BRiDGE: An Intelligent Space for Collaborative Situation Awareness in Safety-Critical Contexts Across the Reality-Virtuality Continuum

Jan-Henrik Schröder
University of Lübeck
Lübeck, Germany

Daniel Schacht
University of Lübeck
Lübeck, Germany

André Calero Valdez
University of Lübeck
Lübeck, Germany

Hans-Christian Jetter
University of Lübeck
Lübeck, Germany

Figure 1: BRiDGE, a collaborative interactive laboratory for safety-critical teams, supports work across the reality-virtuality continuum and across different degrees of automation.

Abstract

Safety-critical work is increasingly collaborative and increasingly shared with autonomous and semi-autonomous subsystems, which raises the question of who holds authority and accountability at any given moment. Such work also draws on a growing range of input and output modalities to make complex information understandable. Supporting both at once, variable levels of control and variable levels of immersion with seamless transitions between them, is, we argue, first of all a question of the space in which a team works. In this work-in-progress position paper we share our view on the design of such a space and present BRiDGE, a collaborative interactive laboratory we have built as an Intelligent Space, an interactive space extended with automation and local intelligences, using Blended Interaction as our design framework. We describe its computing substrate, its sensing, its interaction and output modalities, and its observer space, and we work through its design along the four design domains of Blended Interaction, drawing on collaborative lunar rover teleoperation as a running example. We close by asking who decides in such a setting, arguing that the answer remains the team, supported but not displaced by the room's transitional interfaces and local intelligences.

Mensch und Computer 2026 – Workshopband, Gesellschaft für Informatik e.V., 30. August - 02. September 2026, Duisburg, Germany

https://doi.org/10.18420/muc2026-mci-ws100-246

CCS Concepts

• **Human-centered computing → Computer supported cooperative work**; *Ambient intelligence*.

1 Introduction

Safety-critical work is increasingly collaborative, and it is increasingly carried out together with autonomous and semi-autonomous subsystems that can operate at varying degrees of automation. As these subsystems take on more of the work, one challenge moves to the center: deciding who holds authority at any given moment, and who is therefore accountable for the result [9]. At the same time, the data such teams must grasp is growing more complex, from one- and two-dimensional to three-dimensional, and making it understandable calls for an expanding range of input and output modalities. Supporting this kind of work is no longer a matter of designing one good interface for a fixed setting.

A recent example makes the shape of this development concrete. The MIDA-VIC paradigm [16] envisions collaborative space telerobotics in which several human operators work with a remote robot, and with each other, along two dimensions at once. The first is the *level of control*: depending on the situation, the same operator may command a motor directly, issue a higher-level instruction, or delegate a task to an autonomous process or a language model, so the team needs different degrees of automation side by side. The second is the *level of immersion*: depending on the task, an operator may want a flat two-dimensional overview on a shared display, a less immersive augmented view, or a fully immersive stereoscopic view, so the team needs different output and input modalities side by side. What makes such settings demanding is not either dimension alone but the requirement to move smoothly within and between them as a task unfolds.

We take MIDA-VIC as a representative example of where modern safety-critical systems are heading, and we observe that the smooth movement it calls for cannot be improvised on top of an arbitrary collection of devices. It first requires a platform built to provide it, an environment in which the displays, input devices, and automated services are tuned to one another well enough that an operator can change modality and degree of automation without friction, and in which a team can keep a shared picture of the situation while doing so. To our knowledge, no existing research environment has been built to provide this combination of variable automation, variable immersion, and seamless transitions between them for a co-located team. This is where the present paper enters. It is a work-in-progress position paper guided by the question: What must such an environment provide, and how should it be designed, so that collaborative situation awareness can be established, maintained, and studied across variable levels of control and immersion? We present BRiDGE, a collaborative interactive laboratory in which we build, run, and study such systems. We have designed it as an *Intelligent Space* [20], an interactive space augmented with sensing, computation, and the local machine intelligence that partial automation calls for, and we used Blended Interaction as the framework that guided its design. We return at the end to what such a space implies for who is ultimately responsible for a decision made within it.

2 Related Work

Because we wanted to build the laboratory on a principled foundation, we began by gathering what is known about such systems in general, so that we could later distill it into guiding requirements for the platform itself. We did so selectively rather than through a systematic review. Our position draws on three lines of work: awareness as the central objective of collaborative safety-critical systems, the need for variable levels of automation and immersion, and the notion of Transitional Interfaces together with the study of collaboration within them.

2.1 Awareness as a Design Objective

Research in human factors and human-computer interaction has long sought to reduce the mental, physical, and temporal demands of operating complex systems, so that operators retain the cognitive resources needed to build and maintain good *situation awareness* [7]. Following Endsley [6], situation awareness is the perception of the relevant elements in the environment, the comprehension of their meaning, and the projection of their status into the near future. These three levels map onto how a team's understanding unfolds in time. It perceives the present, understands what it means, and anticipates what is likely to follow. MIDA-VIC [16] illustrates this for telerobotic work, where an operator must perceive telemetry and imagery, comprehend that, for example, high wheel slip indicates soft terrain, and project whether the vehicle will become stuck or whether its battery will last through a planned drive.

Situation awareness as originally conceived describes an individual's understanding. In collaborative settings it is complemented by *workspace awareness*, a team member's up-to-the-moment understanding of which collaborators are present, what those collaborators are doing, where they are working, and how everyone's activities relate [12]. Workspace awareness is the shared counterpart of situation awareness, and coordinated action in a control room depends on it. In command-and-control practice, the representation of the situation that such a team works from is known as a *common operational picture*. We therefore take collective awareness, about the situation and the workspace, as the central objective against which a collaborative environment should be designed. The same design stance is also argued, in the context of collaborative control systems, by Jetter et al. [16].

2.2 Variable Levels of Automation and Immersion

The drive toward automation in safety-critical work is not a matter of replacing operators but of placing the right degree of autonomy on the right function [26]. Automation is not a single switch. Parasuraman et al. [26] distinguish four functions, each of which can sit at a different point on a continuum from manual to fully automatic: acquiring information, analyzing it, selecting an action, and carrying that action out. A team can then raise autonomy where it helps, as in the analysis of information, and keep it low where control must not be ceded, as in the selection of consequential actions. In teleoperation, especially for systems far away on Earth or in space, the need for adjustable automation is sharpened by communication latency, which grows from milliseconds in low Earth orbit to seconds for the Moon and many minutes for Mars [16]. As latency grows, direct manual control must give way to supervisory control. The Mars exploration rovers accordingly offered operators a choice among low-level motor commands, directed driving primitives, and autonomous path selection [21], and on the lower levels human and machine inputs may be shared rather than simply handed over [8].

Adjustable automation has a perceptual and visuospatial counterpart in adjustable immersion. The reality-virtuality continuum [23] spans unmediated reality, augmented reality, and fully virtual environments, and there is no single best point on it for collaborative safety-critical work. Operational control centers today sit almost entirely at the reality end of this continuum. The Mars rovers, for example, have long been commanded from conventional ground workstations [21]. This works because trained operators reconstruct the spatial situation mentally from two-dimensional projections, but that reconstruction consumes the very cognitive resources that situation awareness depends on [7], and its cost grows as the data becomes inherently three-dimensional, as with terrain models, spatial trajectories, and egocentric camera views. Immersive virtual reality can return this reconstruction to perception, improving situation awareness through rich, stereoscopic, depth-faithful views. However, isolating an operator in virtual reality from co-located teammates can undermine workspace awareness and coordination [16]. A practical environment therefore needs several positions available at once: immersive stereoscopic views for spatial tasks, less immersive augmented views that do not isolate operators from the room, and flat two-dimensional overviews on shared displays that support group discussion [16].

The initiative to switch between such levels of control and immersion does not need to rest with the human alone. Under a mixed-initiative paradigm [18], the autonomous partner can itself propose a shift, suggesting a level of control or immersion rather than only responding to one, so that the impulse to move along either dimension may come from either side. In a setting where the question of who decides is central, this two-sided initiative is a notable possibility, and it prompts a further one, namely where such machine initiative could come from. An Intelligent Space, through the local intelligences it can host, is one such place it could originate.

2.3 Transitional Interfaces and Transitional Collaboration

Neither dimension is satisfied by a fixed assignment. In both cases what matters is the ability to move fluidly as a task evolves. For the immersion dimension, this fluid movement along the reality-virtuality continuum is the defining feature of *Transitional Interfaces*, a class of cross-reality systems that let users move freely along the continuum during their work [3, 11]. A central design concern for such interfaces is preserving *interaction continuity* across the transitions, since a change of device and interface paradigm can otherwise break the user's activity [5]. Grasset et al. [11] already noted the collaborative aspect of transitional interfaces, and the difficulty of establishing and maintaining awareness when collaborators occupy different points on the continuum has since been studied directly. Schröder et al. [30] examined dyadic collaboration in a Transitional Interface across desktop, augmented, and virtual reality, and found that workspace awareness can be established across such differences but cannot be taken for granted, and that teams sometimes avoid immersion differences altogether at a cost to performance. The requirement is analogous to a well-known observation in computer-supported cooperative work, namely that teams need to move seamlessly between coupling styles, between phases of tightly coupled work on a shared object and phases of loosely coupled work carried out individually in parallel [25].

3 Designing BRiDGE as an Intelligent Space for Collaboration

To do research and development on interactive systems for safety-critical applications like these, software prototypes and usability labs are not enough on their own. It takes a research environment that supports a wide range of collaboration, interaction, visualization, and sensing techniques in one place. BRiDGE is such an environment. As introduced above, we treat it as an *Intelligent Space*, an interactive space in the sense of Blended Interaction, augmented with the automation and local intelligences that partial autonomy, as in MIDA-VIC, requires [17, 20].

Research environments of this kind have a tradition. Interactive spaces such as i-LAND with its Roomware components [31] and the Stanford Interactive Workspaces [19] pioneered the integration of shared and personal devices for co-located collaboration, later rooms such as the NiCE discussion room [13] refined it for group meetings, and Benyon and Mival [2] carried the idea into Blended Spaces. What these environments share with BRiDGE is the conviction that support for collaboration begins with the room itself, but what none of them provides is the combination argued for above, variable automation and variable immersion with seamless transitions between them, which is the gap BRiDGE is built to fill.

We used Blended Interaction [17] as the framework for its design. Blended Interaction explains what makes some interactive systems feel natural and effortless to use, accounting for this in terms of how well a system blends concepts people already know from the real world with the expressive power of digital computation, which is exactly the property we need here. With so many techniques in one room, the technology must not become a burden but should keep the team's resources on the task itself, and Blended Interaction with its four design domains (1. individual interaction, 2. social interaction and communication, 3. the physical environment, and 4. workflow) gives us a vocabulary for designing toward that.

Methodologically, this is conceptual design work rather than an empirical study, so there was no data collection or analysis. We distilled the lines of work in Section 2 into what the platform must provide and made our decisions within the four design domains as the room took shape. Section 4 retraces those decisions along the same domains to make them explicit and open to critique. The laboratory described here is an early iteration of that design, adjusted during construction and first research use. We have not yet involved domain experts in the process and we return to this limitation in Section 5.

BRiDGE occupies a dedicated laboratory of eight by seven meters, designed for up to four collaboration partners working in parallel (Figure 2). We describe its infrastructure in four parts: the computing substrate that ties it together, the sensing that lets it perceive activity, the interaction and output modalities through which people work and the room responds, and the observer space from which studies are run.

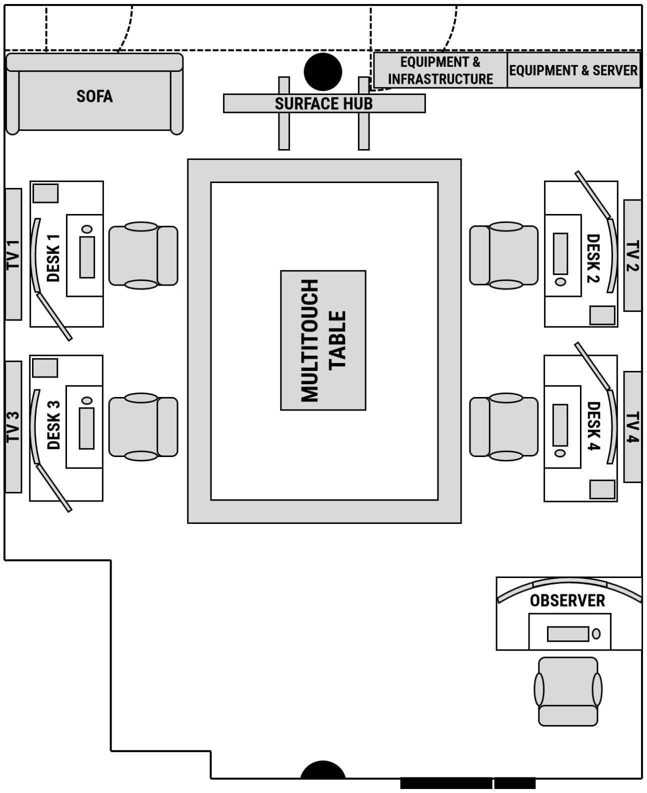

Figure 2: Layout of the laboratory. Four symmetric workstations surround the shared multitouch table and board. The observer space sits apart, near the entrance on the bottom right side.

3.1 Computing Substrate

At the center of BRiDGE is an on-premise server running Proxmox[1], a powerful virtualization platform for servers whose services are reachable through a web interface and an API. On it, virtual machines and containers can easily be created, started, and shut down quickly and in isolation, so new prototypes, such as an implementation of the MIDA-VIC teleoperation scenario, can be brought up and torn down again without disturbing the rest. The server also hosts the room's core services. An Ollama[2] and WhisperLive[3] instance runs large language models locally and powers, among other things, the room's voice assistant and the live transcription of speech. A Home Assistant[4] instance brings the room's many proprietary device APIs under a single interface, so that an assistant or a prototype has one place from which to control lighting, sound, displays, and blinds, and through which the Proxmox virtual machines themselves can be powered up and down. Application-specific servers run alongside these, for example to host the MIDA-VIC teleoperation scenario, to synchronize other large cross-device prototypes, to hold the data behind an information-visualization application, or to provide skills the voice assistant can call. A Samba service supplies a Windows domain and DNS.

The room is networked to keep this traffic local and low-latency. The university uplink enters the server and a router that provides the gateway, DHCP, and the laboratory's wireless network. The router and the Proxmox server both connect to a central laboratory switch, from which every wired device in the room is patched. Head-mounted displays and tablets connect wirelessly.

3.2 Sensing

BRiDGE perceives what happens in it through several complementary sensing systems, which serve both the running collaboration and the later study of it. Spoken activity is captured by four ceiling-mounted microphone arrays, ReSpeaker XMOS XVF-3000 units[5] whose four microphones each let the array estimate the direction of arrival of a sound and whose on-board signal processing improves speech quality. The arrays are attached to a Raspberry Pi[6] and made available as a network audio stream that the voice assistant, the live transcription service, and the observer space can subscribe to. The room also sees through the webcams at the workstations and through Power-over-Ethernet cameras placed around it (Reolink[7]), among them a dome camera that captures the whole laboratory from above in a 360-degree wide angle and can be zoomed in to recover detail.

Spatial position is captured by an OptiTrack[8] optical motion-capture system, whose cameras span their own tracking volume over the room on a dedicated subnet, complemented by lighthouse-based tracking (SteamVR[9]) and the shared spatial anchors of the head-mounted displays. Together these put people, head-mounted displays, tablets, and tagged objects into one common frame of reference. This is also what lets the room record where people are and how they move, which point each person occupies on the reality-virtuality continuum and when they move along it, what their displays show, and, in augmented and virtual reality, where the users are looking, alongside the localized audio and its transcription.

3.3 Interaction and Output Modalities

People work through, and the room responds through, a set of shared and personal devices. Around the room are four workstations, each with a custom-built desktop PC, several monitors, and behind them a ring of wall-mounted televisions. At the center stand the two shared surfaces, a horizontal multi-touch table[10] and a large vertical multi-touch board[11]. Each collaborator also has a personal 8.7 inch tablet PC[12] and a personal Meta Quest 3 head-mounted display[13], and a sofa gives the team a place to step back. Beyond its screens, the room can also act on its occupants through output that is not tied to a display, namely roughly 50 meters of Philips

[1] https://www.proxmox.com
[2] https://ollama.com
[3] https://github.com/hwdsl2/docker-whisper-live
[4] https://www.home-assistant.io

[5] https://wiki.seeedstudio.com/ReSpeaker_Mic_Array_v2.0/
[6] https://www.raspberrypi.com
[7] https://reolink.com
[8] https://www.optitrack.com
[9] https://www.steamvr.com
[10] https://www.interactive-scape.com/de/produkte/scape-movable
[11] https://learn.microsoft.com/de-de/surface-hub/surface-hub-technical-84
[12] https://www.samsung.com/de/support/model/SM-X110NZAEEUB
[13] https://www.meta.com/quest/quest-3

Figure 3: The observer space, screened off near the entrance, with three curved displays for live observation and recording of device streams, camera feeds, and audio.

Hue[14] addressable RGB lighting strips, a 7.1 surround sound system, and motorized blinds that can darken the room completely, all reachable through the Home Assistant interface so a prototype can drive them. We return to how all of these are used in Section 4.

3.4 Observer Space

Set apart from the symmetric working area is the observer space, a small, screened-off area near the entrance from which a study can be watched without intruding on the participants (Figure 3). Three large curved displays give ample room for live observation and recording. From here every device in the room, whether a tablet, a head-mounted display, or a desktop, can be streamed and recorded through OBS Studio[15], together with the camera feeds and the ceiling microphone stream. The room's local language models can transcribe that audio live, using a model such as Faster Whisper[16], so recordings are stored already transcribed.

4 BRiDGE Through the Four Domains of Blended Interaction

Having described the laboratory, we now work through its design along the four design domains of Blended Interaction [17]. (1.) *Individual interaction*: the ability of one person to manipulate an artefact, the basis of all group activity. (2.) *Social interaction and communication*: the coordination of that activity without cutting against established social norms and territoriality. (3.) *The physical environment*: the room, its furniture, and the form factors of its devices. (4.) *Workflow*: the phases and intended outcomes that structure the rest. Unlike the original framework [17], we take workflow last, because in a safety-critical setting it integrates what the other three contribute. We draw concrete examples, where they help, from the lunar teleoperation of MIDA-VIC [16].

4.1 Individual Interaction

Individual interaction is where every collaborative activity begins, in one person's ability to take hold of an artefact and manipulate it [17]. In BRiDGE this happens in personal workspaces, where a collaborator can develop their own ideas, explorations, and partial results before bringing them to the group, and it is what lets a team move between loosely and tightly coupled work rather than having to act together at all times [25]. A personal workspace can be a whole workstation (Figure 4), a region of a shared surface that a collaborator treats as their own corner, or the private view inside a head-mounted display, and augmented and virtual reality let a person set up such a space ad hoc anywhere in the room.

[14]https://www.philips-hue.com
[15]https://obsproject.com
[16]https://github.com/SYSTRAN/faster-whisper

Figure 4: Individual interaction: each participant has a full personal workstation and chooses their own position on the reality-virtuality continuum.

The device we lean on most for individual interaction is the personal tablet, used as the same input instrument across the whole reality-virtuality continuum. Keeping one familiar instrument preserves interaction continuity across the transitions, so a person does not have to relearn how to act each time they change their level of immersion [5]. The tablet works as an ordinary touch device in reality, is enriched with additional digital artefacts when seen through augmented reality, and still takes precise touch input in virtual reality, where its own display, unseen by the immersed user, can show information for bystanders, such as a sign that the person is currently in VR. Because the tablet is tracked by the motion-capture system, its position and orientation can also drive proxemic interactions [1], where what a device offers depends on where it is, which way it faces, and who is holding it. The tablet is thus the default instrument rather than the only one: the Quest controllers remain available where they suit a task better, and personal tangible objects placed on the multi-touch table give each collaborator an identifiable physical handle there if needed.

Figure 5: Social interaction: the shared table and interactive board act as collaboration anchors where operators coordinate face to face.

4.2 Social Interaction and Communication

Collaboration needs the individual activity of the previous domain to be coordinated, and it must be coordinated without cutting against the social norms and territoriality people already rely on [17]. Two shared surfaces support this from different sides. The horizontal multi-touch table lets several people work at it at once and acts as shared group territory, where a shared overview can move a group toward more balanced decisions [4], but content on it has no single up, so its orientation differs for everyone around it. The vertical board has the opposite profile. It gives everyone the same orientation on what it shows and a spatial context for conversation, with information referenced by where it sits on the surface [15], but only a few people can stand at and act on it at a time, and a person working at it tends to leave the others observing [10, 27]. We therefore use the two for different kinds of work, the table for shared, strategic interaction and the board for dashboards and organizational coordination (Figure 5). This is our own division rather than a finding from the literature, though it follows from these orientation and visibility properties.

Shared surfaces of this kind help when the people around them occupy a common sensory space. In a transitional interface that is often not the case. Two collaborators can be physically co-located in the same room and still inhabit two entirely different virtual places and realities, which are not automatically connected to one another. They can be not in the same virtual location, and someone immersed in virtual reality perceives nothing of the room, or of what the others are doing, on their own. Physical proximity does not bring sensory proximity with it [29], so social interaction here depends on explicit awareness cues that rebuild a shared sense of the workspace across the continuum. Earlier work in transitional collaboration designed embodiments for each collaborator. These can show their identity through a color, their location and which context they are in, and the direction of their attention through a visualized gaze cone in virtual reality or a gaze cursor in augmented reality [12, 30]. In the lunar teleoperation of MIDA-VIC, when a health alert and a promising science target arise at once, the operators concerned can gather at the table and board and read each other's focus from these cues even when one of them is working in VR.

Figure 6: Physical environment: a configurable room whose lighting, sound, and blinds are programmable, making controlled study conditions possible. Here the motorized blinds are fully lowered for a simulated shutdown, and blue and orange accent lighting distinguish two teams.

4.3 Physical Environment

The room is the physical environment in which all of this takes place, and its layout is a deliberate design choice [17]. By default it is symmetric around its four workstations (Figure 2). The symmetry gives every collaborator equal access to the shared surfaces, so that no one becomes the de facto owner of the table or board simply by sitting closest to it while others have to cross the room. We believe that this lets collaboration unfold as it naturally would, except where a study deliberately arranges otherwise. Apart from the server rack and the observer station, the furniture is movable and quick to reconnect, so the room can be rearranged from one study to the next. Placing the horizontal table next to the vertical board follows the original Blended Interaction example [17], pairing a surface everyone can reach with one everyone can see.

The physical environment is also instrumented as a variable. Because the room's lighting, sound, and blinds can be driven programmatically (Figure 6), conditions that ordinarily make safety-critical work hard can be introduced deliberately and under control (e.g. ambient noise, sudden loss of light, or simulated shutdown), so that collaboration and decision-making can be studied under something closer to real pressure. A room that hosts the work of a safety-critical team can also reproduce the conditions that strain it.

4.4 Workflow

Individual and social interaction do not happen in a vacuum. In practice they are structured by a workflow, the phases of an operation and its intended outcomes [17], and in a safety-critical setting this is not a detail but the substance of the work, since procedures and coordinated actions are what such teams live by. Workflow is therefore the domain that pulls the others together. The personal workspaces of the individual domain become the places where steps are carried out, the shared surfaces and awareness cues of the social domain make the state of the work legible to everyone.

The vertical board carries much of this as a shared collaboration cockpit, working in two modes. As a dashboard, it keeps the operation's common operational picture in one continuously visible place. Such a shared status display is known to support synchronous, tightly coupled coordination [32]. It can show live views of the team with who is doing what, where each person currently is on the reality-virtuality continuum, and what their devices are showing. As a whiteboard, it turns into a shared surface the team contributes to directly. Each collaborator can send screenshots or configurations of their own view to it and the team can then annotate, group, and discuss these together. A few lightweight coordination mechanisms work alongside this. One person can delegate a data set to another, request the right to present to the group, or, in an escalating situation, pull everyone back to reality at once.

The room's peripheral and ambient output supports the same end more quietly. Addressable light strips behind the displays and along the walls can mark a change of state in the background or escalate visibly when the whole team must be called together, a use of room lighting with a direct precedent in Blended Interaction, where Benyon and Mival [2] used a color wash to signal a room's state. The surround sound system adds spatially placed auditory cues in the tradition of Audio Aura [24] and of the spatial earcons used in transitional collaboration [30], and the ring of wall-mounted televisions can share status with everyone at once. Such peripheral and ambient displays [22, 33, 34] let a team stay aware of the state of an operation without spending on it the attention the primary task needs.

Blended Interaction frames workflow as a tension between enforcing a fixed structure and staying open to unanticipated use [17]. BRiDGE is built for both. The dashboard and its procedures can hold the fixed sequence a safety-critical operation depends on, while the whiteboard keeps room for the open sensemaking that precedes a decision and can rarely be laid out in advance. This places the laboratory in a long line of computer-supported cooperative work on control rooms, going back to the coordination of activity across shared displays in transport control [14].

5 Conclusion, Limitations, and Future Work

We set out from the question of what an environment must provide, and how it should be designed, so that collaborative situation awareness can be established, maintained, and studied across variable levels of control and immersion. We have argued that some safety-critical domains are moving toward collaboration that combines variable levels of control and immersion with seamless transitions between them, that providing these together is first of all a question of a carefully designed space, and that Blended Interaction offers a principled way to design one. We have shared our view of such a space, BRiDGE, as an integrated research environment for safety-critical teams, and worked through its design along the four design domains of Blended Interaction. It also speaks to the question we raised at the outset, who is ultimately responsible for a decision made in such a space. In a setting like this, the decision remains the team's. The room and its local intelligences can support the team, by helping it build and hold a shared picture of the situation and by giving automation a spatial form in the room rather than a separate window, but that support reaches only as far as it does not draw attention from the task itself. Sensemaking can be distributed, with the room and its automated parts as genuine partners, while decision authority and accountability stay with an identifiable person. A machine can take part in a decision but cannot carry responsibility for it. This is the substance of meaningful human control [28], with autonomy in acquiring and analyzing information but not in the selection of consequential actions [26].

As a work-in-progress position paper, this contribution has clear limits. It reports no empirical evaluation, so its claims about the benefits of the design are arguments rather than findings, and the MIDA-VIC paradigm it draws on is conceptual rather than a built system. The design also did not involve domain experts. Empirical work with expert users, both to uncover requirements we may have missed and to evaluate the environment in use, is a natural next step. Whether the account transfers cleanly to other safety-critical domains is plausible but not yet shown.

Our future work proceeds along two lines. The first is the concrete realization of the design possibilities described here as working prototypes within BRiDGE, targeting application domains such as the optimization of airspaces, the detection of hybrid threats, and the maintenance of maritime situational awareness, each chosen to bring out a different facet of the advantages argued for here. The second is the interpretation of the laboratory's integrated sensor data into an understanding of the collaboration itself, so that the space can recognize the state of the work in progress. The longer-term vision that motivates BRiDGE is a space whose components are integrated tightly enough that artificial intelligence can fold new data and new devices into the collaborative process on its own and offer concrete, situated suggestions, so that automation becomes part of the room rather than a chat window beside it, in the spirit of the mixed-initiative collaboration that MIDA-VIC already anticipates.

References

[1] Till Ballendat, Nicolai Marquardt, and Saul Greenberg. 2010. Proxemic Interaction: Designing for a Proximity and Orientation-Aware Environment. In *ACM International Conference on Interactive Tabletops and Surfaces (ITS '10)*. Association for Computing Machinery, New York, NY, USA, 121–130. https://doi.org/10.1145/1936652.1936676

[2] David Benyon and Oli Mival. 2015. Blended Spaces for Collaboration. *Computer Supported Cooperative Work (CSCW)* 24, 2-3 (June 2015), 223–249. https://doi.org/10.1007/s10606-015-9223-8

[3] Mark Billinghurst, Hirokazu Kato, and Ivan Poupyrev. 2001. The magic book: A transitional AR interface. *Computers & Graphics* 25, 5 (2001), 745–753. https://doi.org/10.1016/S0097-8493(01)00117-0

[4] Frederik Brudy, Joshua Kevin Budiman, Steven Houben, and Nicolai Marquardt. 2018. Investigating the Role of an Overview Device in Multi-Device Collaboration. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems (CHI '18)*. Association for Computing Machinery, New York, NY, USA, 1–13. https://doi.org/10.1145/3173574.3173874

[5] Felipe G. Carvalho, Daniela G. Trevisan, and Alberto Raposo. 2012. Toward the design of transitional interfaces: an exploratory study on a semi-immersive hybrid user interface. *Virtual Reality* 16, 4 (Nov. 2012), 271–288. https://doi.org/10.1007/s10055-011-0205-y

[6] Mica R. Endsley. 1995. Toward a Theory of Situation Awareness in Dynamic Systems. *Human Factors* 37, 1 (1995), 32–64. https://doi.org/10.1518/001872095779049543

[7] Mica R. Endsley. 2011. *Designing for Situation Awareness: An Approach to User-Centered Design, Second Edition* (2nd ed.). CRC Press, Inc., USA.

[8] F. Flemisch, D. A. Abbink, M. Itoh, M.-P. Pacaux-Lemoine, and G. Weßel. 2019. Joining the blunt and the pointy end of the spear: towards a common framework of joint action, human–machine cooperation, cooperative guidance and control,

shared, traded and supervisory control. *Cognition, Technology & Work* 21, 4 (Nov. 2019), 555–568. https://doi.org/10.1007/s10111-019-00576-1

[9] Frank Flemisch, Matthias Heesen, Tobias Hesse, Johann Kelsch, Anna Schieben, and Johannes Beller. 2012. Towards a dynamic balance between humans and automation: authority, ability, responsibility and control in shared and cooperative control situations. *Cognition, Technology & Work* 14, 1 (March 2012), 3–18. https://doi.org/10.1007/s10111-011-0191-6

[10] Florian Geyer, Ulrike Pfeil, Anita Höchtl, Jochen Budzinski, and Harald Reiterer. 2011. Designing reality-based interfaces for creative group work. In *Proceedings of the 8th ACM conference on Creativity and cognition*. ACM, Atlanta Georgia USA, 165–174. https://doi.org/10.1145/2069618.2069647

[11] Raphael Grasset, Julian Looser, and Mark Billinghurst. 2006. Transitional interface: concept, issues and framework. In *Proceedings of the 5th IEEE/ACM International Symposium on Mixed and Augmented Reality (ISMAR '06)*. IEEE Computer Society, USA, 231–232. https://doi.org/10.1109/ismar.2006.297819

[12] Carl Gutwin and Saul Greenberg. 2002. A Descriptive Framework of Workspace Awareness for Real-Time Groupware. *Computer Supported Cooperative Work (CSCW)* 11, 3-4 (Sept. 2002), 411–446. https://doi.org/10.1023/A:1021271517844

[13] Michael Haller, Jakob Leitner, Thomas Seifried, James R. Wallace, Stacey D. Scott, Christoph Richter, Peter Brandl, Adam Gokcezade, and Seth Hunter. 2010. The NiCE Discussion Room: Integrating Paper and Digital Media to Support Co-Located Group Meetings. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '10)*. Association for Computing Machinery, New York, NY, USA, 609–618. https://doi.org/10.1145/1753326.1753418

[14] Christian Heath and Paul Luff. 1992. Collaboration and Control: Crisis management and multimedia technology in London Underground Line Control Rooms. *Computer Supported Cooperative Work (CSCW)* 1, 1-2 (March 1992), 69–94. https://doi.org/10.1007/BF00752451

[15] Mikkel R. Jakobsen and Kasper Hornbæk. 2014. Up close and personal: Collaborative work on a high-resolution multitouch wall display. *ACM Transactions on Computer-Human Interaction* 21, 2 (Feb. 2014), 1–34. https://doi.org/10.1145/2576099

[16] Hans-Christian Jetter, Leon Raule, Jens Gerken, and Sören Pirk. 2025. Mixed-Initiative Dynamic Autonomy Through Variable Levels of Immersion and Control (MIDA-VIC): A New Paradigm for Collaborative Robotic Teleoperation in Space Exploration. In *Advancing Human-Computer Interaction for Space Exploration, SpaceCHI 2025, European Astronaut Centre (EAC), Cologne, Germany, June 23-24, 2025 (OASIcs)*, Leonie Bensch, Tommy Nilsson, Martin Nisser, Pat Pataranutaporn, Albrecht Schmidt, and Valentina Sumini (Eds.). Schloss Dagstuhl - Leibniz-Zentrum für Informatik, Dagstuhl, Germany, 22:1–22:10. https://doi.org/10.4230/oasics.spacechi.2025.22

[17] Hans-Christian Jetter, Harald Reiterer, and Florian Geyer. 2014. Blended Interaction: understanding natural human—computer interaction in post-WIMP interactive spaces. *Personal Ubiquitous Comput.* 18, 5 (2014), 1139–1158. https://doi.org/10.1007/s00779-013-0725-4

[18] Shu Jiang and Ronald C. Arkin. 2015. Mixed-Initiative Human-Robot Interaction: Definition, Taxonomy, and Survey. In *2015 IEEE International Conference on Systems, Man, and Cybernetics*. IEEE, Kowloon Tong, Hong Kong, 954–961. https://doi.org/10.1109/SMC.2015.174

[19] Brad Johanson, Armando Fox, and Terry Winograd. 2002. The Interactive Workspaces Project: Experiences with Ubiquitous Computing Rooms. *IEEE Pervasive Computing* 1, 2 (2002), 67–74. https://doi.org/10.1109/MPRV.2002.1012339

[20] Joo-Ho Lee and Hideki Hashimoto. 2002. Intelligent Space — concept and contents. *Advanced Robotics* 16, 3 (2002), 265–280. https://doi.org/10.1163/156855302760121936

[21] Mark Maimone, Jeffrey Biesiadecki, Edward Tunstel, Yang Cheng, and Chris Leger. 2006. Surface navigation and mobility intelligence on the Mars Exploration Rovers. In *Intelligence for Space Robotics*. TSI Press, San Antonio, TX, USA, 45–69.

[22] Tara Matthews, Anind K. Dey, Jennifer Mankoff, Scott Carter, and Tye Rattenbury. 2004. A toolkit for managing user attention in peripheral displays. In *Proceedings of the 17th Annual ACM Symposium on User Interface Software and Technology* (Santa Fe, NM, USA) *(UIST '04)*. Association for Computing Machinery, New York, NY, USA, 247–256. https://doi.org/10.1145/1029632.1029676

[23] Paul Milgram and Fumio Kishino. 1994. A taxonomy of mixed reality visual displays. *IEICE TRANSACTIONS on Information and Systems* 77, 12 (1994), 1321–1329.

[24] Elizabeth D. Mynatt, Maribeth Back, Roy Want, Michael Baer, and Jason B. Ellis. 1998. Designing audio aura. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Los Angeles, California, USA) *(CHI '98)*. ACM Press/Addison-Wesley Publishing Co., USA, 566–573. https://doi.org/10.1145/274644.274720

[25] Thomas Neumayr, Hans-Christian Jetter, Mirjam Augstein, Judith Friedl, and Thomas Luger. 2018. Domino: A Descriptive Framework for Hybrid Collaboration and Coupling Styles in Partially Distributed Teams. *Proc. ACM Hum.-Comput. Interact.* 2, CSCW (2018), 128:1–128:24. https://doi.org/10.1145/3274397

[26] R. Parasuraman, T.B. Sheridan, and C.D. Wickens. 2000. A model for types and levels of human interaction with automation. *IEEE Transactions on Systems, Man, and Cybernetics - Part A: Systems and Humans* 30, 3 (2000), 286–297. https://doi.org/10.1109/3468.844354

[27] Yvonne Rogers and Siân Lindley. 2004. Collaborating around vertical and horizontal large interactive displays: which way is best? *Interacting with Computers* 16, 6 (Dec. 2004), 1133–1152. https://doi.org/10.1016/j.intcom.2004.07.008

[28] Filippo Santoni De Sio and Jeroen Van Den Hoven. 2018. Meaningful Human Control over Autonomous Systems: A Philosophical Account. *Frontiers in Robotics and AI* 5 (Feb. 2018), 15. https://doi.org/10.3389/frobt.2018.00015

[29] Jan-Henrik Schröder and Hans-Christian Jetter. 2023. Towards a Model for Space and Time in Transitional Collaboration. In *IEEE International Symposium on Mixed and Augmented Reality Adjunct, ISMAR 2023, Sydney, Australia, October 16-20, 2023*, Gerd Bruder, Anne-Hélène Olivier, Andrew Cunningham, Evan Yifan Peng, Jens Grubert, and Ian Williams (Eds.). IEEE, Sydney, Australia, 223–227. https://doi.org/10.1109/ismar-adjunct60411.2023.00052

[30] Jan-Henrik Schröder, Daniel Schacht, Niklas Peper, Anita Marie Hamurculu, and Hans-Christian Jetter. 2023. Collaborating Across Realities: Analytical Lenses for Understanding Dyadic Collaboration in Transitional Interfaces. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems (CHI '23)*, Albrecht Schmidt, Kaisa Väänänen, Tesh Goyal, Per Ola Kristensson, Anicia Peters, Stefanie Mueller, Julie R. Williamson, and Max L. Wilson (Eds.). Association for Computing Machinery, New York, NY, USA, 97:1–97:16. https://doi.org/10.1145/3544548.3580879

[31] Norbert A. Streitz, Jörg Geißler, Torsten Holmer, Shin'ichi Konomi, Christian Müller-Tomfelde, Wolfgang Reischl, Petra Rexroth, Peter Seitz, and Ralf Steinmetz. 1999. i-LAND: An Interactive Landscape for Creativity and Innovation. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '99)*. Association for Computing Machinery, New York, NY, USA, 120–127. https://doi.org/10.1145/302979.303010

[32] James R. Wallace, Stacey D. Scott, Eugene Lai, and Deon Jajalla. 2011. Investigating the Role of a Large, Shared Display in Multi-Display Environments. *Computer Supported Cooperative Work (CSCW)* 20, 6 (Dec. 2011), 529–561. https://doi.org/10.1007/s10606-011-9149-8

[33] Mark Weiser and John Seely Brown. 1996. Designing Calm Technology. *PowerGrid Journal* 1, 1 (1996), 75–85. https://ubiq.com

[34] Craig Wisneski, Hiroshi Ishii, Andrew Dahley, Matt Gorbet, Scott Brave, Brygg Ullmer, and Paul Yarin. 1998. Ambient Displays: Turning Architectural Space into an Interface between People and Digital Information. In *Cooperative Buildings: Integrating Information, Organization, and Architecture*, Gerhard Goos, Juris Hartmanis, Jan Van Leeuwen, Norbert A. Streitz, Shin'ichi Konomi, and Heinz-Jürgen Burkhardt (Eds.). Vol. 1370. Springer Berlin Heidelberg, Berlin, Heidelberg, 22–32. https://doi.org/10.1007/3-540-69706-3_4 Series Title: Lecture Notes in Computer Science.