Eidgenössische Technische Hochschule Zürich
Swiss Federal Institute of Technology Zurich

Institute for Dynamic Systems and Control
Institut für Dynamische Systeme und Regelungstechnik

Improving tracking performance by learning from past data

Angela P. Schoellig

Doctoral examination – July 30, 2012

Advisor: Prof. Raffaello D'Andrea // Co-advisor: Prof. Andrew Alleyne

Eidgenössische Technische Hochschule Zürich
Swiss Federal Institute of Technology Zurich

Institute for Dynamic Systems and Control
IDSC
Institut für Dynamische Systeme und Regelungstechnik

Improving tracking performance by learning from past data = *experience*

Angela P. Schoellig

Doctoral examination – July 30, 2012

Advisor: Prof. Raffaello D'Andrea // Co-advisor: Prof. Andrew Alleyne

MOTIVATION

HUMANS learn from experience.

We learn from mistakes and get better through practice.

We constantly adapt to changing environments.

MOTIVATION

AUTOMATED SYSTEMS typically make the *same mistakes* over and over again when performing a task repeatedly.

Why?

Robots of a car assembly line.

MOTIVATION

AUTOMATED SYSTEMS are typically operated using feedback control:

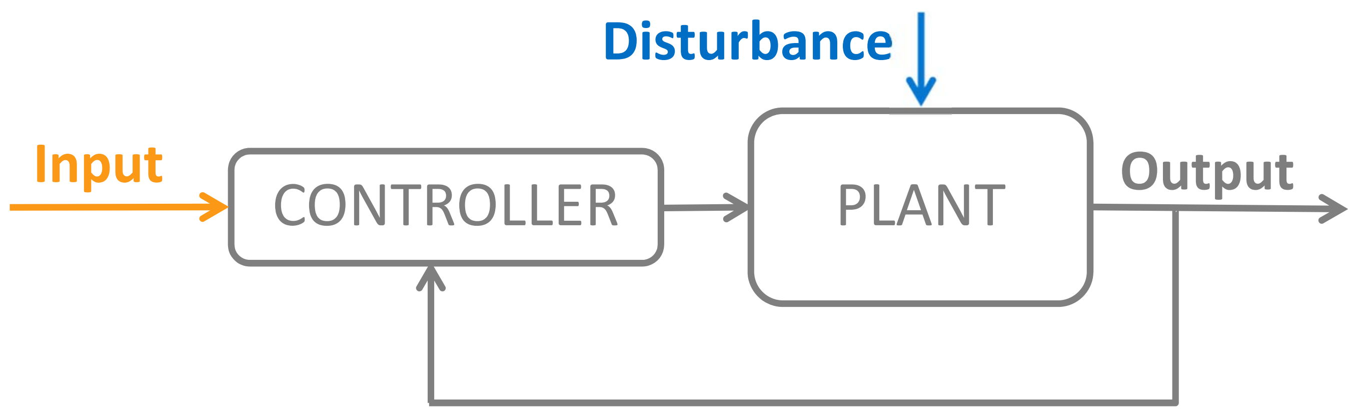

Performance limitations:

- Causality of disturbance correction: "first detect error, then react".
- Model-based controller design; model ≠ real system.

GOAL

Improve the performance over causal, feedback control by learning from previous experiments.

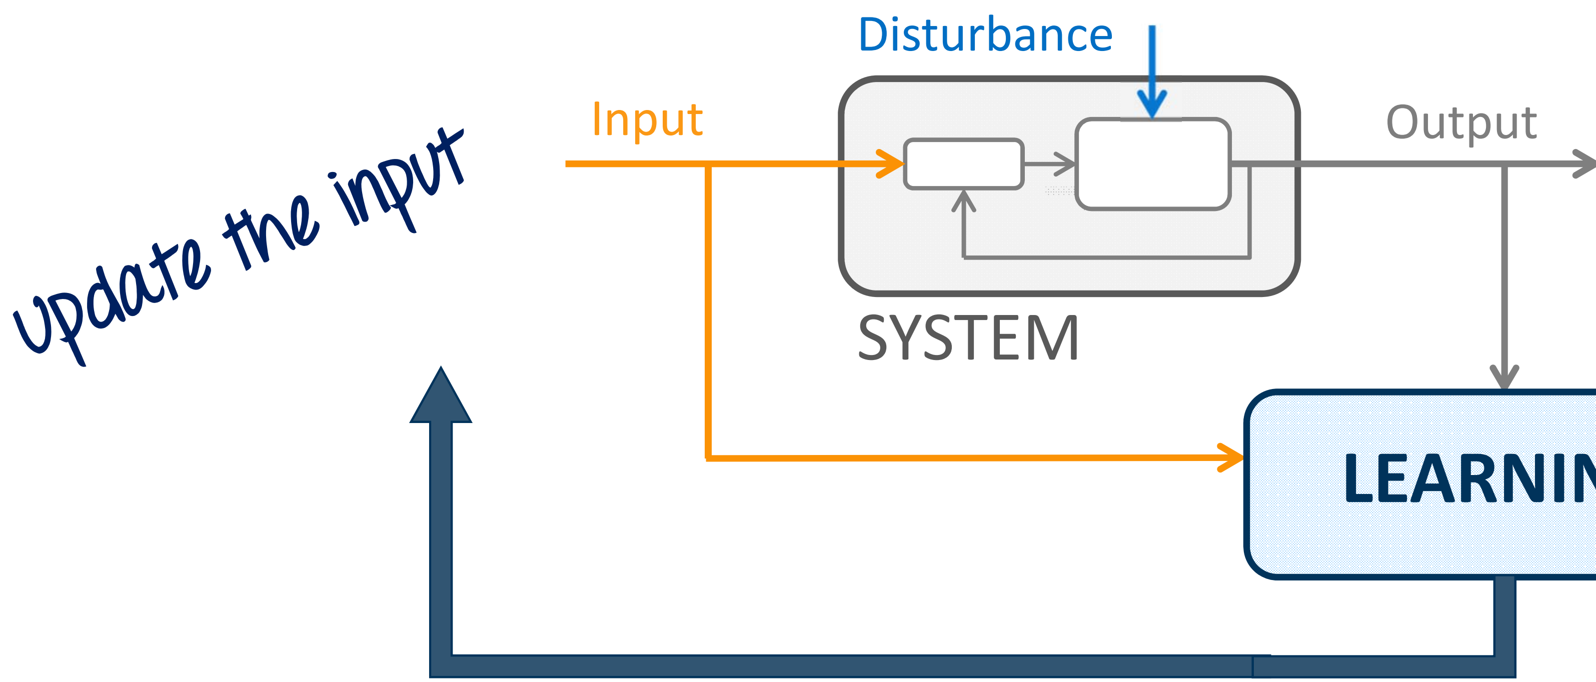

SCOPE OF WORK

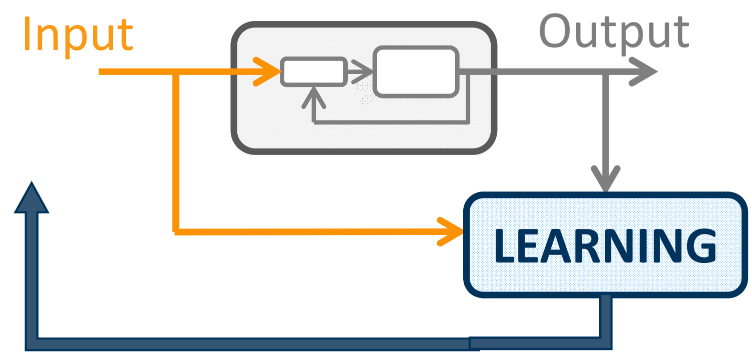

Learning task:

Following a predefined trajectory.

Approach:

- Model-based learning based on *a priori* knowledge of the system dynamics.
- Adaptation of the input.

Potential:

Acausal action, anticipating <u>repetitive</u> disturbances.

OVERVIEW

I. Introduction
   a. Testbed: The Flying Machine Arena
   b. Motivation for learning

II. Project A. Iterative learning for precise trajectory following: single-agent and multi-agent results. Focus of this talk

III. Project B. Learning of feed-forward parameters for rhythmic flight performances

IV. Summary

TESTBED, see www.flyingmachinearena.org

THE TEAM

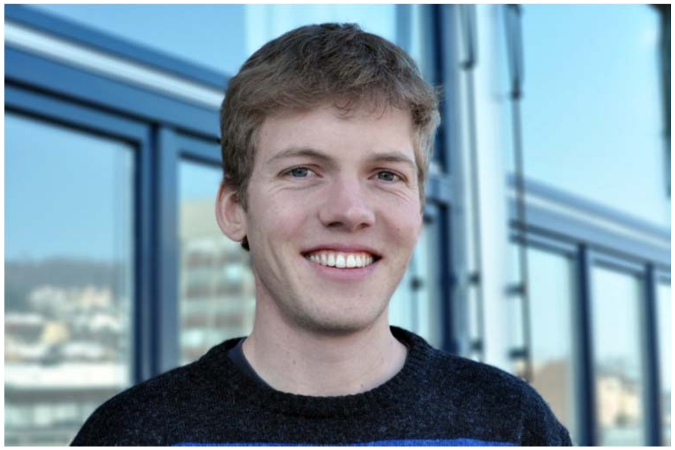
Mark Müller


Markus Hehn


Sergei Lupashin


Federico Augugliaro

THE FLYING MACHINE ARENA

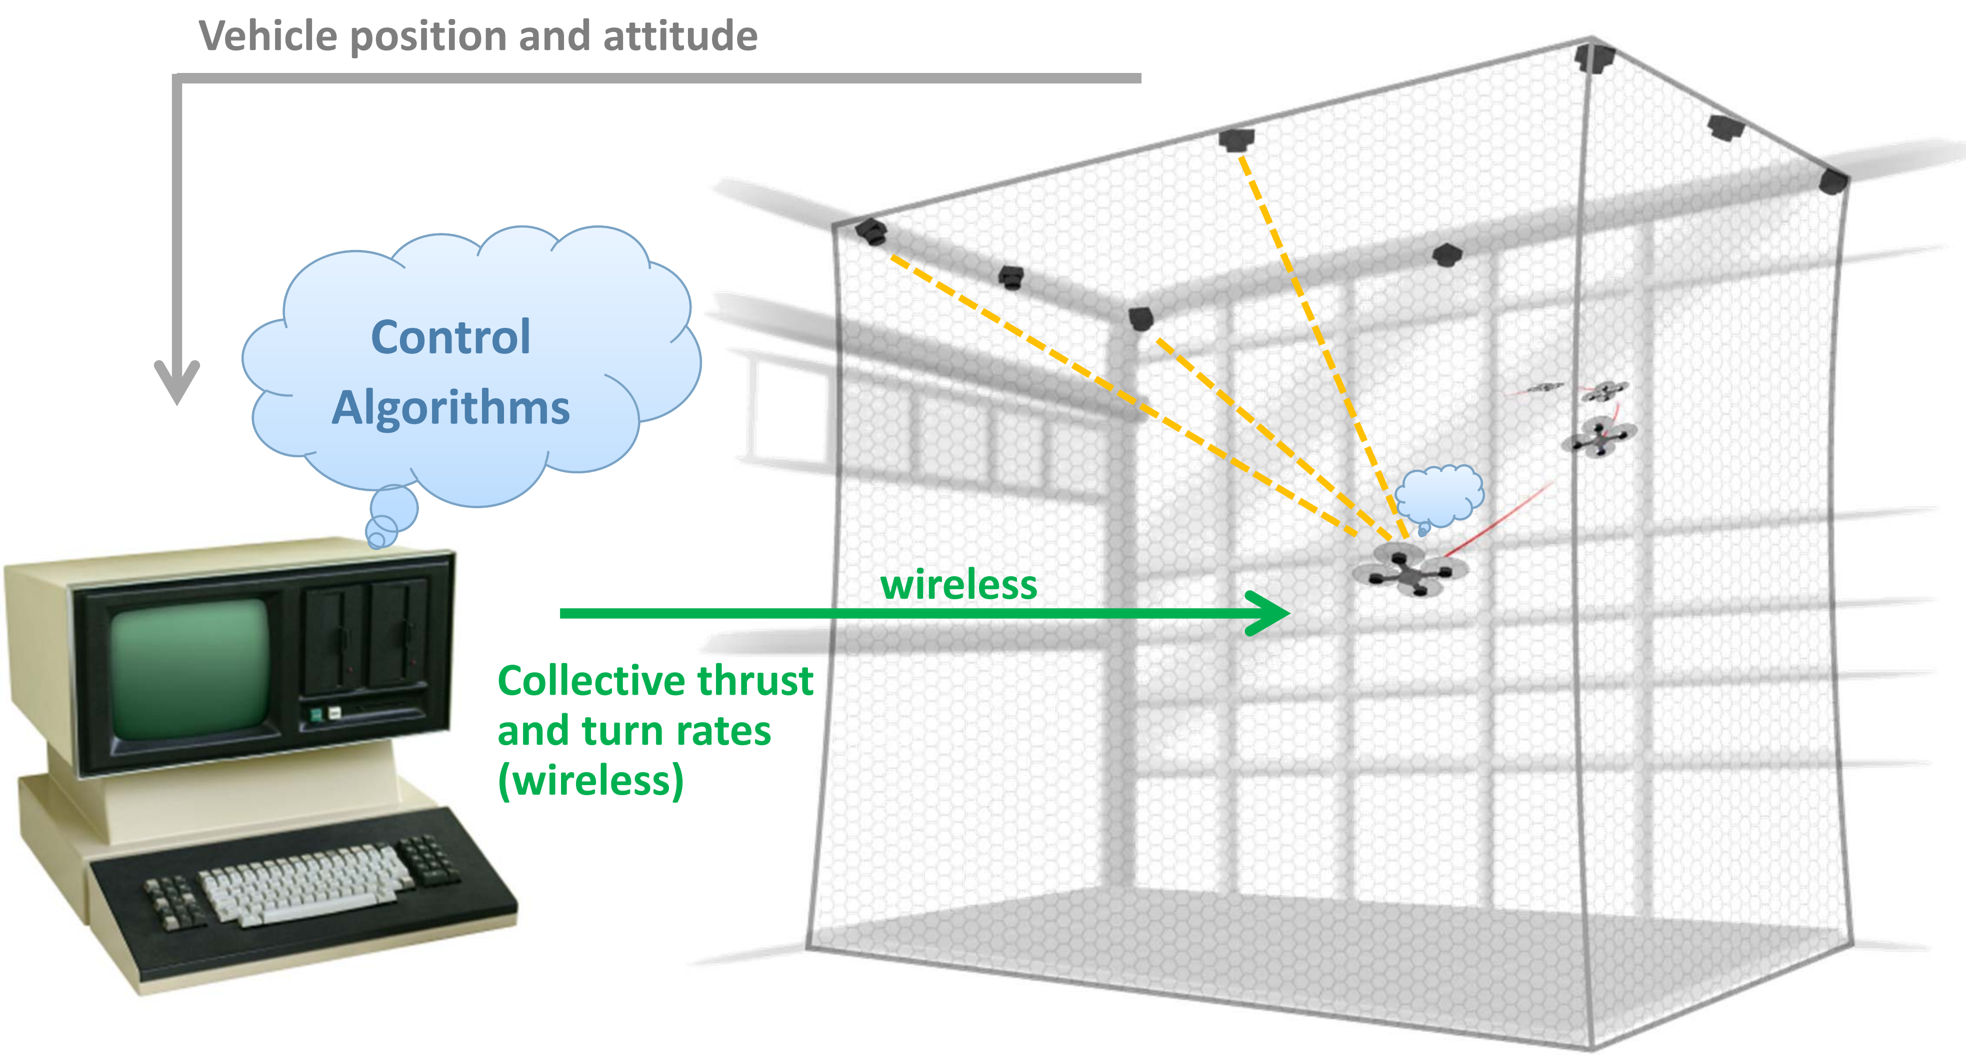

OPERATION

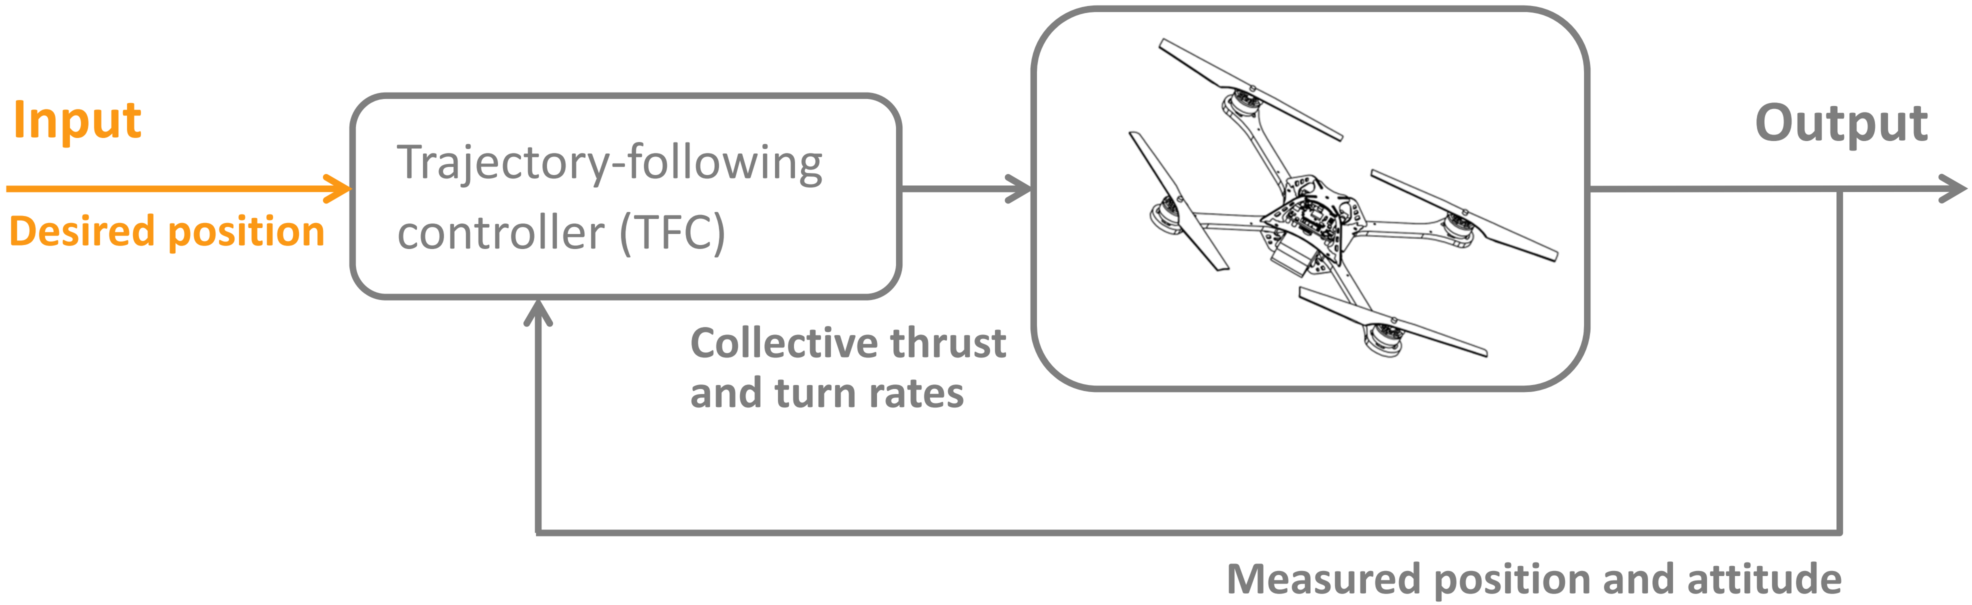
Input
Desired position
Trajectory-following controller (TFC)
Collective thrust and turn rates
Output
Measured position and attitude

MOTIVATION: PROJECT A

Desired motion.

MOTIVATION: PROJECT A

Performance with trajectory-following controller.

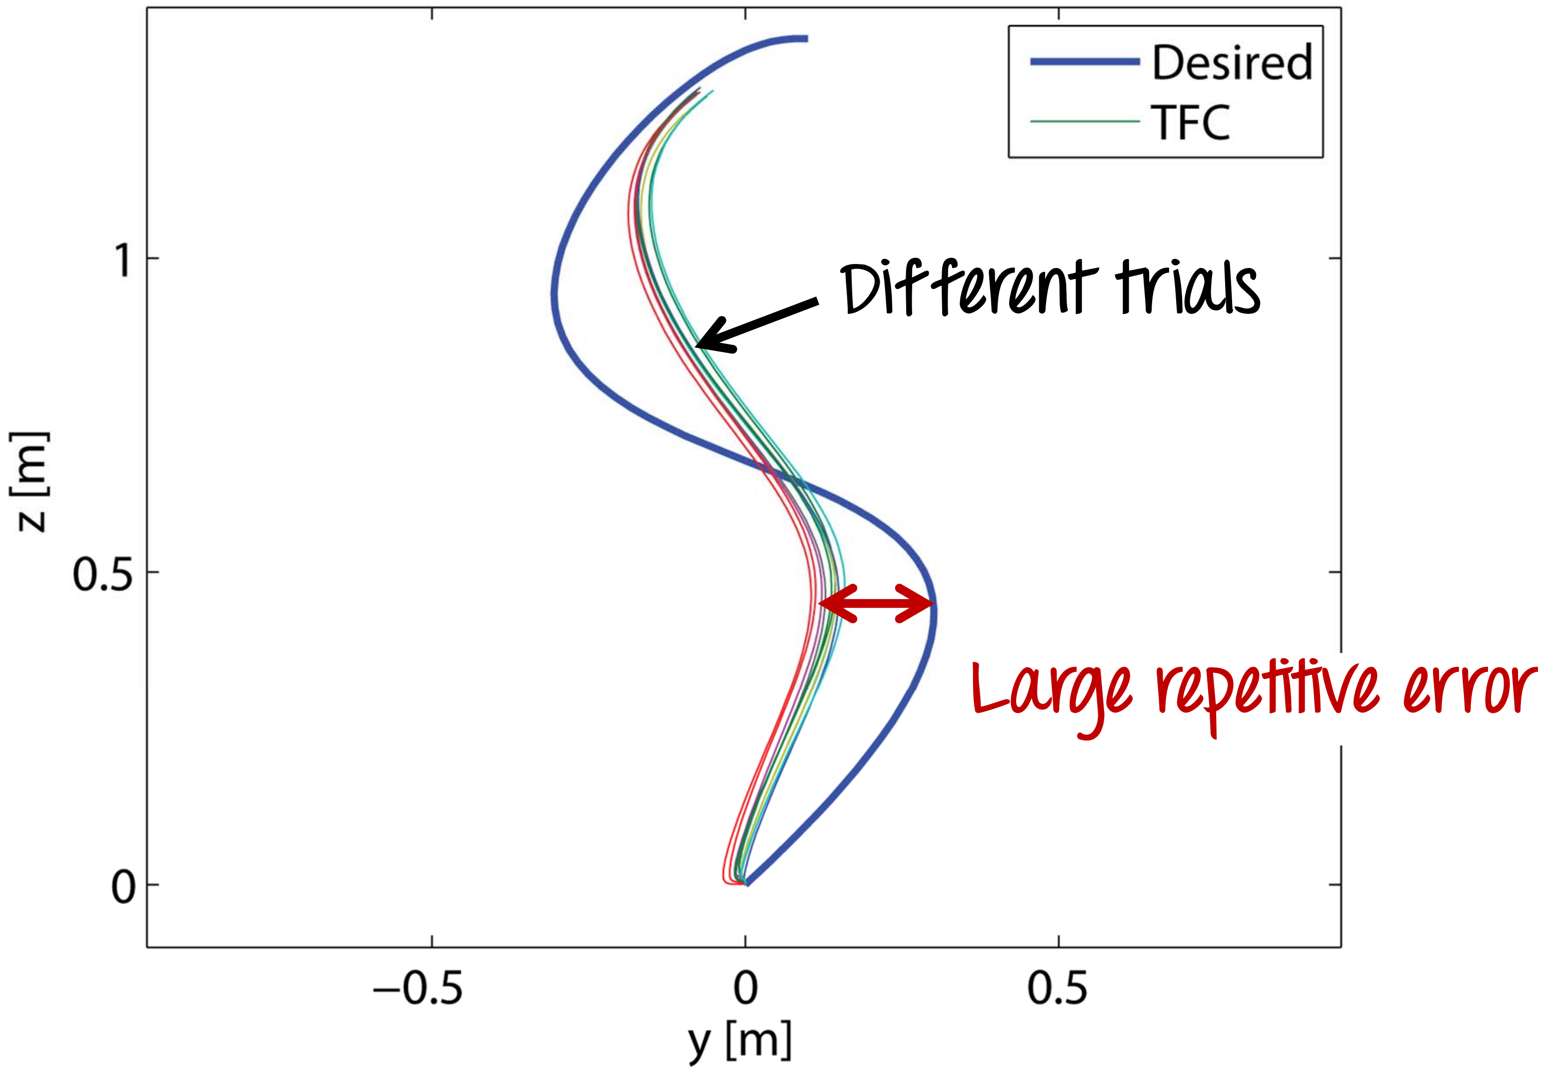

OVERVIEW

I. Introduction

II. Project A. Iterative learning for precise trajectory following
   a. Learning approach
   b. Results

III. Project B. Learning of feed-forward parameters for rhythmic flight performances

IV. Summary

A | PUBLICATIONS

Peer-reviewed publications

Schoellig, A. P. and R. D’Andrea (2009):
“Optimization-based iterative learning control for trajectory tracking.” In *Proceedings of the European Control Conference (ECC)*.

Schoellig, A. P., F. L. Mueller, and R. D’Andrea (2012):
“Optimization-based iterative learning for precise quadrocopter trajectory tracking.” *Autonomous Robots.*

Mueller, F.L., A. P. Schoellig, and R. D’Andrea (2012):
“Iterative learning of feed-forward corrections for high-performance tracking.” To appear in *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*.

Joint work with Fabian L. Mueller (Master student).

A | LEARNING APPROACH

Features: Learning through a <u>repeated operation</u>, updating full input trajectory after each trial.

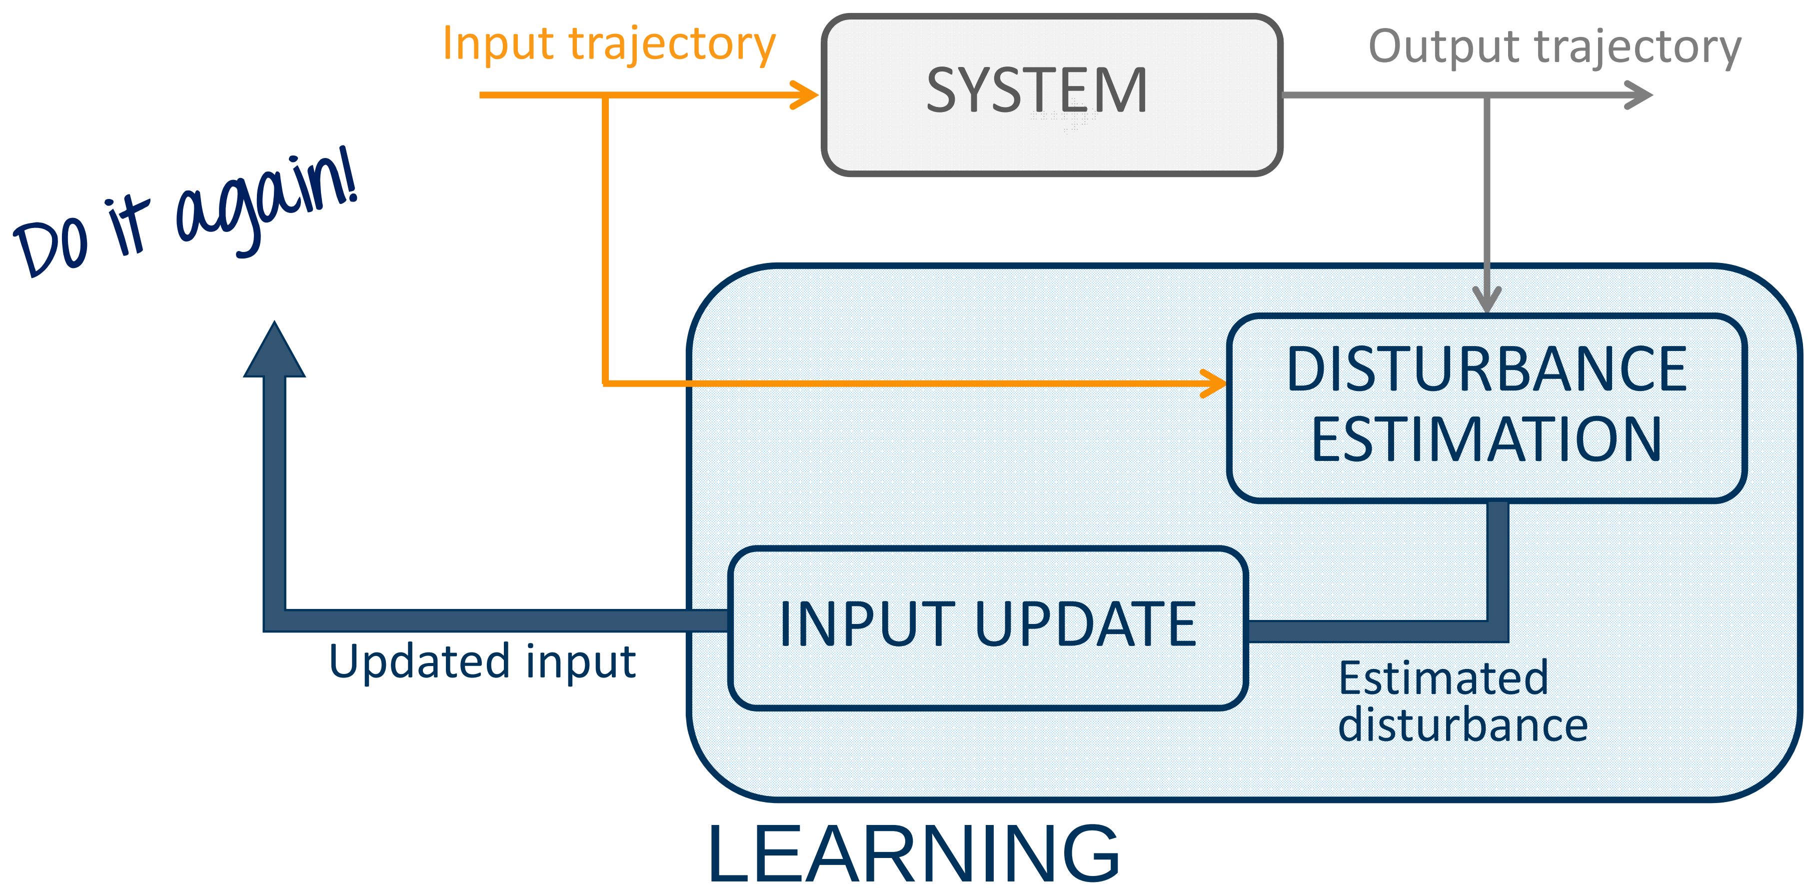

A | LEARNING APPROACH

PREREQUISITES

- **Dynamics model of system**
  (i) in analytical form or
  (ii) in form of a numerical dynamics simulation

- **Desired output trajectory** $y^*(t),\ t \in [0, t_f],$ and corresponding nominal input trajectory $u^*(t)$.
  → $(u^*(t), y^*(t))$ must satisfy the model equations.

RESULT

- Learned input
- Estimated disturbance vector

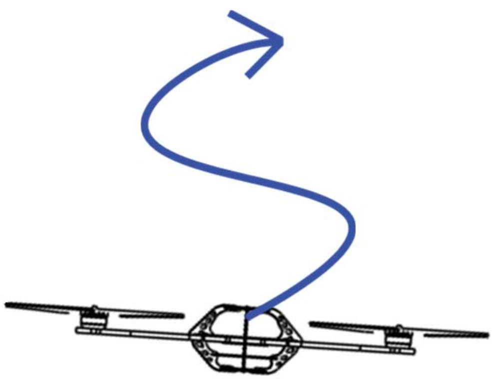

A | LIFTED-DOMAIN REPRESENTATION

Dynamics model of the physical system: $\dot{\check{x}}(t) = f(\check{x}(t), \check{u}(t)), \quad \check{y}(t) = \check{x}(t).$

Consider small **deviations from nominal trajectory.**

$$\tilde{u}(t) = \check{u}(t) - u^*(t), \quad \tilde{x}(t) = \check{x}(t) - x^*(t), \quad \tilde{y}(t) = \check{y}(t) - y^*(t)$$

Linearize and discretize. Linear, time-varying difference equation.

$$\tilde{x}(k+1) = A_D(k)\tilde{x}(k) + B_D(k)\tilde{u}(k), \qquad \tilde{y}(k) = \tilde{x}(k), \qquad k \in \{0, \ldots, N\}.$$

Static mapping. Representing one trial.

$$\underbrace{\begin{bmatrix} \tilde{x}(0) \\ \tilde{x}(1) \\ \tilde{x}(2) \\ \vdots \\ \tilde{x}(N) \end{bmatrix}}_{x} = \underbrace{\begin{bmatrix} 0 & 0 & \cdots & 0 & 0 \\ B_D(0) & 0 & \cdots & 0 & 0 \\ \Phi_{(1,1)}B_D(0) & B_D(1) & \cdots & 0 & 0 \\ \vdots & \vdots & \ddots & \vdots & \vdots \\ \Phi_{(N-1,1)}B_D(0) & \Phi_{(N-1,2)}B_D(1) & \cdots & B_D(N) & 0 \end{bmatrix}}_{F} \underbrace{\begin{bmatrix} \tilde{u}(0) \\ \tilde{u}(1) \\ \tilde{u}(2) \\ \vdots \\ \tilde{u}(N) \end{bmatrix}}_{u}$$

With $\Phi_{(l,m)} = A_D(l)A_D(l+1)\cdots A_D(m), \quad l < m$, and $\tilde{x}(0) = 0$.

A | ITERATION-DOMAIN MODEL

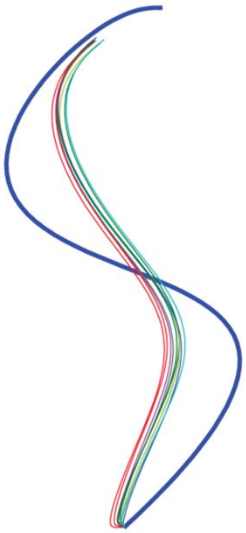

For each trial $j,\ j \in \{1,\, 2,\, \ldots\}$,

$$y_j = F\, u_j + d_j + \mu_j.$$

Recurring disturbance d_j**.**
Unknown. Only small changes between iterations:

$$d_j = d_{j-1} + \omega_{j-1}.$$

$\mu_j, \omega_j -$ trial-uncorrelated, zero-mean Gaussian noise

Noise μ_j**.**
Unknown. Changing from iteration to iteration.

From trial to trial our knowledge about d_j improves.

A | STEP 1: ESTIMATION

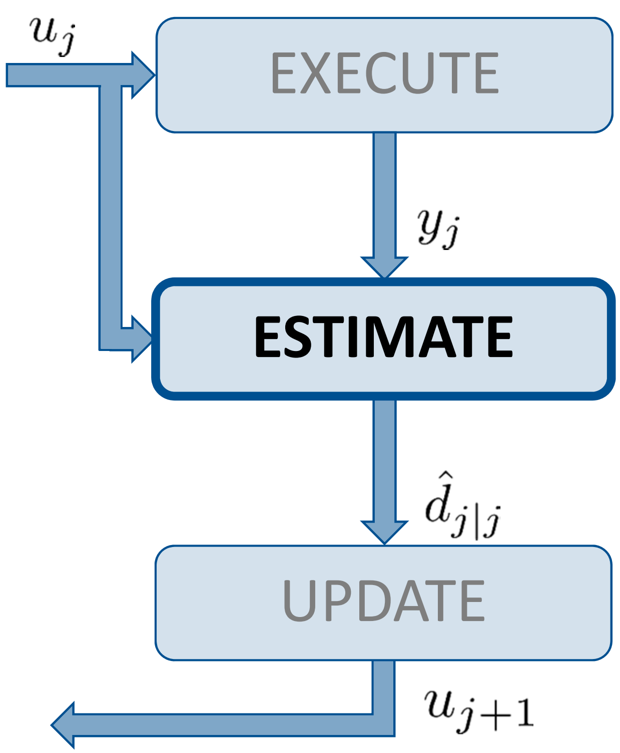

UPDATE OF DISTURBANCE ESTIMATE

via **Kalman filter** in the iteration domain:

estimates the repetitve disturbance d_j
by taking into account all past measurements.

Prediction step:

$$d_j = d_{j-1} + \omega_{j-1}.$$

Measurement update step:

$$y_j = F\,u_j + d_j + \mu_j.$$

Obtain $\hat{d}_{j|j}$.

A | STEP 2: UPDATE

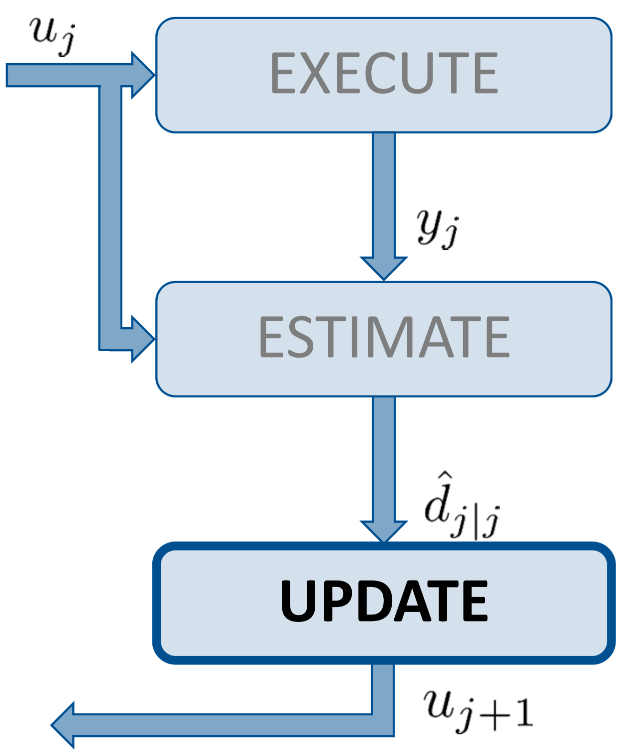

INPUT UPDATE via **convex optimization:**

minimizes the tracking error in the next trial:

$$\mathrm{E}\left[y_{j+1} | \text{all past measurements}\right] = F\, u_{j+1} + \hat{d}_{j|j}.$$

$$\min_{u_{j+1}} \left\| F u_{j+1} + \hat{d}_{j|j} \right\|_p \qquad p \in \{1,\, 2,\, \infty\}$$

subject to

$$u_{\min} \leq u_{j+1} \leq u_{\max}$$
$$x_{\min} \leq x_{j+1} \leq x_{\max}$$

➡ Obtain u_{j+1}.

A | TWO EXPERIMENTAL SCENARIOS

SCENARIO 1	SCENARIO 2
• No feedback from motion capture cameras during task execution	• Camera information is used.
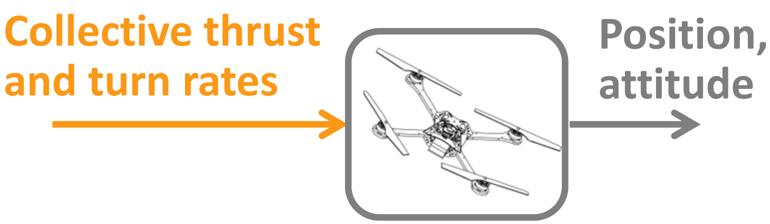	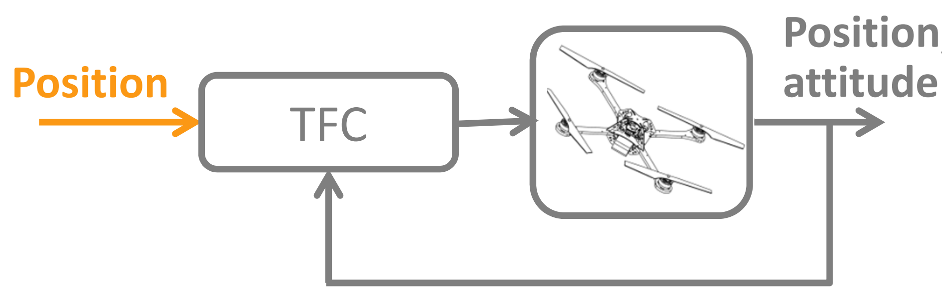
• Analytical model	• Model via numerical simulation
• 2D quadrocopter model	• 3D quadrocopter model
• Constraints on single motor thrusts and turn rates.	

A | SCENARIO 1: state trajectories

S-shaped trajectory.

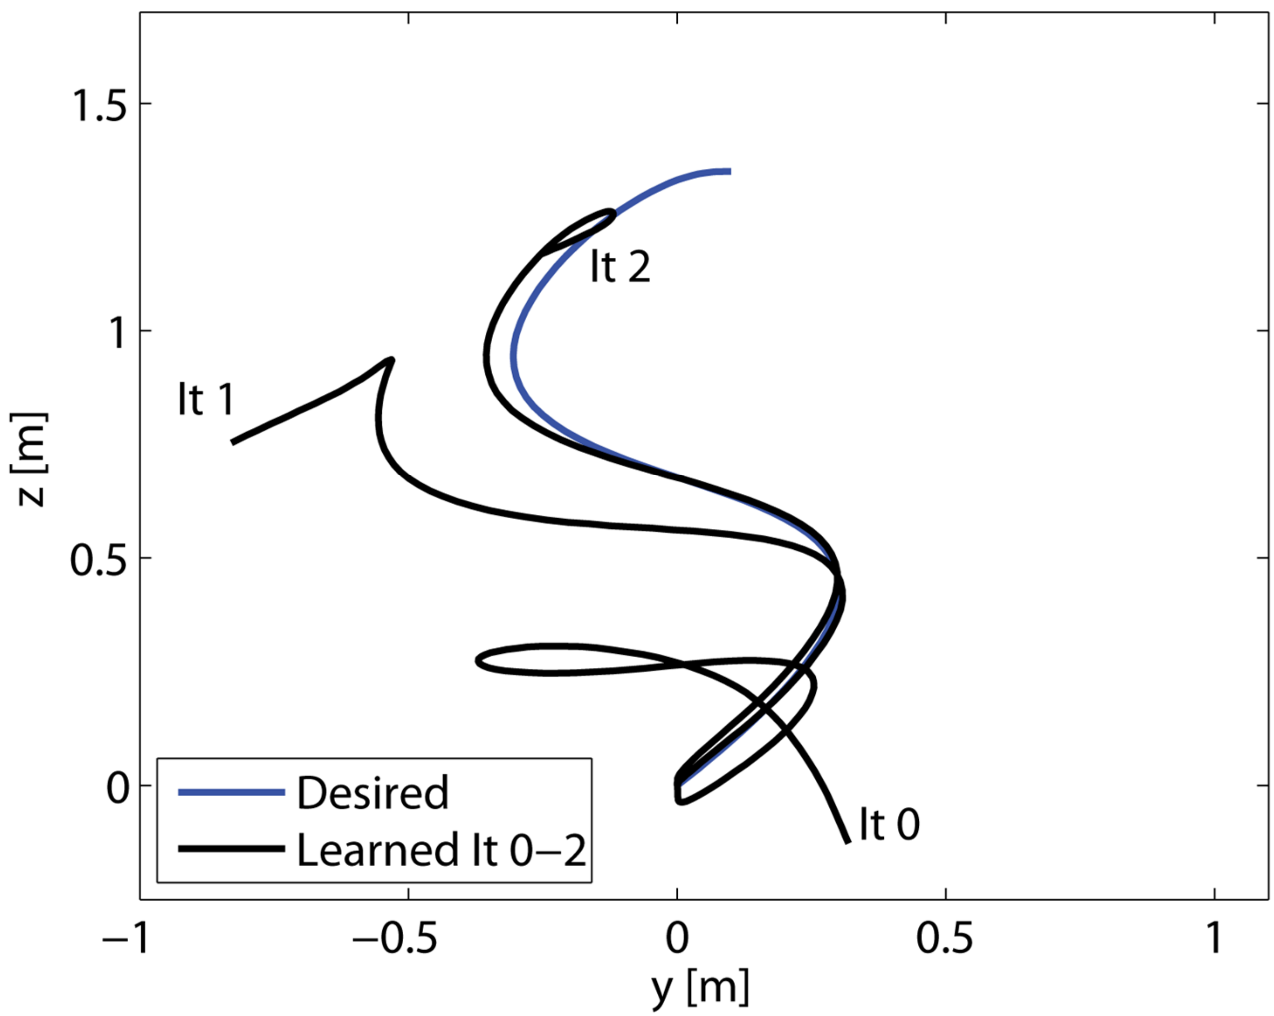

A | SCENARIO 1: input trajectories

S-shaped trajectory.

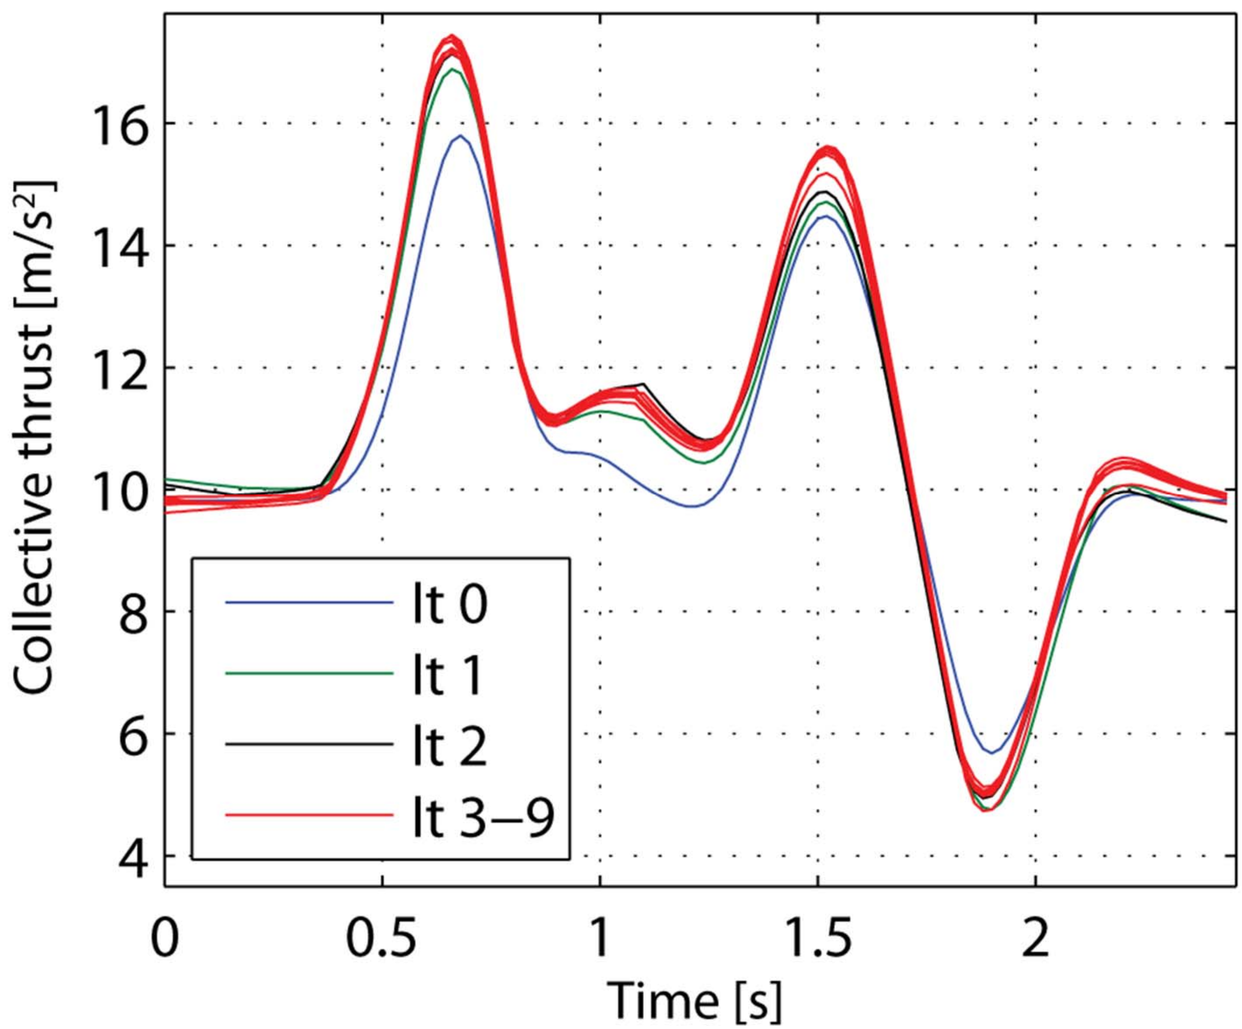

A | SCENARIO 1: state trajectories

S-shaped trajectory.

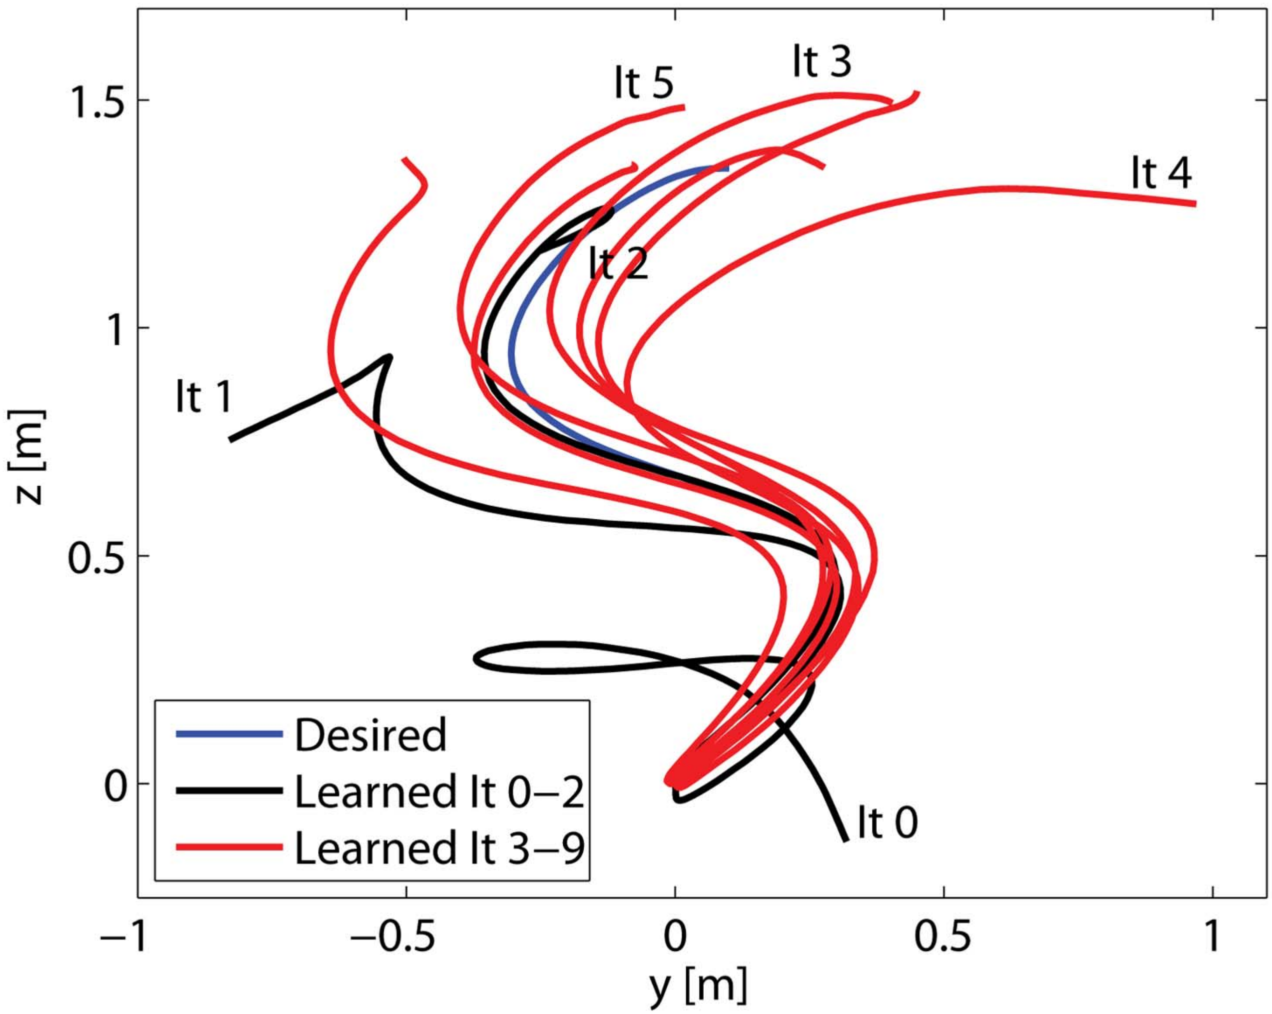

A | TWO EXPERIMENTAL SCENARIOS

<table>
<tr><th>SCENARIO 1</th><th>SCENARIO 2</th></tr>
<tr><td>• No feedback from motion capture cameras during task execution</td><td>• Camera information is used.</td></tr>
<tr><td>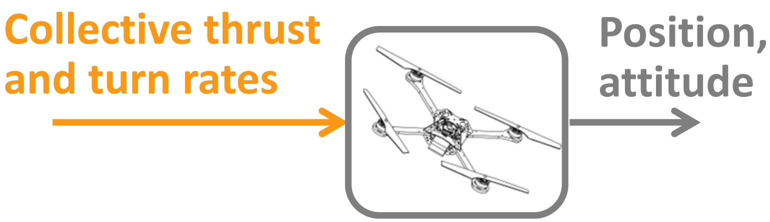
</td><td>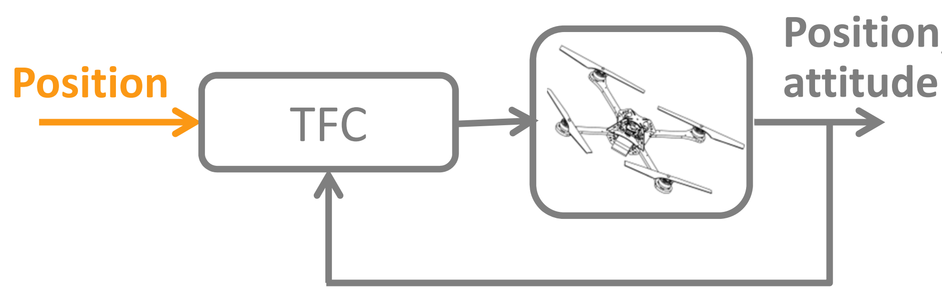
</td></tr>
<tr><td>• Analytical model</td><td>• Model via numerical simulation</td></tr>
<tr><td>• 2D quadrocopter model</td><td>• 3D quadrocopter model</td></tr>
<tr><td colspan="2">• Constraints on single motor thrusts and turn rates.</td></tr>
</table>

A | SCENARIO 2: state trajectories

S-shaped trajectory.

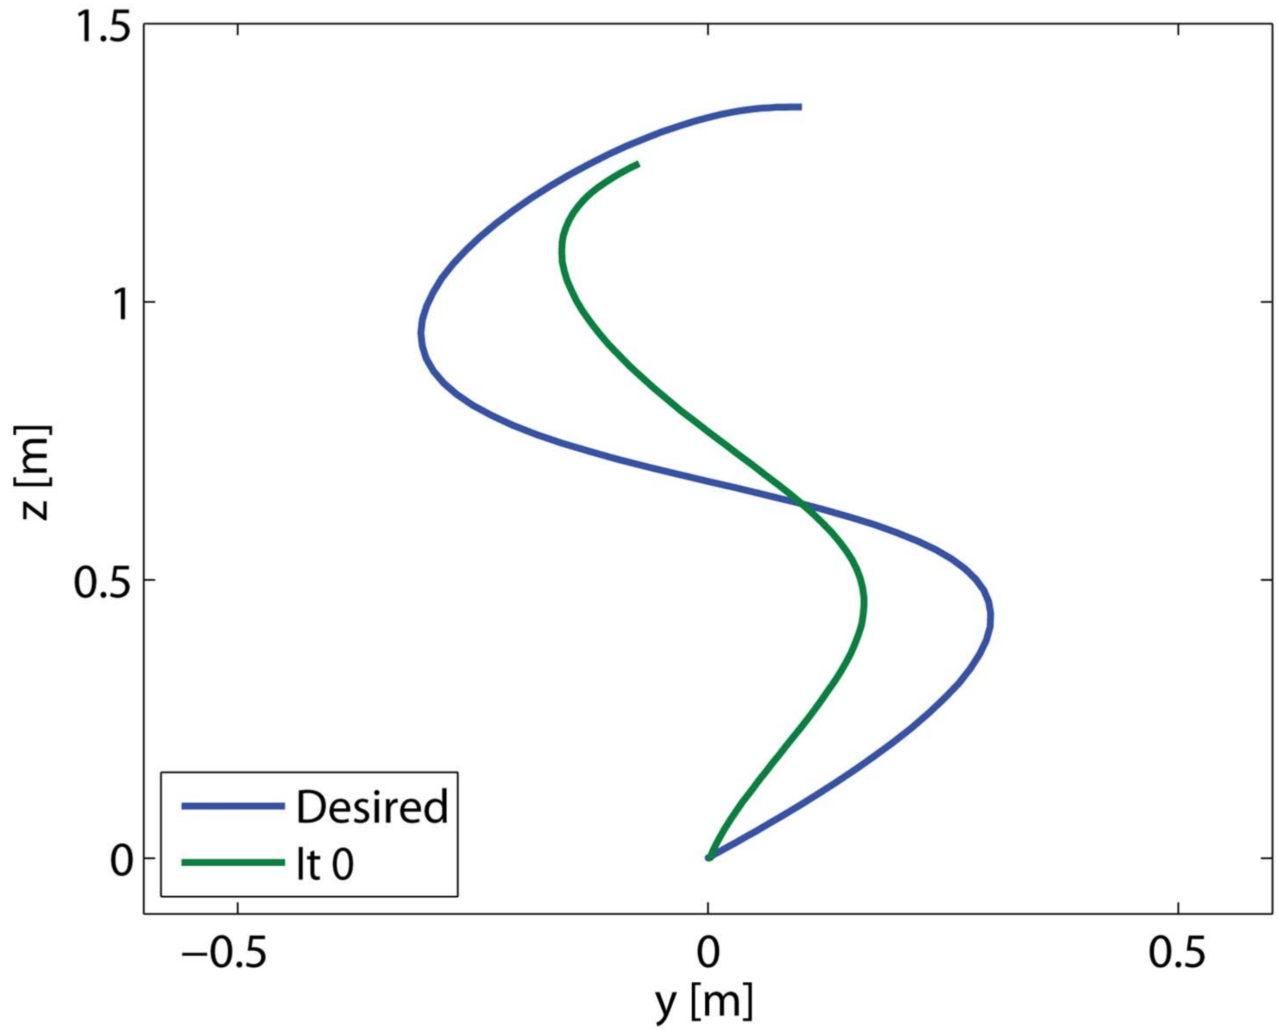

A | SCENARIO 2: state trajectories

S-shaped trajectory.

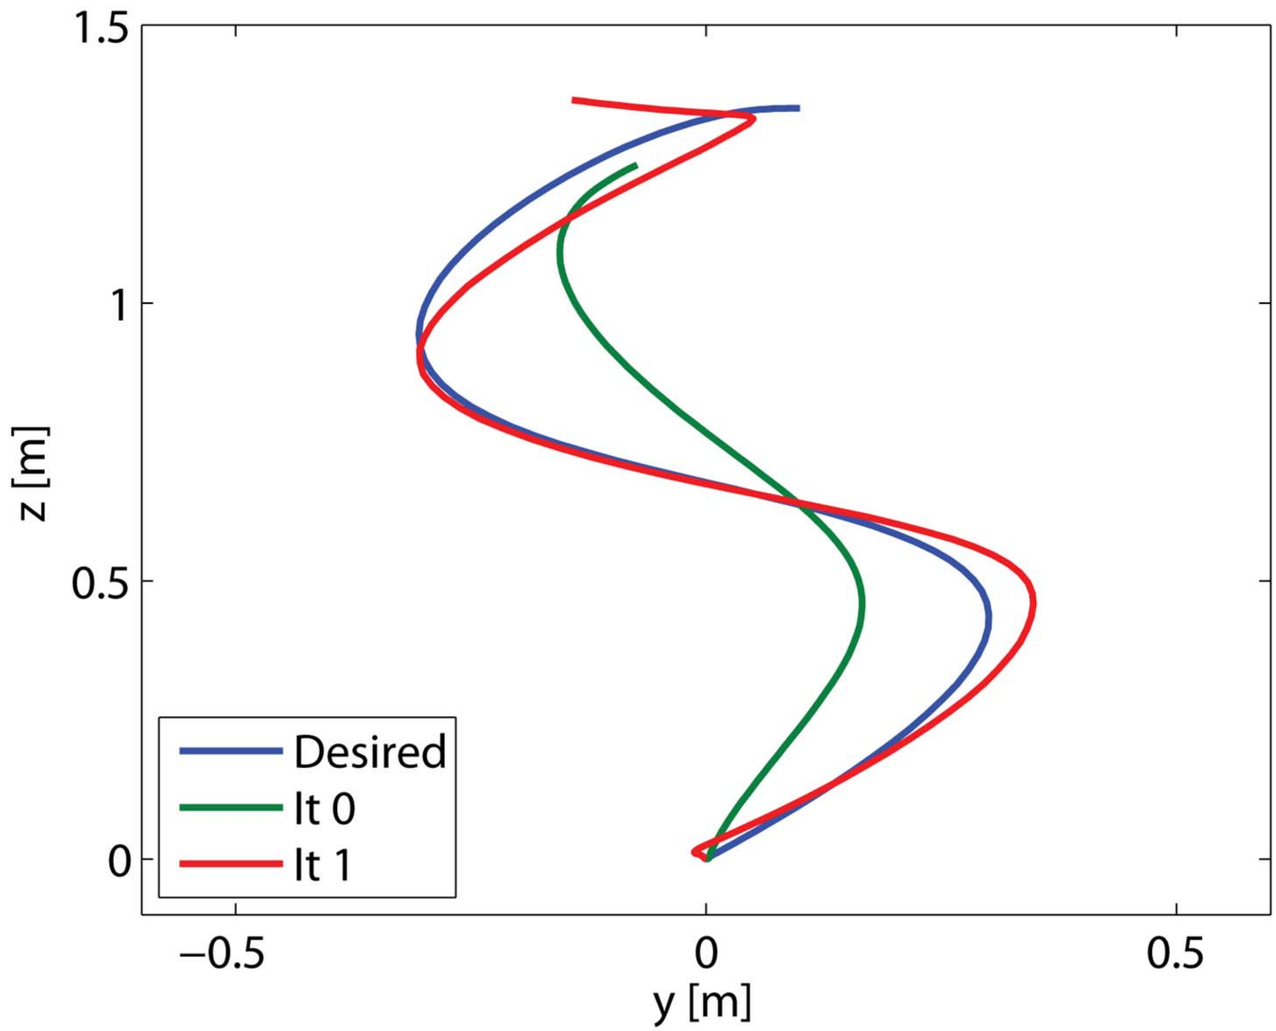

A | SCENARIO 2: state trajectories

S-shaped trajectory.

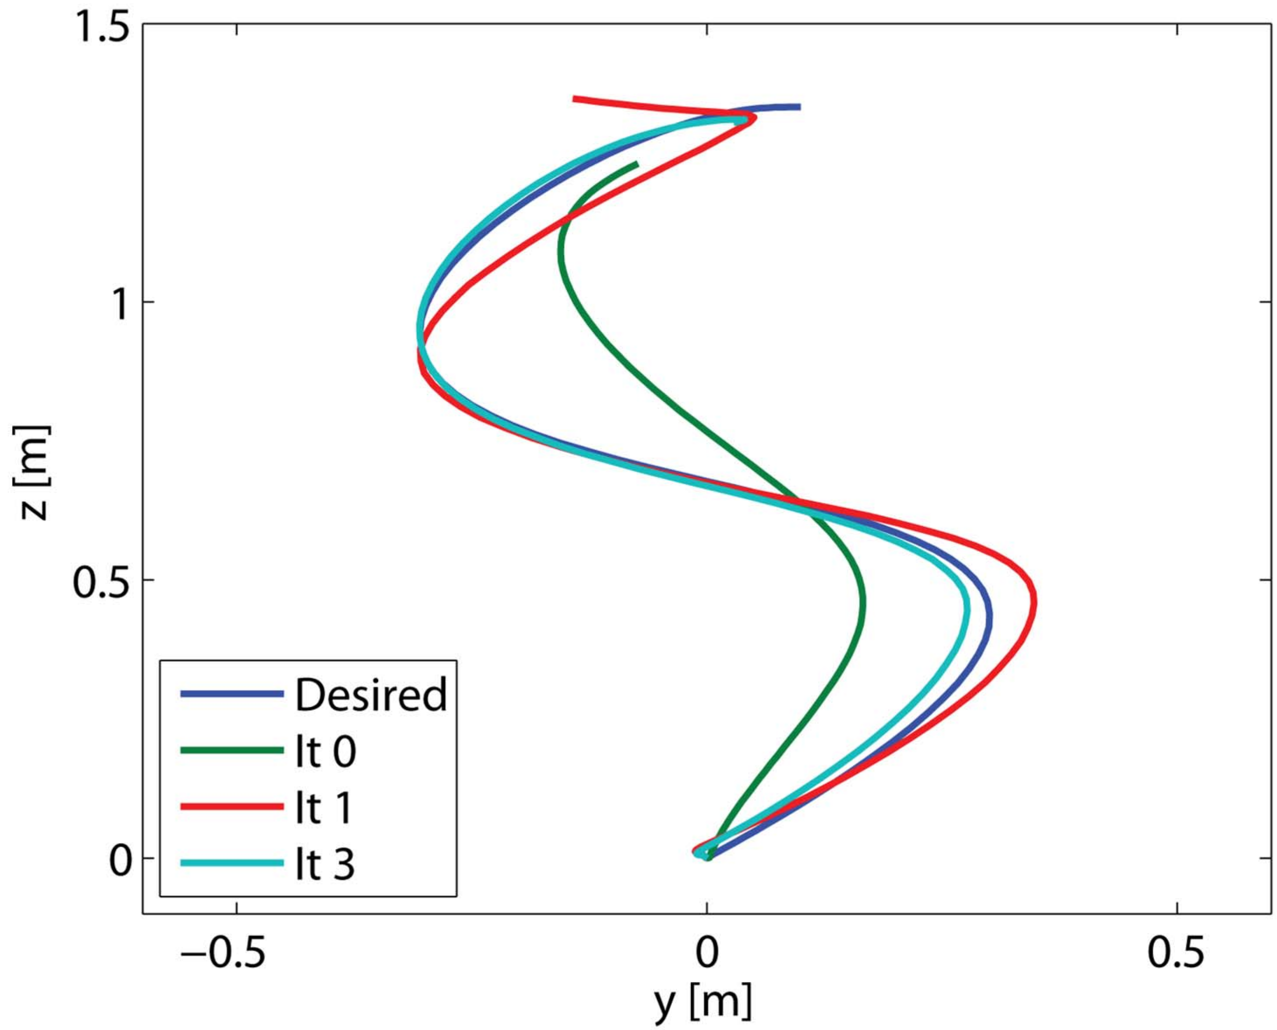

A | SCENARIO 2: state trajectories

S-shaped trajectory.

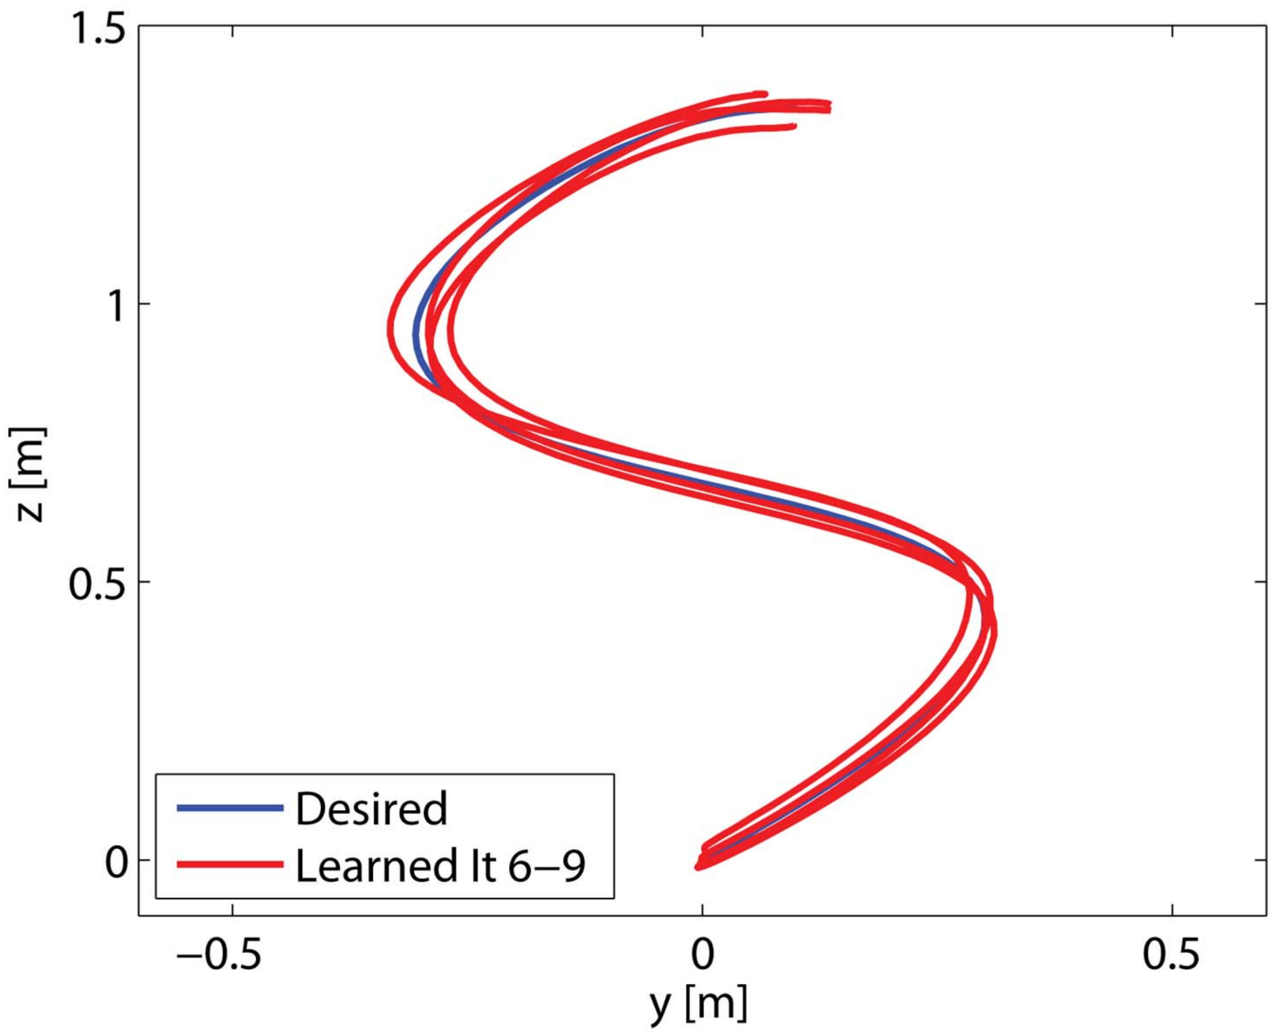

A | SCENARIO 2: error convergence

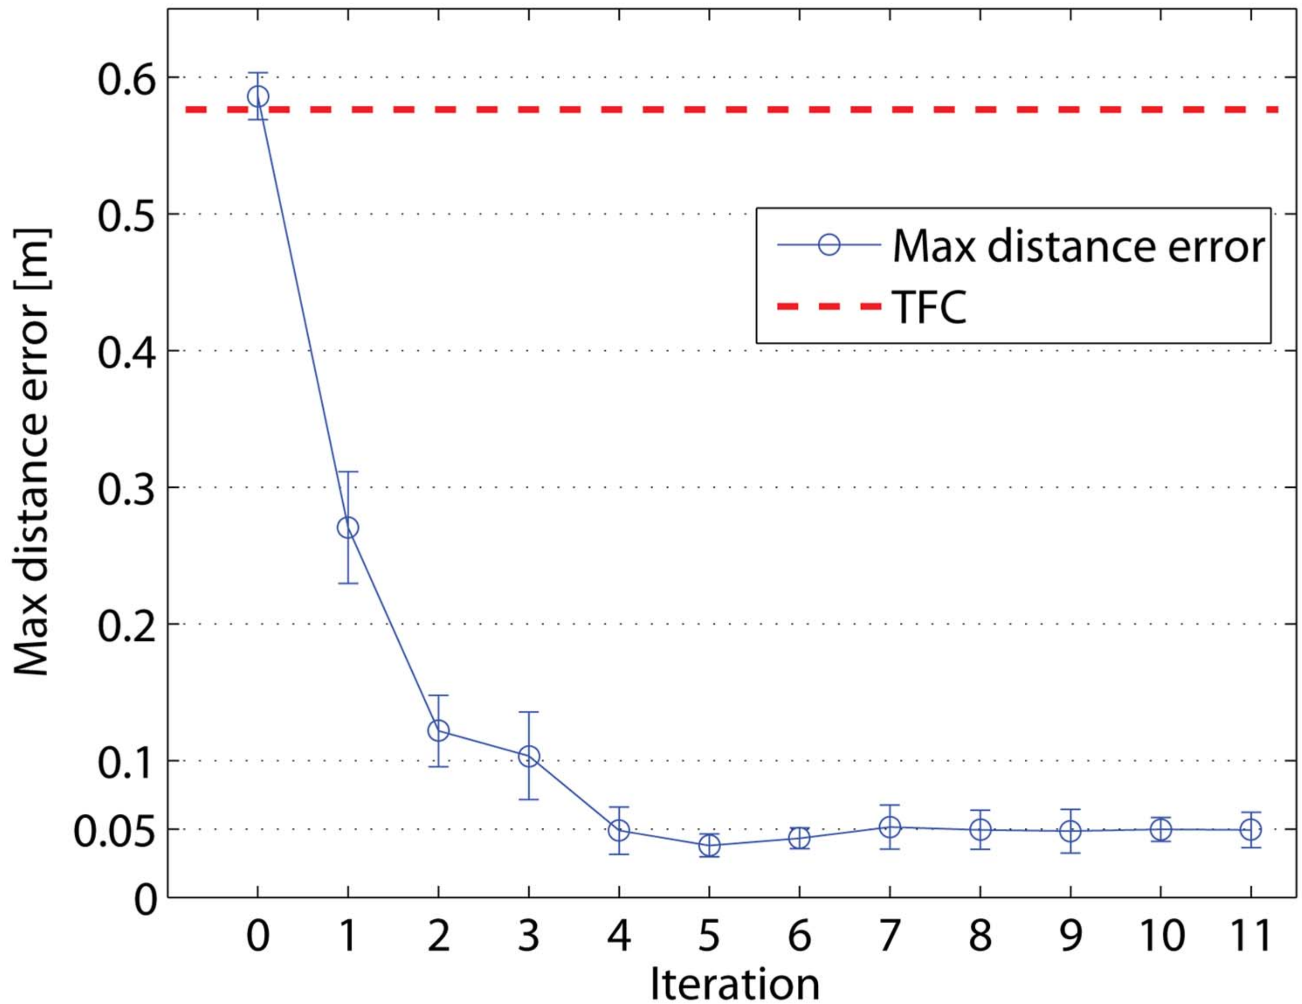

A | SUMMARY

- **Prerequisites:** approximate model of system dynamics.
- **Efficient learning algorithm:** convergence in around 5-10 iterations.
- **Acausal compensation:** outperforms pure feedback control.

Scenario 2: without learning **with learning**

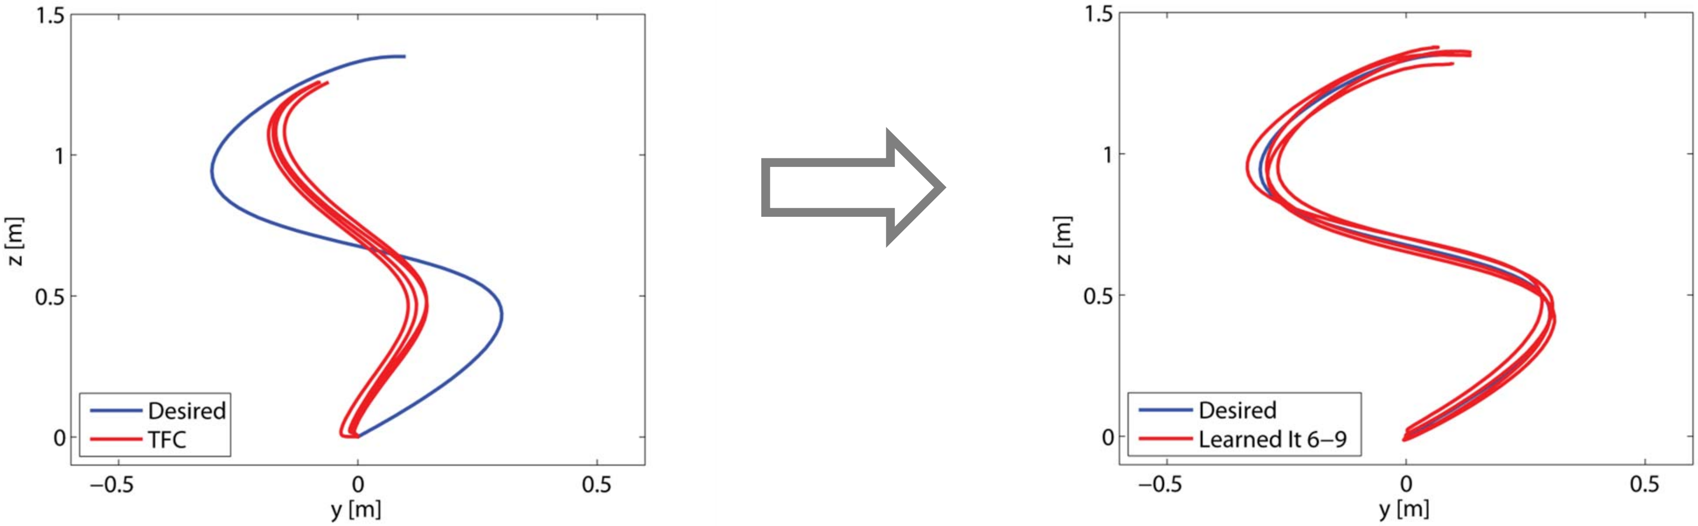

Powerful combination Learning applied to feedback-control systems: compensation for repetitive and non-repetitive disturbances.

VIDEO: http://tiny.cc/SlalomLearning

OVERVIEW

I. Introduction

II. Project A. Iterative learning for precise trajectory following

III. Project B. Learning of feed-forward parameters for rhythmic flight performances
   a. Learning approach
   b. Results

I. Summary

B | PUBLICATIONS

Peer-reviewed publications

<u>Schoellig, A. P.,</u> F. Augugliaro, and R. D’Andrea (2009):
“Synchronizing the motion of a quadrocopter to music.” In *Proceedings of IEEE International Conference on Robotics and Automation (ICRA).*

<u>Schoellig, A. P.,</u> F. Augugliaro, and R. D’Andrea (2010):
“A platform for dance performances with multiple quadrocopters.” In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)–Workshop on Robots and Musical Expressions.*

<u>Schoellig, A. P.,</u> M. Hehn, S. Lupashin, and R. D’Andrea (2011): “Feasibility of motion primitives for choreographed quadrocopter flight.” In *Proceedings of the American Control Conference (ACC)*.

<u>Schoellig, A. P.,</u> C. Wiltsche, and R. D’Andrea (2012):
“Feed-forward parameter identification for precise periodic quadrocopter motions.” In *Proceedings of the American Control Conference (ACC).*

Joint work with Federico Augugliaro (Bachelor/Master student) and Clemens Wiltsche (semester project).

VIDEO: http://tiny.cc/DanceWith3

B | LEARNING APPROACH

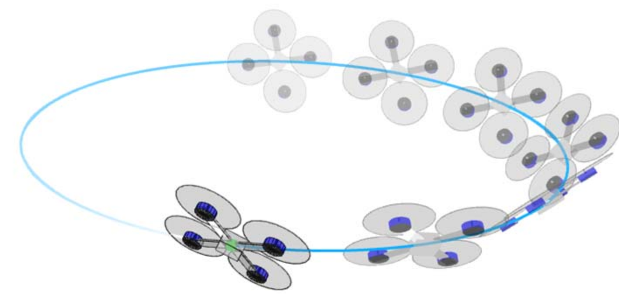

Task: Precise tracking of *periodic* motions.

Features:

- Learning through a dedicated identification routine performed prior to flight performance.
- Adaptation of only a few *input parameters.*

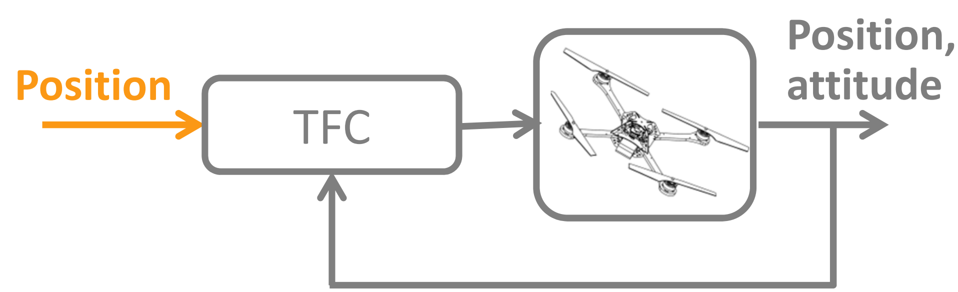

B | LEARNING APPROACH

Amplitude and phase error

PURE FEEDBACK

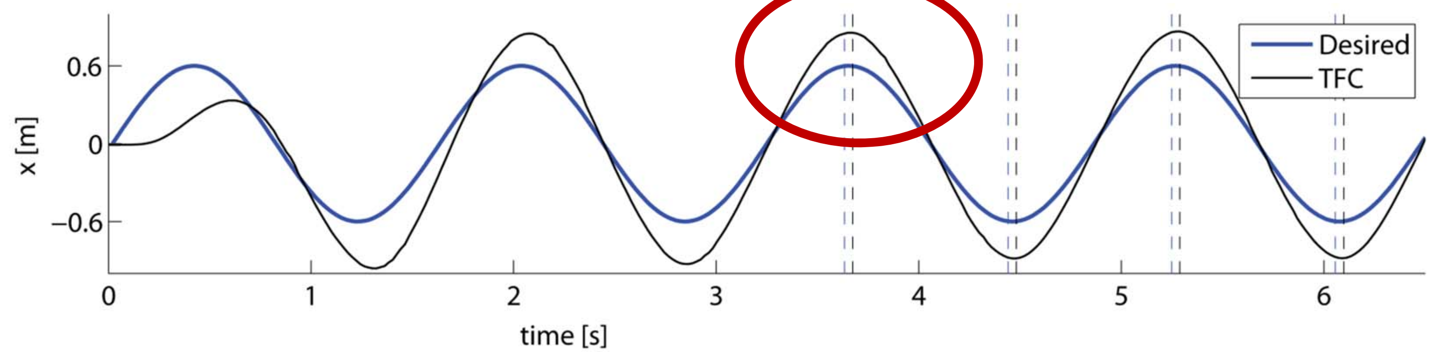

WITH LEARNED CORRECTION FACTORS

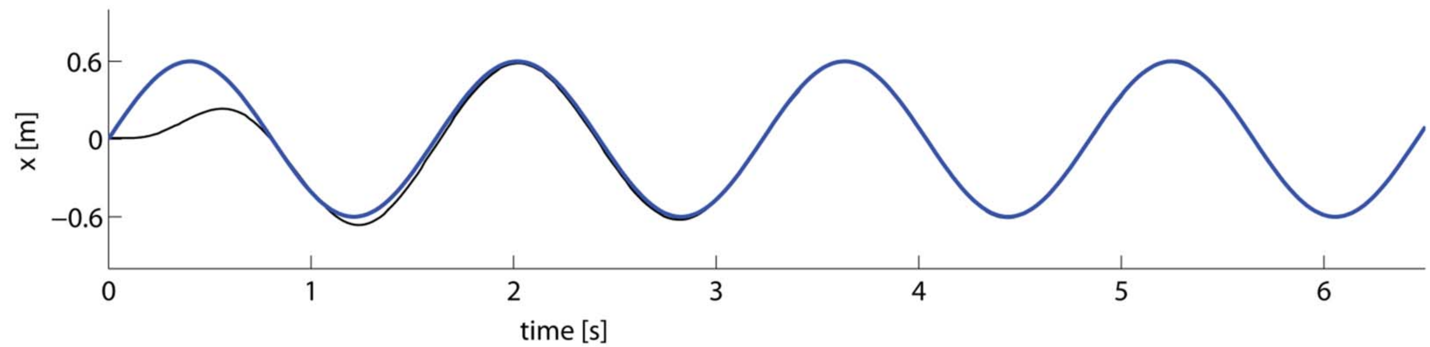

For each directional motion component and frequency, we learn:

(1) amplitude correction factor,

(2) additive phase correction.

VIDEO: http://tiny.cc/Armageddon

OVERVIEW

I. Introduction

II. Project A. Iterative learning for precise trajectory following

III. Project B. Learning of feed-forward parameters for rhythmic flight performances

IV. Summary

SUMMARY

Repetitive error components can be effectively compensated for by learning from past data.

Result is an **improved tracking performance**.

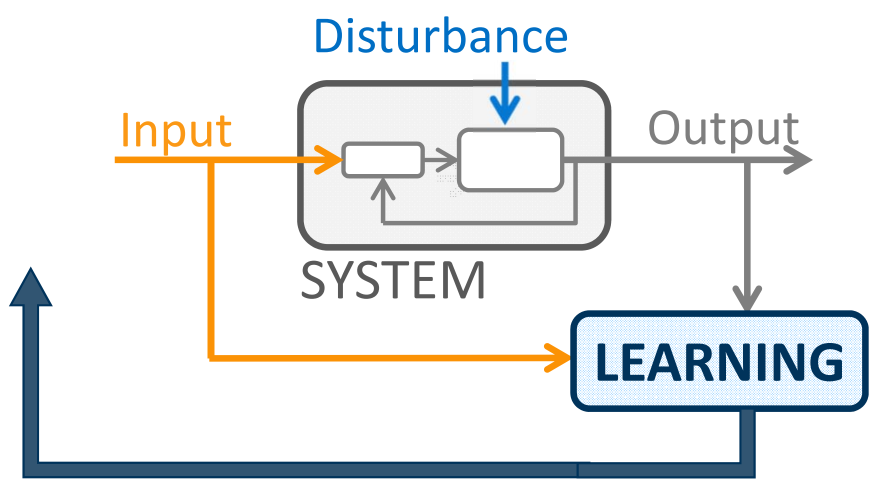

RESEARCH SUPPORT STAFF

Igor Thommen

Carolina Flores

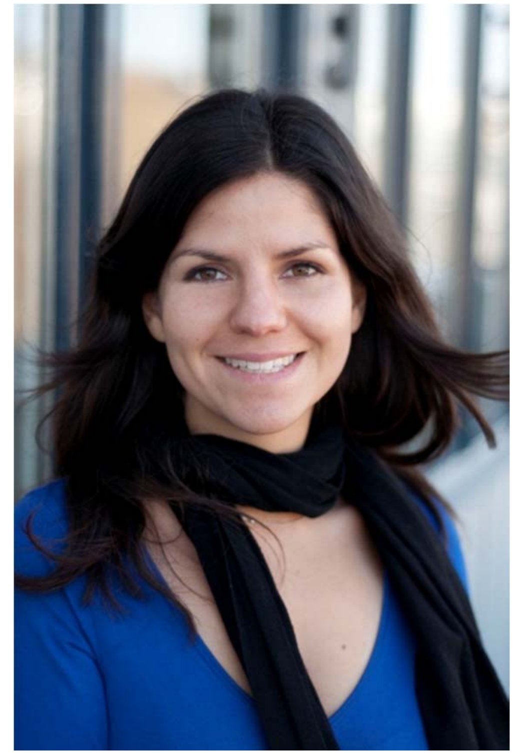

Thank you!

Hans Ulrich Honegger

Marc Corzillius

IT FOLLOWS...

Live demonstration
in the Flying Machine Arena

Eidgenössische Technische Hochschule Zürich
Swiss Federal Institute of Technology Zurich

Institute for Dynamic Systems and Control
IDSC
Institut für Dynamische Systeme und Regelungstechnik

Improving tracking performance by learning from past data

Angela P. Schoellig

Doctoral examination – July 30, 2012

Advisor: Prof. Raffaello D'Andrea // Co-advisor: Prof. Andrew Alleyne