ITU WORKSHOP ON EMBODIED AI FOR MULTIMEDIA TECHNOLOGIES
SESSION 2

Safe Embodied AI in the Real World

Closing the gap between learning and deployment

Prabhat Kumar Mishra
Assistant Professor, Department of Artificial Intelligence, IIT Kharagpur
Hangzhou, China · 16 October 2026



WHERE WE DEPLOY

Three robots, one question



Quadrupeds

Rough terrain, crowds and a chatbot interface. Mistakes land on people.



Under-canopy field robots

Weak GPS under the canopy. Plants must stay undamaged, with little human help.



Aerial drones

Wind, ground effect and lossy radio links, so the vehicle must stay safe on its own.

Can we certify what a learned policy will do when the model, the sensors and the network are all imperfect?

THE DEPLOYMENT GAP

What changes between the lab and the field

Gap	In the lab	In the field	What deployment needs
Dynamics	A simulated or well-identified model	Unmodeled terrain, rolling resistance, wind	Learn online while keeping constraints satisfied
Sensing	Clean state, well-behaved noise	GPS dropouts, estimates that violate known constraints	Estimates should be according to their physical significance
Communication	Perfect links	Packet loss on sensor and command channels due to interference and fading	A practical notion of stability guarantee

Learning has made robots far more capable. These three gaps still need **guarantees.**

OUR APPROACH

Learning proposes. A certified optimizer decides.

Learned module

Deep network, VLA policy or LLM planner

→

Predictive safety layer

Model, physical constraints and stability conditions, solved at every step

→

Robot

A certified action, with its predicted trajectory

Safety certificate

Constraints are enforced at the synthesis stage

Stability

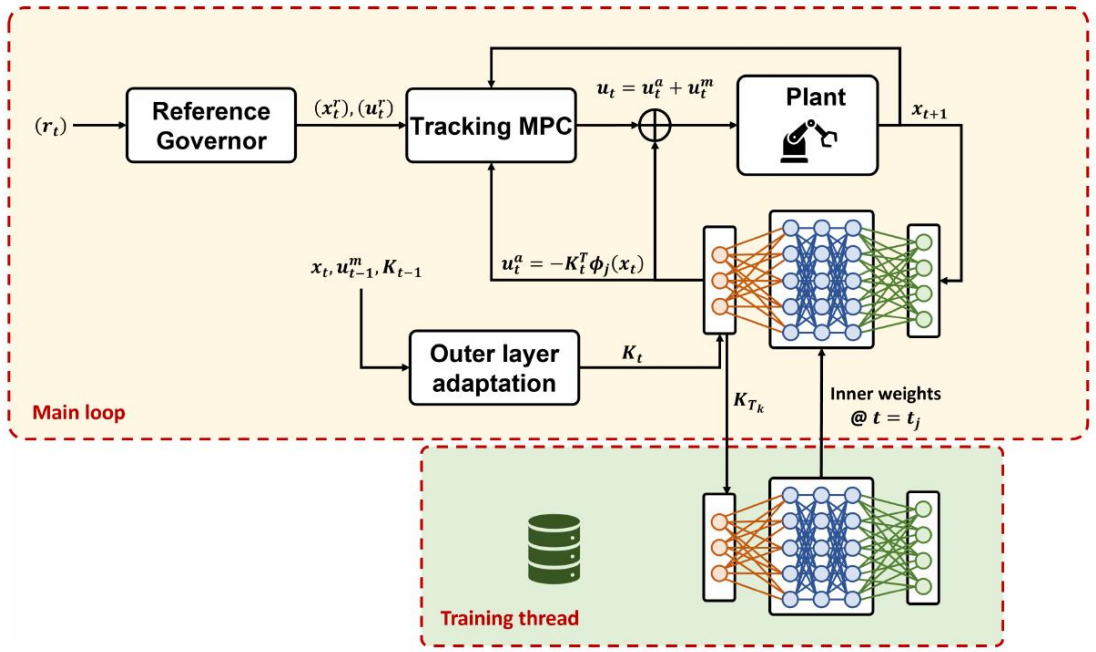
Stability is proven under stated assumptions, not only observed in tests.

Real time

Formulated to run on board, as convex programs wherever possible.

GAP 1 · UNCERTAIN DYNAMICS

Deep MPC: learn online, stay inside the tube



Fast loop

The output layer adapts at every time step.

Slow loop

Hidden layers retrain in a background thread on data collected online.

Tube MPC

Keeps states and inputs within limits while the network is still learning.

CERTIFICATE

Input-to-state stability, and convergence to the origin under stated conditions.

Mishra, Gasparino and Chowdhary, IEEE TAC 2025; implementation notes: Indian Control Conference 2025

GAP 2 · NOISY SENSING

Estimates that respect constraints

Problem

A Kalman filter can return estimates that break known constraints, and the controller then plans from an impossible state.

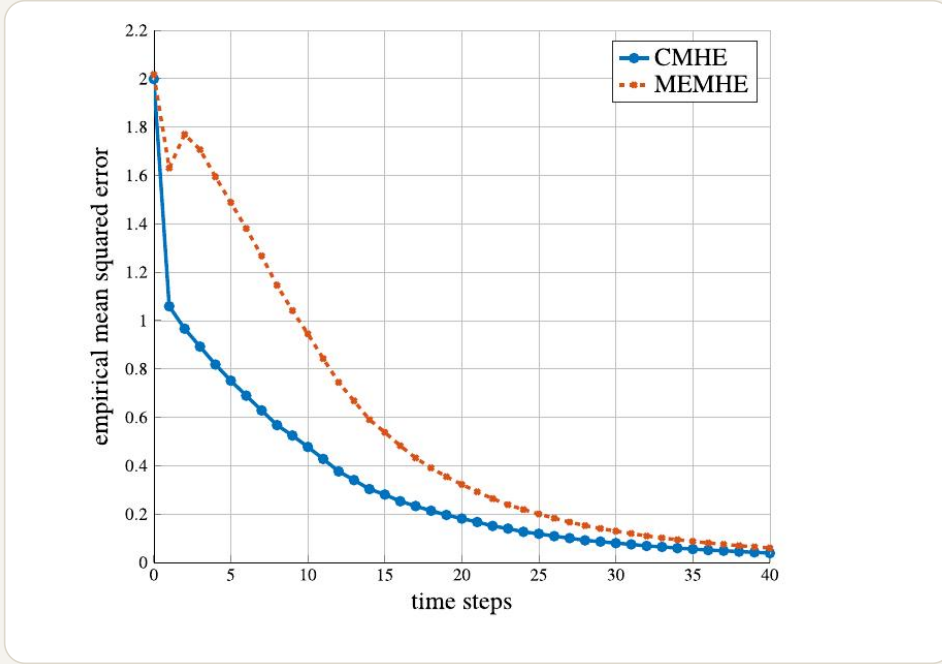
Method

Minimum-variance duality turns constrained estimation into an optimal control problem: a convex QP.

CERTIFICATE

Constrained full-information and moving-horizon estimators are provably stable observers under stated conditions.

Under review: covariance initialization for safe inference, and LQG control from noisy outputs without state estimation.

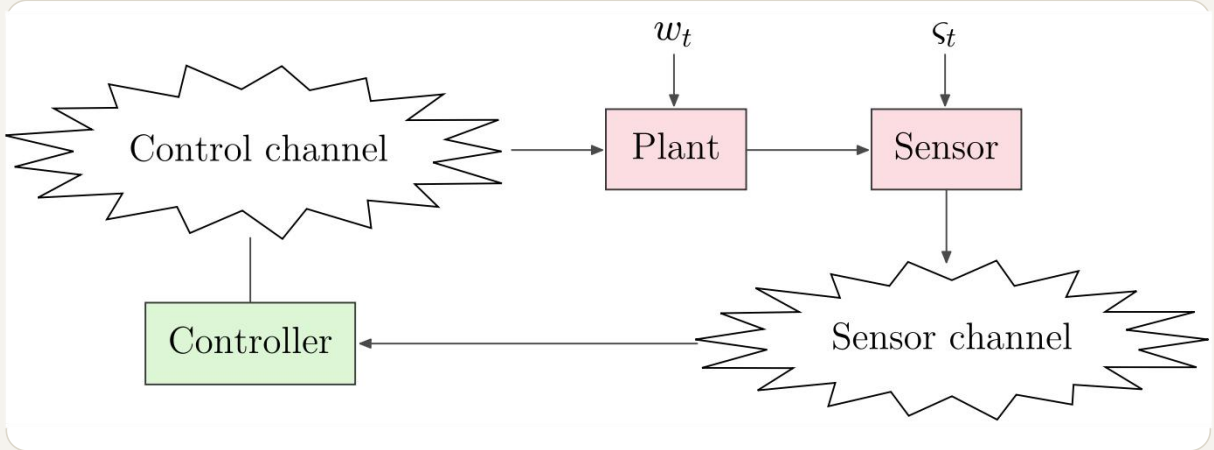


Benchmark batch reactor, 1,000 runs, Gaussian prior: lower error than standard constrained MHE.

Mishra, Chowdhary and Mehta, Minimum Variance Constrained Estimator, Automatica, 2022

GAP 3 · IMPERFECT COMMUNICATION

Control that stays stable over lossy links

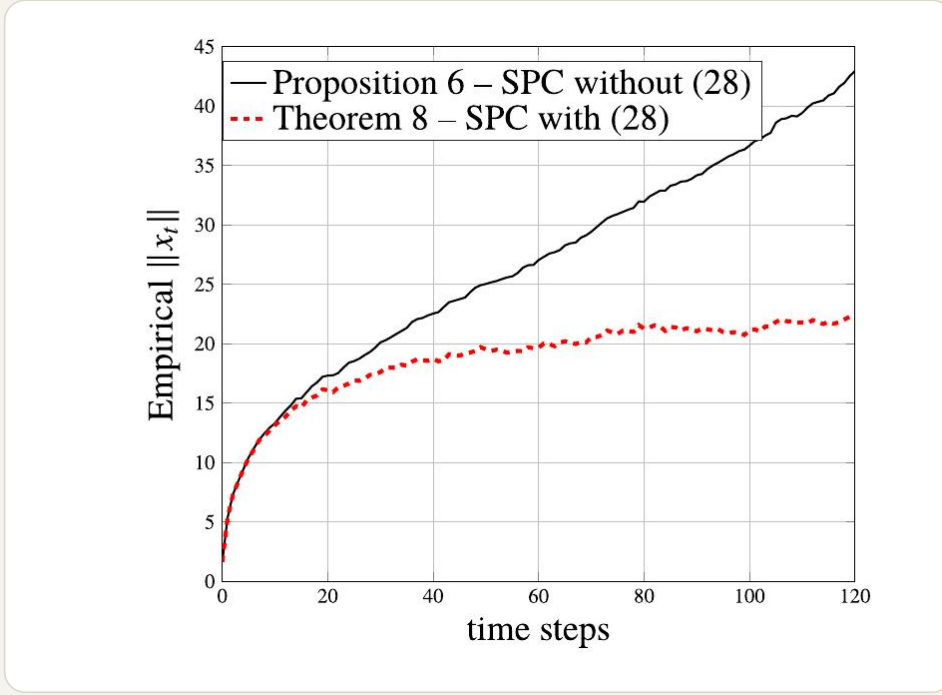


A smart sensor filters on board, a remote estimator fills in lost packets, and the actuator buffers planned future inputs.

CERTIFICATE

A recursively feasible convex QP. States stay mean-square bounded for any input bound and any non-zero delivery probability, for Lyapunov-stable linear plants.

Mishra, Chatterjee and Quevedo, Automatica 2020; IEEE TAC 2018



One example, 500 runs: without the stability constraints the state grows; with them it stays bounded.

LESSONS LEARNT

Four design rules for unreliable networks

If the command channel drops packets

Smart actuator

Learn a feedback policy, not actions, and buffer planned inputs so a lost packet never leaves the robot without a command.

If sensor packets are lost

Smart sensor

Filter on board and send estimates, so the controller receives usable state information.

If on-board compute is limited

Backup controller

Plan remotely, but keep a simple fallback controller running on the robot.

If the goal keeps changing

Reference governor

Reshape the reference so it stays reachable within the robot's limits.

Thank you

Collaborators

Girish Chowdhary, Prashant G. Mehta, Mateus V. Gasparino, Tixian Wang, Mattia Gazzola (UIUC)
Debasish Chatterjee (IIT Bombay)
Daniel E. Quevedo (QUT)
Ujjawal Prasad, Shivang Tiwari, Ahon Bandyopadhyay, Ambalika Raj, Ranjan Sarkar (RHPI Lab)
Aghna Mukherjee (IIT Kharagpur)

Support

ANRF Prime Minister Early Career Research Grant (PMECRG) and SRIC, IIT Kharagpur
Earlier work by ONR MURI, NSF CPS USDA

pkmishra@ai.iitkgp.ac.in