting building integrity [95]. In such scenarios, users may be less interested in directing low-level aspects of robot operation and may instead focus on accomplishing the overall mission objective by synthesizing knowledge products created by robots, such as analyzing new data samples the robot provides, investigating maps the robot creates, or working with a robot to understand and reduce contextual uncertainties or improve data quality. As robot capabilities advance, human-robot interfaces will increasingly need to support such data-centric activities.

Reasoning about data-centric human-robot interaction is also critical for for systems that focus on traditional aspects of robot operation and/or supervision due to limitations in robot autonomy (e.g., many of the systems shown in Figure 1) as humans must understand robot data, often provided as camera feed(s), map overlays, sensor readouts, point clouds, or even mixed and virtual reality displays, to build situational awareness sufficient for directing robots effectively. However, designing robot interfaces that support users in both directing robots and understanding robot-collected data remains a challenge. For instance, Murphy & Tadokoro [94] note that "the End-User interface is the most difficult to build because it requires a working prototype of the robot, an initial interface, and access to high fidelity field conditions, and multiple end users for a domain analysis." We argue that another major reason that such interfaces have historically been, and remain, so difficult to design is that interface researchers and developers are often members of the human-robot interaction (HRI) and field robotics communities who may have little connection with the data and information visualization communities (VIS), a field of research exploring guidelines for helping people quickly and accurately make sense of data. While some HRI practitioners may have training in visualization techniques, we believe there is a general opportunity to substantially improve robot interface design by leveraging knowledge of how to effectively encode data for human users.

In this work, we argue for increased collaboration between experts in robotics and data visualization, motivated by the need for robotic interfaces to increasingly consider and provide direct support for data analysis. Our goal is to begin to bridge the divide between HRI and VIS, recognizing that **there is a considerable potential for natural synergies between HRI, which develops interfaces that enable humans to effectively direct and/or supervise robots while also making use of the data such robots collect, and information visualization, which focuses on designing interfaces that allow people to build targeted knowledge through data.** We argue that robot interfaces need to increasingly consider how users will engage with data provided by robots and thus be designed not only from a robotics-centric, but also a data-centric perspective. In addition, we note that robotics offers rich opportunities for novel VIS research, particularly in identifying practices for dynamic, uncertain, and spatio-temporal data. Capitalizing on such opportunities will require co-innovation and knowledge of the state-of-the-art in both fields. To this end, we briefly review major guidelines and best-practices from the VIS community that are relevant to HRI and propose an initial framework to guide potential parallel developments in HRI and VIS, understand HRI data tasks and activities, and identify potential connections and opportunities for future work in this space of data-centric HRI.

2 BACKGROUND

This paper is written from the perspective of two authors, one from the HRI community and the other from VIS, to help both communities recognize the opportunities each field has to inform the other. We believe there are clear opportunities for mutual benefit; however, while the two fields draw on common intellectual traditions (e.g., human-computer interaction, cognitive science, etc.), to date HRI and VIS have developed largely in isolation from one another. For example, we searched the entire corpus of the IEEE VIS conference from 1990 to 2019 using the term "robot" and found only one paper (by one of the co-authors) exploring analytic systems where users might work collaboratively with robots that are supplying field data [148]. However, robotics was not the primary focus of the paper and no robot was used in the research.

We similarly surveyed IEEE Transactions on Visualization and Computer Graphics (TVCG), the premier journal for visualization research (as well research from other communities such as virtual reality). We identified only three relevant papers: one paper (from more than 20 years ago) on visualizing robot sensors to help developers understand how the sensors work and improve object identification algorithms [140] and two papers combining augmented reality and aerial robots [40, 160]. However, the latter two papers did not originate from, or take advantage of knowledge within, the data visualization community, but rather came from the adjacent mixed and virtual reality communities, although they further demonstrate potential opportunities for better synergies between HRI and related fields (see similarities to recent HRI research in robotics and mixed reality, e.g., [57, 109, 114, 115, 146, 147]).

This survey, paired with the expertise of the authors, indicates that robot data has not been actively investigated in VIS despite research within VIS on developing tools for a range of other domain-focused problems, such as biology, machine learning, and environmental science. But what about the other perspective: is the HRI community well versed in VIS research and best practices? We surveyed all publications from the IEEE/ACM International Conference on Human-Robot Interaction (HRI) using the keyword "visualization," which resulted in 262 papers, many of which focused on developing robotic interfaces that include various aspects of data visualization (e.g., maps, video overlays, sensor displays, etc.). However, we could not identify a single paper that cited any work from the VIS community, meaning that such interfaces may have been designed without considering state-of-the-art principles for effective data visualization. Likewise, we were only able to identify one relevant paper [63] and one Late-Breaking Report (LBR) [118] from HRI that cited visualization work from TVCG.

While a lack of cross-disciplinary citations is not necessarily proof of a lack of knowledge transfer between HRI and VIS (certain principles may be considered "foundational" and thus have no need for direct citations), based on the authors' experiences and our survey of VIS, TVCG, and HRI, we believe that HRI-VIS collaborations are rare and opportunities exist for greater ideological interchange. This paper highlights such opportunities in synergies across HRI and VIS. It is primarily framed for an HRI audience, focusing on how knowledge from VIS may support robot interface design, although §6 outlines how HRI offers potential for innovation in VIS.

2.1 Scope

Our objective is to explore how research in visualization may improve interface design for robotics, recognizing that one of the fundamental activities for human-robot teams is to collect and analyze data. We specifically focus on human-robot teams working collaboratively towards a predefined objective, as in search-and-rescue, collecting data at field sites, telerobotics for space, marine, or terrestrial applications, etc., where some form of technology serves a mediating role in the interaction (i.e., there is an explicit visual interface through which data can be conveyed). Such scenarios may involve collocated (e.g., an emergency responder in the field interacting with a nearby drone through a smartphone or tablet interface) or remote interaction (e.g., an astronaut on board a space station working with a robot deployed on the lunar surface via a traditional laptop or more advanced virtual or mixed reality interface). We do not consider purely social human-robot interactions, interactions that lack any overarching team objective(s), or interactions where there is no mediating technology that could support visual data communication (e.g., a collocated human-robot interaction with exclusively verbal or gestural interaction).

2.2 Human-Robot Interface Design

A full survey of human-robot interfaces is beyond the scope of this work (relevant surveys and design guidelines from robotics can be found in [2, 22, 34, 73, 94, 154, 156] and metrics in [32, 139]). However, we briefly review several major trends in robot interface design to ground our analysis. Existing interfaces primarily support user situational awareness (SA) and user control. HRI has extensively explored various aspects and levels of SA, including categories of SA relevant to human-robot interfaces (e.g., human-robot awareness, robot-human awareness, location awareness, activity awareness, surroundings awareness, status awareness, overall mission awareness, etc.) [36–38] and interactions between SA and robot autonomy levels [3]. HRI has similarly explored control paradigms across the spectrum of teleoperation and supervision (e.g., shared control [14, 35], collaborative control [42], delegation schemes [142], and more exotic systems such as the “adverb pallette” [133], etc.).

Robot interface design frequently investigates how to provide users with information from both robot camera feeds and maps derived through low-level sensors and/or higher-level perception systems. For instance, several early projects examined the relative usefulness of map and video data [98] and explored how both data types might be fused into a single overlay display [15, 29, 99, 154]. Systems continue to adopt this paradigm while leveraging modern graphics capabilities (e.g., [41, 116]). Keyes et al. [78] provides a review of both map-centric and video-centric robot interfaces as part of an iterative interface design for remote robot teleoperation, borrowing not only from certain robot interface design guidelines available at the time [129, 153], but also general user interface design heuristics from human-computer interaction (HCI) [100]. More recently, Murphy & Tadokoro [94] enumerate overlaps between general HCI principles and robotic interface design, but note that HCI principles alone are insufficient for robotics.

Other major aspects in robot interface design research include determining what sensor information may be useful, developing predictive control systems (particularly for high-latency operations [9]), and enabling a cohesive workflow across planning, execution, and live plan adjustments/re-planning [45, 91, 92, 142] (Fig. 1b illustrates this evolution). Several guidelines for robot interface design have been proposed, often based on reflections from robot competitions (e.g., DARPA challenges, Fig. 1c) [71, 103, 153, 154] or real field deployments, such as in disaster response efforts for 9/11 [19] (Fig. 1a) or the Fukushima-Daiichi disaster [76]), with recent work [94] suggesting 32 different guidelines for field robot interfaces (we review the intersection between some of these guidelines and VIS best practices throughout this work). Such guidelines, particularly when combined with paradigms such as coactive design [69] that consider the interrelationships between humans and robots, may help HRI researchers and practitioners reason about various user concerns, such as SA. However, little work has considered principled methods for *how* these interfaces may visualize data to specifically help user analysis and decision-making. Findings and practices from VIS may help address this gap.

2.3 Information Visualization as a Discipline

Visualization research offers principles and design methods for creating interfaces that help people effectively reason with data. To help understand when and how information visualization principles may inform robotic interfaces, we aim to establish a common lexicon and way of structuring problems to translate practices between disciplines (see §3). We start by discussing relevant definitions, knowledge, and guidelines from VIS.

Data analysis is the process of exploring and making sense of data, where data consists of measurable artifacts. This data may be qualitative (e.g., the type of task the robot is performing) or quantitative (e.g., the current battery life of the robot or a sensor measurement). People engage in data analysis to develop *insights*: specific, meaningful conclusions drawn from data [102]. For example, a user may detect a potential survivor using hot spots in a thermal scan or determine that a wall is structurally sound based on visual inspection and sensor measures. Insights collectively help analysts use the data to expand their knowledge of a problem or situation and to inform decisions and actions.

Visualizations are designed to help people develop insights from data. In contrast with fully automated analyses, such as database queries or machine learning, visualization supports users in analyzing data across a variety of questions using a single representation [17]. By offering users the agency to fluidly explore data to answer their questions, insights may “snowball,” allowing users to change questions on the fly. Users can interact with visualizations to reveal new information or focus on different patterns as insights develop.

Visualization offers a valuable tool for HRI as the agency, flexibility, and control over the data analysis process provided by visualizations may support users in developing situational awareness, adjusting operations in dynamic environments, and rapidly and intuitively assessing mission state across multiple factors and measures. However, for visualizations to be effective, they must consider the *tasks*—the specific information people look to draw from data, e.g., finding relevant data values, estimating statistical quantities, or comparing patterns across data sources—that users want to accomplish with their data (see [5, 130] for surveys). Note that the VIS definition of *task* (referred to here as *data tasks* for clarity) differs substantially from how HRI commonly defines *tasks*.

A visualization's design determines the data tasks it best supports. For example, while line graphs convey trends, people are five times more likely to focus on quantities in bar charts [157]; heatmaps efficiently convey summary statistics, while line graphs support value estimation [4]. Tasks can inform designs that optimize for target applications [6]. While robots are commonly used to collect data, the corresponding data tasks are seldom directly enumerated. As a result, robotic interfaces often use visualizations with suboptimal designs. Common examples include:

(1) *Overview First:* Nearly all analyses require first understanding the overall picture provided by the data before gathering specific details [134]. However, many robotic interfaces show all information in the highest available level of detail first, potentially overwhelming the user (Fig. 1e). Others focus only on the immediately available details, removing context for those details (Fig. 1a). Effective interfaces should minimize unnecessary detail while retaining sufficient context, providing additional detail only when requested by the user.
(2) *Color Choices:* Robotic interfaces commonly use rainbow (Fig. 1c) or red-green color schemes (Figure 1e) (e.g., every system surveyed in [103] used one of these two maps). Rainbows are ineffective for many reasons, including inaccurate value estimation, difficulty directing attention to specific information, and artificial "bands" falsely grouping data [12, 143], causing potential inefficiencies and false conclusions that can mislead developers and other users. For example, rainbows can decrease ROI detection by 30% [10]. Red-green schemes more precisely represent values but are inaccessible for colorblind users (8% of people [151]). Further, green-good/red-bad mappings recommended by existing guidelines [94] do not hold for all cultures [68]. Visualization tools (e.g., ColorBrewer [54], Colorgorical [50], Color Crafter [137]), and guidelines [135] (e.g., use lightness to encode numeric values [119]) can readily guide improved color choices in robotic interfaces in as little as one line of code.
(3) *Comparing Data:* Existing guidelines encourage interfaces that tile a set of visualizations that each answer a single question [94]; however, too many *juxtaposed* visualizations may make it harder to reason across data [66]. Interfaces may alternatively choose to superimpose datasets (i.e., overlay data on a common space as in [99]) or explicitly encode (i.e., compute and visualize) relevant relationships between datasets [48]. Robotic interfaces have experimented with different forms of comparison (Fig. 1a, 1b, 1d). Prior studies of teleoperation interfaces suggest that the right design depends on the goals of the interaction [15, 99, 154]. Best practices in composite visualization [67], dashboard design [126], and visual comparison [47] may help illuminate trade-offs between designs based on the available data and users' goals.

While visualization research offers guidelines for how to support data tasks in isolation, data in HRI typically informs a variety of both data and robotic tasks. These tasks may interact with each other in complex ways and may also interact with external aspects, such as domain standards or aspects of the data distributions [80]. Visualization offers a suite of methods for designing data visualizations that support key tasks and integrate contextual knowledge about a user's data, constraints, and goals [6, 132]. Integrating data visualization principles into robot interfaces through such methods can improve HRI by better facilitating *sensemaking*.

Sensemaking, a key principle within VIS, defines the process of how humans work with information, including that extracted directly from data and relevant context or expert knowledge, to generate conclusions or actions [110, 124]. Sensemaking asserts that we can use information to reason about the current state of the world and use that reasoning to build awareness and inform action. One common outcome of sensemaking is *decision making*, which occurs when people must choose between a set of options (e.g., determine which building to investigate) or courses of action (e.g., the best route to explore a building). When people use visualizations to engage in decision making, they use patterns in data to form knowledge that provides holistic context to their decision. For example, they may reason about where they *predict* fire will spread to choose how to deploy limited resources, how *certain* they are that a hotspot on a map represents a person rather than a sensor error, or how much *risk* there is to a particular structure [106]. Some decisions use only the information presented in the visualization (i.e., they are exclusively based on the data), while others require integrating data with expertise or contextual information to reason across a broad body of factors.

We believe that the design of robot interfaces can be improved by explicitly considering and designing for sensemaking and decision making processes, including determining the information and data tasks necessary for accomplishing a given objective. To help interface designers better reason about the use and presentation of data in robotics, we present a framework that may help HRI researchers take this more data-centric view while retaining traditional considerations regarding robot control and supervision.

3 TOWARDS DATA-CENTRIC HRI

We propose a framework that emphasizes the data-oriented processes within human-robot interactions to help researchers reason about the design of new HRI interfaces and inspire deeper collaborations between visualization and robotics. Our framework is structured around data flow among human(s) and robot(s) as it relates to accomplishing a shared core **Objective**. Each individual team member carries out various **Activities**, formed through sequences of **Actions**, in service to the **Objective**. We visually illustrate our framework for team data flow in Fig. 2 and detail each framework component in Table 1.

In this framework, data can flow between humans and robots in two main ways: (1) at the **Action** level, where a human might query a robot regarding a particular sensor reading/group of readings or direct the robot to perform a **Robot Action** (e.g., explore a certain region) or a robot might query a human to take advantage of human perceptual or decision-making processes (e.g., is it safe for the robot to move forward) as in collaborative control paradigms [42], and (2) from the robot **Autonomous Processes** to the robot interface for use in the human **Data Analysis Process**. HRI has traditionally focused on the first type of information flow (supporting humans in directing **Robot Actions**). Generating data for the second type of flow has been a major focus of traditional robotics (e.g., improving SLAM, supporting autonomous reasoning, etc.), although

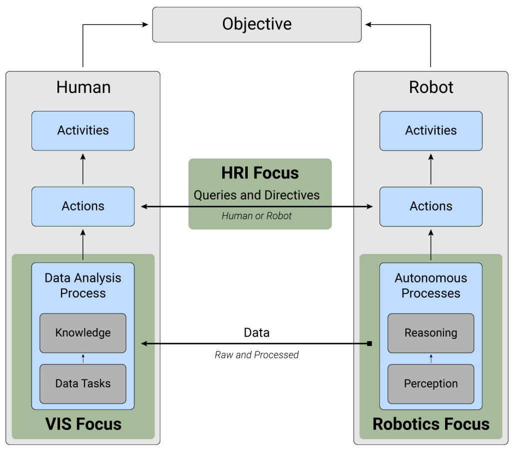

Figure 2: We present a framework based on human-robot data flow that visualizes the traditional focii of HRI, robotics, and VIS research and highlights Data Analysis Processes as a critical consideration for robot interface design.

such research is often motivated from the perspective of improving robot autonomy rather than improving human-robot teamwork. Our framework highlights that such data will directly feed into the human's **Data Analysis Process** (i.e., sensemaking, a traditional focus of VIS) to generate new knowledge and insights that inform more effective **Human Actions** towards the team **Objective**.

Our framework bridges perspectives in HRI, robotics, and VIS to understand how crossovers between the fields can inform systems that allow robots and humans to collaboratively achieve a given objective. Our goal is for this framework to complement (not replace), other models from HRI (e.g., coactive design [69] considerations, the GEDIS framework for evaluating robot interfaces [111, 159], video game-based frameworks for characterizing interaction design [117], etc.) and related areas within HCI and cognitive science (e.g., distributed [60] and situated cognition [23], the human action cycle [101], etc.). Each of these models views interactions at different levels of abstraction; our framework aims to specifically highlight the role of the **Data Analysis Process** in HRI.

4 HRI DATA TASKS

A core component of our framework is recognizing that users regularly engage with various **Data Tasks** to build knowledge through the **Data Analysis Process** to inform human **Actions**. To identify these **Data Tasks**, we surveyed papers across HRI and field robotics, covering a diverse set of domains including search-and-rescue [8, 34, 49, 73, 75, 150], emergency/disaster response [19, 95, 131], terrestrial [77, 94, 99, 109, 114, 115, 154], marine [13, 27, 65, 136], and space [18, 44, 81, 90, 97] exploration/search/environmental data collection, DARPA robotics challenges [69, 71, 88, 103, 104], health systems [28], unmanned aerial systems [21, 29, 45, 70, 91, 159], agricultural robotics [1], and large robot teams [123]. We analyzed each paper using our framework, decomposing the stated human-robot interactions into the various framework components to understand what data insights and knowledge users would need to accomplish their intended objectives (regardless of whether the interface was explicitly designed to support such analysis).

This process revealed seven **Data Tasks** commonly conducted using existing interfaces (c.f., Fig. 3). These tasks provide a direct bridge to relevant best practices, techniques, and findings from VIS that may inform more effective data-centric interfaces. While not exhaustive, this list reflects common themes we observed across HRI scenarios that we can connect to techniques and practices in VIS to inform future interfaces. We use Fig. 1f, a notional redesign of Fig. 1e created using our framework, as a running example of how these tasks can inform interface design in a rich and complex HRI scenario (see Appendix A for details on this redesign).

Each of these tasks requires that users estimate specific statistics or data values of interest. While we do not explicitly discuss statistical estimation tasks below, they are a universal consideration. Visualization offers quantified design insight into how to best encode data for tasks such as estimating individual values [24], assessing trends [30], approximating averages or variance [4], and inferring correlation [53]. For example, people are 25% more accurate at averaging values in heat maps than in line graphs, but 45% more accurate in identifying maxima using line graphs [4]. Such component statistics often form the basis for more complex tasks, including the seven discussed below. However, we caution that VIS is a rich, multifaceted design problem, much as HRI itself. As

Table 1: Details of each framework component depicted in Figure 2

Objective: a set of common goals, contexts, and criteria for success determined by the domain and known by the team in advance			
Agent		**Human(s)**	**Robot(s)**
Activities		Sequence of **Human Actions** taken in service to the larger **Objective** (e.g., searching a particular set of coordinates within an overall mission framework or monitoring robot health) as informed by individual human sub-goals and user roles.	Sequence of linked **Robot Actions** taken in service to the larger **Objective** (e.g., a robot performing a search activity or an inspection activity).
Actions		Set of specific acts and decisions, informed by the **Data Analysis Process**, that a human performs, which may involve the robot directly (e.g., tasking the robot to collect data or manipulate an object), indirectly (e.g., deciding whether to investigate an area further or move on), or not at all (e.g., radioing to another human in the field to communicate an insight).	Set of specific acts, informed by the robot's **Autonomous Processes**, that a robot performs (e.g., collecting a soil sample or querying a human to determine if it is safe to proceed); in the traditional robotics sense/plan/act paradigm (i.e., perception/cognition/actuation), this represents act.
Data Analysis Process The steps a human collaborator takes to make sense of robot data, which may be about the robot (e.g., robot battery level) or collected by the robot (e.g., environmental readings). In essence, this represents the data sensemaking process.	**Knowledge Formation**	Synthesizing patterns and statistics from data into insights that expand the human's knowledge and understanding of the **Objective** and drive **Actions.**	
	Data Tasks	Foraging for relevant information in the data to answer questions about the **Objective**, which can drive **Knowledge Formation.**	
Autonomous Processes The set of low-level robot functional primitives that enable more complex robot actions.	**Reasoning**		The robot's inference and belief estimation capabilities (i.e., planning/cognition sub-systems).
	Perception		The robot's sensing, localization, mapping, and object detection capabilities (i.e., sense sub-systems).

visualizations are often used because users need to achieve multiple tasks at once, empirical VIS results should scaffold reasoning about design/task trade-offs, rather than provide algorithmic guidance.

Find Relevant Information: Users must efficiently detect key information about a given **Objective**. For example, they may need to locate a door handle to turn it [71] or determine areas of safe air quality. Interfaces can draw users' attention to potentially relevant information by manipulating the *salience* of that data [16, 56, 89]. For instance, an interface may make incoming data more opaque and stale data more transparent. Interfaces can also make critical data "pop-out." However, providing too much information can make key data harder to find, a phenomena known as *visual clutter* [122], which visualization techniques like aggregation or filtering can address [39]. For example, drawing trajectories in multirobot systems can make it hard to assess any individual trajectory (Fig. 1e). Techniques like *edge bundling* [61] may simplify trajectory collections to emphasize patterns across robots. Interfaces can then provide precise information about any single trajectory on demand. Alternatively, if **Human Actions** require assessing individual trajectories, managing the salience of key trajectories through bolding or related techniques and using encodings that readily distinguish robots, such as differently nameable colors [50], can help manage clutter and make focusing on individual robots easier (Fig. 1f).

Synthesize Data Across Sources: Robot data is frequently *multi-dimensional*: it contains many variables often from many sources (e.g., a sensor suite or measures of robot state) or even within any single source (e.g., position, color, and time in camera feeds). Visualizations can help people combine data across sources to generate insights, providing context for holistic decision making and allowing users to rapidly answer complex questions. Interfaces can support this combination by, for example, explicitly allowing users to compare data using a suite of visual comparison techniques, including *juxtaposition* (putting data side-by-side; Fig. 1a, 1c, 1d, & 1f), *superposition* (layering data on a common set of axes; Fig.1b), or explicitly computing and visualizing key relationships between data sources. These methods of supporting comparison across variables offer trade-offs in precision, clarity, ease of use, and other factors that interface designers can select between based on the users' needs (see [47, 48, 66] for discussions).

Most robot interfaces and interface design guidelines encourage juxtaposed visualizations [94]. While overreliance on juxtaposed views can introduce clutter and inhibit comparisons by pushing charts further apart [149], dashboard design practices from VIS may inform effective multiview interfaces [126]. Such interfaces can follow a suite of best practices to support people in readily connecting related information across displays, such as using consistent scales, visual channels (e.g., size, position, color), and mappings and not duplicating encodings across unrelated data [113].

Develop and Maintain Awareness: The human visual system allows people to make sense of complex visual information at a glance [43]. Robotic interfaces can use these capabilities to help users develop and maintain situational awareness (SA) through global views of mission data, including data about the state of the environment [99], the state of the robot [19], and robot capabilities [57]. While tables (Fig. 1e) require actively reading and comparing information, visualizations can rapidly summarize key relationships even without active attention [11, 144]. However, interfaces can inhibit this awareness by providing too much detail (introducing clutter), by decontextualizing information (e.g., zooming into the current operational state while providing no context for how the current data fits into the broader environment [72]), or by using ineffective cues (e.g., relying on text at the periphery of the display). Visualization offers a suite of techniques for providing *overviews* that summarize key information in large and complex datasets and that contextualize relevant information in data about the overall mission [96, 127]. Overviews do not often show all available information at once: effective summaries distill information into concise representations that help users develop a sense of the state of an operation and detect locations of interest to examine in detail on-demand. For example, our redesign (Fig. 1f left panel) allows operators to maintain awareness of the state of the AUV formation by visually summarizing critical aspects of motion profiles and resources (e.g., battery, disk, synchronization), revealing specific details about target robots (i.e., robot F_54, purple) on-demand.

Relevant visualization techniques include *overview+detail*, where one visualization provides a concise global overview and another shows details about the active environment, and *focus+context*, where details are shown in the context of an overview [25]. For example, C2I [8] (Fig. 1d) efficiently manages complexity using overview+detail visualizations coupled with detail-on-demand interactions, adding mission details during planning through pop-ups. *Minimaps* frequently provide overview+detail in robotic interfaces (Fig. 1d); however, such maps are often as large or larger than the detail view, making it difficult for users to know where to attend. HRI interfaces seldom provide global overviews of nonspatial data, such as sensor readings. A few interfaces surrounded a camera feed with directional data to provide focus+context representations (e.g., [108]); however, extending these principles to other datatypes (e.g., *periphery plots* [93]) may further enhance SA using abstract data.

Monitor Data Quality: Robots frequently collect data in locations that may be unsafe or unreachable for humans, such as damaged buildings [95], radioactive sites [76], and space [18, 44, 81, 90, 97]. Such locations often correlate with environments where errors in data collection are common. For example, a sensor may become miscalibrated [19], images degraded [22], or wireless reception lost [103]. Data visualizations can help users rapidly identify data quality errors by making data visible as it is collected [148]. For example, *heatmaps* can show where data has or has not been collected to assess coverage (e.g., Fig. 1f, grey regions on the right show completed coverage). Comparing data across *juxtaposed line graphs* can reveal calibration issues between sensors. While analyzing data quality is an open challenge in VIS [87], prior results offer a wealth of techniques for representing uncertain [74, 107, 125] or even missing data [138] that may guide robot interfaces in better informing user actions and decisions. We can also design complementary views that allow users to accommodate variations in quality across data sources (e.g., supplementing low-quality video with sensor readings or robot confidence in obstacle detection [22] or peripherally monitoring data updates as in the bar chart of sync times in Fig. 1f).

Identify Anomalies: Anomalies in robot operation (e.g., high frequency control signals that can cause mechanical failure [71]) or collected data (e.g., spikes in temperature readings) are often events of interest for users. Sometimes these anomalies can be automatically detected using predetermined thresholds; however, failures are often difficult or impossible to detect autonomously [103]. Visualization tools can be designed to enable people to rapidly detect outliers [30] or relevant differences in patterns [31]. Further, by comparing across visualizations (e.g., whether different sensors show correlated changes) or to common frames of reference (e.g., threshold bounds drawn on line graphs or expected bounds on movement [71]), users may more rapidly identify and characterize anomalies and use these observations to drive appropriate action. Interfaces can accommodate anomaly detection using designs that preserve *data provenance* (historical patterns) through techniques like *data sedimentation* [64], that support comparison against known thresholds (e.g., reference lines showing acceptable upper and lower bounds) or that emphasize relevant patterns (e.g., encoding all data samples as lines when noise matters or using smoothed lines or color when mean performance matters [4]).

Make Predictions: Robotic interfaces often support users in making predictions, such as determining whether a robot can safely traverse a doorway [70] or estimating the current mission state when signal is lost or delayed [9, 79]. Prediction tasks are most obvious in control and supervisory **Activities** where they directly guide operator decisions; however, prediction can also guide more collaborative **Activities**, such as estimating a wildland fire's spread from data collected by the robot to adjust operational plans [33]. While HRI has designed predictive interfaces (e.g., [120, 146, 147]), integrating data into predictive reasoning must involve effective visualization of data both over time and with uncertainty. Innovations in uncertainty visualization for temporal geospatial data, such as hurricanes [86], may inform predictive interface design. For example, representing potential motion trajectories using a *bounding contour* causes people to overestimate the likelihood of trajectories at the contour's edges, whereas showing all possible trajectories shifts predictions towards the modal trajectory [106].

Assess Risks: Robots are often deployed in mission-critical domains and pose risks to the success of the operation and/or damage to the environment, robot itself, or collocated people [94]. Data collected by robots may also help evaluate other risks relevant to the **Objective** (e.g., building inspection data might inform users about structural integrity [95]). To assess risk, users combine data about current robot and environment state with human knowledge about potential consequences of actions to estimate an internal cost function for different outcomes. For example, Fig. 1f shows available disk (% circle fill) in the context of the remaining planned path (dotted lines) to help operators anticipate if the system has sufficient resources to complete data collection.

The ways that interfaces present data influence user perceptions of risk [106]. For example, people often value salient data more highly than less visible data when assessing risk. Robot interfaces should consider how to effectively direct user attention towards the most relevant factors for estimating risk. Roldan et al. [120] draw attention to predicted risk factors in data using a "spotlight" and indicate robots at risk using smoke. While increasing salience may improve risk assessment, smoke and other overlays may occlude important details and make target objects less salient [89].

Data Task	VIS Design Concepts	Navigation	Exploration	Manipulation	Inspection & Search	Debugging
Find Relevant Information *Readily direct attention towards data that matters for the current activity*	Data Salience Clutter	●	●	●	●	●
Synthesize Data Across Sources *Compare and fuse different pieces of data into a larger insight*	Comparison Multiple Views	●	●	●	●	●
Develop and Maintain Awareness *Foster and revise situational awareness through data*	Overview+Detail Focus+Context	●	●		●	●
Monitor Data Quality *Assess the coverage and validity of data*	Missing Data Uncertainty		●		●	
Identify Anomalies *Determine unusual or unexpected patterns in data*	Data Provenance Statistical Estimation		●	●		●
Make Predictions *Use data to estimate future states of robots, environment, and mission*	Uncertainty Temporal Data	●	●	●	●	●
Assess Risks *Estimate the costs of different courses of action*	Direct Attention Value Estimation	●	●	●		

Figure 3: We identify seven common Data Tasks employed in HRI. These tasks help connect VIS principles for effective interface design to a variety of Human and Robot Activities.

5 HRI ACTIVITIES

We can use these tasks to help characterize major domain applications (i.e., the **Objectives** and Human/Robot **Activities** in our framework) within HRI. We briefly discuss four such **Activities** in the context of our framework and common related **Data Tasks**:

Environment Navigation and Exploration: Almost all robot deployments involve elements of environment navigation (directing a robot to some known goal) and exploration (investigating an area to find a goal or to survey a region). In navigation **Activities**, a primary **Data Task** is to *find relevant information*, where interfaces should help users by indicating known targets or factors impacting the robot's abilities to successfully reach a target. Users must *synthesize data across sources* to understand the relation between the robot's planned and current trajectory (Fig. 1b) and to successfully navigate the environment (e.g., combining RGB camera and point cloud data to assess the 3D geometry of the environment; Fig. 1d). Interfaces should allow users to *develop and maintain awareness* of the operational environment to respond to potential changes in the environment, *make predictions* about the robot's current path and state (e.g., does the robot have sufficient power to reach the target?), and *assess risks* in control decisions (e.g., is cutting through the canyon worth the risk of the UAV crashing?). Wayfinding [62] may be of particular relevance to such navigation activities.

Exploration builds on these tasks to help the human-robot team successfully reconnoitre an environment. However, analysts must also *monitor data quality* to understand how much of the target environment has been successfully explored and *identify anomalies* to detect areas for further investigation (e.g., gaining more information about unexpected temperature readings).

Manipulation: One of the key advantages robots offer as physically situated agents is the ability to manipulate the physical environment. While some aspects of manipulation may be automated, manipulation **Activities** that involve human guidance or direction may particularly rely on the **Data Task** of *finding relevant information* to direct the user's attention towards key objects or mechanisms while minimizing distracting information that may complicate assessment or distract the user. For example, colorful depth visualizations applied to background objects provided in many interfaces (e.g., Fig. 1c) may distract or impede assessments of a foreground object [103]. During manipulations, users must *synthesize data across sources* to understand the current state of the robot and manipulated object, *identify anomalies* that may indicate failure patterns (e.g., high-frequency signals [71]), *make predictions* as to what commands are most likely to result in successful execution, and *assess risks* of potential failure modes (e.g., what are the risks of applying too much or too little force to an object [103]).

Inspection and Search: Robots can assist humans by investigating and collecting data on environments as well as objects, structures, and people within environments (e.g., inspecting buildings after earthquakes [95], conducting geological [65] or archaeological [27] surveys, or assisting with search-and-rescue emergency response [34]). Such activities may involve aspects of environment navigation, exploration, and manipulation, as described above, but represent a fundamentally different type of **Activity** due to the focus on collecting contextual data (rather than primarily spatial data, as in environment exploration), usually about specific targets of interest (e.g., buildings, ruins, geologic points of interest, survivors, etc.). In such activities, users must be able to *find relevant data* about the focus of their inspection (e.g., structural elements of a building) and *synthesize data* across both contextual (e.g., temperature or force sensors) and spatial sources to, for example, decide whether a wall is compromised. Users must *develop and maintain awareness* of the state of the entire space being assessed and *monitor data quality* to ensure that contextual data provides an accurate and complete assessment of the focus of the inspection (e.g., the fieldsite or structure). To conserve time or resources, users may often wish to *make predictions* about the likelihood of success to adjust plans to ensure adequate and thorough coverage. *Assessing risks* associated with operational decisions (e.g., what is the risk that rescuers will miss a survivor by not investigating further [19]; Fig. 1a) should be emphasized in emergency response interfaces.

Debugging and Recovery from Error: Across many domains, users may need to engage in **Activities** involving recognizing that something went wrong (either with the overall operation, the robot, or both), determining what specifically went wrong, assessing the severity of an error, and reasoning over potential solutions. Prior HRI interfaces for examining programmatic state flows [46, 112, 128] and for robot debugging [26, 82] may be particularly relevant and directly informed by VIS research on network visualization [58] and visual debugging techniques [59]. In this context, key **Data Tasks** include *finding relevant information* related to an error, *synthesize information across sources* and *identify relevant anomalies* to determine what caused to the error, *develop and maintain awareness* of the system state to characterize the error and its magnitude, and *make predictions* about how potential resolutions may influence the system and the likelihood of recurrence.

6 DISCUSSION AND OPPORTUNITIES

We argue for closer collaboration between HRI and VIS, focusing on how knowledge from VIS may inform data-centric robot interface design. This interdisciplinary crossover can happen at both a low-level (e.g., helping designers reason about color choices) and a higher level (e.g., helping developers understand sensemaking and decision making). One recent example of the value in such cross-pollination is the MOSAIC Viewer [7], a visualization interface for multirobot systems developed using a design study approach [132]. This interface demonstrates how managing level of detail across data dimensions and views can emphasize key aspects of robot state and belief. MRS operators perceived that the interface increased speed, trust, and understanding in detecting anomalous behaviors.

While our analysis and framework here focus on how visualization can benefit HRI, we note that HRI also offers novel opportunities for visualization researchers. For example, most visualizations are designed for static data (i.e., data that is already available to users). While some tools have explored streaming data [64, 84] and progressive computations [141, 158], visualization currently has limited insights into designing for dynamic data at the volume and variety collected by robots or for considering how data importance may vary as the operational context, data type, and data age change. Many robotics applications require combining several forms of information-rich, yet visually complex, image and spatial data, such as LiDAR, IR images, and camera feeds. While image-based visualizations exist [51, 83, 121, 152, 155], understanding how to fuse complex image data with other forms of information and how to best analyze multiple visual streams simultaneously to support real-time decision making remain open challenges.

There are also many potential directions for mutually advancing HRI and VIS beyond those detailed here. For example, both VIS [20] and HRI [55, 145] researchers are exploring how to effectively communicate the reasoning processes used in AI algorithms. Similarly, dynamic data physicalization using robot swarms is being simultaneously investigated by both VIS [85] and HRI [52, 105].

Our work promotes the need for greater collaboration between HRI and visualization. Such cross-disciplinary collaboration will lead to mutually beneficial innovations that help create data-centric HRI interfaces for more effectively reasoning about and acting on the vast quantities of information produced by robots. Our data-centric framework aims to scaffold the design of such interfaces by highlighting the importance of **Data Tasks**, establishing connections across robotics and visualization, and identifying VIS best practices relevant to effective robot interface design. However, our framework is currently limited in its ability to provide specific guidance regarding objective trade-offs in potential designs. We hope that future research will enable our framework to evolve in its ability to provide actionable strategies for achieving required outcomes in interface design, yet caution that VIS, like HRI, is a complex design problem: purely prescriptive guidance will not always be possible or desired. Rather, we hope this paper serves as a call to action for both communities to recognize potential synergies and explore co-innovation at the rich intersection of HRI and VIS.

7 ACKNOWLEDGEMENTS

This work was supported by the NSF under Award #1764092.

REFERENCES

[1] George Adamides, Christos Katsanos, Ioannis Constantinou, Georgios Christou, Michalis Xenos, Thanasis Hadzilacos, and Yael Edan. 2017. Design and development of a semi-autonomous agricultural vineyard sprayer: Human–robot interaction aspects. *Journal of Field Robotics* 34, 8 (2017), 1407–1426.

[2] Julie A Adams. 2002. Critical considerations for human-robot interface development. In *Proceedings of 2002 AAAI Fall Symposium*. 1–8.

[3] Julie A Adams. 2007. Unmanned vehicle situation awareness: A path forward. In *Human systems integration symposium*. 31–89.

[4] Danielle Albers, Michael Correll, and Michael Gleicher. 2014. Task-driven evaluation of aggregation in time series visualization. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. 551–560.

[5] Robert Amar, James Eagan, and John Stasko. 2005. Low-level components of analytic activity in information visualization. In *IEEE Symposium on Information Visualization, 2005. INFOVIS 2005*. IEEE, 111–117.

[6] Robert Amar and John Stasko. 2004. A knowledge task-based framework for design and evaluation of information visualizations. In *IEEE Symposium on Information Visualization*. IEEE, 143–150.

[7] Suyun Bae, Federico Rossi, Joshua Vander Hook, Scott Davidoff, and Kwan-Liu Ma. 2021. A Visual Analytics Approach to Debugging Cooperative, Autonomous Multi-Robot Systems' Worldviews. *IEEE Transactions on Visualization and Computer Graphics* (2021).

[8] Haris Balta, Janusz Bedkowski, Shashank Govindaraj, Karol Majek, Pawel Musialik, Daniel Serrano, Kostas Alexis, Roland Siegwart, and Geert De Cubber. 2017. Integrated data management for a fleet of search-and-rescue robots. *Journal of Field Robotics* 34, 3 (2017), 539–582.

[9] Jonathan Bohren, Chris Paxton, Ryan Howarth, Gregory D Hager, and Louis L Whitcomb. 2016. Semi-autonomous telerobotic assembly over high-latency networks. In *ACM/IEEE International Conference on Human-Robot Interaction (HRI)*. 149–156.

[10] Michelle Borkin, Krzysztof Gajos, Amanda Peters, Dimitrios Mitsouras, Simone Melchionna, Frank Rybicki, Charles Feldman, and Hanspeter Pfister. 2011. Evaluation of artery visualizations for heart disease diagnosis. *IEEE Transactions on Visualization and Computer Graphics* 17, 12 (2011), 2479–2488.

[11] Michelle A Borkin, Azalea A Vo, Zoya Bylinskii, Phillip Isola, Shashank Sunkavalli, Aude Oliva, and Hanspeter Pfister. 2013. What makes a visualization memorable? *IEEE Transactions on Visualization and Computer Graphics* 19, 12 (2013), 2306–2315.

[12] David Borland and Russell M Taylor Ii. 2007. Rainbow color map (still) considered harmful. *IEEE computer graphics and applications* 27, 2 (2007), 14–17.

[13] Gerald Brantner and Oussama Khatib. 2020. Controlling Ocean One: Human–robot collaboration for deep-sea manipulation. *Journal of Field Robotics* (2020).

[14] Connor Brooks and Daniel Szafir. 2019. Balanced information gathering and goal-oriented actions in shared autonomy. In *ACM/IEEE International Conference on Human-Robot Interaction (HRI)*. 85–94.

[15] David J Bruemmer, David I Gertman, Curtis W Nielsen, Douglas A Few, and William D Smart. 2007. Supporting complex robot behaviors with simple interaction tools. In *Human Robot Interaction*.

[16] Zoya Bylinskii, Nam Wook Kim, Peter O'Donovan, Sami Alsheikh, Spandan Madan, Hanspeter Pfister, Fredo Durand, Bryan Russell, and Aaron Hertzmann. 2017. Learning visual importance for graphic designs and data visualizations. In *Proceedings of the 30th Annual ACM symposium on user interface software and technology*. 57–69.

[17] Stuart K Card and Jock Mackinlay. 1997. The structure of the information visualization design space. In *Proceedings of VIZ'97: Visualization Conference, Information Visualization Symposium and Parallel Rendering Symposium*. IEEE, 92–99.

[18] Roberto Carlino, Jonathan Barlow, Jose Benavides, Maria Bualat, Aric Katterhagen, Yunji Kim, Ruben Garcia Ruiz, Trey Smith, and Andres Mora Vargas. [n.d.]. Astrobee Free Flyers: Integrated and Tested. Ready for Launch! *NASA Technical Reports* ([n. d.]).

[19] Jennifer Casper and Robin R. Murphy. 2003. Human-robot interactions during the robot-assisted urban search and rescue response at the world trade center. *IEEE Transactions on Systems, Man, and Cybernetics, Part B (Cybernetics)* 33, 3 (2003), 367–385.

[20] Angelos Chatzimparmpas, Rafael M Martins, Ilir Jusufi, and Andreas Kerren. 2020. A survey of surveys on the use of visualization for interpreting machine learning models. *Information Visualization* (2020), 1473871620904671.

[21] Jessie YC Chen and Michael J Barnes. 2008. Robotics operator performance in a military multi-tasking environment. In *ACM/IEEE International Conference on Human-Robot Interaction (HRI)*. 279–286.

[22] J. Y. C. Chen, E. C. Haas, and M. J. Barnes. 2007. Human Performance Issues and User Interface Design for Teleoperated Robots. *IEEE Transactions on Systems, Man, and Cybernetics, Part C (Applications and Reviews)* 37, 6 (2007), 1231–1245.

[23] William J Clancey. 1997. *Situated cognition: On human knowledge and computer representations*. Cambridge university press.

[24] William S Cleveland and Robert McGill. 1984. Graphical perception: Theory, experimentation, and application to the development of graphical methods. *Journal of the American statistical association* 79, 387 (1984), 531–554.

[25] Andy Cockburn, Amy Karlson, and Benjamin B Bederson. 2009. A review of overview+ detail, zooming, and focus+ context interfaces. *ACM Computing Surveys (CSUR)* 41, 1 (2009), 1–31.

[26] THJ Collett and Bruce A MacDonald. 2006. Developer oriented visualisation of a robot program. In *Proceedings of the ACM SIGCHI/SIGART conference on Human-robot interaction*. 49–56.

[27] Giuseppe Conte, Silvia Maria Zanoli, David Scaradozzi, and Andrea Caiti. 2009. Robotics techniques for data acquisition in underwater archeology. *International Journal of Mechanics and Control (JoMaC)* 10, 1 (2009), 45–51.

[28] François Conti, Jaeheung Park, and Oussama Khatib. 2014. Interface design and control strategies for a robot assisted ultrasonic examination system. In *Experimental Robotics*. 97–113.

[29] Joseph Cooper and Michael A Goodrich. 2008. Towards combining UAV and sensor operator roles in UAV-enabled visual search. In *ACM/IEEE International Conference on Human-Robot Interaction (HRI)*. 351–358.

[30] Michael Correll and Jeffrey Heer. 2017. Regression by eye: Estimating trends in bivariate visualizations. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*. 1387–1396.

[31] Michael Correll, Mingwei Li, Gordon Kindlmann, and Carlos Scheidegger. 2018. Looks good to me: Visualizations as sanity checks. *IEEE Transactions on Visualization and Computer Graphics* 25, 1 (2018), 830–839.

[32] Jacob W Crandall and Mary L Cummings. 2007. Developing performance metrics for the supervisory control of multiple robots. In *Proceedings of the ACM/IEEE international conference on Human-robot interaction*. 33–40.

[33] Daniel Crawl, Jessica Block, Kai Lin, and Ilkay Altintas. 2017. Firemap: A dynamic data-driven predictive wildfire modeling and visualization environment. *Procedia Computer Science* 108 (2017), 2230–2239.

[34] Jeffrey Delmerico, Stefano Mintchev, Alessandro Giusti, Boris Gromov, Kamilo Melo, Tomislav Horvat, Cesar Cadena, Marco Hutter, Auke Ijspeert, Dario Floreano, et al. 2019. The current state and future outlook of rescue robotics. *Journal of Field Robotics* 36, 7 (2019), 1171–1191.

[35] Anca D Dragan and Siddhartha S Srinivasa. 2013. A policy-blending formalism for shared control. *The International Journal of Robotics Research* 32, 7 (2013), 790–805.

[36] Jill L Drury, Brenden Keyes, and Holly A Yanco. 2007. LASSOing HRI: analyzing situation awareness in map-centric and video-centric interfaces. In *ACM/IEEE International Conference on Human-Robot Interaction (HRI)*. 279–286.

[37] Jill L Drury, Laurel Riek, and Nathan Rackliffe. 2006. A decomposition of UAV-related situation awareness. In *Proceedings of the ACM SIGCHI/SIGART conference on Human-robot interaction*. 88–94.

[38] Jill L Drury, Jean Scholtz, and Holly A Yanco. 2003. Awareness in human-robot interactions. In *IEEE International Conference on Systems, Man and Cybernetics: System Security and Assurance*, Vol. 1. 912–918.

[39] Geoffrey Ellis and Alan Dix. 2007. A taxonomy of clutter reduction for information visualisation. *IEEE Transactions on Visualization and Computer Graphics* 13, 6 (2007), 1216–1223.

[40] Okan Erat, Werner Alexander Isop, Denis Kalkofen, and Dieter Schmalstieg. 2018. Drone-augmented human vision: Exocentric control for drones exploring hidden areas. *IEEE Transactions on Visualization and Computer Graphics* 24, 4 (2018), 1437–1446.

[41] François Ferland, François Pomerleau, Chon Tam Le Dinh, and François Michaud. 2009. Egocentric and exocentric teleoperation interface using real-time, 3D video projection. In *Proceedings of the ACM/IEEE international conference on Human robot interaction*. 37–44.

[42] Terrence Fong, Charles Thorpe, and Charles Baur. 2001. *Collaborative control: A robot-centric model for vehicle teleoperation*. Vol. 1. Carnegie Mellon University, The Robotics Institute.

[43] Steven L Franconeri. 2013. The nature and status of visual resources. (2013).

[44] Daniel Gaines, Gary Doran, Michael Paton, Brandon Rothrock, Joseph Russino, Ryan Mackey, Robert Anderson, Raymond Francis, Chet Joswig, Heather Justice, et al. 2020. Self-reliant rovers for increased mission productivity. *Journal of Field Robotics* 37, 7 (2020), 1171–1196.

[45] Jack Gale, John Karasinski, and Steve Hillenius. 2018. Playbook for UAS: UX of Goal-Oriented Planning and Execution. In *International Conference on Engineering Psychology and Cognitive Ergonomics*. 545–557.

[46] Dylan F Glas, Takayuki Kanda, and Hiroshi Ishiguro. 2016. Human-robot interaction design using Interaction Composer eight years of lessons learned. In *2016 ACM/IEEE International Conference on Human-Robot Interaction (HRI)*. 303–310.

[47] Michael Gleicher. 2017. Considerations for visualizing comparison. *IEEE Transactions on Visualization and Computer Graphics* 24, 1 (2017), 413–423.

[48] Michael Gleicher, Danielle Albers, Rick Walker, Ilir Jusufi, Charles D Hansen, and Jonathan C Roberts. 2011. Visual comparison for information visualization. *Information Visualization* 10, 4 (2011), 289–309.

[49] Shashank Govindaraj, Keshav Chintamani, Jeremi Gancet, Pierre Letier, Boris van Lierde, Yashodhan Nevatia, Geert De Cubber, Daniel Serrano, Miguel Esbri Palomares, Janusz Bedkowski, et al. 2013. The icarus project-command, control

and intelligence (c2i). In *IEEE International Symposium on Safety, Security, and Rescue Robotics (SSRR)*. 1–4.

[50] Connor C Gramazio, David H Laidlaw, and Karen B Schloss. 2016. Colorgorical: Creating discriminable and preferable color palettes for information visualization. *IEEE Transactions on Visualization and Computer Graphics* 23, 1 (2016), 521–530.

[51] Yi Gu, Chaoli Wang, Jun Ma, Robert J Nemiroff, and David L Kao. 2015. iGraph: a graph-based technique for visual analytics of image and text collections. In *Visualization and Data Analysis 2015*, Vol. 9397. International Society for Optics and Photonics, 939708.

[52] Darren Guinness, Annika Muehlbradt, Daniel Szafir, and Shaun K Kane. 2019. RoboGraphics: Dynamic Tactile Graphics Powered by Mobile Robots. In *The International ACM SIGACCESS Conference on Computers and Accessibility*. 318–328.

[53] Lane Harrison, Fumeng Yang, Steven Franconeri, and Remco Chang. 2014. Ranking visualizations of correlation using weber's law. *IEEE Transactions on Visualization and Computer Graphics* 20, 12 (2014), 1943–1952.

[54] Mark Harrower and Cynthia A Brewer. 2003. ColorBrewer. org: an online tool for selecting colour schemes for maps. *The Cartographic Journal* 40, 1 (2003), 27–37.

[55] Bradley Hayes and Julie A Shah. 2017. Improving robot controller transparency through autonomous policy explanation. In *ACM/IEEE International Conference on Human-Robot Interaction (HRI)*. 303–312.

[56] Christopher Healey and James Enns. 2011. Attention and visual memory in visualization and computer graphics. *IEEE Transactions on Visualization and Computer Graphics* 18, 7 (2011), 1170–1188.

[57] Hooman Hedayati, Michael Walker, and Daniel Szafir. 2018. Improving collocated robot teleoperation with augmented reality. In *Proceedings of the 2018 ACM/IEEE International Conference on Human-Robot Interaction*. 78–86.

[58] Ivan Herman, Guy Melançon, and M Scott Marshall. 2000. Graph visualization and navigation in information visualization: A survey. *IEEE Transactions on Visualization and Computer Graphics* 6, 1 (2000), 24–43.

[59] Jane Hoffswell, Arvind Satyanarayan, and Jeffrey Heer. 2016. Visual debugging techniques for reactive data visualization. In *Computer Graphics Forum*, Vol. 35. 271–280.

[60] James Hollan, Edwin Hutchins, and David Kirsh. 2000. Distributed cognition: toward a new foundation for human-computer interaction research. *ACM Transactions on Computer-Human Interaction (TOCHI)* 7, 2 (2000), 174–196.

[61] Danny Holten and Jarke J Van Wijk. 2009. Force-directed edge bundling for graph visualization. In *Computer graphics forum*, Vol. 28. Wiley Online Library, 983–990.

[62] Haikun Huang, Ni-Ching Lin, Lorenzo Barrett, Darian Springer, Hsueh-Cheng Wang, Marc Pomplun, and Lap-Fai Yu. 2017. Automatic optimization of wayfinding design. *IEEE Transactions on Visualization and Computer Graphics* 24, 9 (2017), 2516–2530.

[63] Curtis M Humphrey and Julie A Adams. 2008. Compass visualizations for human-robotic interaction. In *Proceedings of the 3rd ACM/IEEE international conference on Human robot interaction*. 49–56.

[64] Samuel Huron, Romain Vuillemot, and Jean-Daniel Fekete. 2012. Towards visual sedimentation. In *IEEE Transactions on Visualization and Computer Graphics*. IEEE.

[65] AV Inzartsev, Alexander Kamorniy, Lev Kiselyov, Yury Matviyenko, Nicolay Rylov, Roman Rylov, and Yury Vaulin. 2010. Integrated positioning system of autonomous underwater robot and its application in high latitudes of arctic zone. *Motion Control. Vienna: InTech* (2010), 229–244.

[66] Nicole Jardine, Brian D Ondov, Niklas Elmqvist, and Steven Franconeri. 2019. The Perceptual Proxies of Visual Comparison. *IEEE Transactions on Visualization and Computer Graphics* 26, 1 (2019), 1012–1021.

[67] Waqas Javed and Niklas Elmqvist. 2012. Exploring the design space of composite visualization. In *2012 IEEE Pacific Visualization Symposium*. IEEE, 1–8.

[68] Feng Jiang, Su Lu, Xiang Yao, Xiaodong Yue, and Wing Tung Au. 2014. Up or down? How culture and color affect judgments. *Journal of Behavioral Decision Making* 27, 3 (2014), 226–234.

[69] Matthew Johnson, Jeffrey M Bradshaw, Paul J Feltovich, Catholijn M Jonker, M Birna Van Riemsdijk, and Maarten Sierhuis. 2014. Coactive design: Designing support for interdependence in joint activity. *Journal of Human-Robot Interaction* 3, 1 (2014), 43–69.

[70] Matthew Johnson, John Carff, and Jerry Pratt. 2012. Coactive design for human-MAV team navigation. In *International micro aerial vehicle conference and competition*.

[71] Matthew Johnson, Brandon Shrewsbury, Sylvain Bertrand, Tingfan Wu, Daniel Duran, Marshall Floyd, Peter Abeles, Douglas Stephen, Nathan Mertins, Alex Lesman, et al. 2015. Team IHMC's lessons learned from the DARPA robotics challenge trials. *Journal of Field Robotics* 32, 2 (2015), 192–208.

[72] Susanne Jul and George W Furnas. 1998. Critical zones in desert fog: aids to multiscale navigation. In *Proceedings of the 11th annual ACM symposium on User interface software and technology*. 97–106.

[73] M Waleed Kadous, Raymond Ka-Man Sheh, and Claude Sammut. 2006. Effective user interface design for rescue robotics. In *Proceedings of the 1st ACM SIGCHI/SIGART conference on Human-Robot Interaction (HRI)*. 250–257.

[74] Alex Kale, Matthew Kay, and Jessica Hullman. 2019. Decision-making under uncertainty in research synthesis: Designing for the garden of forking paths. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*. 1–14.

[75] Walter Kasper. 2016. Team monitoring and reporting for robot-assisted usar missions. In *IEEE International Symposium on Safety, Security, and Rescue Robotics (SSRR)*. 246–251.

[76] Shinji Kawatsuma, Mineo Fukushima, and Takashi Okada. 2012. Emergency response by robots to Fukushima-Daiichi accident: summary and lessons learned. *Industrial Robot: An International Journal* (2012).

[77] Alonzo Kelly, Nicholas Chan, Herman Herman, Daniel Huber, Robert Meyers, Pete Rander, Randy Warner, Jason Ziglar, and Erin Capstick. 2011. Real-time photorealistic virtualized reality interface for remote mobile robot control. *The International Journal of Robotics Research* 30, 3 (2011), 384–404.

[78] Brenden Keyes, Mark Micire, Jill L Drury, Holly A Yanco, and Daisuke Chugo. 2010. Improving human-robot interaction through interface evolution. In *Human-robot interaction*. InTech, 183–202.

[79] Won Soo Kim, Paul S Schenker, Antal K Bejczy, Stephen Leake, and Stanford Ollendorf. 1993. Advanced operator interface design with preview/predictive displays for ground-controlled space telerobotic servicing. In *Telemanipulator Technology and Space Telerobotics*, Vol. 2057. 96–107.

[80] Younghoon Kim and Jeffrey Heer. 2018. Assessing effects of task and data distribution on the effectiveness of visual encodings. In *Computer Graphics Forum*, Vol. 37. Wiley Online Library, 157–167.

[81] Tobias Klamt, Max Schwarz, Christian Lenz, Lorenzo Baccelliere, Domenico Buongiorno, Torben Cichon, Antonio DiGuardo, David Droeschel, Massimiliano Gabardi, Malgorzata Kamedula, et al. 2020. Remote mobile manipulation with the centauro robot: Full-body telepresence and autonomous operator assistance. *Journal of Field Robotics* 37, 5 (2020), 889–919.

[82] Tijn Kooijmans, Takayuki Kanda, Christoph Bartneck, Hiroshi Ishiguro, and Norihiro Hagita. 2006. Interaction debugging: an integral approach to analyze human-robot interaction. In *Proceedings of the ACM SIGCHI/SIGART conference on Human-robot interaction*. 64–71.

[83] Ashnil Kumar, Falk Nette, Karsten Klein, Michael Fulham, and Jinman Kim. 2014. A visual analytics approach using the exploration of multidimensional feature spaces for content-based medical image retrieval. *IEEE journal of biomedical and health informatics* 19, 5 (2014), 1734–1746.

[84] Ove Daae Lampe and Helwig Hauser. 2011. Interactive visualization of streaming data with kernel density estimation. In *2011 IEEE pacific visualization symposium*. IEEE, 171–178.

[85] Mathieu Le Goc, Charles Perin, Sean Follmer, Jean-Daniel Fekete, and Pierre Dragicevic. 2018. Dynamic composite data physicalization using wheeled micro-robots. *IEEE Transactions on Visualization and Computer Graphics* 25, 1 (2018), 737–747.

[86] Le Liu, Lace Padilla, Sarah H Creem-Regehr, and Donald H House. 2018. Visualizing uncertain tropical cyclone predictions using representative samples from ensembles of forecast tracks. *IEEE Transactions on Visualization and Computer Graphics* 25, 1 (2018), 882–891.

[87] Shixia Liu, Gennady Andrienko, Yingcai Wu, Nan Cao, Liu Jiang, Conglei Shi, Yu-Shuen Wang, and Seokhee Hong. 2018. Steering data quality with visual analytics: The complexity challenge. *Visual Informatics* 2, 4 (2018), 191–197.

[88] Pat Marion, Maurice Fallon, Robin Deits, Andrés Valenzuela, Claudia Pérez D'Arpino, Greg Izatt, Lucas Manuelli, Matt Antone, Hongkai Dai, Twan Koolen, et al. 2017. Director: A user interface designed for robot operation with shared autonomy. *Journal of Field Robotics* 34, 2 (2017), 262–280.

[89] Laura E Matzen, Michael J Haass, Kristin M Divis, Zhiyuan Wang, and Andrew T Wilson. 2017. Data visualization saliency model: A tool for evaluating abstract data visualizations. *IEEE Transactions on Visualization and Computer Graphics* 24, 1 (2017), 563–573.

[90] Mark Micire, Terrence Fong, Ted Morse, Eric Park, Chris Provencher, Ernest Smith, Vinh To, R Jay Torres, DW Wheeler, and David Mittman. 2013. Smart spheres: a telerobotic free-flyer for intravehicular activities in space. In *AIAA Space 2013 Conference and Exposition*. 5338.

[91] Christopher Miller, Harry Funk, Peggy Wu, Robert Goldman, John Meisner, and Marc Chapman. 2005. The Playbook approach to adaptive automation. In *Proceedings of the Human Factors and Ergonomics Society Annual Meeting*, Vol. 49. 15–19.

[92] Christopher A Miller, Harry B Funk, Michael Dorneich, and Stephen D Whitlow. 2002. A playbook interface for mixed initiative control of multiple unmanned vehicle teams. In *IEEE Digital Avionics Systems Conference*, Vol. 2. 7E4–7E4.

[93] Bryce Morrow, Trevor Manz, Arlene E Chung, Nils Gehlenborg, and David Gotz. 2019. Periphery Plots for Contextualizing Heterogeneous Time-Based Charts. In *2019 IEEE Visualization Conference (VIS)*. IEEE, 1–5.

[94] Robin R Murphy and Satoshi Tadokoro. 2019. User Interfaces for Human-Robot Interaction in Field Robotics. In *Disaster Robotics*. 507–528.

[95] Francesca Negrello, Alessandro Settimi, Danilo Caporale, Gianluca Lentini, Mattia Poggiani, Dimitrios Kanoulas, Luca Muratore, Emanuele Luberto, Gaspare Santaera, Luca Ciarleglio, et al. 2018. Humanoids at work: The walk-man robot in a postearthquake scenario. *IEEE Robotics & Automation Magazine* 25, 3 (2018), 8–22.

[96] F Nguyen, X Qiao, J Heer, and J Hullman. 2020. Exploring the Effects of Aggregation Choices on Untrained Visualization Users' Generalizations From Data. In *Computer Graphics Forum*. Wiley Online Library.

[97] Laurent A Nguyen, Maria Bualat, Laurence J Edwards, Lorenzo Flueckiger, Charles Neveu, Kurt Schwehr, Michael D Wagner, and Eric Zbinden. 2001. Virtual reality interfaces for visualization and control of remote vehicles. *Autonomous Robots* 11, 1 (2001), 59–68.

[98] Curtis W Nielsen and Michael A Goodrich. 2006. Comparing the usefulness of video and map information in navigation tasks. In *Proceedings of the ACM SIGCHI/SIGART conference on Human-robot interaction*. 95–101.

[99] Curtis W Nielsen, Michael A Goodrich, and Robert W Ricks. 2007. Ecological interfaces for improving mobile robot teleoperation. *IEEE Transactions on Robotics* 23, 5 (2007), 927–941.

[100] Jakob Nielsen. 1994. *Usability engineering*. Morgan Kaufmann.

[101] Donald A Norman. 1984. Stages and levels in human-machine interaction. *International journal of man-machine studies* 21, 4 (1984), 365–375.

[102] Chris North. 2006. Toward measuring visualization insight. *IEEE computer graphics and applications* 26, 3 (2006), 6–9.

[103] Adam Norton, Willard Ober, Lisa Baraniecki, David Shane, Anna Skinner, and Holly Yanco. 2018. Perspectives on Human-Robot Team Performance from an Evaluation of the DARPA Robotics Challenge. In *The DARPA Robotics Challenge Finals: Humanoid Robots To The Rescue*. 631–666.

[104] Kyohei Otsu, Scott Tepsuporn, Rohan Thakker, Tiago Stegun Vaquero, Jeffrey A Edlund, William Walsh, Gregory Miles, Tristan Heywood, Michael T Wolf, and Ali-Akbar Agha-Mohammadi. 2020. Supervised Autonomy for Communication-degraded Subterranean Exploration by a Robot Team. In *2020 IEEE Aerospace Conference*. 1–9.

[105] Ayberk Özgür, Séverin Lemaignan, Wafa Johal, Maria Beltran, Manon Briod, Léa Pereyre, Francesco Mondada, and Pierre Dillenbourg. 2017. Cellulo: Versatile handheld robots for education. In *ACM/IEEE International Conference on Human-Robot Interaction (HRI*. 119–127.

[106] Lace M Padilla, Sarah H Creem-Regehr, Mary Hegarty, and Jeanine K Stefanucci. 2018. Decision making with visualizations: a cognitive framework across disciplines. *Cognitive research: principles and implications* 3, 1 (2018), 29.

[107] Lace M Padilla, Grace Hansen, Ian T Ruginski, Heidi S Kramer, William B Thompson, and Sarah H Creem-Regehr. 2015. The influence of different graphical displays on nonexpert decision making under uncertainty. *Journal of Experimental Psychology: Applied* 21, 1 (2015), 37.

[108] Ramviyas Parasuraman, Sergio Caccamo, Fredrik Båberg, Petter Ögren, and Mark Neerincx. 2017. A New UGV Teleoperation Interface for Improved Awareness of Network Connectivity and Physical Surroundings. *J. Hum.-Robot Interact.* 6, 3 (2017), 48–70. https://doi.org/10.5898/JHRI.6.3.Parasuraman

[109] Savannah Paul, Christopher Reardon, Tom Williams, and Hao Zhang. 2020. Designing augmented reality visualizations for synchronized and time-dominant human-robot teaming. In *Virtual, Augmented, and Mixed Reality (XR) Technology for Multi-Domain Operations*, Vol. 11426. International Society for Optics and Photonics, 1142607.

[110] Peter Pirolli and Stuart Card. 2005. The sensemaking process and leverage points for analyst technology as identified through cognitive task analysis. In *Proceedings of international conference on intelligence analysis*, Vol. 5. McLean, VA, USA, 2–4.

[111] Pere Ponsa and Marta Díaz. 2007. Creation of an ergonomic guideline for supervisory control interface design. In *International Conference on Engineering Psychology and Cognitive Ergonomics*. 137–146.

[112] David Porfirio, Allison Sauppé, Aws Albarghouthi, and Bilge Mutlu. 2018. Authoring and verifying human-robot interactions. In *Proceedings of the ACM Symposium on User Interface Software and Technology*. 75–86.

[113] Zening Qu and Jessica Hullman. 2017. Keeping multiple views consistent: Constraints, validations, and exceptions in visualization authoring. *IEEE Transactions on Visualization and Computer Graphics* 24, 1 (2017), 468–477.

[114] Christopher Reardon, Jason Gregory, Carlos Nieto-Granda, and John G Rogers. 2020. Enabling Situational Awareness via Augmented Reality of Autonomous Robot-Based Environmental Change Detection. In *International Conference on Human-Computer Interaction*. 611–628.

[115] Christopher Reardon, Kevin Lee, John G Rogers, and Jonathan Fink. 2019. Communicating via Augmented Reality for Human-Robot Teaming in Field Environments. In *2019 IEEE International Symposium on Safety, Security, and Rescue Robotics (SSRR)*. 94–101.

[116] Aurélien Reveleau, François Ferland, Mathieu Labbé, Dominic Létourneau, and François Michaud. 2015. Visual representation of interaction force and sound source in a teleoperation user interface for a mobile robot. *Journal of Human-Robot Interaction* 4, 2 (2015), 1–23.

[117] Justin Richer and Jill L Drury. 2006. A video game-based framework for analyzing human-robot interaction: characterizing interface design in real-time interactive multimedia applications. In *Proceedings of the ACM SIGCHI/SIGART conference on human-robot interaction*. 266–273.

[118] Kris Rogers, Janet Wiles, Scott Heath, Kristyn Hensby, and Jonathon Taufatofua. 2016. Discovering patterns of touch: a case study for visualization-driven analysis in human-robot interaction. In *2016 11th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*. 499–500.

[119] Bernice E Rogowitz and Alan D Kalvin. 2001. The" Which Blair Project": a quick visual method for evaluating perceptual color maps. In *Proceedings Visualization, 2001. VIS'01*. IEEE, 183–556.

[120] Juan Jesús Roldán, Elena Peña-Tapia, Andrés Martín-Barrio, Miguel A Olivares-Méndez, Jaime Del Cerro, and Antonio Barrientos. 2017. Multi-robot interfaces and operator situational awareness: Study of the impact of immersion and prediction. *Sensors* 17, 8 (2017), 1720.

[121] Ork Rooij, Jarke van Wijk, and Marcel Worring. 2010. Mediatable: Interactive categorization of multimedia collections. *IEEE Computer Graphics and Applications* 30, 5 (2010), 42–51.

[122] Ruth Rosenholtz, Yuanzhen Li, and Lisa Nakano. 2007. Measuring visual clutter. *Journal of vision* 7, 2 (2007), 17–17.

[123] Karina A Roundtree, Jason R Cody, Jennifer Leaf, H Onan Demirel, and Julie A Adams. 2019. Visualization design for human-collective teams. In *Proceedings of the Human Factors and Ergonomics Society Annual Meeting*, Vol. 63. 417–421.

[124] Daniel M Russell, Mark J Stefik, Peter Pirolli, and Stuart K Card. 1993. The cost structure of sensemaking. In *Proceedings of the INTERACT'93 and CHI'93 conference on Human factors in computing systems*. 269–276.

[125] Jibonananda Sanyal, Song Zhang, Gargi Bhattacharya, Phil Amburn, and Robert Moorhead. 2009. A user study to compare four uncertainty visualization methods for 1d and 2d datasets. *IEEE Transactions on Visualization and Computer Graphics* 15, 6 (2009), 1209–1218.

[126] Alper Sarikaya, Michael Correll, Lyn Bartram, Melanie Tory, and Danyel Fisher. 2018. What do we talk about when we talk about dashboards? *IEEE Transactions on Visualization and Computer Graphics* 25, 1 (2018), 682–692.

[127] Alper Sarikaya, Michael Gleicher, and Danielle Albers Szafir. 2018. Design factors for summary visualization in visual analytics. In *Computer Graphics Forum*, Vol. 37. Wiley Online Library, 145–156.

[128] Allison Sauppé and Bilge Mutlu. 2014. Design patterns for exploring and prototyping human-robot interactions. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. 1439–1448.

[129] Jean Scholtz, Jeff Young, Jill L Drury, and Holly A Yanco. 2004. Evaluation of human-robot interaction awareness in search and rescue. In *IEEE International Conference on Robotics and Automation (ICRA)*, Vol. 3. 2327–2332.

[130] Hans-Jörg Schulz, Thomas Nocke, Magnus Heitzler, and Heidrun Schumann. 2013. A design space of visualization tasks. *IEEE Transactions on Visualization and Computer Graphics* 19, 12 (2013), 2366–2375.

[131] Max Schwarz, Tobias Rodehutskors, David Droeschel, Marius Beul, Michael Schreiber, Nikita Araslanov, Ivan Ivanov, Christian Lenz, Jan Razlaw, Sebastian Schüller, et al. 2017. NimbRo Rescue: Solving disaster-response tasks with the mobile manipulation robot Momaro. *Journal of Field Robotics* 34, 2 (2017), 400–425.

[132] Michael Sedlmair, Miriah Meyer, and Tamara Munzner. 2012. Design study methodology: Reflections from the trenches and the stacks. *IEEE Transactions on Visualization and Computer Graphics* 18, 12 (2012), 2431–2440.

[133] Meher T Shaikh and Michael A Goodrich. 2017. Design and evaluation of adverb palette: A gui for selecting tradeoffs in multi-objective optimization problems. In *ACM/IEEE International Conference on Human-Robot Interaction (HRI*. 389–397.

[134] Ben Shneiderman. 1996. The eyes have it: A task by data type taxonomy for information visualizations. In *Proceedings 1996 IEEE symposium on visual languages*. IEEE, 336–343.

[135] Samuel Silva, Beatriz Sousa Santos, and Joaquim Madeira. 2011. Using color in visualization: A survey. *Computers & Graphics* 35, 2 (2011), 320–333.

[136] Enrico Simetti, Giovanni Indiveri, and António M Pascoal. 2020. WiMUST: A cooperative marine robotic system for autonomous geotechnical surveys. *Journal of Field Robotics* (2020).

[137] Stephen Smart, Keke Wu, and Danielle Albers Szafir. 2019. Color Crafting: Automating the Construction of Designer Quality Color Ramps. *IEEE Transactions on Visualization and Computer Graphics* 26, 1 (2019), 1215–1225.

[138] Hayeong Song and Danielle Albers Szafir. 2018. Where's my data? evaluating visualizations with missing data. *IEEE Transactions on Visualization and Computer Graphics* 25, 1 (2018), 914–924.

[139] Aaron Steinfeld, Terrence Fong, David Kaber, Michael Lewis, Jean Scholtz, Alan Schultz, and Michael Goodrich. 2006. Common metrics for human-robot interaction. In *Proceedings of the ACM SIGCHI/SIGART conference on Human-robot interaction*. 33–40.

[140] Mark R Stevens, J Ross Beveridge, and Michael E Goss. 1997. Visualizing Multisensor Model-Based Object Recognition. In *IEEE Transactions on Visualization and Computer Graphics*.

[141] Charles D Stolper, Adam Perer, and David Gotz. 2014. Progressive visual analytics: User-driven visual exploration of in-progress analytics. *IEEE Transactions*

on Visualization and Computer Graphics 20, 12 (2014), 1653–1662.
[142] Daniel Szafir, Bilge Mutlu, and Terrence Fong. 2017. Designing planning and control interfaces to support user collaboration with flying robots. *The International Journal of Robotics Research* 36, 5-7 (2017), 514–542.
[143] Danielle Albers Szafir. 2018. The good, the bad, and the biased: five ways visualizations can mislead (and how to fix them). *interactions* 25, 4 (2018), 26–33.
[144] Danielle Albers Szafir, Steve Haroz, Michael Gleicher, and Steven Franconeri. 2016. Four types of ensemble coding in data visualizations. *Journal of vision* 16, 5 (2016), 11–11.
[145] Aaquib Tabrez, Shivendra Agrawal, and Bradley Hayes. 2019. Explanation-based reward coaching to improve human performance via reinforcement learning. In *ACM/IEEE International Conference on Human-Robot Interaction (HRI)*. 249–257.
[146] Michael Walker, Hooman Hedayati, Jennifer Lee, and Daniel Szafir. 2018. Communicating robot motion intent with augmented reality. In *Proceedings of the 2018 ACM/IEEE International Conference on Human-Robot Interaction*. 316–324.
[147] Michael E Walker, Hooman Hedayati, and Daniel Szafir. 2019. Robot teleoperation with augmented reality virtual surrogates. In *2019 14th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*. IEEE, 202–210.
[148] Matt Whitlock, Keke Wu, and Danielle Albers Szafir. 2019. Designing for mobile and immersive visual analytics in the field. *IEEE Transactions on Visualization and Computer Graphics* 26, 1 (2019), 503–513.
[149] Christopher D Wickens and C Melody Carswell. 1995. The proximity compatibility principle: Its psychological foundation and relevance to display design. *Human factors* 37, 3 (1995), 473–494.
[150] Grant A Wilde and Robin R Murphy. 2018. User interface for unmanned surface vehicles used to rescue drowning victims. In *IEEE International Symposium on Safety, Security, and Rescue Robotics (SSRR)*. 1–8.
[151] Bang Wong. 2011. Color blindness. *Nature Methods* 8, 6 (2011), 441–442.
[152] Xiao Xie, Xiwen Cai, Junpei Zhou, Nan Cao, and Yingcai Wu. 2018. A semantic-based method for visualizing large image collections. *IEEE Transactions on Visualization and Computer Graphics* 25, 7 (2018), 2362–2377.
[153] Holly A Yanco, Jill L Drury, and Jean Scholtz. 2004. Beyond usability evaluation: Analysis of human-robot interaction at a major robotics competition. *Human–Computer Interaction* 19, 1-2 (2004), 117–149.
[154] Holly A Yanco, Brenden Keyes, Jill L Drury, Curtis W Nielsen, Douglas A Few, and David J Bruemmer. 2007. Evolving interface design for robot search tasks. *Journal of Field Robotics* 24, 8-9 (2007), 779–799.
[155] Jing Yang, Jianping Fan, Daniel Hubball, Yuli Gao, Hangzai Luo, William Ribarsky, and Matthew Ward. 2006. Semantic image browser: Bridging information visualization with automated intelligent image analysis. In *2006 IEEE Symposium On Visual Analytics Science And Technology*. IEEE, 191–198.
[156] Sierra N Young and Joshua M Peschel. 2020. Review of Human–Machine Interfaces for Small Unmanned Systems With Robotic Manipulators. *IEEE Transactions on Human-Machine Systems* 50, 2 (2020), 131–143.
[157] Jeff Zacks and Barbara Tversky. 1999. Bars and lines: A study of graphic communication. *Memory & cognition* 27, 6 (1999), 1073–1079.
[158] Emanuel Zgraggen, Alex Galakatos, Andrew Crotty, Jean-Daniel Fekete, and Tim Kraska. 2016. How progressive visualizations affect exploratory analysis. *IEEE Transactions on Visualization and Computer Graphics* 23, 8 (2016), 1977–1987.
[159] Wenjuan Zhang, David Feltner, James Shirley, David Kaber, and Manida S Neubert. 2020. Enhancement and Application of a UAV Control Interface Evaluation Technique: Modified GEDIS-UAV. *ACM Transactions on Human-Robot Interaction (THRI)* 9, 2 (2020), 1–20.
[160] Stefanie Zollmann, Christof Hoppe, Tobias Langlotz, and Gerhard Reitmayr. 2014. Flyar: Augmented reality supported micro aerial vehicle navigation. *IEEE Transactions on Visualization and Computer Graphics (TVCG* 20, 4 (2014), 560–568.