# Stable Shared Control based on Models learned from Physical Demonstration


**PI (lead): Todd Murphey, *Northwestern University***
**Co-PI: Brenna Argall, *Northwestern University***
**PI: Magnus Egerstedt, *Georgia Institute of Technology***


## *How much should a person be allowed to interact with a controlled machine?*

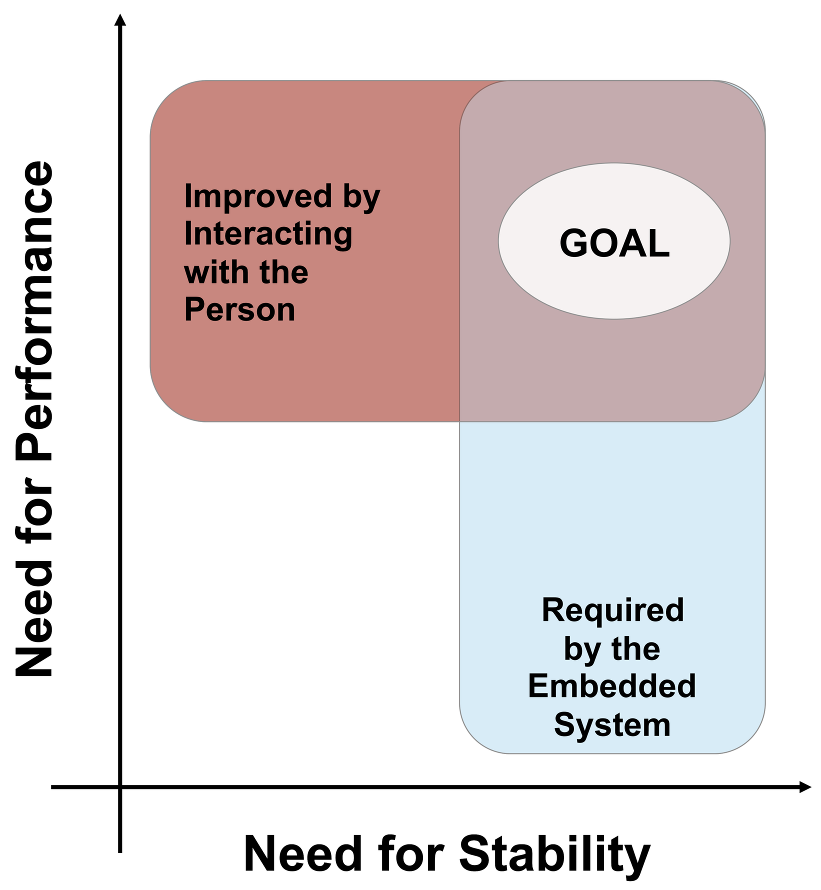


**Figure 1:** Our goal is to algorithmically resolve the tension between the need for stability and the need for performance.

**Aim:** Balancing the ability of a person to direct a cyber-physical system, against the system's representation of its own capabilities and limitations.

**Why?** Physical interactions with cyber-physical systems need to be understood in terms of ***shared autonomy***, where the embedded software and the human together have to interface directly with the system dynamics.

**How?** Develop a model of the ***joint human-robot system*** that describes the dynamics of the physical system and the interaction with a human operator.

For cyber-physical systems, an understanding of *the system and control dynamics* is of crucial importance. The *automated system* must be capable of controlling the physical device when necessary, but also permissive of controls provided by the *human operator* when safety and stability are not an issue.

## Mutually Controlled Motion

**Idea:** Derive control behaviors via optimal control, while...

Engage the **human operator** for low-level interactions which

can be evaluated and constrained based on optimal policies

**Challenge:** The operator may destabilize the system.

This risk changes from operator to operator.

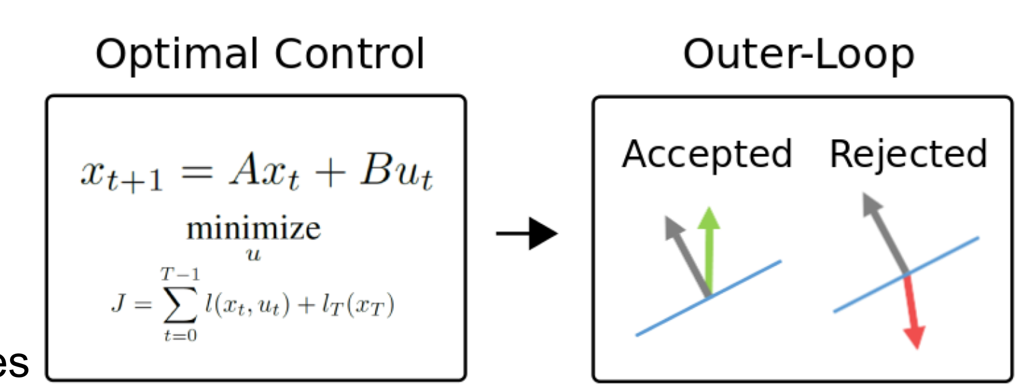


**Figure 2:** Our goal is to use optimal control techniques to provide outer-loop stabilization to the low-level input provided by human operators.

## Individual and Generalizable Shared Control Models

**Solution:** Learn models of **joint human-robot system** directly from demonstration data, in order to...

derive **individual supervisory** behaviors via **optimal control** which...

decide on-line **how much control to cede** to the operator during physical interaction

**Result:** A computable notion of optimality and stability → The system assesses the safety of the instructi

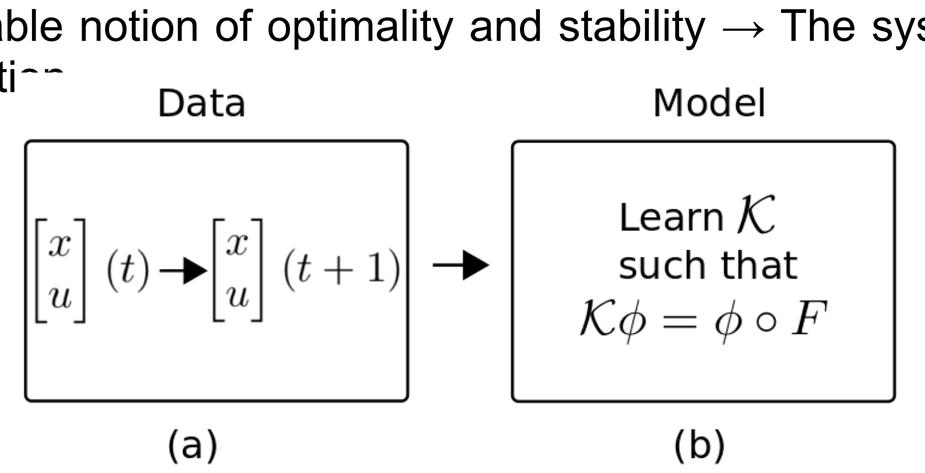


**Figure 3:** Our goal is to use techniques from machine learning to learn actionable models of the joint-human robot system that can be used to evaluate the joint system and produce control policies.

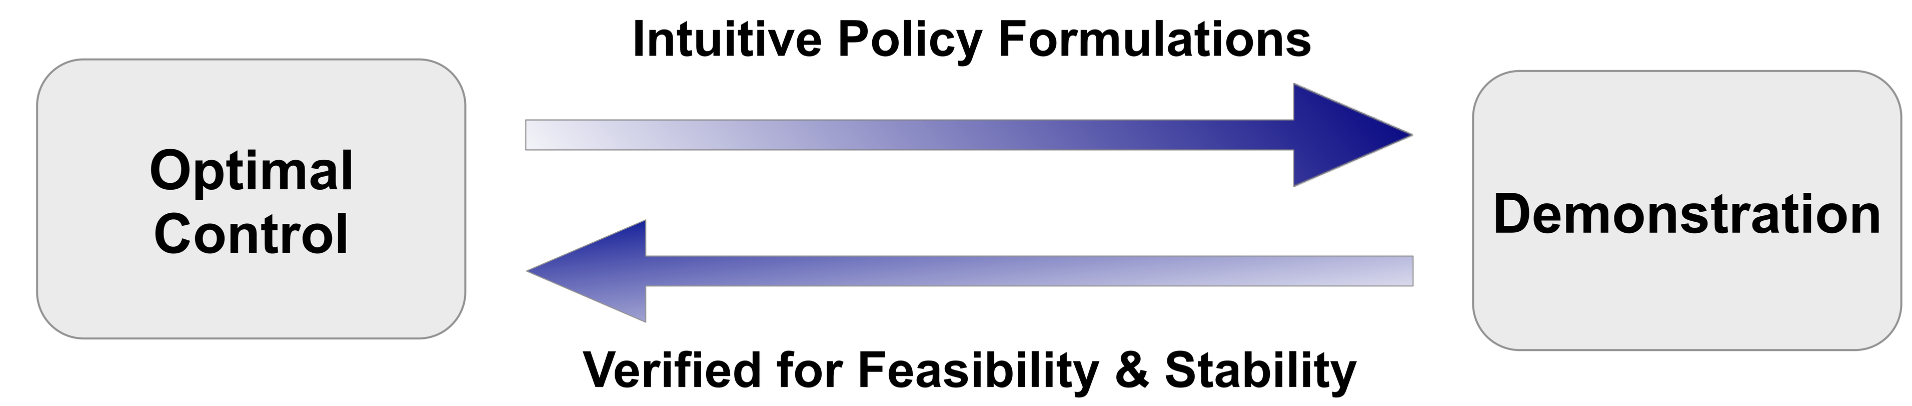


**Figure 4:** Respective benefits of two procedures for deriving controllers: Optimal Control and Learning from Demonstration.

## Assessments

**Platforms:** Lunar lander with non-linear dynamics → easily destabilized.

**Goal:** By the end of this project, we will be able to learn a joint model of the human and robot interaction. This model can be used to compute optimal control policies which will in turn be used as outer-loop control and stabilization. Our shared control paradigm allows the human to remain in full control of the dynamic system except in extreme cases when the autonomous agent detects unsafe control input. We also begin to evaluate the individuality of the learned models.

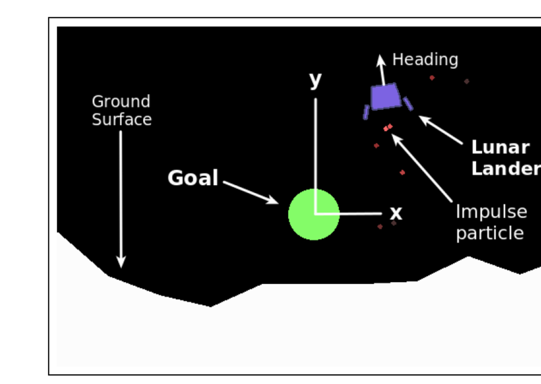


Fig. 2: Simulated lunar lander system. The green circle is the visualization of the desired goal location. The red dots represent an engine firing.

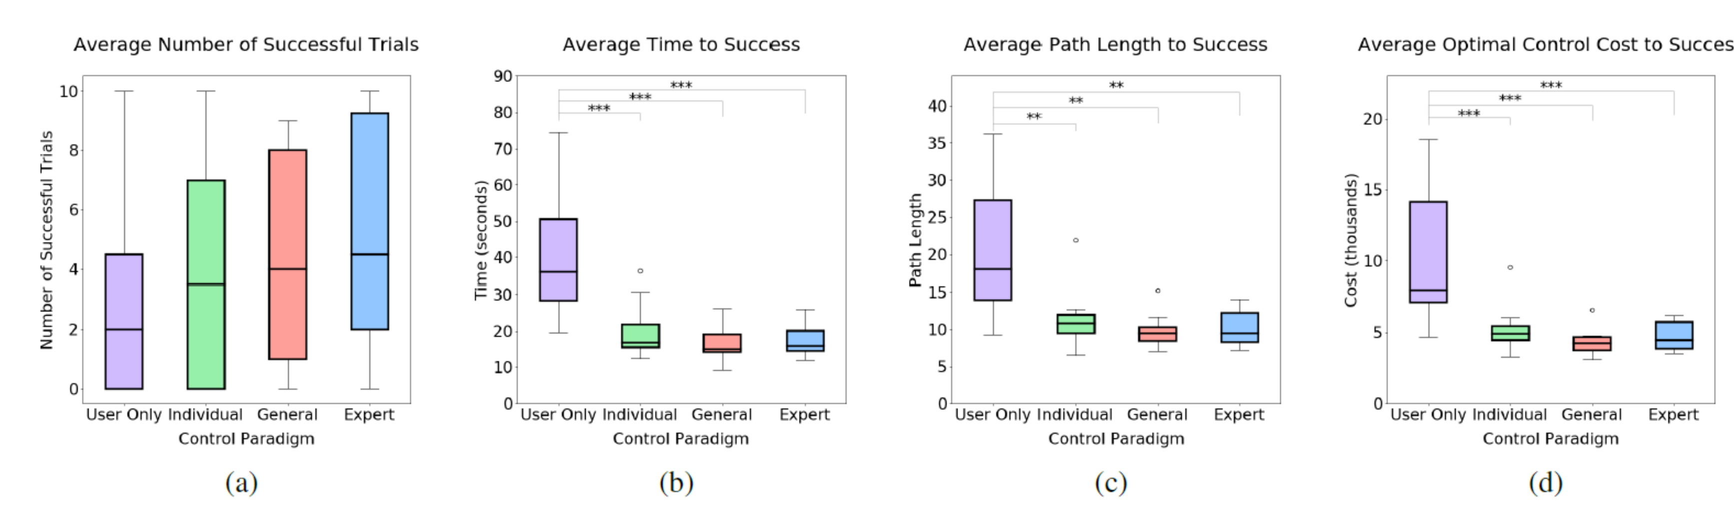


Fig. 3: (a) Number of successful trials under each control paradigm. Despite a visual trend, we find no statistically significant difference between the user-only control paradigm and the shared control paradigms. (b) Average time to successfully complete a trial under each control paradigm. (c) Average path length taken to successfully complete a trial under each control paradigm. (d) Average total optimal control cost computed during a successful trial under each paradigm. For all three metrics (b-d), post-hoc pair-wise t-tests using Holm-Bonferonni corrected alpha values find statistically significant differences between the user-only paradigm and each shared control paradigm (time: $p < 0.005$, path length: $p < 0.01$, cost: $p < 0.005$). Additionally, our post-hoc t-tests find no statistically significant difference between any of the shared control paradigms.

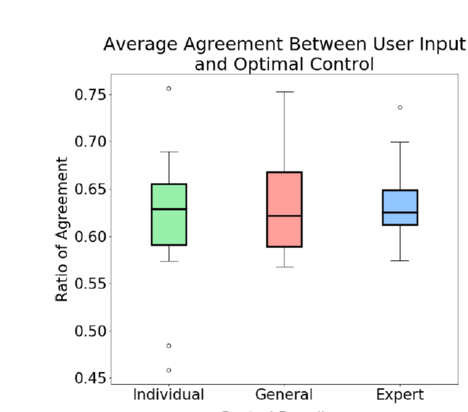


Fig. 4: Average agreement.

## Relevance to Cyber-Physical Systems

**Impact:** Cyber-physical systems for which (i) control authority is shared between the human and machine, (ii) the machine automation is adaptable by and able to receive instruction from a human who is not an automation expert, (iii) there are physical, possibly destabilizing, interactions between the human and machine.

**Domains:** Immediately impacted: Rehabilitation, assistive devices, and human augmentation. Near-term impact: Manufacturing, which will soon involve skilled workers working side-by-side with robots and teaching robots tasks. More broadly: Non-mechanical but highly interconnected systems, such as air traffic control and power grid management. Such systems are often too complex to understand completely, yet the operator still must provide instruction that is feasible for the system to reliably execute.

**Relationship to CPS Needs:** *Interaction and potential interference among CPS and humans*, by explicitly reasoning about when to cede control authority to a human operator, and when to request instruction for stability assistance. *Cross-disciplinary collaborative research*, by building a synergy between the areas of data-driven machine learning and formal control theory. *Jointly modeling the interaction of both cyber and physical components*, by taking steps to quantify the level of understanding needed by the human to provide effective corrections, and by explicitly computing the system's understanding of the consequences of physical or interaction during instruction. *Incorporating CPS science into education*, by incorporating CPS-centric coverage in the Control of Mobile Robotics MOOC taught by co-PI Egerstedt.