# Two Heads Are Better Than One: A Dimension Space for Unifying Human and Artificial Intelligence in Shared Control

Gabriele Cimolino
gabriele.cimolino@queensu.ca
School of Computing, Queen's University
Kingston, Ontario, Canada

T.C. Nicholas Graham
nicholas.graham@queensu.ca
School of Computing, Queen's University
Kingston, Ontario, Canada

## ABSTRACT

*Shared control* is an emerging interaction paradigm in which a human and an AI partner collaboratively control a system. Shared control unifies human and artificial intelligence, making the human's interactions with computers more accessible, safe, precise, effective, creative, and playful. This form of interaction has independently emerged in contexts as varied as mobility assistance, driving, surgery, and digital games. These domains each have their own problems, terminology, and design philosophies. Without a common language for describing interactions in shared control, it is difficult for designers working in one domain to share their knowledge with designers working in another. To address this problem, we present a dimension space for shared control, based on a survey of 55 shared control systems from six different problem domains. This design space analysis tool enables designers to classify existing systems, make comparisons between them, identify higher-level design patterns, and imagine solutions to novel problems.

## CCS CONCEPTS

• **Human-centered computing → Systems and tools for interaction design**; **Accessibility**; **Collaborative interaction**.

## KEYWORDS

Shared Control, Human-Machine Cooperation, Design Space Analysis

**ACM Reference Format:**
Gabriele Cimolino and T.C. Nicholas Graham. 2022. Two Heads Are Better Than One: A Dimension Space for Unifying Human and Artificial Intelligence in Shared Control. In *CHI Conference on Human Factors in Computing Systems (CHI '22), April 29-May 5, 2022, New Orleans, LA, USA.* ACM, New York, NY, USA, 21 pages. https://doi.org/10.1145/3491102.3517610

## 1 INTRODUCTION

Human-AI shared control is an emerging interaction paradigm in which systems are controlled through the tightly-coupled cooperation of a human and an artificially intelligent (AI) agent. While the human user controls the system to perform their tasks, the AI provides assistance, making the humans' interactions more accessible, safe, precise, reliable, creative, and playful. For example, digital games use shared control to improve their accessibility to players with motor impairments [13–17, 23, 40, 63, 80]. Teleoperated robots and unmanned aerial vehicles (UAVs) share control to make complicated manipulation tasks easier [19–22, 24, 32, 37, 42, 44, 45, 81]. Surgical robots filter out tremulous movements and guide the surgeon's control of their implements [36, 43, 47, 48, 61, 77, 96, 107]. Creativity support systems make sketching and playing musical instruments easier using motors or electrical muscle stimulation [50, 51, 60, 84, 89, 95, 97, 99, 108]. However, each of these domains has its own assumptions about the roles that human users and the AI agent play, and uses different designs to facilitate these interactions. For example, smart power wheelchairs are typically designed to help people who have difficulty using a joystick to avoid collisions with obstacles in their environments [18, 25, 26, 34, 56–58, 90]. Since the user is unable to perform this task on their own, smart power wheelchair AI supervises the human user and amends their control when needed. In contrast, semi-autonomous vehicles expect the human user to supervise the AI's driving activities and take over control when necessary [29, 33, 46, 53, 91].

When the human and AI cooperate harmoniously, shared control can blur the lines between human and AI action. For example, Kasahara et al. modified a device called *Wired Muscle* [69] that helps humans perform difficult timing tasks using electrical muscle stimulation (EMS) [49]. If the human is tasked with catching a falling ruler, then an EMS forearm strap stimulates their muscles to close their hand with superhuman timing. However, the human might not feel in control if the device reacts before they have noticed the ruler is falling, so it delays its response to maximize the human's sense of agency. This can improve users' reaction timing while allowing them to retain the sense of being fully in control. By carefully designing the human's interactions with the AI, shared control systems can unify human and artificial intelligence, enabling each actor to complement and extend the abilities of the other.

Even though many technologies use shared control, the styles of interaction that they afford can differ vastly across domains. Designers working in one domain may be unaware of how shared control is used by others and may be unaware of solutions that have proven useful for solving similar problems. The terminology frequently differs, making it difficult for design insights discovered in one domain to be transferred to others. In order to understand how commonly used design patterns overcome a problem and to make comparisons between designs, we need a common language for describing the design space of shared control systems. We need tools that help designers to better communicate their ideas about how interactions between human users and AI agents should be


*CHI '22, April 29-May 5, 2022, New Orleans, LA, USA*

ACM ISBN 978-1-4503-9157-3/22/04…$15.00
https://doi.org/10.1145/3491102.3517610

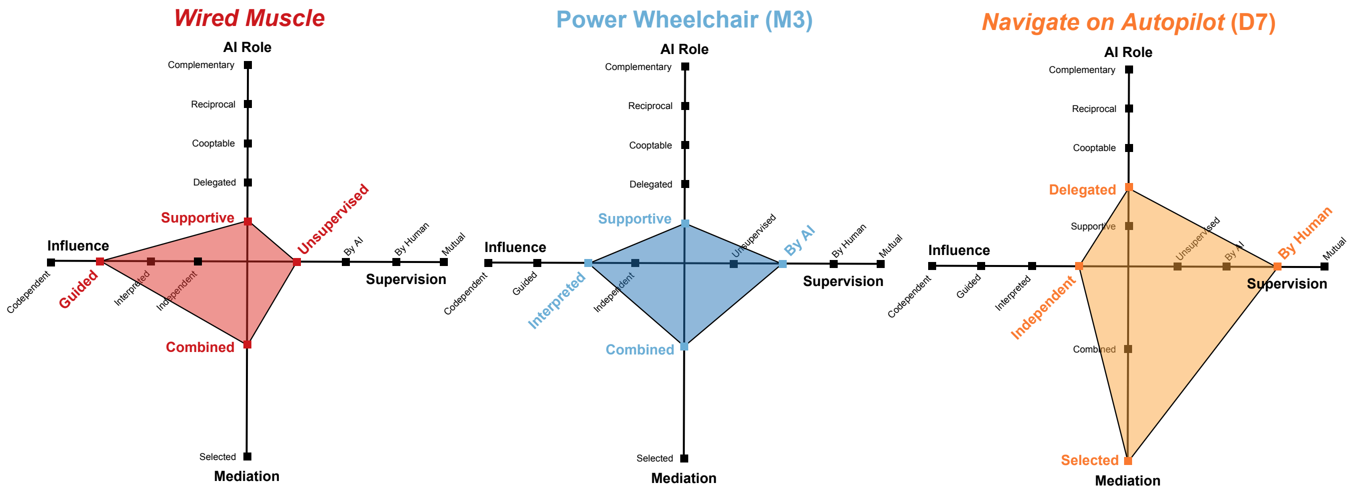


**Figure 1: Kiviat diagrams capturing three shared control systems using our dimension space. Each is plotted in four dimensions. *Wired Muscle*, on the left, supports the human in catching a ruler; the AI supports the human, guiding the human's actions. The M3 power wheelchair supervises the human's movement and helps to avoid collisions. Tesla's *Navigate on Autopilot* enables drivers to delegate driving tasks to the AI under the driver's supervision. The details of the dimension space are presented in Section 4**

.

structured in shared control, to identify gaps in the approaches used in one domain, and take inspiration from approaches used in others.

This paper makes two contributions. We present the first design space analysis tool for classifying human-AI shared control systems. As a secondary contribution, we survey the use of shared control in six problem domains. Our *dimension space* is defined along four axes: *AI Role*, *Supervision*, *Influence*, and *Mediation*. These dimensions were constructed through an inductive process of discovering distinctions between systems' designs that have meaningful consequences for the human-AI interactions they afford, described in Section 3. Figure 1 depicts the classification of three example systems using our dimension space, which is presented in Section 4. In Section 5, we use this tool to identify common design patterns in six domains using shared control (i.e., *games*, *semi-autonomous driving*, *creativity support*, *mobility assistance*, *telerobotics*, and *surgery*), and compare the interactions that these systems afford. We then use the dimension space to identify opportunities for design. In Section 5.2 we identify under-explored gaps in the design of human-AI shared control systems, propose novel designs that leverage existing design knowledge in new ways, and suggest directions for future research. We begin, in the next section, by defining what we mean by shared control, introducing the six domains we have surveyed, and summarizing prior reflections on the design of shared control systems.

# 2 BACKGROUND

It is important to distinguish shared control from other uses of AI agents. For example, when used in a car, Apple's Maps application suggests destinations based on the contents of the user's calendar. This example is *not* shared control, because it is not tightly-coupled, and because the AI agent does not share a control interface with the user; it makes suggestions, but does not actually control the car. Also note that in this paper, we adopt the established use of the term shared control [1, 22, 27, 67, 68, 94] to designate the partnership of human and AI. This term can be confusing to an HCI audience, who might consider human-human cooperation in a shared text editor such as Google Docs to be an example of shared control. In this paper, however, the term shared control refers to the special case of *human-AI* shared control.

In this section, we describe what shared control is, where it is used, and how it can be understood. This motivates the need for tools such as our dimension space that help in understanding the range of possible designs for shared control and suggesting new designs that have not yet been considered. We first provide a definition for shared control. We then introduce six application domains in which shared control is used, explain how shared control overcomes domain-specific problems, and describe how systems from these domains afford different styles of human-AI interaction. Finally, we summarize existing knowledge about the shared control design space and explain how our dimension space addresses gaps in our understanding.

## 2.1 Human-AI Shared Control

Shared control systems are characterized by tightly-coupled collaboration between a human and an AI agent. In this partnership, two *actors*, one human and one AI, control the same system through a *common interface* [22, 67, 68]. For example, power wheelchairs are often controlled by a joystick, which can be difficult to accurately manipulate by persons with fine-motor deficits [81]. In a smart power wheelchair, an AI agent works with the user of the wheelchair, adapting the user's commands to avoid collisions and to better navigate to a desired destination [26]. Both actors control the system directly and in real time. This unifies what the human wants to do and what the AI wants to do into a single command, as per the adage "*two heads are better than one*."

More technically, the power wheelchair might expose an interface that enables users to move *forward* and *backward*, and rotate to the *left* or *right*. These are the system's *inputs* [22, 67, 68], which actors control using a sequence of commands called a control *signal* [67, 68]. A human user of this power wheelchair might use their control signal to make the wheelchair move forward along a sidewalk, turn right at a corner, and move forward again down another street. So when control is shared between the human user and an AI agent, each actor sends a control signal to the interface's inputs, which is used to construct a new *shared control signal*. This creates a tightly-coupled form of cooperation between human and AI that unifies their control of an interactive system. For example, a shared control power wheelchair agent might detect that the human has erroneously commanded their wheelchair to move forward into traffic and instead command the wheelchair to turn right and follow the sidewalk.

This example describes a style of interaction that is typical of power wheelchairs [57, 58] but also applicable to other interactive systems. The AI listens to the human's commands and interprets them to infer what command they really intended. Unlike other systems that *combine* actors' commands, it then *selects* its own command to control the wheelchair when it believes the human to be in error. This same approach is used in digital games to help players with disabilities to aim more precisely [40], in cars to decide when to assist with lane keeping [85], and in remotely controlled robots to decide when it is unsafe to execute the human's commands [44, 45]. This design pattern is similar to the interleaved controller switching of mixed-initiative systems [5, 39], as it provides a style of interaction in which users can expect the AI to control the system on its own. And as we shall see, this is only one of several design patterns used in shared control.

From a technical perspective, all that matters is how control signals are created and how the human and AI's control signals are unified. Architecture or block diagrams and kinematics or dynamics models can tell us a lot about what the AI does and how actors' control signals are put together, but they cannot tell us what it is like to interact with them. For example, the distinction described in the last paragraph, wherein actors' control signals are either *Selected* or *Combined*, corresponds to the Mediation axis of our dimension space (Section 4.4). This simple binary classification enables designers to discuss human-AI interactions in shared control using language, rather than systems of equations. It enables us to reason about the design of shared control systems more abstractly, to compare available design choices, and to explain how low-level design choices afford different interactions. Therefore, insights into the design of shared control surgical robots, for example, can inform the design of sketching assistants, and vice versa. In the rest of this section, we describe how different domains use shared control and summarize existing high-level design knowledge.

## 2.2 Shared Control Domains

In our search for dimensions of variability that have meaningful consequences for human users' interactions in shared control, we examined 55 systems and categorized them according to their intended uses. We found that shared control is used in playing digital games, driving motor vehicles, creating artifacts or performances, getting around using mobility aids, operating robots at a distance, and performing surgical operations. Others have identified similar categories (e.g., {*automotive*, *robot-assisted surgery*, *brain-machine interfaces*, *learning*} [1] or {*robotic wheelchairs*, *telerobotics*, *autopilots*, *intelligent vehicles*, *autonomous manufacturing*} [94]). In this section, we describe six domains using shared control, what they have in common with other domains, and what properties distinguish these domains from the others. Systems referenced in the rest of this paper can be found in Table 1 according to an ID tag that we have assigned; for example, G1 is the game *Zac - O Esquilo* [63].

*2.2.1 Digital Games.* *Digital games* (henceforth simply "games") have become an enormously popular pastime. However, some players are disadvantaged or even excluded from the play of games due to differences in physical ability [109]. Shared control is used in some games to make them accessible and to level the playing field between persons with different abilities [14].

There are many ways that games can be made more accessible to players with motor disabilities [109], and among them *player balancing* and *input automation* may use shared control [14]. Player balancing helps weaker players compete with stronger players [10]; for example, many first-person shooter games designed for home consoles provide *aim-assist* to make aiming with an analog stick easier [103]. The AI in *Gekku Aim* (G6) aims directly at the closest opponent when the player shoots, stepping in at the last moment to override the player's command.

In contrast, steering assistance in *Mario Kart 8 Deluxe* (G5) steers the player's kart when they are in danger of driving off course. This enables disabled players to overcome gameplay challenges that would be too difficult otherwise. Alternatively, disabled players who are unable to control some aspects of play at all can be assisted using input automation. In this approach, an AI agent automates control of inaccessible game inputs, enabling players to play using whatever inputs they can control. For example, acceleration automation in *Mario Kart 8 Deluxe* makes the player's kart accelerate constantly, as though the player is holding down the "A" button.

Assisting human users who have difficulty controlling a system is common to all six of the domains we surveyed, although automating control of inaccessible inputs is unique to games. Most games using input automation are classified as *one-switch games*—games designed to be played using a single button or switch [109]—and were developed to make existing games more accessible to players with motor disabilities. For example, *Alienated* (G2) is a one-switch clone of the popular arcade game *Space Invaders*. An AI agent automates control of shooting lasers at enemies and pressing the button makes the player's avatar move right, if it was moving left, or left, if it was moving right. In this way, shared control overcomes a pervasive problem encountered by players with motor disabilities, extending access to the fully-featured games enjoyed by their non-disabled peers [14].

*2.2.2 Semi-Autonomous Driving.* Fully autonomous vehicles are fast approaching, but in the meantime drivers have been invited to share control with driving automation systems. Here, both the human and AI control the car through its interface of accelerating,

**Table 1: Our corpus of shared control systems according to their classifications in our dimension space. More intense colors indicate greater extension along an axis. Italicized names are those designated by the system's designers.**

| Domain | ID | Name | AI Role | Supervision | Influence | Mediation |
|---|---|---|---|---|---|---|
| Games | G1 | *Zac - O Esquilo* [63] | Complementary | Unsupervised | Independent | Combined |
| | G2 | *Alienated* [80] | Complementary | Unsupervised | Independent | Combined |
| | G3 | *Partial Automation* [14] | Complementary | Unsupervised | Independent | Combined |
| | G4 | *Imaginary Pong for One* [59] | Complementary | Unsupervised | Guided | Combined |
| | G5 | *Mario Kart 8 Deluxe* [23] | Delegated | By AI | Independent | Combined |
| | G6 | *Gekku Aim* [40] | Supportive | By AI | Interpreted | Selected |
| | G7 | Racing game [13] | Supportive | By AI | Interpreted | Combined |
| | G8 | Missile Command clone [15–17] | Cooptable | By Human | Independent | Selected |
| Semi-Autonomous Driving | D1 | Lane-keeping assistant [93] | Supportive | Unsupervised | Guided | Combined |
| | D2 | Lane-keeping assistant [91] | Supportive | By Human | Guided | Combined |
| | D3 | Lane-keeping assistant [46] | Supportive | Mutual | Codependent | Combined |
| | D4 | *Lane Keeping Aid* [104] | Supportive | Mutual | Codependent | Combined |
| | D5 | Lane-keeping assistant [85] | Supportive | By AI | Interpreted | Selected |
| | D6 | Automated driving system [53] | Delegated | By Human | Independent | Selected |
| | D7 | *Navigate on Autopilot* [98] | Delegated | By Human | Independent | Selected |
| | D8 | *Hotzenplotz Interface* [33] | Cooptable | Unsupervised | Independent | Selected |
| | D9 | H-Mode car [29] | Cooptable | By Human | Codependent | Combined |
| Creativity Support | C1 | *PossessedHand* [95] | Supportive | Unsupervised | Guided | Combined |
| | C2 | *EMS Air Guitar* [97] | Supportive | Unsupervised | Guided | Combined |
| | C3 | *dePENd* [108] | Supportive | By AI | Guided | Combined |
| | C4 | *Haptic Intelligentsia* [99] | Supportive | By AI | Guided | Combined |
| | C5 | *FreeD* [110] | Supportive | By AI | Guided | Combined |
| | C6 | Sketching assistant [50] | Supportive | Mutual | Guided | Combined |
| | C7 | Image classifier [84] | Cooptable | By Human | Guided | Selected |
| | C8 | *Tanagra* [89] | Delegated | By Human | Independent | Selected |
| | C9 | *Muscle-Plotter* [60] | Delegated | Unsupervised | Independent | Combined |
| | C10 | *TransPen & MimeoPad* [51] | Complementary | Unsupervised | Independent | Combined |
| | C11 | *Origin* [83] | Complementary | Unsupervised | Independent | Combined |
| | C12 | Digital Airbrush [88] | Complementary | By AI | Codependent | Combined |
| Mobility Assistance | M1 | Mobility robot [105] | Supportive | By AI | Interpreted | Combined |
| | M2 | Power wheelchair [56] | Supportive | By AI | Interpreted | Combined |
| | M3 | Power wheelchair [18] | Supportive | By AI | Interpreted | Combined |
| | M4 | Power wheelchair [26] | Supportive | By AI | Interpreted | Combined |
| | M5 | Power wheelchair [57] | Supportive | By AI | Interpreted | Selected |
| | M6 | Power wheelchair [58] | Supportive | By AI | Interpreted | Selected |
| | M7 | Power wheelchair [90] | Supportive | By AI | Independent | Selected |
| | M8 | Power wheelchair [25] | Supportive | By AI | Independent | Combined |
| | M9 | Mobility robot [34] | Complementary | By AI | Independent | Combined |
| Telerobotics | R1 | Telemanipulation robot [81] | Supportive | By AI | Independent | Combined |
| | R2 | Telemanipulation robot [20–22] | Supportive | Mutual | Independent | Combined |
| | R3 | Unmanned aerial vehicle [24] | Supportive | Unsupervised | Interpreted | Combined |
| | R4 | Unmanned aerial vehicle [44] | Supportive | By AI | Interpreted | Selected |
| | R5 | Rescue robot [45] | Supportive | By AI | Interpreted | Selected |
| | R6 | Unmanned aerial vehicle [32] | Supportive | By Human | Independent | Combined |
| | R7 | Telemanipulation robot [42] | Supportive | Mutual | Codependent | Combined |
| | R8 | *DJI Mavic Air 2* [19] | Reciprocal | Mutual | Interpreted | Combined |
| | R9 | Unmanned aerial vehicle [37] | Complementary | Unsupervised | Independent | Combined |
| Surgery | S1 | *Steady-Hand* [96] | Supportive | Unsupervised | Codependent | Combined |
| | S2 | *Micron* [61] | Supportive | Unsupervised | Codependent | Combined |
| | S3 | *Craniostar* [47] | Supportive | By AI | Codependent | Combined |
| | S4 | *Supervisory Steady-Hand* [48] | Delegated | By Human | Codependent | Combined |
| | S5 | Comanipulation robot [107] | Complementary | Unsupervised | Codependent | Combined |
| | S6 | Force amplifier [77] | Supportive | Unsupervised | Guided | Combined |
| | S7 | *Acrobot* [43] | Supportive | By AI | Guided | Combined |
| | S8 | Comanipulation robot [36] | Supportive | By AI | Guided | Combined |

braking, and steering. These systems come in two flavors with similar acronyms: advanced driving assistance systems (ADAS) and automated driving systems (ADS) [71]. ADAS assist drivers in performing *primary* [54, 79] or *operational* [64, 71] driving tasks, such as accelerating, braking, and steering. For example, lane-keeping assistants (D1-5) help drivers to stay in their lane by steering towards the lane's center when the vehicle deviates. This leads to tightly-coupled interaction, where the human and AI agent can be literally moving the steering wheel at the same time, the human with their hands and the agent with a motor. ADAS' assistance helps drivers to drive more safely when they are fatigued or distracted, and the AI's intervention can help them to recognize their mistakes, for example changing lanes without signalling (D4).

In contrast, ADS relieve drivers of driving tasks entirely (D6-7). For example, Tesla's Navigate on Autopilot (D7) fully automates highway driving from on-ramp to off-ramp. The human is required to drive before the feature is initiated and after it disengages, as well as supervise the AI at all times so that they can take over in case of an emergency. ADS systems remove much of the tedium from driving, since they relieve drivers of their tasks by delegating them to the AI, and have the potential of increasing the safety and accessibility of driving.

Much like systems from other domains, player balancing features in games, driver support systems work to prevent disaster when control is difficult, although some are able to support drivers even further. They may use *haptic shared control*, which provides the additional benefit of communicating the AI intentions to the driver via force feedback [2]. For example, Volvo's Lane Keeping Aid (D4) rotates and vibrates the steering wheel when the vehicle unexpectedly crosses over lane markings, thereby informing the driver of their situation and amending their control. This same approach is used in surgical robots to improve surgeons' awareness of the tissue they are cutting [78]. Otherwise, driving automation systems relieve drivers of their tasks entirely. This use case is unique to semi-autonomous vehicles, since driving is the only application domain we surveyed in which it is preferable for the AI to do all the work. In contrast, input automation controls inaccessible inputs to make games more accessible, not to make play less tedious.

*2.2.3 Creativity Support.* Shared control can also support humans in creative endeavors such as making sketches or sculptures (C3-6,C9,C10), musical performances (C1-2), and digital artifacts such as levels for a digital game (C8) or labelled datasets (C7). For example, *dePENd* (C3) is a sketching assistant that helps designers to draw shapes by guiding a pen across a sheet of paper using a magnet. By drawing two short lines at two points on the sheet, users can instruct dePENd to drag the pen from the second point to the first. The AI therefore supervises the human while they sketch; however, the user is free to wiggle the pen, to create wavy lines, or periodically lift the pen to create dotted lines. The user requests assistance by modifying the artifact and dePENd responds to these changes. Both human and AI control the pen at the same time in the real-time task of drawing. The collaboration is tightly-coupled, where movements of the pen immediately convey intent.

Systems can also assist users in real time performances, such as playing an instrument. For example, PossessedHand (C1) uses EMS to share control of the user's body while playing a traditional Japanese instrument called the *koto*. Fourteen electrodes are placed on the user's forearm and EMS causes the user's hand to play along with a piece of sheet music. In this way, PossessedHand demonstrates to users how the koto is played and may facilitate learning to play without assistance. These systems were designed to overcome technical barriers experienced by creators, guiding their performance of difficult tasks, and are therefore especially useful to novices.

Shared control in creativity support is often explicitly didactic. PossessedHand, for example, was designed for users who had never played the koto before. In contrast, surgical robots guide surgeons' control of their implements to supervise and support them, but not to teach them how to perform an operation. This makes creativity support unique among domains we surveyed in that, although it is not yet known whether using such a system can yield long term improvement in the performance of creative acts, creativity support AI can be designed to show novice users how a task should be done.

*2.2.4 Mobility Assistance.* Many people have mobility disabilities, as did an estimated 9.6% of Canadians in 2017 [92], and may need the assistance of a device, such as a walker or wheelchair, to get around. Mobility assistance devices enable many users to navigate independently, but others may find using their wheelchair's joystick too difficult. These users may be able to use an alternative device, such as a sip-n-puff or head controller, but these devices may themselves be too difficult or tiring to use without assistance [25, 57]. Therefore, smart power wheelchairs share control with their users, helping them to avoid obstacles and navigate smoothly. For example, Soh & Demiris designed a smart power wheelchair (M7) that learns and mimics the *hand-over-hand* assistance occupational therapists provide for novice power wheelchair users.

Mobility assistance devices use shared control to assist users who have difficulty navigating independently. Therefore, their designs embed the assumption that users need to be supervised to navigate safely. All of the devices we surveyed help users to avoid collisions (M1-9) and many help with navigating smoothly (M1-4, M9). For example, Ezeh et al. compared two smart power wheelchairs [25], representing radically different forms of shared control. The first (M4) implements *linear blending* (i.e. *policy blending* [22] or *direct blending* [70]), a form of shared control that weights and sums the human user and AI agent's control signals. The path planner generates a control signal that smoothly steers the wheelchair away from obstacles and uses the average of the human and AI's control signals to control the wheelchair. In so doing, the linear blending wheelchair supervises the user and refines their control signal. In contrast, the second (i.e. M8) smart power wheelchair implements *probabilistic shared control* [25, 26], a form of shared control that learns from the human to infer their intentions [100]. The path planner uses its model of the human to select a command that it judges the human will find most agreeable and that best satisfies obstacle avoidance and smoothness constraints. The probabilistic approach used by this wheelchair illustrates another assumption embedded in the design of mobility assistance devices using shared control: since the human is expected to err, it is better to interpret their intentions than obey their commands. In this way, many systems use the AI to supervise and override the human's control (M1-6), preventing disaster when the human is unable to navigate safely.

Much like player balancing in games and ADAS in cars, shared control mobility assistance devices use the AI to supervise the human's actions and assist when their actions are dangerous. However, these systems typically interpret the human's commands in a way that is unique to their problem domain, and may be inappropriate in other situations. For example, M8 gives the AI more authority than the human user. The AI interprets the human's intentions and navigates on their behalf. In contrast, surgical robots have been designed to filter surgeons' tremors, but it would be inappropriate for these robots to operate autonomously because the surgeon is considered the ultimate authority [78]. Differences in the design of these systems expose designers' assumptions about each actors' capabilities and how their interactions should be structured. Therefore, shared control of mobility assistance devices takes the form of a strict hierarchy, with the supervisory AI interpreting the human's commands and acting on their behalf from above.

*2.2.5 Telerobotics.* The first shared control systems were designed to make dangerous and difficult work easier for humans [22]. In particular, teleoperated robots can perform manipulation tasks on the human's behalf (R1,R2,R7) or explore environments that are unsafe for humans (R3-6,R8,R9). For example, Rakhimkul et al. created a robot arm (R1) for users with motor disabilities that identifies objects in its environment and changes its pose to make it easier for the human to pick them up. Therefore, telemanipulation systems use shared control to make human work and activities of daily living safer and easier. Otherwise, teleoperated robots can work in places that are inaccessible or unsafe for humans. For example, recreational drones—miniature unmanned aerial vehicles—such as DJI's Mavic Air 2 (i.e. R8), or those created for research (R3-6,R9), could assist search and rescue services in locating persons missing in dangerous or difficult environments.

Teleoperated robots can automate tasks at a distance and in environments that would be unsafe for humans. But, as has been demonstrated time and time again, automation does not replace human work, it changes human work [8, 75]. And so, telerobotics comes with its own set of problems that shared control is used to overcome. For example, telemanipulation AI can share control of a robot arm that the user controls using a joystick (R1) or skeleton tracking system (R2). Using these input devices may be tiring, imprecise, or difficult, so the AI assists to make the human's control less arduous. In some cases, working remotely also introduces issues related to signal interference and time delays that shared control can alleviate. For example, the Mavic Air 2 has a Failsafe Return to Home feature that causes it to fly back to a specified location when the human's control signal is lost for more than three seconds and Engel et al. created quadrocopter AI (R3) that overcomes unstable behaviour caused by short time delays. Shared control likely cannot overcome problems caused by extreme time delays, such as those experienced by robots on Mars [101], but it can account for short lapses in real-time control. Therefore, shared control is used in telerobotics to overcome some issues inherent to controlling robots remotely, making them more reliable.

Shared control telerobots leverage AI to make the human's tasks easier and safer, but in ways unlike systems in other domains. For example, shared control telerobots expect interruptions in the human's control signal, but fewer than half of the systems we surveyed used these data to construct the AI's control signal (R3-6, R8), as is typically done by mobility assistance AI. Therefore, telerobots are designed to provide a form of shared control characterized by the human and AI's real-time and joint performance of their tasks. Rather than supervising the human and trying to infer their intentions, telerobots share control by performing the same tasks as the human and providing continuous support, working in the background in case the human needs assistance.

*2.2.6 Surgery.* Surgery is a risky and delicate task that sometimes requires more precision than human surgeons can provide. As explained by Jakopec et al., total knee replacement surgery leads to large misalignment of knee prostheses, which may necessitate revision surgery in over one third of cases [43]. It is surgeons' need for precision that drives the development of shared control surgical robots [78]. These are grounded robots (i.e., affixed to a static structure, such as an operating table) or ungrounded robots (i.e., mounted on the surgeon's body or tools) that share control of the surgeon's implements. For example, the grounded Acrobot (S7) guides the surgeon's control of an orthopaedic cutter system for milling patients' bones during total knee replacement surgery, eliminating unwanted deviations in prosthesis alignment. In this way, shared control surgical robots overcome imprecision in surgeons' control of their implements, making their work safer and more effective.

The imprecision that surgical robots are typically designed to overcome is called *tremulous motion*. These are involuntary and high-frequency hand movements that may be unnoticeable in non-surgical contexts but disabling for surgeons, as they can be orders of magnitude larger than than some of the smaller bodily structures surgeons routinely manipulate [61]. Due to their distinctively high frequencies, surgical robot designers rely on a signal processing analogy in which shared control is used to "*filter*" [48, 61, 78, 96] tremors. The ungrounded Micron system (S2), for example, implements a low-pass filter that removes high-frequency components from the surgeon's control signal, using only the low-frequency components to drive a piezoelectric manipulator that moves the tool's tip. This can give the surgeon the sensation of being in complete control while the assistant removes tremulous motion that the surgeon may not even be aware of.

These systems enable surgeons to operate with superhuman precision and their motors can confer superhuman awareness as well. For example, instead of reducing noise, the force amplifying device developed by Payne et al. (i.e. S6) amplifies forces at the tip of the surgeon's implement. This enables the surgeon to determine an appropriate amount of pressure to apply, overcoming the reduced kinaesthetic feedback surgeons experience during minimally invasive surgery [77]. A similar approach was used in Acrobot, which uses force feedback to prevent surgeons from cutting beyond predefined regions. In this way, surgical robots can share control of the surgeon's implements to guide their control, informing them of their mistakes and preventing them from making them.

Designers' expectations of surgeons' capabilities make surgery unique among domains we surveyed. The human user in these systems is a surgeon who is expected to be less precise than the AI, but far more knowledgeable and prudent. Games, mobility assistance devices, and telerobots expect that human users may be unable to provide their intended command, so they step in to control the system on the user's behalf. Some creativity support systems guide users' performance and even go so far as to force the user to take action when they may not have intended to. In contrast, surgical robots expect their users to err, but never act on the human's behalf. They have been designed to filter out a specific form of noise in the human's control signal, leveraging shared control to discover the *true signal* that it suggests. Force feedback can inform surgeons about the tissue they are cutting and help them to recognize when they have made a mistake, but it is never used to force the surgeon's hand. Therefore, shared control is used by surgical robots to support surgeons' control of their implements as a subordinate, but not as an equal or superior.

We see from this presentation that human-AI shared control has arisen in numerous contexts, often independently and without knowledge of its use elsewhere. Each is characterized by the human and the agent sharing the same interface to the system – sometimes through software, and sometimes through physical embodiment of the agent through magnets, motors, or even electrical stimulation of the user's muscles. This usage of the system is tightly-coupled, oftxen involving simultaneous control of the system by both human and agent. Our review has revealed some of the difficult decisions that are faced by designers of shared control systems. For example, should the human or the agent have primary control? What forms of supervision should be present between human and agent? How and when should control transfer between the two actors? These are questions that our dimension space will help address, first by precisely characterizing the forms of interaction exhibited by specific systems, and then by identifying design patterns that can be used to drive design and analyze design choices.

We turn our focus now to the design of shared control systems more generally. In the next section, we summarize what is known about the design of shared control systems and explain how the design space analysis tool proposed in this work addresses gaps in our knowledge. The dimension space itself is presented in Section 4 and its uses are demonstrated in Section 5.

## 2.3 Shared Control Design Space Analysis

Design space analysis is "*a perspective on design which emphasizes the role and representation of design rationale.*" [62] It recognizes that a design process produces not only a reified artifact but also an abstract design space of possible options. This notion of design space analysis is closely related to Alexander's notion of *pattern languages* for describing *design patterns* [3, 4]. The connection is most obvious in Alexander's own words: "*The real work of any process of design lies in this task of making up the language, from which you can later generate the one particular design.*" [3] A design pattern is "*a rule which describes what you have to do to generate the entity which it defines*" [3] and they can be composed to express novel designs using pattern languages.

In this work, we attempt to elucidate a design space of shared control systems and propose a design space analysis tool through which it can be understood. The tool we construct in this paper is called a *dimension space*; these are used to structure a design space, classifying and comparing systems along different *dimensions*. For example, one could construct a dimension space to understand the properties of interactive systems in a physical environment according to the system's role and the physicality of its manifestation [35], or to classify and compare musical devices according to their required expertise and number of degrees of freedom [11]. Dimension spaces enable designers to describe a design space, explore the design choices available to them, and communicate their design rationale to others. When they describe meaningful differences between designs, they can help designers to realize possibilities that they may not have conceived of otherwise. To the best of our knowledge, no such tool exists for reasoning about the design of human-AI shared control systems. In the rest of this section, we describe several analogous design spaces that have influenced our thinking. We illustrate that existing tools for understanding and designing shared control systems do not provide adequate understanding of the human-AI interactions afforded by different design choices.

*2.3.1 Verplank Notions.* Sheridan and Inagaki's extensions to Verplank's roles of automation can be used to describe the roles AI play in shared control. The AI can *extend* the human's capabilities, *relieve* them of burdens, or *partition* and perform part of their functions [41, 86, 87]. These ideas have been used to describe human-machine cooperation for over four decades and have been hugely influential in designers' thinking. However, their expressiveness is limited, as they can only describe *what* roles the automation plays and not *how* it plays those roles. Furthermore, the roles of automation are potentially problematic when used as a dimension for design, because each is defined in different terms and are therefore not mutually exclusive. For example, input automation in games *extends* the player's capabilities by controlling inaccessible inputs, it *relieves* them of this burden, and it *partitions* their function by performing inaccessible parts. The historical significance of the roles of automation indicates that understanding the AI's role in shared control is important for understanding how it cooperates with the human.

*2.3.2 Design Frameworks.* Many recent works [27, 30, 55, 65, 66, 72, 74, 75] have iteratively constructed a human-machine interaction model composed of layers of shared and cooperative control, assistance, and automation [73]. As explained by Pacaux-Lemoine & Flemisch, human-machine cooperation involves a human and AI agent communicating with each other and controlling a system, via a Common Work Space [66, 72]. Their framework can express what types of tasks each agent performs [75], on what levels they communicate [72], and whether they control the system directly, but it cannot express how agents' control signals are composed to create a shared control signal. In a similar vein, Abbink et al. have proposed a design framework for shared control systems [1]. It can express how human and AI perform a hierarchy of tasks and communicate using signals, signs, and symbols [82] at each level, but says little about how a particular design choice might influence users' experiences. These design frameworks provide an especially technical account of shared control, and thereby provide little intuition for how designers' choices affect users' experiences of sharing control. These authors' focus on precisely how human and AI communicate with each other indicates that understanding how actors influence each other is important for understanding interactions in shared control.

*2.3.3 Pattern Languages.* There are at least two pattern languages that can be applied to shared control, although neither is specific to it. Baltzer et al. proposed an interaction pattern language for human-machine cooperation defined in terms of the problem a pattern addresses, the solution that overcomes the problem, the consequences of using the solution, and example systems that implement it [9]. This tool can express how and why a particular design overcomes a problem, but as van Diggelen & Johnson have previously argued [102], pattern languages reliant on natural language descriptions lack uniformity and therefore make comparing designs difficult. Since this pattern language provides no guidance regarding the design space of shared control systems, designers using it could easily get lost. van Diggelen & Johnson have proposed their own pattern language for human-agent team design patterns, defined in discrete terms[102]. They considered the types of work actors perform, whether the work is physical or cognitive, actors' spatial distribution, and how actors communicate. Using these elements of team work, the authors have demonstrated how higher-level design patterns, such as human supervisory control [86], could be expressed. This pattern language enables designers to describe existing human-agent team patterns, and envision new ones, but it cannot describe how teams share control. It indicates that understanding how actors supervise each other is necessary for understanding how they cooperate, but more specificity is needed to understand how design choices determine users' experiences.

Flemisch et al. have argued that shared control is the *sharp end* of human-machine cooperation [27, 28]. Like a single person wielding a spear, the human and AI cooperatively hold the *blunt end*, communicating their goals and plans, and strike using the *sharp end*, jointly controlling the system using their common interface. It is therefore cooperation at the sharp end (i.e. controlling a common interface) that is unique to shared control. The blunt end (i.e. communication between human and AI), however, is common to all forms of human-machine cooperation and, therefore, much more is known about it. This is the gap that our dimensions space fills. For the first time, designers are able to express how design choices at the sharp end determine how human users interact with their AI partners. We provide a structure to the design space of shared control systems and describe how different design choices afford different interactions and experiences to users. In the next section, we describe our surveying method and how the dimension space was constructed.

## 3 METHOD

In order to better understand how human-agent interactions are structured in shared control, we set out to create a dimension space, which is shown in Figure 1 and will be presented fully in Section 4. We compared descriptions, diagrams, use cases, and user evaluations of shared control systems to identify dimensions of variability that have meaningful consequences for the interactions they afford. We surveyed the literature using an approach similar to a scoping review [6]. Our goal was to discover the boundaries of shared control and identify problem domains in which it is used, so that designers working in one domain could understand design insights from other domains. We followed an iterative process in which we (1) searched online databases using keywords found in relevant papers, (2) identified papers that described human-AI shared control systems that were somehow novel, and (3) mined these papers for new keywords and references to potentially relevant systems. This process is described in Table 2 and the resulting 55 works consulted are summarized in Table 1.

From the outset, we recognized that shared control is an interaction paradigm used in disparate domains and that there were differences in terminology between these domains. Most importantly, "*shared control*" is a term used primarily in robotics [22] and so non-robotics uses of shared control may not be identified as such. Therefore, an exhaustive search using fixed search terms could unintentionally exclude relevant, albeit unconventional, areas of the literature. Furthermore, we knew a priori of some domains in which human-AI shared control was used (i.e., digital games, driving automation, and robotics), we did not know whether there were more or how they related to each other. We therefore followed the iterative approach shown in Table 2, repeating our search as new keywords were discovered.

As we built our corpus, we grouped systems according to the problems that they were designed to overcome and called these groups *domains*. For example, *mobility assistance* systems help users to get around while *creativity support* systems support users in creative endeavors. This process was repeated until we had exhausted all of our keywords and identified samples representing the variability in designs used in each domain.

Table 2 summarizes our search process. In Phase 1, we performed a sequence of searches on Google Scholar and the ACM Digital Library for an expanding set of keywords, starting with "shared control" and "shared-control", and ending with the list of keywords shown in the table. In this first phase, 58 papers were examined and 13 were found to be relevant, i.e., presenting systems deploying Human-AI Shared Control. The second phase focused on two series that were found to be especially relevant to human-AI shared control: the IEEE *International Conference on Systems, Man, and Cybernetics* (SMC) and the International Federation of Automated Control (IFAC) *Symposium on Analysis, Design, and Evaluation of Human-Machine Systems* (HMS). We searched the proceedings of both conferences, SMC on IEEE Xplore and HMS on ScienceDirect, using the keywords identified in phase 1. These searches yielded 41 potentially relevant systems, of which 12 were considered relevant and included in our corpus. The third phase addressed newly published research [16, 17, 50]), which expanded the search to creativity support systems, and also identified the Human-Agent Interaction (HAI) conference as a venue for shared control research. 57 systems were examined in phase 3, of which 18 were found to be relevant. We then repeated phase 1 with the complete set of keywords identified in phases 1 through 3. In this phase 8 systems were examined and 7 were added to our corpus. Finally, 4 especially notable creativity support systems were examined during the review process, of which 3 (i.e., C5 & C11-12) were included.

As we built our corpus, we described the systems we encountered and recorded how these systems' designs differed. By considering what made systems similar or different informally and subjectively, we constructed dimensions of variability that enabled us to classify

**Table 2: The systems examined, topics encountered, and search terms used in each phase of our survey.**

| Phases | Keywords | Systems |
|---|---|---|
| **Phase 1** (Google Scholar, ACM Digital Library) | player balancing, input automation, crowd control semi-autonomous vehicles, haptic shared control, drones smart power wheelchairs, mobility assistance robots, human-supervisory control, aviation automation, smart homes, teleoperated robots, human-adaptive mechatronics, surgical training | (58 found, 45 excluded) G1, G3 D2, D3 M2, M4-7, M9 R2, R6-8 |
| **Phase 2** (IEEE Xplore, ScienceDirect) | human-machine cooperation, cooperative control, adaptive automation, adaptive cruise control, lane keeping assistants, air traffic control, quadrocopters, quadrotors, comanipulation robots | (41 found, 29 excluded) D1, D4, D5, D7, D9 M1, M3, M8 R1, R3, R4 S8 |
| **Phase 3** (Google Scholar, ACM Digital Library) | interactive machine learning, gaze control, mixed-initiative, interactive fabrication, sketching assistants, electrical muscle stimulation, human-robot teams, human-AI teams, | (57 found, 40 excluded) G2, G4-8 D6, D8 C1-4,C6-10 R5, R9 |
| **Phase 1 Bis** (Google Scholar ACM Digital Library) | surgical robots, *all of the above* | (8 found, 1 excluded) S1-7 |

systems more formally and objectively. At the end of each phase of our survey (Table 2), we chose the dimensions that we believed were most important for explaining how human-AI interactions in shared control are designed, and classified all of the systems in our corpus using these dimensions. When our dimensions classified systems as inappropriately similar or needlessly different, we split or combined axes to form new concepts. This process of expanding our dimension space to capture notions inspired by the literature, and contracting our dimension space by removing and combining axes, was repeated until we converged on a set of dimensions describing the most important similarities and differences for explaining human users' interactions with the systems we surveyed.

## 4 DIMENSION SPACE

The dimension space has four axes: *AI Role*, *Supervision*, *Influence*, and *Mediation*. They are depicted as Kiviat diagrams in Figure 2 (left), as is conventional for dimension spaces (e.g., [11, 35, 38]). *AI Role* describes the different ways tasks might be shared between human and the AI agent. *Supervision* specifies how actors monitor and correct each other. *Influence* captures how an actor chooses its own actions in response to those of the other actor. *Mediation* describes how actors' commands are unified through combination or selection. As suggested by Figure 2 (right), these dimensions correspond to different aspects of actors' interactions in shared control. *AI Role* is about actors' tasks, while *Mediation* is about their commands. *Influence* is about being aware of and responding to the other's actions, while *Supervision* is about taking action to correct the other. These dimensions have been constructed so that they can describe all human-AI shared control systems of which we are aware, and so that any position in the dimension space describes a non-empty set of plausible systems.

We designed this dimension space to provide the properties of *totality*, *orthogonality*, and *mutual exclusion*, which we define here. Each dimension was designed to be *total*, meaning that systems belong to at least one of the categories along each axis. This property holds of every system in our corpus. The dimensions are designed to be *orthogonal*, meaning that a system's classification in any dimension is independent of its classification along the others. Not all dimension spaces are orthogonal (e.g., [11]), but this property is desirable for our purpose because every position in the dimension space describes a valid system that designers could then implement. Figure 2 (right) shows how each dimension corresponds to a different aspect of the human and AI agent's interactions in shared control. There is no reason, for example, why a specific role would constrain choices around supervision, influence, or mediation. Finally, we designed the dimensions to be *mutually exclusive*, meaning that a system occupies at most one point in the design space. This property holds of all systems in our corpus. By aspiring to these properties, the dimension space can be used to classify systems in which one human shares control with one AI agent, and can be used to identify common design patterns between sets of systems. In this section, we describe the four dimensions of the dimension space: *AI Role*, *Supervision*, *Influence*, and *Mediation*.

### 4.1 AI Role—*Who Does What and When?*

The AI Role dimension answers questions of the form "*who does what and when?*", a phrase borrowed from Inagaki [41]. Users interact with a system for some purpose; they have some set of tasks that they are performing. When control is shared, the AI can support the human by cooperatively performing these tasks at the same time, the AI can take over control of the human's tasks to perform them at a different time, or it can perform its own tasks that are separate

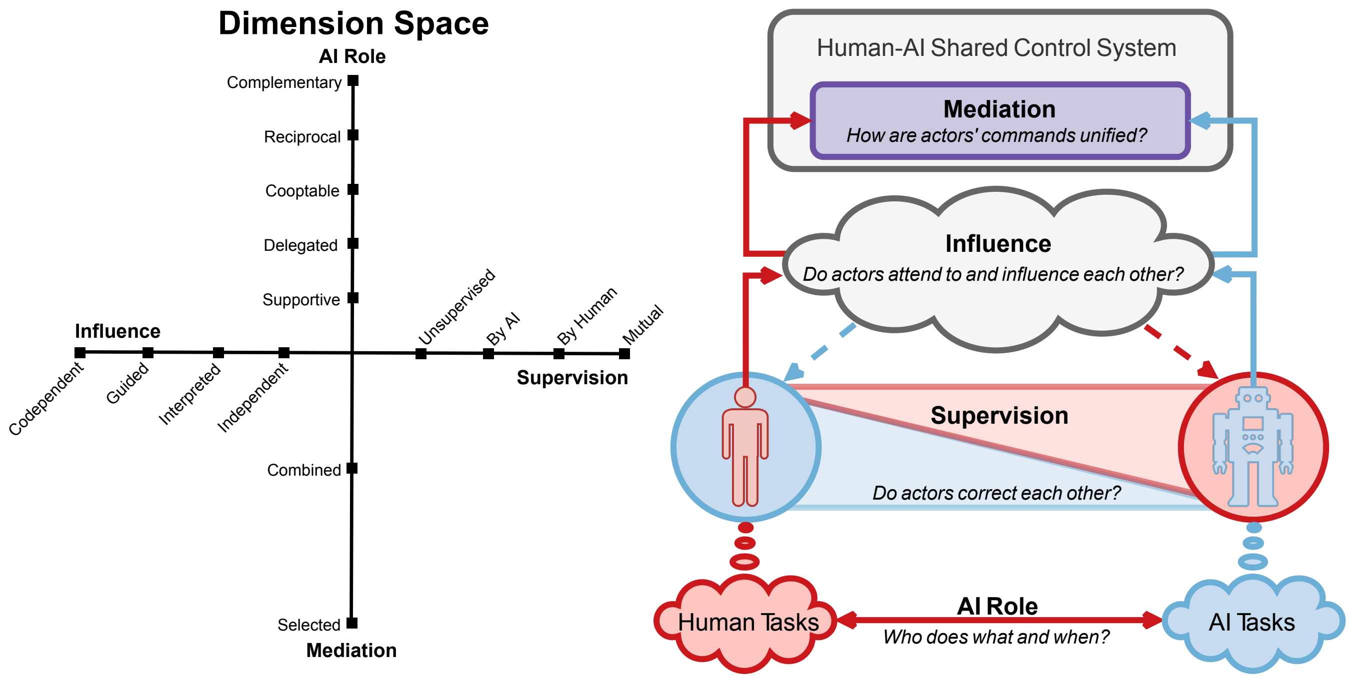


**Figure 2: On the left, our dimension space represented as Kiviat diagram axes. Representing systems this way communicates their differences visually, leveraging the spatial arrangement of each classification to show how different they are. Therefore, the area that a diagram covers is not meaningful, although its shape is. On the right, our dimension space represented as a block diagram. Each actor has a set of tasks (i.e., clouds at the bottom) that the human (i.e., red figure on the left) may delegate or coopt to or from the AI (i.e., blue figure on the right). Actors may supervise (i.e., vision cones emanating from the figures) or influence (i.e., dashed lines) each other, but they always share control of the system (i.e., solid lines).**

from the human's. This dimension captures how tasks are shared among actors and how the human assigns tasks. It tells designers how the human's interactions with the AI are related to the tasks they perform and how these tasks are shared between them. In all of the systems we surveyed, it is the human who allocates tasks and never the AI. We therefore define five types of AI role: *Supportive*, *Delegated*, *Cooptable*, *Reciprocal*, and *Complementary*, which we define in this section.

#### 4.1.1 ***Supportive****: The AI assists the human with a subset of the human's tasks*

The human is primarily in control when the AI performs a *Supportive* role. The human is engaged in all parts of their overall task while the AI performs some tasks at the same time to help the user if they run into trouble. For example, Gurnel et al. created a comanipulation robot (S8) that helps surgeons to more precisely insert needles into tissue samples using virtual fixtures. While the surgeon moves the needle, five haptic guides exert forces on the needle to guide the surgeon towards the desired needle position and orientation. Systems designed such that the AI performs a *Supportive* role assist humans with their tasks while the human retains primary control.

#### 4.1.2 ***Delegated****: Supportive & The human can hand over a subset of their tasks to the AI*

AI agents that perform a *Delegated* role relieve humans of specific tasks and perform them on their behalf. For example, human users may command Muscle-Plotter (C9) to do physical simulations using car designs that they have sketched and then graph the results using EMS. Digital game level designers can instruct Tanagra (C8) to fill in missing level geometry after they have made a change to the level. *Delegated* AI performs tasks that the human could do but prefers not to, or takes over control of tasks that the AI can perform more precisely.

#### 4.1.3 ***Cooptable****: Supportive & The human can take over a subset of the AI's tasks*

In contrast to the *Delegated* role, which enables humans to delegate their tasks to AI, the *Cooptable* role enables humans to take over control of the AI's tasks. For example, Flemisch et al. designed a semi-autonomous vehicle (D9) according to the H-Metaphor [31, 94], which suggests that drivers' interactions with autonomous vehicles should resemble riders' interactions with horses. In horseback riding, riders can *loosen* their grip on the reins, thereby shifting control authority to the horse, or *tighten* their grip, to seize authority. In this way, the vehicle autonomously drives around a racetrack, albeit less skillfully than a human might, while the driver subtly takes control, tightening the rein, to assist the AI. In sum, AI agents performing a *Cooptable* role are primarily in control of their tasks, while the human is able to take over the AI's tasks when they see fit.

#### 4.1.4 ***Reciprocal***: *Delegated & Cooptable*

When the AI performs a *Reciprocal* role, the human can delegate tasks to the AI and take over tasks that the AI is controlling. For example, the Mavic Air 2 (R8) is a recreational drone with an array of operational modes. *Active Track* enables pilots to instruct their drone to follow them and record video, delegating this task to the AI. The pilot can then take back control by turning the feature off, but only so long as the drone receives their control signal. Should the Mavic Air 2 lose the human's control signal, due to interference from the environment, the *Failsafe Return to Home* feature causes the drone to return to a predetermined location. Once their connection is reestablished, the pilot is free to disable this feature and coopt the AI's control of the drone's navigation. Thus *Reciprocal* shared control AI confers the benefits of both *Delegated* and *Cooptable*. The AI relieves users of their tasks and enables them to amend the AI's control of its own tasks.

#### 4.1.5 ***Complementary***: *Supportive & The AI has its own tasks that the human never performs*

A *Complementary* design partitions the human's task and allocates some parts to the AI in their entirety. This role can reduce the human's control burden when controlling the system is complicated, and can make systems more accessible when the human is unable to perform some tasks at all. For example, Shaper's Origin (C11) is a woodworking router that helps users to cut along a reference path more precisely. When using a typical router, moving the device's frame also moves the cutting tool, causing it to cut. In contrast, Origin separates control of the frame, which is moved by the human, from the tool, which is moved within the frame by the AI. In this way, the human is tasked with positioning the device while the AI is tasked with cutting. These systems relieve humans of part of their task, either because the AI can do it better or because the human is unable to do it at all.

### 4.2 Supervision—*Do Actors Correct Each Other?*

A shared control system's design embeds assumptions about the supervisory relationships among actors. For example, the SAE's levels of automation assume that human drivers of semi-autonomous vehicles, who until recently have been solely responsible for all primary driving tasks, are more capable drivers than the AI. They are therefore responsible for supervising the AI agent's control and overriding it when necessary [71]. In this way, one actor intervenes to prevent the other actor's mistakes. For example, a driver using Navigate on Autopilot (D7) may not know exactly what the AI is doing, only that they don't like it, and choose to take back control. Many systems assume a one-way supervisory relationship, with either the human supervising the AI (G8, D2, D6-7, D9, C7-8, R6, S4) or the AI supervising the human (G5-7, D5, C3-5, C12, M1-9, R1, R4-5, S3, S7-8). Some systems make no such assumptions (G1-4, D1, D8, C1-2, C9-11, R3, R9, S1-2, S5-6), while others afford mutual supervision (D3-4, C6, R2, R7-8) where AI and human each supervise the other. Our Supervision dimension captures the supervisory responsibilities of actors according to these four categories: *Unsupervised*, *By AI*, *By Human*, and *Mutual*.

#### 4.2.1 ***Unsupervised***: *Neither actor supervises the other*

When both actors are *Unsupervised*, neither actor controls the system to prevent the other from making mistakes. Their intentions may still conflict, such as when a smart power wheelchair steers to the right to avoid an obstacles while the human steers to the left, but not because one believes the other to be in error. For example, EMS Air Guitar (C2) helps humans to strum an imaginary guitar along with music. The system has no sensors, and is therefore unable to react to anything, so it stimulates the human's arm muscles without supervising their movements.

#### 4.2.2 ***By AI***: *The AI supervises the human*

Many AI agents are designed to monitor the human and take action when the human errs. For example, every mobility assistance system we surveyed (M1-9) uses AI to supervise the human and amend or override their control to prevent collisions. This form of supervision is applicable when the AI is able to detect that the human has erred; for example, the craniotomy tool Craniostar (S3) tracks a reference path across a patient's skull and steers the drill towards the path when it deviates.

#### 4.2.3 ***By Human***: *The human supervises the AI*

Conversely, shared control systems can be designed such that the human user supervises the AI. For example, Supervisory Steady-Hand (S4) performs microinjections fully autonomously, but requires a surgeon to supervise its performance and verify that it succeeded.

#### 4.2.4 ***Mutual***: *Both actors supervise the other*

Finally, shared control systems can be designed such that both actors supervise each other. For example, the Mavic Air 2 supervises the human to help with collision avoidance during manual operation, and the human needs to supervise the AI when it flies autonomously. In lane assistance driving systems (D3-4), the AI supervises the human's steering, taking over when the car drifts out of lane; the human in turn supervises this correction, and can override it by applying extra force to the steering wheel.

### 4.3 Influence—*Do Actors Attend to and Influence Each Other?*

To cooperate effectively, both human and AI may need to monitor what the other is doing to determine how they should control the system together. In some systems, the AI's control signal is a function of the human's. For example, the racing game AI of Cechanowicz et al. (G7) only assists the player when it detects that they are steering. In other systems, the human responds to the AI's control signal, which they perceive via force feedback (D1, C3-6), EMS (C1-2), or a visual display (C7). For example, Haptic Intelligentsia (C4) helps users to construct sculptures out of glue by guiding their control of a hot glue gun. When the user moves the gun outside a predefined 3D volume, known to the AI but not to the user, a haptic device pushes the gun back towards the volume's surface. This informs the user of the volume's location and helps them to decide where to put more glue. In this way, users of these systems are made aware of the AI's commands and may infer and react to its intentions. When both the human and AI each respond to the other's control signal, they can communicate with each other and negotiate how the system is controlled. We term the way in which human and AI interpret and respond to the other's control signal Influence, which we have observed as being *Independent*, *Interpreted*, *Guided*, and *Codependent*.

#### 4.3.1 ***Independent****: Neither actor influences the other*

In many systems, there is no obvious benefit to actors influencing each others' control. For example, in the *Space Invaders* clone *Alienated* (G2) the player controls the avatar's movement while the AI controls its shooting. Since movement does not affect shooting, the AI's decision to fire is determined by the game's state, rather than the player's actions. Neither actor is directly aware of the other's commands.

#### 4.3.2 ***Interpreted****: The AI's actions are influenced by the human's actions*

When the AI interprets the human's control signal, it may be able to infer and help the human achieve their goals. For example, *Gekku Aim* (G6) is a 2D racing game designed for children with cerebral palsy who have difficulty aiming at targets. When the player aims near an enemy and fires a shot, the AI steps in at the last moment and aims directly at the enemy closest to where the player was pointing. In this way, the AI may interpret users' control signals to infer their intentions and then take action to help the human realize them.

#### 4.3.3 ***Guided****: The human's actions are influenced by the AI's actions*

Several systems use EMS or force feedback to guide the human's control. For example, Kianzad et al. created a physically assisted sketching system (C6) that pushes the user's pen away from bounding lines they have drawn on paper, helping them to stay within a predefined region. This form of Influence can guide the human's control not only by physically pushing or pulling the system's hardware but also by improving the human's awareness. For example, the force-amplifying surgical robot created by Payne et al. (S6) improves kinaesthetic feedback from the tip of the surgeon's instruments, which is otherwise not present in minimally invasive surgery. This enables surgeons to sense what the AI senses and use their heightened awareness to operate more precisely.

#### 4.3.4 ***Codependent****: Both actors' actions are influenced by each other's actions*

When both actors are aware of the other's control signal, they may infer the other's intentions to cooperate better. If both actors control the same parts of the system's interface, they may engage in a negotiation to determine their form of control. For example, Shilkrot et al. created a digital airbrush (C12) that uses force feedback on its trigger to help users paint a reference image. When the human pushes down on the trigger to paint over an already painted area, the AI pushes back. Therefore, both actors influence each other by communicating their intentions through the trigger. This specific form of *Codependent* shared control, in which control is negotiated using physical forces, has been called *haptic shared control* by Abbink et al. [2]. It is especially applicable to systems where actors need to maintain awareness of the other's control in real time, such as semi-autonomous vehicles (D3-4, D9) and surgical robots (S1-5).

### 4.4 Mediation—*How are Actors' Commands Unified?*

Sometimes, the actions that the human and the AI agent perform may conflict. For example, the image classifier of Seno et al. (C7) may incorrectly label an object as a car while the human correctly labels it as a horse. The Mediation dimension captures how such conflicts are resolved. More technically, this dimension describes how actors' commands are used to control the system.

#### 4.4.1 ***Combined****: Actors' commands are combined to construct the shared control signal*

A system's mediation of actors' control signals is said to be *Combined* if the shared control signal is always a combination of both actors' control signals. For example, a human driver may turn their steering wheel to change lanes while their car's lane-keeping AI turns the wheel in the opposite direction, negating the human's action because they did not signal. In this example, actors' commands are *Combined* physically to create a new command. Both actors contribute to this new command and may control the system more forcefully, turning the wheel further in this case, to override the other actor's contribution. This approach is used by several lane-keeping assistants (D1-4).

#### 4.4.2 ***Selected****: One actor's command is selected as the shared control signal*

Alternatively, a system might select one actor's control signal as the preferred control signal, in which case its mediation is said to be *Selected*. These systems continuously check if the human or the AI should be solely in control. For example, Sentouh et al. created a lane-keeping assistant (D5) that monitors the human driver's commands and selectively ignores them when they might violate stabilization and lane-keeping constraints. Instead of combining both actors' commands, this assistant decides which actor should steer and *selects* their command to control the vehicle. In some *Selected* systems, actors' commands are both *Combined* and *Selected* at different times. For example, G8 combines the player's shooting commands with the AI's aiming commands, but selects the player's aiming commands when they take control of that input. This type of Mediation has important consequences for the human-agent interactions systems afford. *Selected* systems enable actors to remove the other from the control loop, which means that one actor can act autonomously and without interference.

## 5 ANALYSIS

As we have shown, our dimension space describes human-agent interactions in shared control along four dimensions: *AI Role*, *Supervision*, *Influence*, and *Mediation*. The dimension space enables designers to classify systems and discover patterns in the designs used within a problem domain. Our dimension space can help designers to identify common assumptions about actors' roles, responsibilities, competencies, and ways of cooperating and then imagine how these interactions might otherwise be structured. In this section, we apply the dimension space to the systems from which it was derived to demonstrate how it can elucidate common design patterns. We then apply it to our entire corpus to identify larger design patterns used across domains.

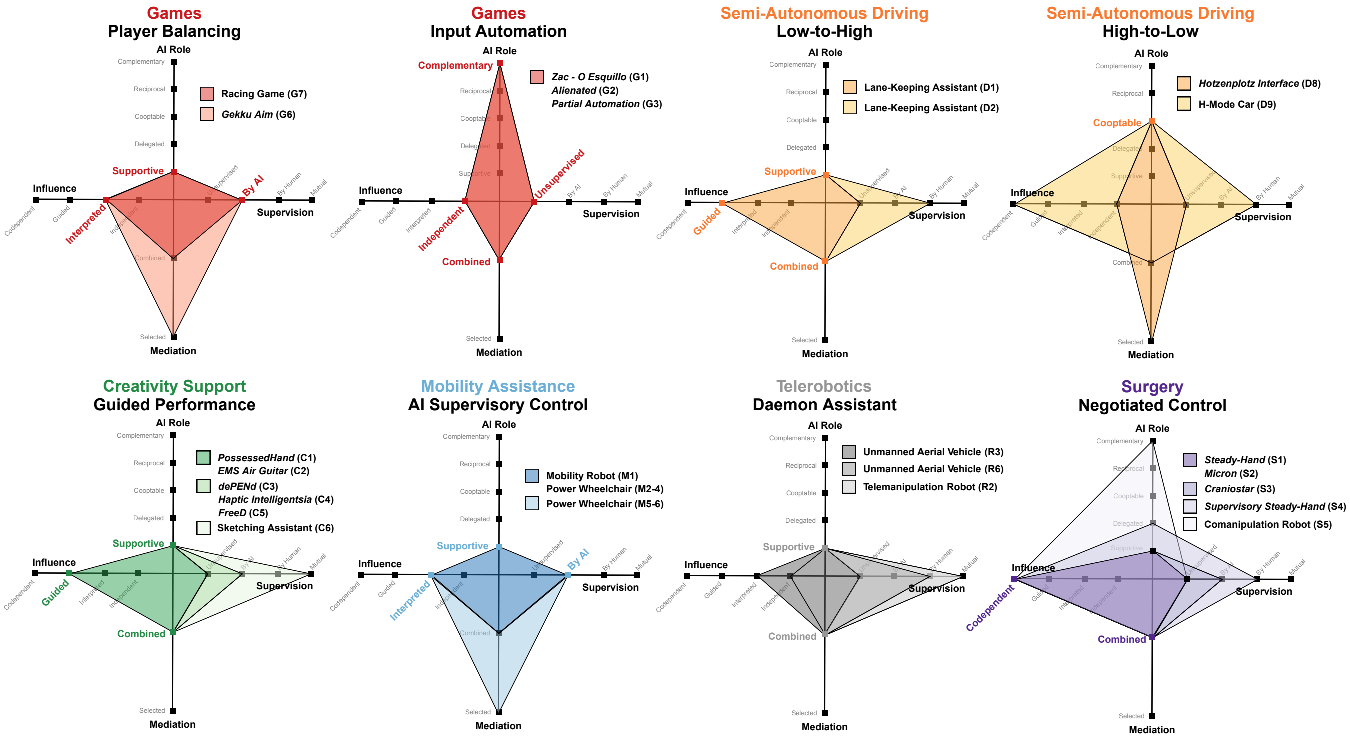


**Figure 3: Design patterns found in each of the six domains we surveyed. Each pattern is defined using at least one of the dimension space's four axes. For example, *guided performance* (bottom left corner in green) provides a *Supportive*, *Guided*, and *Combined* form of shared control, indicated by the bold green axis labels. This pattern is exhibited in six systems (C1-6); however, individual systems vary with respect to their types of Supervision, which are shown in the diagram using different shades of green.**

## 5.1 Domain-Specific Design Patterns

Designers can classify a shared control system by assigning it to one category in each dimension in the dimension space. When multiple systems from the same problem domain are classified, patterns emerge, providing insight into to the types of interactions they afford. For example, human users of mobility assistance systems are typically supervised *by the AI* who *supports* them in avoiding collisions and *interprets* their commands to infer their intentions. However, they may either *combine* or *select* actors' control signals. Therefore, mobility assistance devices exhibit a *Supportive*, *By AI*, and *Interpreted* design pattern, defined along three dimensions (leaving the category in the fourth dimension open). We have called this pattern *AI supervisory control* in mobility assistance devices, but in games it is already known as *player balancing*. Therefore, similarities in the patterns used in multiple domains may reveal correspondences in the interactions systems afford and users' experiences of them. Design insights discovered in one domain may be directly applicable in designing systems for another.

Of course, identifying design patterns is only the first step. Once we know how a set of design choices overcomes domain specific problems, we can better understand how similar designs might overcome other problems. We can describe what about a pattern makes it effective and imagine how the strengths of multiple patterns might be combined to provide novel forms of shared control.

In this section, we demonstrate how our dimension space enables designers to classify, compare, and understand shared control solutions to similar problems by identifying common design patterns within each of the six domains we surveyed. We present eight patterns drawn from dimension space plots of our 55 surveyed systems, shown in Figure 3.

*5.1.1 Digital Games.* Player balancing enables weaker players to compete with stronger opponents. For example, in *Gekku Aim* (G2) players shoot projectiles at each other in a 2D play area and aim their shots using a gamepad's analog stick. The game assists players with deficits in manual dexterity by aiming directly at an opponent when their aim is misaligned but close. This makes the game easier to play for players who experience difficulty. As shown in Figure 3, *Player Balancing* is a *Supportive*, *By AI*, and *Interpreted* design pattern. The AI's role is to *support* the player's activities. To do so, the AI *supervises* the player and *interprets* their aiming actions, modifying this control signal to provide the player with improved aim. The pattern is agnostic with respect to Mediation strategy. Figure 3 shows two games whose aim assistance algorithms follow this pattern.

A second design pattern for games, *Input Automation*, is also shown in Figure 3. This pattern delegates control of inaccessible game inputs to an AI agent. For example, in *Zac - O Esquilo*, a one-switch clone of *Frogger*, the player presses a button when they want their avatar to move [63]. The AI then selects the movement direction it deems most appropriate and the avatar moves in that direction. While the player controls both *when* and *where* their avatar moves in the original game, *Zac - O Esquilo* partitions the player's task by asking them only to choose *when* to move. Games like *Zac - O Esquilo* use input automation to make games nearly universally accessible by adhering to a *Unsupervised*, *Complementary*, *Independent*, and *Combined* design pattern. Actors perform different tasks simultaneously and therefore do not directly supervise each other (Unsupervised Supervision) or influence each other (Independent Influence). The AI's Role is to complement the player's actions. The player and AI's inputs do not conflict and are therefore Combined as a Mediation strategy.

Input automation can enable players with radically different physical abilities to play the same games, by broadening their accessibility to players they were not designed for. However, players may experience *automation confusion* [14] if they do not understand what the AI is doing, because it performs a *Complementary* role, or become frustrated when it does things they do not like, because their control signals are *Independent*.

*5.1.2 Semi-Autonomous Driving.* Early semi-autonomous driving systems, such as lane-keeping assistants, sought to make simple driving tasks easier. These systems represent a transition from *low-to-high* automation, introducing automation into manual driving. These systems follow a *Supportive*, *Combined*, and *Guided* pattern that enables the AI to perform operational level driving tasks at the same time as the human (Figure 3). The AI *supports* drivers as they both steer to stay in the lane. AI in low-to-high automation uses force feedback to *guide* the driver's steering, which makes lane keeping easier by *combining* steering forces from both actors. The pattern is agnostic as to Supervision strategy.

As more automated features have been integrated into cars, designers have imagined novel interaction metaphors and design philosophies for semi-autonomous driving. Consequently, a *high-to-low* automation pattern has begun to emerge, introducing manual control into automated driving. For example, the *H-Metaphor* suggests that semi-autonomous vehicles should cooperate with drivers in the same way horses cooperate with riders, in that they can act autonomously or allow the human to take control [31, 94]. Otherwise, designers may see automating driving tasks as "*an amputation*" [52] of the driver's task and seek to provide them with playful ways to engage in driving [33, 106]. High-to-low systems put the human in a position of greater authority and put the AI in a *Cooptable* role. Systems adopting this approach differ in their choices of Influence, Supervision, and Mediation approaches, indicating that the design of high-to-low systems is still under active exploration.

*5.1.3 Creativity Support.* Several of the creativity support systems we surveyed were designed to help novices sketch (C3, C6, C9-10) or play a musical instrument (C1-2). They use EMS and force feedback to guide creators' control of their tools, enabling them to perform skilled actions that may be too difficult without assistance. These systems represent a form of shared control that we have called *guided performance* (Figure 3), in which the human is *Guided* by the AI which performs a *Supportive* role while actors' control signals are *Combined*. Example systems use different forms of Supervision.

Since guided performance systems are designed for novice users, the AI's guidance is intended to help them overcome technical barriers to creative endeavors. For example, sketching assistants help humans to draw straight lines (C3) and stay within them (C6), both of which are tasks that the human may learn to do on their own with more experience. Therefore, this pattern helps novices to get acquainted with their craft by removing technical burdens and enabling creators to create freely.

*5.1.4 Mobility Assistance.* Mobility assistance is one of the more popular applications of shared control and was one of the more homogeneous domains we surveyed. Since the user is expected to have difficulty controlling these devices, the AI is put in the role of supervisor and is often permitted to override the user's control when it sees fit. These systems interpret the human's control signal to infer their intentions and may ignore their commands if executing them is dangerous. Therefore, the human's control signal is *Interpreted* as they are supervised *By AI*, which plays a *Supportive* role. These systems have the AI supervise the human, sometimes preventing them from controlling the device at all, and therefore adhere to a design pattern that we have called *AI supervisory control*. The pattern does not specify a Mediation strategy.

Unfortunately, little has been reported about users' experiences of interacting with these devices, since many evaluations have not included disabled users (e.g. [18, 26, 57, 58, 90, 105]) and authors are typically more focused on the technical challenges of avoiding collisions. However, what is known indicates that these devices can help disabled users to navigate more safely [25, 34]. If interacting with these systems is anything like interacting with games that use player balancing, which exhibit the same design pattern, then users may find the AI's intervention helpful but potentially intrusive.

*5.1.5 Telerobotics.* As explained by Jiang & Odom, the history of human-robot team design marks an ideological shift from seeing robots as subordinate tools to seeing them as equal partners [45]. Regardless of whether it is posing a robot arm to make manipulation easier or stabilizing a drone to counteract signal interference, teleoperated robots are being designed to overcome technical and human factors problems by leveraging their unique strengths. For example, Dragan & Srinivasa created a telemanipulation robot to assist remote operators in picking up objects [20–22]. They point out that human operators have a better understanding of the task than the robot (e.g., knowing that caution is needed around breakable objects), although they have difficulty controlling the robot precisely and can become fatigued over time. In contrast, an AI partner is tireless and can control a robot perfectly precisely. In order to account for the deficiencies of one actor by leveraging the strengths of the other, the AI predicts which object the user is trying to grab and assists them in doing so safely. However, unlike *AI supervisory control*, these systems do not necessarily have the AI supervise the human, since the human is assumed to have greater authority in some situations. Therefore, teleoperated robot systems typically adhere to a *Supportive* and *Combined* pattern, which we have called *daemon assistant* (Figure 3), in which the AI works in the background to *support* the human by *combining* their commands.

*5.1.6 Surgery.* Surgical robots help surgeons to perform delicate and precise operations by providing haptic guidance and filtering tremulous movement. Because the surgeon is the designated expert in the operating room, the AI agents in these systems are seldom capable of performing tasks on their own, although some examples can be found (S4-5). These systems also vary with regards to actors' supervisory responsibilities. The surgeon is authoritative, but not infallible, so some systems let the AI supervise (S3, S7-8) while others are designed for neither actor to supervise the other (S1-2, S5-6). Surgical robots therefore adhere to a *Combined* and *Codependent* design pattern, which we have called *negotiated control* (Figure 3). Since actors sense the other's actions via the forces they exert on the surgeon's implements, actors negotiate how the system should be controlled. Each responds to the other's movements, which move the surgeon's tools when *combined*, so their action is *codependent*. If the AI believes the surgeon to be in error, the surgeon may discover their mistake when the AI counteracts their movement using force feedback. This approach can provide surgeons with superhuman precision with their implements and superhuman awareness of their workspace.

## 5.2 Design Patterns Across Domains

Having shown how human-agent interactions are structured in various domains using shared control, we now apply our dimension space to our entire corpus to determine which designs are used across domains. In so doing, we elucidate commonalities in the design of these systems and describe the forms of shared control they provide. These patterns tell us not only how shared control is used across multiple domains but also what problems these domains have in common. Because designers working in the domains we surveyed may be unaware of how shared control is used by others, these patterns represent more general forms of shared control that have been discovered independently multiple times. The broader perspective we take in this section enables use to identify the most popular design patterns and understand how common design choices overcome the most pervasive problems in controlling interactive systems.

To identify these patterns, we used Artur & Minghim's *Subspaces Explorer* system, shown in Figure 4, which uses correlation analyses to embed and cluster data in an alternative RadViz space [7]. The color coding and spatial arrangement of systems plotted this way are explained in the figure caption. We then selected homogeneous clusters and combined them to create patterns, such that each had a similar number of systems that adhere to it (i.e., 7 to 16). The cluster in the center of the RadViz plot in Figure 4 contains the outliers that did not belong to a more cohesive pattern, so this cluster was not selected. In this section, we describe the four higher-level design patterns we identified.

*5.2.1 Vigilant Savior.* The *vigilant savior* pattern has the AI step in to override the user's control in dangerous situations. The human is supervised *By AI*, which performs a *Supportive* role. When the AI detects that the human is in need of assistance, the human's control signal is *Interpreted* by the AI whose control signal is *Selected* as the shared control signal. For example, Jiang et al. created a UAV (R4) that predicts whether the human's command would put it into an unsafe state and overrides their control if it would. This pattern is a more specific form of the player balancing and AI supervisory control patterns described in the last section. However, this pattern has also been used in semi-autonomous vehicles for lane-keeping (D5) and in teleoperated robots to take over control when the human performs poorly (R5). In this way, vigilant savior systems assume a deficiency in the human's capabilities and override the human's control to prevent them from making mistakes, saving the day.

*5.2.2 Supportive Patron.* Like the vigilant savior, the *supportive patron* pattern has the AI supervise the human and support them with their tasks as needed. However, this pattern uses *Combined* control and the AI does not necessarily interpret to the human's commands. Rather, the AI is *Supportive* of the human, who is supervised *By AI*, and their control signals are *Combined* to perform the same tasks at the same time. The pattern is used across a range of Influence styles. For example, Deng et al. created a smart power wheelchair (M3) that refines the human's control by blending it with the control signal of an autonomous path planner that moves away from obstacles. Should the user command their wheelchair to move towards an obstacle, blending nullifies the user's command and prevents the collision. When the wheelchair gets too close to an obstacle, the planner adjusts the user's course to give them the space they need to operate the wheelchair safely. This is similar to how the comanipulation robot of Gurnel et al. (S8) pushes and rotates the surgeon's implements to guide them towards a desired position and orientation. Instead of stepping in to replace the human's control signal, this pattern's AI uses shared control to make the human's tasks easier by performing the tasks at the same time.

*5.2.3 Compromise Negotiator.* Many AI agents both interpret the human's control signal and use their own to guide the human. These provide a form of *Codependent* and *Combined* shared control that Abbink et al. have previously called haptic shared control [2]. For example, the haptic shared control vehicle presented by Johns et al. (D3) enables both actors to steer and communicate their intentions using the steering wheel (Codependent Influence). Both the human and AI negotiate the shared control signal by simultaneously applying forces to the system's physical interface (Combined Mediation). However, as described in Section 5.1.3, the pattern encompasses other ways in which AI's control signal can be communicated to the human. The *compromise negotiator* pattern encompasses these haptic shared control systems, but is instead defined in terms of the Mediation and Influence of actor's control signals. The AI interprets and guides the human's control to negotiate a shared control signal that both actors find agreeable.

As seen in Figure 4, numerous systems with differing AI Role and Supervision styles follow the Compromise Negotiator pattern. For example, the digital airbrush of Shilkrot et al. (C12) selects which color to paint and supervises the user, negotiating how much paint to apply. The H-Mode car (D9) is supervised by the human, who coopts the AI's tasks, negotiating how quickly the vehicle accelerates and turns. Many surgical robots (i.e., S1-5) use this pattern to negotiate how the surgeon handles their implements.

*5.2.4 Equal Partner.* When AI agents share control to perform their own tasks that are *Complementary* to the human's, they can extend the human's capabilities by performing tasks that the human cannot, or partition the human's tasks to make the human's job

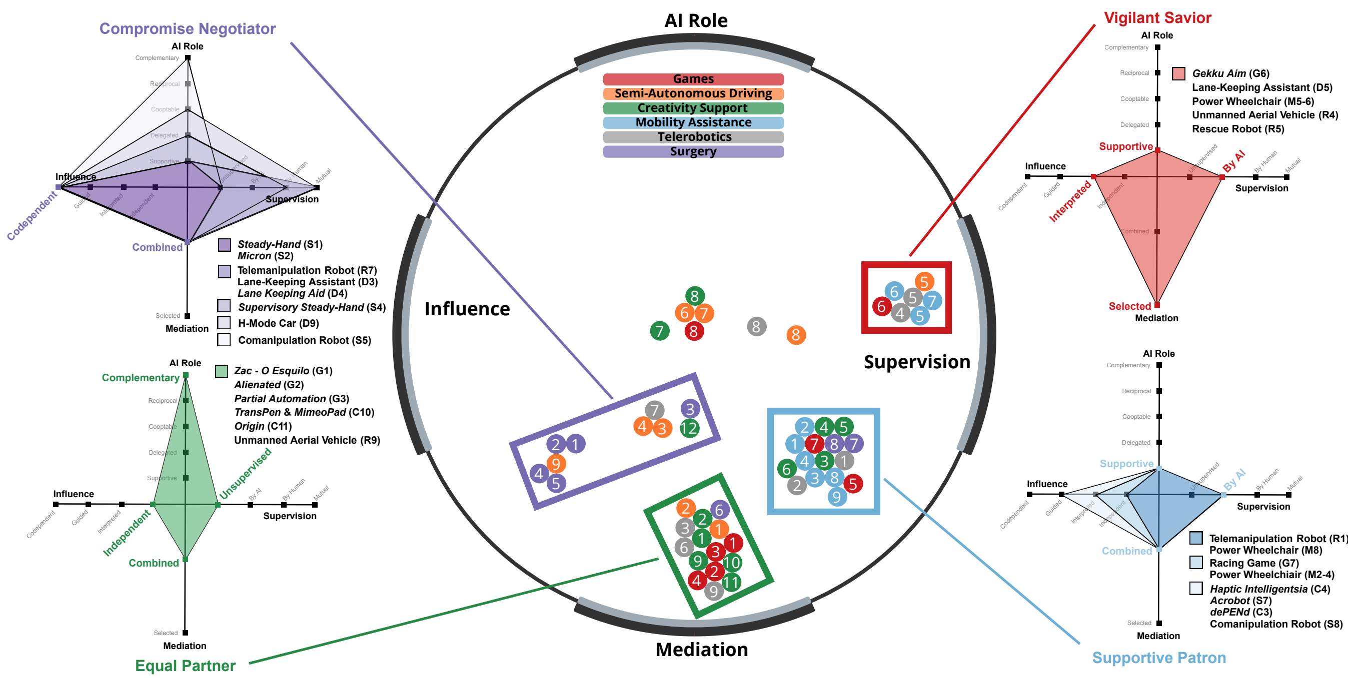


**Figure 4: The *data view* RadViz plot generated by the *Subspaces Explorer* tool and higher-level design patterns used across multiple domains. In the data view in the center, each of the systems in our corpus are represented as circles colored according to their domains and plotted according to their similarities with other systems as well as the axes along which they are similar. For example, the cluster in the top right corner use a distinctive combination of Influence and Mediation, while the cluster at the bottom places the AI in a unique AI Role. Systems found in the center of the plot did not exhibit a cohesive pattern, so this cluster was not selected. Selected patterns are encased in colored rectangles that each correspond to a diagram of the same color on the right and left.**

easier. For example, a mobility assistance robot could brake to catch the human if they fall (M9) or an endomicroscopy robot could rotate a scanner while the human translates it across a sample (S5). Since the AI performs tasks that the human does not, their control signals are often *Independent*, *Combined* with equal control authority, and *Unsupervised*. This form of shared control frames the AI as an *equal partner* who simultaneously performs its own tasks. It has been used in each of the domains we surveyed, save for semi-autonomous driving.

## 6 FUTURE OF SHARED CONTROL

Human-AI shared control can personalize the control of systems to the abilities of users, by leveraging the abilities of one actor to extend the abilities of the other. It unifies human and artificial intelligence, enabling humans to play and create without barriers, get around more safely, and work more effectively. The systems we surveyed demonstrate how shared control can improve human users' interactions with computers in tasks as diverse as surgery and sketching. We turn our attention now to the future of shared control and demonstrate how our dimension space might help designers to explore the design choices available to them.

We have shown how our dimension space supports analysis of shared control systems at different scales. Applying the dimension space to individual systems (as in Figure 1) enables designers to understand and compare the specific human-agent interactions they afford. By using the dimension space to classify a more diverse sample of systems (as in Section 5.1), designers can gain insights into how broadly applicable design patterns can address problems encountered in their own domain. However, there is still one further use case for our dimension space that we demonstrate in this section. Once known solutions are classified, designers can use the dimension space to imagine how other types of Supervision, AI Role, Mediation, and Influence might shape the interactions that systems afford. In this section, we provide examples of designs generated with aid of the dimension space.

We propose novel designs for shared control systems in three of the domains we surveyed. Our proposals, depicted in Figure 5, are largely speculative and there may be concrete human factors or technical problems that preclude their creation. Section 6.1 describes how the Equal Partner pattern, modified to provide *Interpreted* Influence, might make driving more accessible to persons with disabilities. Section 6.2 describes how *Cooptable* AI, similar to the *High-to-Low* pattern used in semi-autonomous driving, might overcome an emerging problem in the performing arts. Section 6.3 describes how a *Guided* telemanipulation system might help users to understand the AI's assistance.

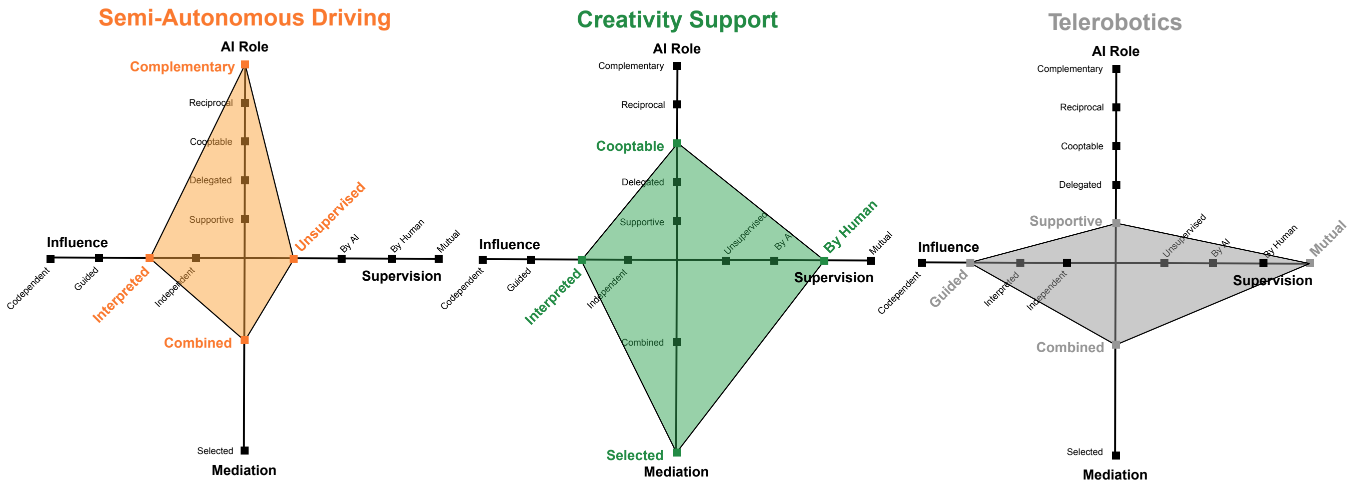


**Figure 5: Kiviat diagrams representing the three systems described in Section 6.**

## 6.1 Semi-Autonomous Driving

Designers of semi-autonomous vehicles tend to approach the problem from one of two perspectives: removing manual control from driving (low-to-high automation) or introducing manual control into autonomous driving (high-to-low automation). However, both assume a one-dimensional shift in control authority between human and AI. These approaches are appropriate for persons without disabilities, but the advent of autonomous vehicles has also sparked discussion regarding how they might better serve persons with disabilities. Brewer & Kameswaran asked persons with vision impairments how automation might enable them to drive and discovered that they wanted autonomous vehicles to present drivers with a "*spectrum of desired control*" [12]. They wanted to personalize their control of the vehicle based on their abilities, rather than designers' expectations of their abilities.

We did not encounter any systems that partition driving tasks to make driving more accessible to drivers with disabilities. For example, many people with spinal cord injury find aspects of driving inaccessible. They may customize their vehicles with hand controls instead of pedals or pedal controls instead of a steering wheel, but for some there may be no configuration of assistive technologies that make all tasks safe and accessible. Using shared control, designers could automate different aspects of the driver's task separately and drivers with spinal cord injury could perform whichever tasks are accessible to them. The AI could perform an entirely *Complementary* role, filling in for the driver by performing tasks that they cannot. Using *Interpreted* Influence, the AI could look for clues in the driver's commands that indicate what the driver wants to do next. For example, the AI could change lanes when the driver accelerates towards slower vehicles ahead. In accordance with the equal partner pattern, in which the AI performs a *Complementary* role, it may be appropriate for neither actor to supervise. Just as input automation controls inaccessible inputs to make games more accessible, AI performing a *Complementary* role could make driving more accessible to drivers with motor disabilities.

## 6.2 Creativity Support

*The Under Presents: Tempest* is a production of William Shakespeare's *The Tempest* staged live in virtual reality (VR). A lone actor, of the thespian variety in this case, plays multiple roles throughout the performance by controlling virtual character models, although they can only play one part at a time. Were the performance to require that multiple parts be played simultaneously, then multiple actors would be needed. Instead, productions could be scaled up dramatically by sharing control of each virtual character with an AI agent that performs scripted behaviours autonomously. The cast could supervise these agents and selectively coopt their roles. For example, should an audience member try interacting with a virtual character without any scripted dialogue, a cast member might take control of this character and improvise. While the character is being controlled, the AI could monitor the cast member's movement, to infer which actions they are doing, and continue doing those actions when control is relinquished. In this way, these *Cooptable* agents are supervised *By Human* and the human actor's commands are *Interpreted* when they are *Selected* to control the character.

## 6.3 Telerobotics

Telerobots are operated at a distance, so human users may have lower awareness of the robots' environment than the robot itself. For example, operators of a robot arm may be unsure whether the arm's gripper has successfully picked up an object using visual feedback alone [76]. For this reason, haptic feedback has been used in telemanipulation systems to improve operators' telepresence. However, we encountered no telerobotic systems that use haptic feedback to share control. This is unfortunate because controlling these robots may be difficult [21, 22, 81] and unsafe in some environments [76]. Sharing control may help operators to avoid making mistakes (e.g., R1-2), but without adequate awareness of how the AI has amended their commands, operators may find control confusing. Instead, telerobots could be *Supportive* of operators and inform them of the AI's commands, such that their supervision of each other is *Mutual* and *Guided* using force feedback. For example, systems in which actors' joystick commands are *Combined* could position the user's joystick to reflect the *Combined* command the system received. Users could sense that their command is dangerous, not because the robot's sensors indicate they might hit an object but because the AI has already prevented it. This approach may improve operators' awareness of the robot's environment, actions, and intentions while also providing the improved task performance afforded by haptic feedback.

Throughout this work, we have described many systems that solve similar problems in significantly different ways using shared control. We have also shown how designing novel uses of shared control benefits from our existing design knowledge and exposes new problems that we know little about. It has only been in the last decade that shared control has emerged as a way of controlling interactive systems beyond the relatively narrow design space of robots. These recent uses of shared control in games and creativity support suggest a future in which shared control is used to improve all sorts of human activities. As shown by the speculative examples presented in this section, the dimension space presented in this paper provides designers with the language to express how shared control might overcome new problems. We are now equipped to go beyond the domains where shared control has proven useful, using this dimension space to guide our exploration.

## 7 CONCLUSION

Shared control is being used in radically different ways to improve our lives. It can extend the accessibility of interactive systems to users with motor disabilities, it can help human users to perform tasks more safely, and it can facilitate engagement in creative and playful activities. But our knowledge of how to design shared control systems is often specific to an application domain, and some systems are not identified as using shared control by their designers. This makes it difficult for designers using different approaches to share their design knowledge or take inspiration from solutions used by others. Solutions may look different when they are similar (e.g., sketching and lane-keeping assistants) or they may look similar when they are different (e.g., surgical and teleoperated robots). Therefore, designers may not see whether solutions used in other domains can be used to overcome problems in their own. To break down barriers between domains using shared control, we need a common language for describing the design space of shared control systems.

In this paper we presented a *dimension space* for shared control defined along four axes: *AI Role*, *Supervision*, *Influence*, and *Mediation*. It enables designers to classify the human-AI interactions a system affords, make comparisons with other systems, and imagine novel approaches to shared control that have never been tried before. Using this simple language, designers may be able to better communicate their design ideas to others, allowing their knowledge of the shared control design space to be more easily shared. With this dimension space as our guide, shared control promises a future in which humans' interactions with computers are more accessible, safe, creative, and playful.

## REFERENCES


[1] David A. Abbink, Tom Carlson, Mark Mulder, Joost C. F. de Winter, Farzad Aminravan, Tricia L. Gibo, and Erwin R. Boer. 2018. A Topology of Shared Control Systems—Finding Common Ground in Diversity. *IEEE Transactions on Human-Machine Systems* 48, 5 (Oct. 2018), 509–525. https://doi.org/10.1109/THMS.2018.2791570 Conference Name: IEEE Transactions on Human-Machine Systems.

[2] David A. Abbink, Mark Mulder, and Erwin R. Boer. 2012. Haptic shared control: smoothly shifting control authority? *Cognition, Technology & Work* 14, 1 (March 2012), 19–28. https://doi.org/10.1007/s10111-011-0192-5

[3] Christopher Alexander. 1979. *The timeless way of building.* Oxford University Press, New York.

[4] Christopher Alexander, Sara Ishikawa, and Murray Silverstein. 1977. *A pattern language: towns, buildings, construction.* Oxford University Press, New York.

[5] J.E. Allen, C.I. Guinn, and E. Horvtz. 1999. Mixed-initiative interaction. *IEEE Intelligent Systems and their Applications* 14, 5 (Sept. 1999), 14–23. https://doi.org/10.1109/5254.796083 Conference Name: IEEE Intelligent Systems and their Applications.

[6] Edoardo Aromataris and Zachary Munn (Eds.). 2020. *JBI Manual for Evidence Synthesis.* JBI. https://doi.org/10.46658/JBIMES-20-01

[7] Erasmo Artur and Rosane Minghim. 2019. A novel visual approach for enhanced attribute analysis and selection. *Computers & Graphics* 84 (Nov. 2019), 160–172. https://doi.org/10.1016/j.cag.2019.08.015

[8] Lisanne Bainbridge. 1983. Ironies of automation. *Automatica* 19, 6 (1983), 775–779. https://doi.org/10.1016/0005-1098(83)90046-8

[9] Marcel Caspar Attila Baltzer, Daniel López, and Frank Flemisch. 2019. Towards an interaction pattern language for human machine cooperation and cooperative movement. *Cognition, Technology & Work* 21, 4 (Nov. 2019), 593–606. https://doi.org/10.1007/s10111-019-00561-8

[10] Scott Bateman, Regan L. Mandryk, Tadeusz Stach, and Carl Gutwin. 2011. Target assistance for subtly balancing competitive play. In *Proceedings of the 2011 annual conference on Human factors in computing systems - CHI '11.* ACM Press, Vancouver, BC, Canada, 2355. https://doi.org/10.1145/1978942.1979287

[11] D. Birnbaum, R. Fiebrink, J. Malloch, and M. M. Wanderley. 2005. Towards a dimension space for musical devices. In *Proceedings of the 2005 conference on New interfaces for musical expression (NIME '05).* National University of Singapore, SGP, 192–195.

[12] Robin N. Brewer and Vaishnav Kameswaran. 2018. Understanding the Power of Control in Autonomous Vehicles for People with Vision Impairment. In *Proceedings of the 20th International ACM SIGACCESS Conference on Computers and Accessibility (ASSETS '18).* Association for Computing Machinery, New York, NY, USA, 185–197. https://doi.org/10.1145/3234695.3236347

[13] Jared E. Cechanowicz, Carl Gutwin, Scott Bateman, Regan Mandryk, and Ian Stavness. 2014. Improving player balancing in racing games. In *Proceedings of the first ACM SIGCHI annual symposium on Computer-human interaction in play (CHI PLAY '14).* Association for Computing Machinery, New York, NY, USA, 47–56. https://doi.org/10.1145/2658537.2658701

[14] Gabriele Cimolino, Sussan Askari, and T.C. Nicholas Graham. 2021. The Role of Partial Automation in Increasing the Accessibility of Digital Games. In *Proceedings of the Annual Symposium on Computer-Human Interaction in Play (CHI PLAY '21).* ACM, New York, NY, USA.

[15] Sylvain Daronnat. 2020. Human-Agent Trust Relationships in a Real-Time Collaborative Game. In *Extended Abstracts of the 2020 Annual Symposium on Computer-Human Interaction in Play (CHI PLAY '20).* Association for Computing Machinery, New York, NY, USA, 18–20. https://doi.org/10.1145/3383668.3419953

[16] Sylvain Daronnat, Leif Azzopardi, and Martin Halvey. 2020. Impact of agents' errors on performance, reliance and trust in human-agent collaboration. 1–5. https://www.hfes2020.com/

[17] Sylvain Daronnat, Leif Azzopardi, Martin Halvey, and Mateusz Dubiel. 2020. Impact of Agent Reliability and Predictability on Trust in Real Time Human-Agent Collaboration. In *Proceedings of the 8th International Conference on Human-Agent Interaction (HAI '20).* Association for Computing Machinery, New York, NY, USA, 131–139. https://doi.org/10.1145/3406499.3415063

[18] Xiaoyan Deng, Zhu Liang Yu, Canguang Lin, Zhenghui Gu, and Yuanqing Li. 2020. Self-adaptive shared control with brain state evaluation network for human-wheelchair cooperation. *Journal of Neural Engineering* 17, 4 (July 2020), 045005. https://doi.org/10.1088/1741-2552/ab937e Publisher: IOP Publishing.

[19] DJI. [n. d.]. Mavic Air 2 - Up Your Game. Retrieved 2022-01-10 from https://www.dji.com/ca/mavic-air-2?site=brandsite&from=nav

[20] Anca D. Dragan, Kenton C.T. Lee, and Siddhartha S. Srinivasa. 2013. Legibility and predictability of robot motion. In *2013 8th ACM/IEEE International Conference on Human-Robot Interaction (HRI).* 301–308. https://doi.org/10.1109/HRI.2013.6483603 ISSN: 2167-2148.

[21] Anca D. Dragan and Siddhartha S. Srinivasa. 2012. Assistive teleoperation for manipulation tasks. In *Proceedings of the seventh annual ACM/IEEE international conference on Human-Robot Interaction (HRI '12).* Association for Computing Machinery, Boston, Massachusetts, USA, 123–124. https://doi.org/10.1145/

2157689.2157716

[22] Anca D Dragan and Siddhartha S Srinivasa. 2013. A policy-blending formalism for shared control. *The International Journal of Robotics Research* 32, 7 (June 2013), 790–805. https://doi.org/10.1177/0278364913490324

[23] Nintendo EAD. 2017. *Mario Kart 8 Deluxe.* Game [Nintendo Switch]. Nintendo, Kyoto, Japan..

[24] Jakob Engel, Jürgen Sturm, and Daniel Cremers. 2014. Scale-aware navigation of a low-cost quadrocopter with a monocular camera. *Robotics and Autonomous Systems* 62, 11 (Nov. 2014), 1646–1656. https://doi.org/10.1016/j.robot.2014.03.012

[25] Chinemelu Ezeh, Pete Trautman, Louise Devigne, Valentin Bureau, Marie Babel, and Tom Carlson. 2017. Probabilistic vs linear blending approaches to shared control for wheelchair driving. In *2017 International Conference on Rehabilitation Robotics (ICORR).* 835–840. https://doi.org/10.1109/ICORR.2017.8009352 ISSN: 1945-7901.

[26] Chinemelu Ezeh, Pete Trautman, Catherine Holloway, and Tom Carlson. 2017. Comparing shared control approaches for alternative interfaces: A wheelchair simulator experiment. In *2017 IEEE International Conference on Systems, Man, and Cybernetics (SMC).* 93–98. https://doi.org/10.1109/SMC.2017.8122584

[27] F. Flemisch, D. Abbink, M. Itoh, M-P. Pacaux-Lemoine, and G. Weßel. 2016. Shared control is the sharp end of cooperation: Towards a common framework of joint action, shared control and human machine cooperation. *IFAC-PapersOnLine* 49, 19 (Jan. 2016), 72–77. https://doi.org/10.1016/j.ifacol.2016.10.464

[28] 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

[29] Frank Flemisch, Johann Kelsch, Christan Löper, Anna Schieben, Julian Schindler, and Matthias Heesen. 2008. Cooperative Control and Active Interfaces for Vehicle Assitsance and Automation. München. https://elib.dlr.de/57618/

[30] Frank Ole Flemisch, Klaus Bengler, Heiner Bubb, Hermann Winner, and Ralph Bruder. 2014. Towards cooperative guidance and control of highly automated vehicles: H-Mode and Conduct-by-Wire. *Ergonomics* 57, 3 (March 2014), 343–360. https://doi.org/10.1080/00140139.2013.869355 Publisher: Taylor & Francis _eprint: https://doi.org/10.1080/00140139.2013.869355.

[31] Frank O. Adams Flemisch. 2003. *The H-Metaphor as a Guideline for Vehicle Automation and Interaction.* Technical Report. https://ntrs.nasa.gov/search.jsp?R=20040031835

[32] Antonio Franchi, Cristian Secchi, Markus Ryll, Heinrich H. Bulthoff, and Paolo Robuffo Giordano. 2012. Shared Control : Balancing Autonomy and Human Assistance with a Group of Quadrotor UAVs. *IEEE Robotics Automation Magazine* 19, 3 (Sept. 2012), 57–68. https://doi.org/10.1109/MRA.2012.2205625

[33] Anna-Katharina Frison, Philipp Wintersberger, Andreas Riener, and Clemens Schartmüller. 2017. Driving Hotzenplotz: A Hybrid Interface for Vehicle Control Aiming to Maximize Pleasure in Highway Driving. In *Proceedings of the 9th International Conference on Automotive User Interfaces and Interactive Vehicular Applications (AutomotiveUI '17).* Association for Computing Machinery, New York, NY, USA, 236–244. https://doi.org/10.1145/3122986.3123016

[34] Milad Geravand, Christian Werner, Klaus Hauer, and Angelika Peer. 2016. An Integrated Decision Making Approach for Adaptive Shared Control of Mobility Assistance Robots. *International Journal of Social Robotics* 8, 5 (Nov. 2016), 631–648. https://doi.org/10.1007/s12369-016-0353-z

[35] T. C. Nicholas Graham, Leon A. Watts, Gaëlle Calvary, Joëlle Coutaz, Emmanuel Dubois, and Laurence Nigay. 2000. A dimension space for the design of interactive systems within their physical environments. In *Proceedings of the 3rd conference on Designing interactive systems: processes, practices, methods, and techniques (DIS '00).* Association for Computing Machinery, New York, NY, USA, 406–416. https://doi.org/10.1145/347642.347799

[36] Hadrien Gurnel, Maud Marchal, Laurent Launay, Luc Beuzit, and Alexandre Krupa. 2019. Design of haptic guides for pre-positioning assistance of a comanipulated needle. In *2019 IEEE International Conference on Systems, Man and Cybernetics (SMC).* 478–485. https://doi.org/10.1109/SMC.2019.8914395 ISSN: 2577-1655.

[37] Marco Antonio Gutierrez, Luis Fernando D'Haro, and Rafael Banchs. 2016. A Multimodal Control Architecture for Autonomous Unmanned Aerial Vehicles. In *Proceedings of the Fourth International Conference on Human Agent Interaction (HAI '16).* Association for Computing Machinery, New York, NY, USA, 107–110. https://doi.org/10.1145/2974804.2980522

[38] Ian Hattwick and Marcelo Wanderley. 2012. A Dimension Space for Evaluating Collaborative Musical Performance Systems. In *Proceedings of the International Conference on New Interfaces for Musical Expression.* University of Michigan, Ann Arbor, Michigan. https://doi.org/10.5281/zenodo.1178281

[39] Eric Horvitz. 1999. Mixed-Initiative Interaction. *IEEE Intelligent Systems* (Sept. 1999), 14–24. https://www.microsoft.com/en-us/research/publication/mixed-initiative-interaction/ Edition: IEEE Intelligent Systems.

[40] Susan Hwang, Adrian L. Jessup Schneider, Daniel Clarke, Alexander Macintosh, Lauren Switzer, Darcy Fehlings, and T.C. Nicholas Graham. 2017. How Game Balancing Affects Play: Player Adaptation in an Exergame for Children with Cerebral Palsy. In *Proceedings of the 2017 Conference on Designing Interactive Systems - DIS '17.* ACM Press, Edinburgh, United Kingdom, 699–710. https://doi.org/10.1145/3064663.3064664

[41] Toshiyuki Inagaki. 2003. Adaptive automation: Sharing and trading of control. In *Handbook of cognitive task design.* Lawrence Erlbaum Associates Publishers, Mahwah, NJ, US, 147–169. https://doi.org/10.1201/9781410607775.ch8

[42] Vahid Izadi, Arjun Yeravdekar, and Amirhossein Ghasemi. 2019. Determination of Roles and Interaction Modes in a Haptic Shared Control Framework. American Society of Mechanical Engineers Digital Collection. https://doi.org/10.1115/DSCC2019-9042

[43] Matjaž Jakopec, Simon J. Harris, Ferdinando Rodriguez y Baena, Paula Gomes, Justin Cobb, and Brian L. Davies. 2002. Preliminary Results of an Early Clinical Experience with the Acrobot™ System for Total Knee Replacement Surgery. In *Medical Image Computing and Computer-Assisted Intervention — MICCAI 2002 (Lecture Notes in Computer Science)*, Takeyoshi Dohi and Ron Kikinis (Eds.). Springer, Berlin, Heidelberg, 256–263. https://doi.org/10.1007/3-540-45786-0_32

[44] Jingjing Jiang and Alessandro Astolfi. 2015. Shared-control for a UAV operating in the 3D space. In *2015 European Control Conference (ECC).* 1633–1638. https://doi.org/10.1109/ECC.2015.7330771

[45] Shu D. Jiang and Jonathan Odom. 2018. Toward Initiative Decision-Making for Distributed Human-Robot Teams. In *Proceedings of the 6th International Conference on Human-Agent Interaction (HAI '18).* Association for Computing Machinery, New York, NY, USA, 286–292. https://doi.org/10.1145/3284432.3284467

[46] Mishel Johns, Brian Mok, David Sirkin, Nikhil Gowda, Catherine Smith, Walter Talamonti, and Wendy Ju. 2016. Exploring shared control in automated driving. In *2016 11th ACM/IEEE International Conference on Human-Robot Interaction (HRI).* 91–98. https://doi.org/10.1109/HRI.2016.7451738 ISSN: 2167-2148.

[47] Gavin Kane, Georg Eggers, Robert Boesecke, Jörg Raczkowsky, Heinz Wörn, Rüdiger Marmulla, and Joachim Mühling. 2009. System Design of a Hand-Held Mobile Robot for Craniotomy. In *Medical Image Computing and Computer-Assisted Intervention – MICCAI 2009 (Lecture Notes in Computer Science)*, Guang-Zhong Yang, David Hawkes, Daniel Rueckert, Alison Noble, and Chris Taylor (Eds.). Springer, Berlin, Heidelberg, 402–409. https://doi.org/10.1007/978-3-642-04268-3_50

[48] Ankur Kapoor, Rajesh Kumar, and Russell H. Taylor. 2003. Simple Biomanipulation Tasks with "Steady Hand" Cooperative Manipulator. In *Medical Image Computing and Computer-Assisted Intervention - MICCAI 2003 (Lecture Notes in Computer Science)*, Randy E. Ellis and Terry M. Peters (Eds.). Springer, Berlin, Heidelberg, 141–148. https://doi.org/10.1007/978-3-540-39899-8_18

[49] Shunichi Kasahara, Jun Nishida, and Pedro Lopes. 2019. Preemptive Action: Accelerating Human Reaction using Electrical Muscle Stimulation Without Compromising Agency. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems.* Association for Computing Machinery, New York, NY, USA, 1–15. https://doi.org/10.1145/3290605.3300873

[50] Soheil Kianzad, Yuxiang Huang, Robert Xiao, and Karon E. MacLean. 2020. Phasking on Paper: Accessing a Continuum of PHysically Assisted SKetchING. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems (CHI '20).* Association for Computing Machinery, New York, NY, USA, 1–12. https://doi.org/10.1145/3313831.3376134

[51] Hyunjung Kim, Seoktae Kim, Boram Lee, Jinhee Pak, Minjung Sohn, Geehyuk Lee, and Woohun Lee. 2008. Digital rubbing: playful and intuitive interaction technique for transferring a graphic image onto paper with pen-based computing. In *CHI '08 Extended Abstracts on Human Factors in Computing Systems (CHI EA '08).* Association for Computing Machinery, New York, NY, USA, 2337–2342. https://doi.org/10.1145/1358628.1358680

[52] Holger Klapperich and Marc Hassenzahl. 2016. Hotzenplotz: Reconciling Automation with Experience. In *Proceedings of the 9th Nordic Conference on Human-Computer Interaction (NordiCHI '16).* Association for Computing Machinery, New York, NY, USA, 1–10. https://doi.org/10.1145/2971485.2971532

[53] Moritz Körber, Eva Baseler, and Klaus Bengler. 2018. Introduction matters: Manipulating trust in automation and reliance in automated driving. *Applied Ergonomics* 66 (Jan. 2018), 18–31. https://doi.org/10.1016/j.apergo.2017.07.006

[54] Jung Min Lee and Da Young Ju. 2018. Classification of Human-Vehicle Interaction: User Perspectives on Design. *Social Behavior and Personality: an international journal* 46, 7 (July 2018), 1057–1070. https://doi.org/10.2224/sbp.6242

[55] M. P. Lemoine, S. Debernard, I. Crevits, and P. Millot. 1996. Cooperation between humans and machines: First results of an experiment with a multi-level cooperative organisation in air traffic control. *Computer Supported Cooperative Work (CSCW)* 5, 2 (June 1996), 299–321. https://doi.org/10.1007/BF00133661

[56] Qinan Li, Weidong Chen, and Jingchuan Wang. 2011. Dynamic shared control for human-wheelchair cooperation. In *2011 IEEE International Conference on Robotics and Automation.* 4278–4283. https://doi.org/10.1109/ICRA.2011.5980055 ISSN: 1050-4729.

[57] Ana C. Lopes, Urbano Nunes, Luis Vaz, and Luís Vaz. 2010. Assisted navigation based on shared-control, using discrete and sparse human-machine interfaces.

In *2010 Annual International Conference of the IEEE Engineering in Medicine and Biology*. 471–474. https://doi.org/10.1109/IEMBS.2010.5626221 ISSN: 1558-4615.

[58] Ana C. Lopes, Gabriel Pires, Luís Vaz, and Urbano Nunes. 2011. Wheelchair navigation assisted by Human-Machine shared-control and a P300-based Brain Computer Interface. In *2011 IEEE/RSJ International Conference on Intelligent Robots and Systems*. 2438–2444. https://doi.org/10.1109/IROS.2011.6094748 ISSN: 2153-0866.

[59] Pedro Lopes, Alexandra Ion, Willi Mueller, Daniel Hoffmann, Patrik Jonell, and Patrick Baudisch. 2015. Proprioceptive Interaction. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems (CHI '15)*. Association for Computing Machinery, New York, NY, USA, 939–948. https://doi.org/10.1145/2702123.2702461

[60] Pedro Lopes, Doăa Yüksel, François Guimbretière, and Patrick Baudisch. 2016. Muscle-plotter: An Interactive System based on Electrical Muscle Stimulation that Produces Spatial Output. In *Proceedings of the 29th Annual Symposium on User Interface Software and Technology (UIST '16)*. Association for Computing Machinery, New York, NY, USA, 207–217. https://doi.org/10.1145/2984511.2984530

[61] Robert A. MacLachlan, Brian C. Becker, Jaime Cuevas Tabares, Gregg W. Podnar, Louis A. Lobes, and Cameron N. Riviere. 2012. Micron: An Actively Stabilized Handheld Tool for Microsurgery. *IEEE Transactions on Robotics* 28, 1 (Feb. 2012), 195–212. https://doi.org/10.1109/TRO.2011.2169634 Conference Name: IEEE Transactions on Robotics.

[62] Allan Maclean and Diane McKerlie. 1995. Design space analysis and use representations. In *Scenario-based design: envisioning work and technology in system development*. John Wiley & Sons, Inc., USA, 183–207.

[63] Lucas Medeiros and Flavio Coutinho. 2015. Developing an Accessible One-Switch Game for Motor Impaired Players. In *Proceedings of SBGames 2015*. Teresina, PI, Brazil, 236–239.

[64] John A. Michon. 1985. A Critical View of Driver Behavior Models: What Do We Know, What Should We Do? In *Human Behavior and Traffic Safety*, Leonard Evans and Richard C. Schwing (Eds.). Springer US, Boston, MA, 485–524. https://doi.org/10.1007/978-1-4613-2173-6_19

[65] P. Millot and A. Kamoun. 1988. An Implicit Method for Dynamic Task Allocation Between Man and Computer in Supervision Posts of Automated Processes. *IFAC Proceedings Volumes* 21, 5 (June 1988), 77–82. https://doi.org/10.1016/S1474-6670(17)53885-9

[66] Patrick Millot and Marie-Pierre Pacaux-Lemoine. 2013. A Common Work Space for a mutual enrichment of Human-machine Cooperation and Team-Situation Awareness. *IFAC Proceedings Volumes* 46, 15 (Jan. 2013), 387–394. https://doi.org/10.3182/20130811-5-US-2037.00061

[67] Selma Musić and Sandra Hirche. 2016. Classification of human-robot team interaction paradigms**The research leading to these results has received funding from the European Union Seventh Framework Programme FP7/2007- 2013 under grant agreement no. 601165 of the project: WEARHAP -WEARable HAPtics for humans and robots. *IFAC-PapersOnLine* 49, 32 (Jan. 2016), 42–47. https://doi.org/10.1016/j.ifacol.2016.12.187

[68] Selma Musić and Sandra Hirche. 2017. Control sharing in human-robot team interaction. *Annual Reviews in Control* 44 (Jan. 2017), 342–354. https://doi.org/10.1016/j.arcontrol.2017.09.017

[69] Jun Nishida, Shunichi Kasahara, and Kenji Suzuki. 2017. Wired muscle: generating faster kinesthetic reaction by inter-personally connecting muscles. In *ACM SIGGRAPH 2017 Emerging Technologies (SIGGRAPH '17)*. Association for Computing Machinery, New York, NY, USA, 1–2. https://doi.org/10.1145/3084822.3084844

[70] Yoojin Oh, Shao-Wen Wu, Marc Toussaint, and Jim Mainprice. 2020. Natural Gradient Shared Control. In *2020 29th IEEE International Conference on Robot and Human Interactive Communication (RO-MAN)*. 1223–1229. https://doi.org/10.1109/RO-MAN47096.2020.9223465 ISSN: 1944-9437.

[71] On-Road Automated Driving (ORAD) committee. 2021. *Taxonomy and Definitions for Terms Related to Driving Automation Systems for On-Road Motor Vehicles*. Standard J3016_202104. SAE International. 41 pages. https://doi.org/10.4271/J3016_202104

[72] M. P. Pacaux-Lemoine and S. Debernard. 2000. A Common Work Space to Support the Air Traffic Control. *IFAC Proceedings Volumes* 33, 12 (June 2000), 75–78. https://doi.org/10.1016/S1474-6670(17)37280-4

[73] Marie-Pierre Pacaux-Lemoine and Frank Flemisch. 2019. Layers of shared and cooperative control, assistance, and automation. *Cognition, Technology & Work* 21, 4 (Nov. 2019), 579–591. https://doi.org/10.1007/s10111-018-0537-4

[74] Marie-Pierre Pacaux-Lemoine and Itoh Makoto. 2015. Towards Vertical and Horizontal Extension of Shared Control Concept. In *2015 IEEE International Conference on Systems, Man, and Cybernetics*. 3086–3091. https://doi.org/10.1109/SMC.2015.536

[75] 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 (May 2000), 286–297. https://doi.org/10.1109/3468.844354 Conference Name: IEEE Transactions on Systems, Man, and Cybernetics - Part A: Systems and Humans.

[76] Tae Mun Park, Seung Yeon Won, Sang Ryong Lee, and Gabor Sziebig. 2016. Force feedback based gripper control on a robotic arm. In *2016 IEEE 20th Jubilee International Conference on Intelligent Engineering Systems (INES)*. 107–112. https://doi.org/10.1109/INES.2016.7555102

[77] Christopher. J. Payne, Win Tun Latt, and Guang-Zhong Yang. 2012. A new hand-held force-amplifying device for micromanipulation. In *2012 IEEE International Conference on Robotics and Automation*. 1583–1588. https://doi.org/10.1109/ICRA.2012.6225306 ISSN: 1050-4729.

[78] Christopher J. Payne, Khushi Vyas, Daniel Bautista-Salinas, Dandan Zhang, Hani J. Marcus, and Guang-Zhong Yang. 2021. Shared-Control Robots. In *Neurosurgical Robotics*, Hani J. Marcus and Christopher J. Payne (Eds.). Springer US, New York, NY, 63–79. https://doi.org/10.1007/978-1-0716-0993-4_4

[79] Bastian Pfleging, Maurice Rang, and Nora Broy. 2016. Investigating user needs for non-driving-related activities during automated driving. In *Proceedings of the 15th International Conference on Mobile and Ubiquitous Multimedia (MUM '16)*. Association for Computing Machinery, New York, NY, USA, 91–99. https://doi.org/10.1145/3012709.3012735

[80] William Pilgrim. 2008. *Alienated*. Game [Windows]. OneSwitch.org.uk, Essex, UK..

[81] Sanzhar Rakhimkul, Anton Kim, Askarbek Pazylbekov, and Almas Shintemirov. 2019. Autonomous Object Detection and Grasping Using Deep Learning for Design of an Intelligent Assistive Robot Manipulation System. In *2019 IEEE International Conference on Systems, Man and Cybernetics (SMC)*. 3962–3968. https://doi.org/10.1109/SMC.2019.8914465 ISSN: 2577-1655.

[82] J. Rasmussen. 1983. Skills, rules, and knowledge; signals, signs, and symbols, and other distinctions in human performance models. *IEEE Transactions on Systems, Man, and Cybernetics* SMC-13, 3 (May 1983), 257–266. https://doi.org/10.1109/TSMC.1983.6313160

[83] Alec Rivers, Ilan E. Moyer, and Frédo Durand. 2012. Position-correcting tools for 2D digital fabrication. *ACM Transactions on Graphics* 31, 4 (July 2012), 88:1–88:7. https://doi.org/10.1145/2185520.2185584

[84] Takuma Seno, Kohei Okuoka, Masahiko Osawa, and Michita Imai. 2018. Adaptive Semi-autonomous Agents via Episodic Control. In *Proceedings of the 6th International Conference on Human-Agent Interaction (HAI '18)*. Association for Computing Machinery, New York, NY, USA, 377–379. https://doi.org/10.1145/3284432.3287192

[85] Chouki Sentouh, Anh-Tu Nguyen, Jrme Floris, and Jean-Christophe Popieul. 2018. Multiple Controller Switching Concept for Human-Machine Shared Control of Lane Keeping Assist Systems. In *2018 IEEE International Conference on Systems, Man, and Cybernetics (SMC)*. 2730–2735. https://doi.org/10.1109/SMC.2018.00466 ISSN: 2577-1655.

[86] Thomas B. Sheridan. 1992. *Telerobotics, automation, and human supervisory control*. MIT Press, Cambridge, Mass.

[87] Thomas B. Sheridan and William L. Verplank. 1978. *Human and Computer Control of Undersea Teleoperators*. Technical Report. MASSACHUSETTS INST OF TECH CAMBRIDGE MAN-MACHINE SYSTEMS LAB. https://apps.dtic.mil/sti/citations/ADA057655 Section: Technical Reports.

[88] Roy Shilkrot, Pattie Maes, Joseph A. Paradiso, and Amit Zoran. 2015. Augmented Airbrush for Computer Aided Painting (CAP). *ACM Transactions on Graphics* 34, 2 (March 2015), 19:1–19:11. https://doi.org/10.1145/2699649

[89] Gillian Smith, Jim Whitehead, and Michael Mateas. 2010. Tanagra: a mixed-initiative level design tool. In *Proceedings of the Fifth International Conference on the Foundations of Digital Games (FDG '10)*. Association for Computing Machinery, New York, NY, USA, 209–216. https://doi.org/10.1145/1822348.1822376

[90] Harold Soh and Yiannis Demiris. 2015. Learning assistance by demonstration: smart mobility with shared control and paired haptic controllers. *Journal of Human-Robot Interaction* 4, 3 (Dec. 2015), 76–100. https://doi.org/10.5898/JHRI.4.3.Soh

[91] Micah Steele and R. Brent Gillespie. 2001. Shared Control between Human and Machine: Using a Haptic Steering Wheel to Aid in Land Vehicle Guidance. *Proceedings of the Human Factors and Ergonomics Society Annual Meeting* 45, 23 (Oct. 2001), 1671–1675. https://doi.org/10.1177/154193120104502323 Publisher: SAGE Publications Inc.

[92] Stuart Morris, Gail Fawcett, Laurent Brisebois, and Jeffrey Hughes. 2018. *A demographic, employment and income profile of Canadians with disabilities aged 15 years and over, 2017*. Technical Report. Statistics Canada. https://www150.statcan.gc.ca/n1/pub/89-654-x/89-654-x2018002-eng.htm

[93] Joshua P. Switkes, Eric J. Rossetter, Ian A. Coe, and J. Christian Gerdes. 2006. Handwheel Force Feedback for Lanekeeping Assistance: Combined Dynamics and Stability. *Journal of Dynamic Systems, Measurement, and Control* 128, 3 (Sept. 2006), 532–542. https://doi.org/10.1115/1.2229256

[94] K.A. Tahboub. 2001. Natural and manmade shared-control systems: an overview. In *Proceedings 2001 ICRA. IEEE International Conference on Robotics and Automation (Cat. No.01CH37164)*, Vol. 3. 2655–2660 vol.3. https://doi.org/10.1109/ROBOT.2001.933023 ISSN: 1050-4729.

[95] Emi Tamaki, Takashi Miyaki, and Jun Rekimoto. 2011. PossessedHand: techniques for controlling human hands using electrical muscles stimuli. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '11)*. Association for Computing Machinery, New York, NY, USA, 543–552. https://doi.org/10.1145/1978942.1979018

[96] Russell Taylor, Patrick Jensen, Louis Whitcomb, Aaron Barnes, Rajesh Kumar, Dan Stoianovici, Puneet Gupta, ZhengXian Wang, Eugene deJuan, and Louis Kavoussi. 1999. A Steady-Hand Robotic System for Microsurgical Augmentation. In *Medical Image Computing and Computer-Assisted Intervention – MICCAI'99 (Lecture Notes in Computer Science)*, Chris Taylor and Alain Colchester (Eds.). Springer, Berlin, Heidelberg, 1031–1041. https://doi.org/10.1007/10704282_112

[97] Shan-Yuan Teng, Rong-Hao Liang, Jo-Hsi Tang, Bing-Yu Chen, and Yi-Chi Liao. [n. d.]. Playing Air Guitar by Electrical Muscle Stimulation. *AUI 2016: 1 st Asian Workshop on User Interface* ([n. d.]), 2.

[98] Tesla. 2018. Introducing Navigate on Autopilot. Retrieved 2022-01-10 from https://www.tesla.com/blog/introducing-navigate-autopilot

[99] TIEL. [n. d.]. HAPTIC INTELLIGENTSIA – TIEL. Retrieved 2022-01-10 from https://tiel.ch/project/haptic-intelligentsia/

[100] Pete Trautman. 2015. Assistive Planning in Complex, Dynamic Environments: A Probabilistic Approach. In *2015 IEEE International Conference on Systems, Man, and Cybernetics*. 3072–3078. https://doi.org/10.1109/SMC.2015.534

[101] E. Tunstel, M. Maimone, A. Trebi-Ollennu, J. Yen, R. Petras, and R. Willson. 2005. Mars Exploration Rover mobility and robotic arm operational performance. In *2005 IEEE International Conference on Systems, Man and Cybernetics*, Vol. 2. 1807–1814 Vol. 2. https://doi.org/10.1109/ICSMC.2005.1571410 ISSN: 1062-922X.

[102] Jurriaan van Diggelen and Matthew Johnson. 2019. Team Design Patterns. In *Proceedings of the 7th International Conference on Human-Agent Interaction (HAI '19)*. Association for Computing Machinery, New York, NY, USA, 118–126. https://doi.org/10.1145/3349537.3351892

[103] Rodrigo Vicencio-Moreira, Regan L. Mandryk, Carl Gutwin, and Scott Bateman. 2014. The Effectiveness (or Lack Thereof) of Aim-Assist Techniques in First-Person Shooter Games. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Toronto, Ontario, Canada) *(CHI '14)*. Association for Computing Machinery, New York, NY, USA, 937–946. https://doi.org/10.1145/2556288.2557308

[104] Volvo. 2020. Lane Keeping Aid | Lane Keeping Aid | Driver support | XC60 2021 Early | Volvo Support. Retrieved 2022-01-10 from https://www.volvocars.com/en-ca/support/manuals/xc60/2020w17/driver-support/lane-keeping-aid/lane-keeping-aid

[105] Huanran Wang and Xiaoping P. Liu. 2014. Adaptive Shared Control for a Novel Mobile Assistive Robot. *IEEE/ASME Transactions on Mechatronics* 19, 6 (Dec. 2014), 1725–1736. https://doi.org/10.1109/TMECH.2014.2299213 Conference Name: IEEE/ASME Transactions on Mechatronics.

[106] Gesa Wiegand, Kai Holländer, Katharina Rupp, and Heinrich Hussmann. 2020. The Joy of Collaborating with Highly Automated Vehicles. In *12th International Conference on Automotive User Interfaces and Interactive Vehicular Applications (AutomotiveUI '20)*. Association for Computing Machinery, New York, NY, USA, 223–232. https://doi.org/10.1145/3409120.3410643

[107] Piyamate Wisanuvej, Petros Giataganas, Konrad Leibrandt, Jindong Liu, Michael Hughes, and Guang-Zhong Yang. 2017. Three-dimensional robotic-assisted endomicroscopy with a force adaptive robotic arm. In *2017 IEEE International Conference on Robotics and Automation (ICRA)*. 2379–2384. https://doi.org/10.1109/ICRA.2017.7989276

[108] Junichi Yamaoka and Yasuaki Kakehi. 2013. dePENd: augmented handwriting system using ferromagnetism of a ballpoint pen. In *Proceedings of the 26th annual ACM symposium on User interface software and technology (UIST '13)*. Association for Computing Machinery, New York, NY, USA, 203–210. https://doi.org/10.1145/2501988.2502017

[109] Bei Yuan, Eelke Folmer, and Frederick C. Harris. 2010. Game accessibility: a survey. *Universal Access in the Information Society* 10, 1 (June 2010), 81–100. https://doi.org/10.1007/s10209-010-0189-5

[110] Amit Zoran and Joseph A. Paradiso. 2013. FreeD: a freehand digital sculpting tool. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '13)*. Association for Computing Machinery, New York, NY, USA, 2613–2616. https://doi.org/10.1145/2470654.2481361