## Research Statement | Thomas Langerak

thomas.langerak@inf.ethz.ch | thomaslangerak.nl

# Research Vision

The goal of my research is to improve the **interaction between intelligent systems (e.g., robots, semi-autonomous vehicles, and digital assistants) and users through human-in-the-loop optimal control**. I believe that Contextual and Embodied Artificial Intelligence (AI) causes the boundaries between the physical and digital worlds to merge, necessitating technological advances that **prioritize human-centric control of intelligent systems**. My research aims to enable a future where humans can interact with intelligent systems instinctively, unobtrusively, and with minimal effort required to master.

My work focuses on developing computational approaches to advance this human-machine collaboration. **Vital for this collaboration is the control strategy of the intelligent system**: given a context and user behavior how should the machine act for collaboration towards a, sometimes unknown, goal? Crucially, the payoff might happen in the future (e.g., breaking in a vehicle now; might prefer an accident in the future), hence the system should take the **optimal action given the expected reward** over an (in)finite horizon given a goal and constraints. Prediction based optimal control is not trivial, due to uncertainty of the world. In collaborative tasks, this becomes even more challenging; as **the system will need to predict the human future states**. I believe that the integration of **cognitively-plausible computational models of human behavior** in optimal control strategies is the way forward.

# Research Contributions

Traditionally humans used interfaces to directly manipulate variables, like using a hammer to drive a nail. More recently, this approach has evolved: we interact with interfaces that communicate with intelligent agents to manipulate variables for us, such as controlling a Nest thermostat to adjust the temperature. **With the rise of Contextual and Embodied AI, the user interacts through the variable with the agent**. For instance, steering a semi-autonomous vehicle. In this case, the vehicle is both the interface (the user uses it to communicate with the autonomous agent) and the variable (both the user and agent update its acceleration). Thereby, making the variable both the target of manipulation and the interface for interaction. This convergence of variables and interfaces introduces **challenges in rethinking interface design and balancing user autonomy with system automation.**

My research follows these two connected threads of work: 1) the design, creation and evaluation of **joint interfaces**, interfaces on which both a human and artificial agent jointly interact; and 2) the exploration of **cooperative control strategies** that combine **user models with optimal control strategies.**

## Interface + Variable = Joint Interface

Traditionally, **an interface is defined as a shared boundary for information exchange between separate computer system components**. A joint interface retains these properties but also **serves as a variable that the user seeks to control (right),** such as a semi-autonomous vehicle. I am not arguing that all variables are interfaces or vice versa. I am saying that these historically separate elements have moved closer over time, due to contextual AI, and that there is now a partial overlap.

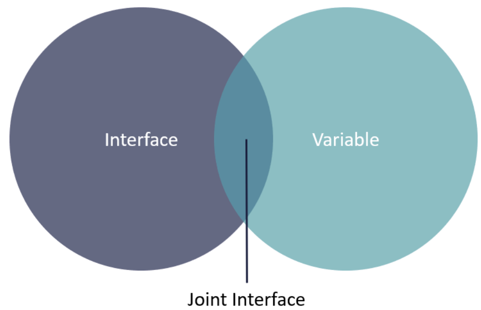


I have explored this overlap in three projects; in which I specifically focused on physical input and output haptic interfaces for Virtual and Augmented Reality. In all my projects an intelligent agent and a user act (using force) on the same interface (a pen). Tool-based haptic in- and output interfaces form an interesting scenario of a joint interface. On one hand, the user and agent both act on the tool/interface (through force) where the variable is the tool's spatial position. On the other hand, the tool position is the proxy for a second task-specific variable (e.g., controlling a car in a game). Using the haptic interface we can explore the nuances of the joint interface.

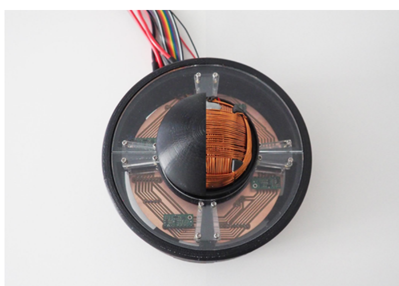

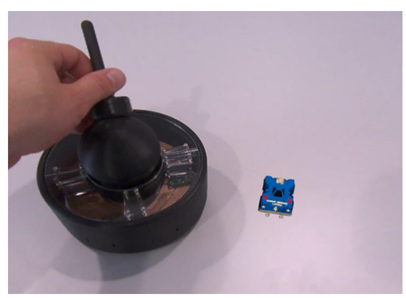

Omni [Langerak2020a] (above) is a haptic device that is used as a controller for Virtual Reality applications, such as gaming, design, and haptic rendering of deformable materials. A symmetric electromagnet is used in combination with a dipole magnet model and a control law to deliver dynamically adjustable forces onto a hand-held tool. Here the haptic feedback is elicited on the tool, while the user interacts with the tool. Omni showed that **physical joint interfaces provide more intuitive and accurate user inputs.**

[Langerak2020b] extends Omni to include sensing of the tool position via the magnetic field. Where Omni required external tracking devices; the spatial haptic capabilities of Omni 2.0 are enabled by a novel gradient-based method to reconstruct the 3D position of the permanent magnet in midair using the measurements from eight off-the-shelf hall sensors that are integrated into the base. In [Langerak2022] we extend this to a deep learning method that improves tracking latency, frequency, and accuracy. Omni 2.0 teaches us that by **incorporating advanced sensing technologies directly within our tools, we**

**can significantly reduce reliance on external devices, paving the way for more seamless and integrated human-computer interactions.**

## Cooperative Control Strategies

A major implication of the contextualization and embodiment of AI for interfaces and variables: **rather than users having full ownership over the interface, ownership is now shared**. For instance, until recently, users would write code on their own. However, nowadays, both the user and CoPilot can write code together, simultaneously. This raises questions about who can edit and manipulate the code and to what extent, at any given time. **We need to balance user autonomy with system automation.**

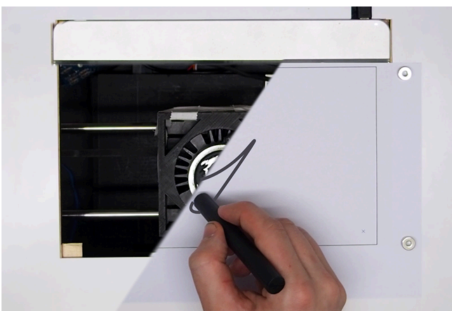

One of the core limitations of omni is that the magnetic field rapidly declines over distance, thereby limiting the working area. In [Langerak2020c], we overcame this by mounting a cylindrical electromagnet on a biaxial linear stage (left). However, this approach complicates the control dynamics. If the magnetic actuator is too close and active, the feedback perceived by the user is too strong, causing the user to lose any sense of autonomy. However, if the magnet is too far away, no haptic feedback is perceived, rendering the system useless. **In our work, we demonstrate that by incorporating user and predictive models of world state transitions into a control strategy, we can guide the user while still preserving their autonomy.**

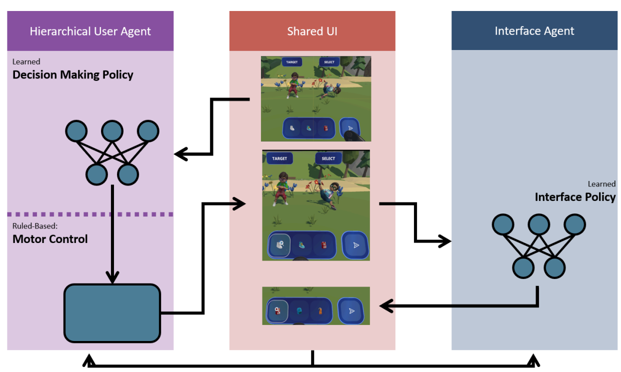


In [Langerak2023] (right) we formulate the **interaction with the joint interface as a multi-agent reinforcement learning problem**. A user interacts with a graphic interface, while an intelligent agent adapts the interface based on observations of the user's actions. In our formulation, a user agent mimics a real user and learns to interact with an interface via point-and-click actions. Simultaneously, an interface agent learns interface adaptations, by observing the user agent's behavior, to maximize the user agent's efficiency. **Our work demonstrates how we can learn intelligent cooperative control strategies, by treating interaction as a multi-agent game.**

# Future Research Agenda

In my future research, I plan to focus on three distinct aspects and opportunities to realize the vision of human-centric control over intelligent systems.

## Human Sensing & Inference

I believe that the impending transformation of the interface will **integrate the world at large—including the users themselves—into the interface itself**. For instance, consider a scenario in which a robot and a human collaborate to clean an apartment. The user should **not need to explicitly instruct the robot** on what to clean; instead, the robot should be capable of inferring this from the state of the world and the user's actions.

This requires two main components. Firstly, **accurate human state estimation** is essential. This includes **not only the physical** state in the world, which can be captured through computer vision and other modalities, but **also latent states** such as expertise, fatigue, and intent. Secondly, **accurate sensing of the world** state and reasoning is required. The system should be aware of its environment and capable of predicting future states from the current context, such as understanding that an apple will fall if released.

## Human Behavioral Models

However, sensing the human state is insufficient. To apply current state estimations to intelligent control, we need to **predict future human—both physical and latent—states.** This requires behavior models that capture the complex dynamics of human behavior.

To solve this problem, we need **advances in data-driven Reinforcement Learning, such as imitation learning (IL) and inverse reinforcement learning (IRL), combined with existing cognitive models.** Reinforcement Learning is uniquely suited, as it can result in approximate optimal behavior given bounds (cognitive models). If we assume computational rationality, that is humans act optimal given their constraint resources, and we have accurate cognitive models then reinforcement learning should convert to human-like policies . However, RL suffers from an exploration problem. By incorporating data-driven prior, we can model more complex behaviors. By integrating these three approaches, I aim to develop methods that approximate human strategies, are generalizable, and computationally feasible. These models will not only enhance our understanding of human behavior but also improve how we design interfaces and interaction paradigms between humans and machines.

## Cooperative Control

Finally, we need to **integrate the human sensing, inference, and behavioral models into intelligent control strategies**, paving the way for more intuitive and seamless interactions between humans and machines that respect and enhance human decision-making processes.

I foresee a future in **hierarchical and multi-agent reinforcement learning**, where predictive human models can be efficiently integrated into control strategies. Exactly how this integration will occur remains an open question. The development of such systems promises to significantly advance our ability to design intelligent systems that truly augment human abilities, marking a significant milestone in human-computer interaction.

# References

**[Langerak2023]** MARLUI: Multi-Agent Reinforcement Learning for Goal-Agnostic Adaptive UIs. Thomas Langerak, Sammy Christen, Mert Albaba, and Christoph Gebhardt, Chirstian Holz, and Otmar Hilliges. *Under Review*

**[Langerak2022]** Robust Real-Time Tracking of Axis-Symmetric Magnets via Neural Networks.
Mengfan Wu*, Thomas Langerak*, Otmar Hilliges, and Juan Zarate. *Under Review*

**[Langerak2020a]** Contact-free nonplanar haptics with a spherical electromagnet. Juan Jose Zarate*, Thomas Langerak*, Bernhard Thomaszewski, and Otmar Hilliges. *In 2020 IEEE Haptics Symposium (HAPTICS), 2020*

**[Langerak2020b]** Omni: Volumetric sensing and actuation of passive magnetic tools for dynamic haptic feedback. Thomas Langerak*, Juan José Zárate*, David Lindlbauer, Christian Holz, and Otmar Hilliges. *In Proceedings of the 33rd Annual ACM Symposium on User Interface Software and Technology, 2020*

**[Langerak2020c]** Optimal control for electromagnetic haptic guidance systems.
Thomas Langerak, Juan José Zárate, Velko Vechev, David Lindlbauer, Daniele Panozzo, and Otmar Hilliges. *In Proceedings of the 33rd Annual ACM Symposium on User Interface Software and Technology, 2020*