Sensitivity of Joint Estimation in Multi-Agent Iterative Learning Control

Angela Schoellig and Raffaello D'Andrea

Institute for Dynamic Systems and Control

ETH Zurich, Switzerland

OUR FOCUS

- Group of similar agents
- Individual agents learn to perform a **single-agent task**
- The task: learn to follow a trajectory

- Does sharing information speed up simultaneous learning?

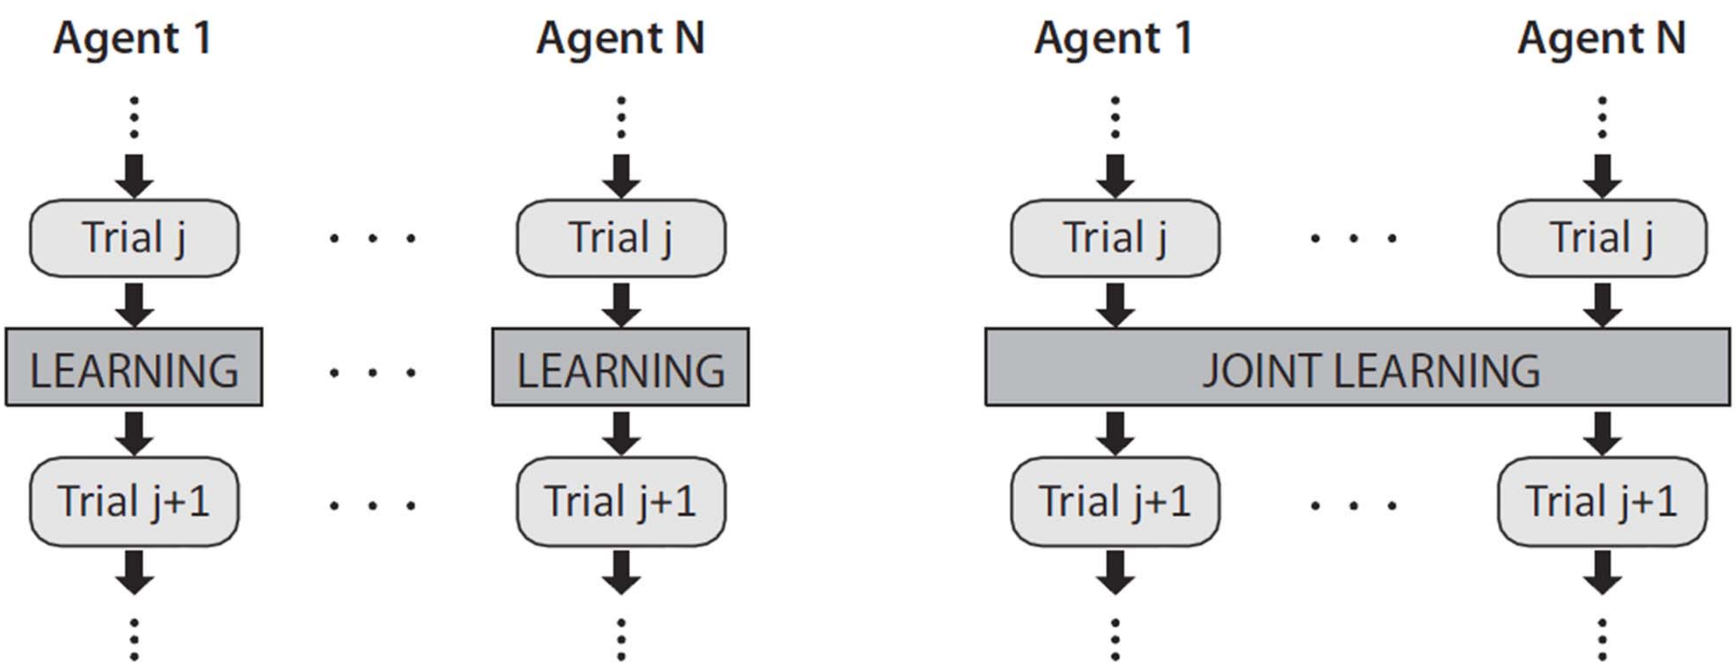

AGENTS ARE ABLE TO LEARN...

Trajectory tracking with a quadrocopter.

Full-length video. www.tiny.cc/QuadroLearnsTrajectory
[Schoellig and D'Andrea, ECC 2009]
[Schoellig, Mueller and D'Andrea, submitted to Autonomous Robots]

CAN AGENTS BENEFIT FROM EACH OTHER...

...when learning the same task?

PROBLEM STATEMENT

Group of similar agents.

Same nominal dynamics

$$\dot{\mathrm{x}}(t) = \mathrm{f}(\mathrm{x}(t), \mathrm{u}(t))$$

$$\mathrm{y}(t) = \mathrm{x}(t)$$

Physical model of real-world system.

Performing the same task.

$$(u^*(t),\, x^*(t),\, y^*(t)) \quad t \in [0, t_f]$$

GOAL OF LEARNING: Follow the desired trajectory.

Repeated and simultaneous operation.

⇨ **LEARNING OF OPEN-LOOP CONTROL CORRECTIONS.**

Q1: Is an individual agent able to learn faster when performing a task simultaneously with a group of similar agents?

LIFTED-DOMAIN REPRESENTATION

Linearize. Small deviations from nominal trajectory.

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

Discretize. Linear, time-varying difference equations.

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

Lifted-system representation. Static mapping representing one execution.

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

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

SIMILAR BUT NOT IDENTICAL...

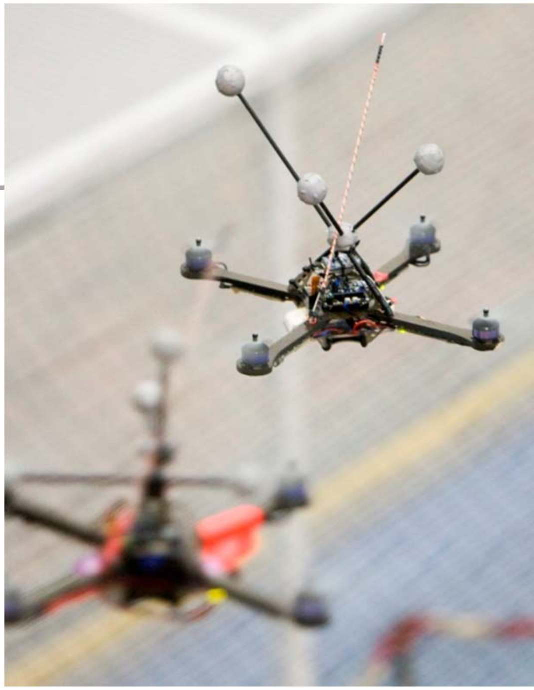

For *trial* $j,\ j \in \{1, 2, \ldots\},$ and *agent* $i,\ i \in \{1, 2, \ldots, N\},$

$$\begin{aligned} x_j^i &= F\, u_j^i + d^i + \xi_j^i \\ y_j^i &= x_j^i \qquad\qquad + \mu_j^i \end{aligned}$$

Repetitive disturbance. Unknown. Constant over iterations.

$$d^i = d^{\text{common}} + d^{i,\text{ind}}$$

$$d^{\text{common}} \sim \mathcal{N}(0, \Sigma^{\text{common}})$$
$$d^{i,\text{ind}} \sim \mathcal{N}(0, \Sigma^{\text{ind}})$$

Agents differ in the unknown part. SIMILARITY ASSUMPTION.

Noise. Unknown. Uncorrelated between iterations.

Over iterations our knowledge on d^{common} and $d^{i,\text{ind}}$ changes...

HOW DOES A SINGLE AGENT LEARN?

$$x_j^i = F\,u_j^i + d^i + \xi_j^i$$
$$y_j^i = x_j^i \qquad\qquad + \mu_j^i$$

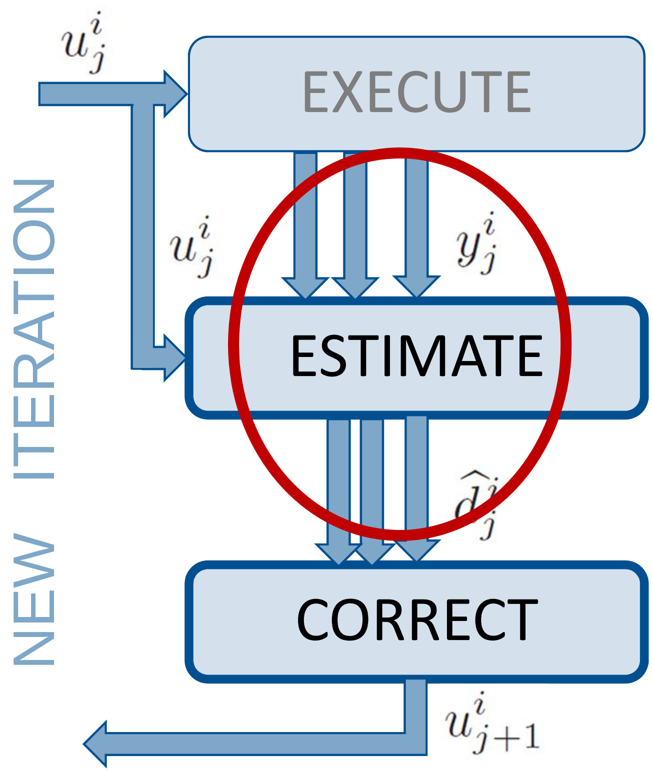

(1) Estimate the repetitive disturbance d^i by taking into account all past measurements.

Obtain $\widehat{d}_j^i$.

(2) Correct for $\widehat{d}_j^i$ by updating the input.

"Minimize" $x_{j+1}^i \approx F u_{j+1}^i + \widehat{d}_j^i$.

For example, $u_{j+1}^i = \underset{u}{\operatorname{argmin}} \left\| F u + \widehat{d}_j^i \right\|$

Can the disturbance estimate be improved by taking into account the measurements of the other agents?

FOCUS: ESTIMATION PROBLEM

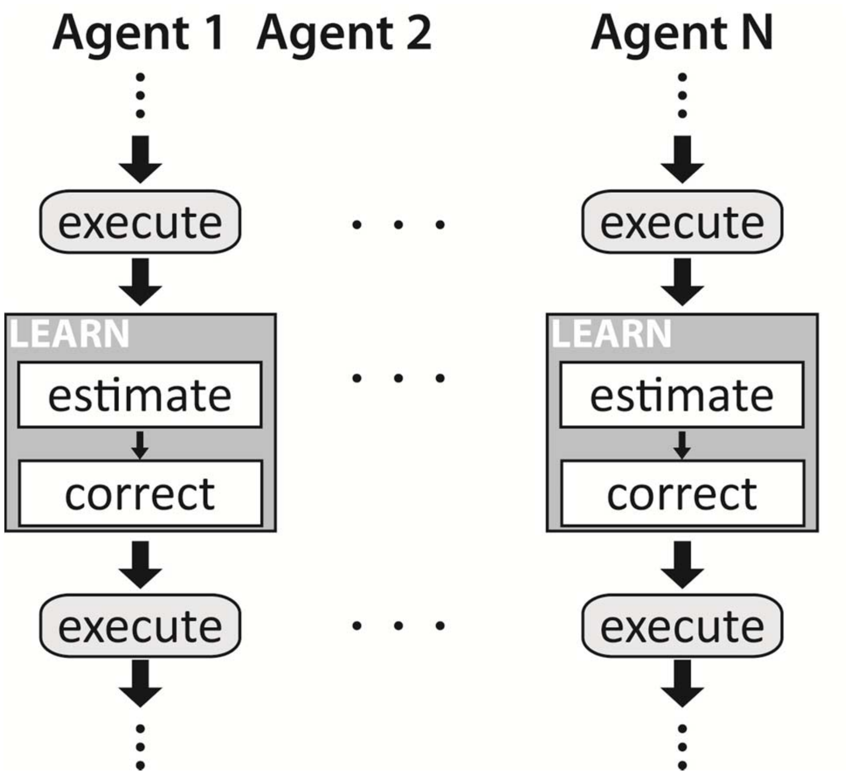

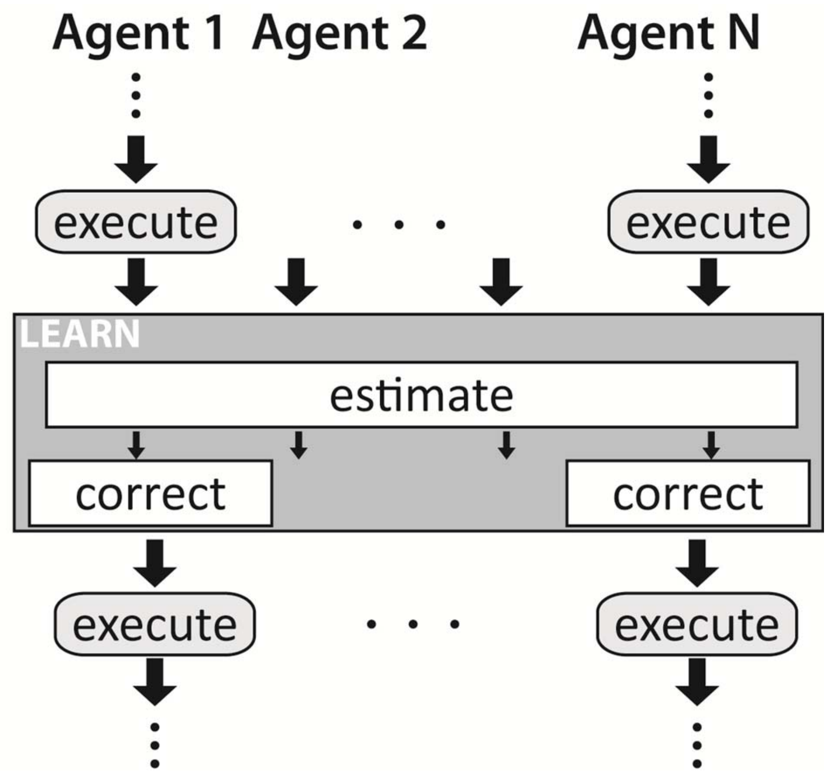

INDEPENDENT ESTIMATION vs. JOINT ESTIMATION

REDUCE MODEL

DYNAMICS

$$x_j^i = \cancel{F u_j^i} + d^i + \xi_j^i$$
$$y_j^i = x_j^i \qquad\qquad + \mu_j^i$$

- neglect deterministic part
- assume independence of vector entries

$$y_j^i = d^{\text{common}} + d^{i,\text{ind}} + v_j^i \in \mathbb{R}$$

with
$$d^{\text{common}} \sim \mathcal{N}(0,\, \sigma^{\text{common}})$$
$$d^{i,\text{ind}} \sim \mathcal{N}(0,\, \sigma^{\text{ind}})$$
$$v_j^i \sim \mathcal{N}(0,\, 1)$$

MEASUREMENT AND PROCESS NOISE

$$v_j^i = \xi_j^i + \mu_j^i$$

with $\xi_j^i \sim \mathcal{N}(0, \sigma^{\text{proc}})$, $\quad 0 \leq \sigma^{\text{proc}} \leq 1$
$\mu_j^i \sim \mathcal{N}(0, 1 - \sigma^{\text{proc}})$

JOINT ESTIMATION

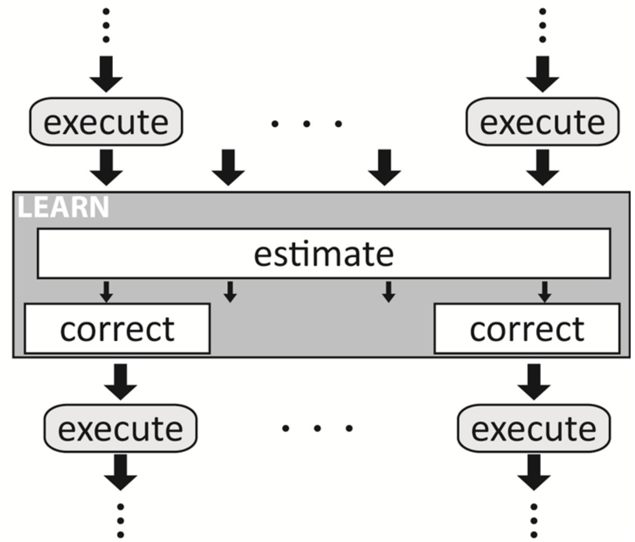

Kalman filter for the joint problem.

Estimation objective: $D = [\, d^{\text{common}},\, d^1,\, \ldots,\, d^N \,]^{\mathrm{T}} \;\in \mathbb{R}^{N+1}$

System equation:
$$D_j = D_{j-1}$$
$$Y_j = [\, \mathbf{0},\; I \,]\; D_j + V_j$$

Initial condition: $\widehat{D}_0 = \mathbf{0}, \quad P_0 = \left[p_0^{(k,l)}\right], \quad p_0^{(k,l)} = \begin{cases} \sigma^{\text{common}} + \sigma^{\text{ind}} & \text{for} \quad k = l \geq 1 \\ \sigma^{\text{common}} & \text{otherwise}\,. \end{cases}$

SIMILARITY ASSUMPTION.

LEMMA: We obtain covariance matrix in closed form. *(Proof by induction)*

$$P_j = \left[p_j^{(k,l)}\right] = \text{FUNCTION}\left(j, N, \sigma^{\text{common}}, \sigma^{\text{ind}}\right)$$

Special case: independent estimation $P_j\big|_{N=1}$

COMPARISON

JOINT LEARNING BENEFIT METRIC:

ratio of state covariances of independent vs. joint estimation

$$R = \frac{p_j^{(1,1)}\big|_{N=1} + \sigma^{\text{proc}}}{p_j^{(1,1)} + \sigma^{\text{proc}}}$$

If R > 1, joint learning is beneficial.

The VARIANCE OF THE STATE ESTIMATE is a measure for the learning performance (=experimental outcome).

$$E\left[(x_j^i - \widehat{x}_j^i)^2\right] = E\left[(d^i + \xi_j^i - \widehat{d}_j^i)^2\right] \quad \text{with} \quad \widehat{x}_j^i = \widehat{d}_j^i$$

$$= p_j^{(1,1)} + \sigma^{\text{proc}}$$

RESULT

$$x_j^i = F\,u_j^i + d^i + \xi_j^i$$
$$y_j^i = x_j^i \qquad + \mu_j^i$$

Performance increase due to joint estimation:

THEOREM 1: Pure Process Noise

$$1 \leq R^{\text{proc}} \leq \frac{1+j}{j}$$

limit case for $N \to \infty, \;\; \sigma^{\text{common}} \to \infty, \;\; \sigma^{\text{ind}} \to 0$

THEOREM 2: Pure Measurement Noise

$$1 \leq R^{\text{meas}} \leq N$$

limit case for $\sigma^{\text{common}} \to \infty, \;\; \sigma^{\text{ind}} \to 0$

[Schoellig, Alonso-Mora and D'Andrea; CDC 2010, accepted AJC]

SUMMARY

Under the given assumptions, **joint estimation**...

- improves the performance of an individual agent
- the benefit is *only* significant if

(1) agents are highly similar AND
(2) process noise is negligible AND
(3) common disturbance large compared to the measurement noise

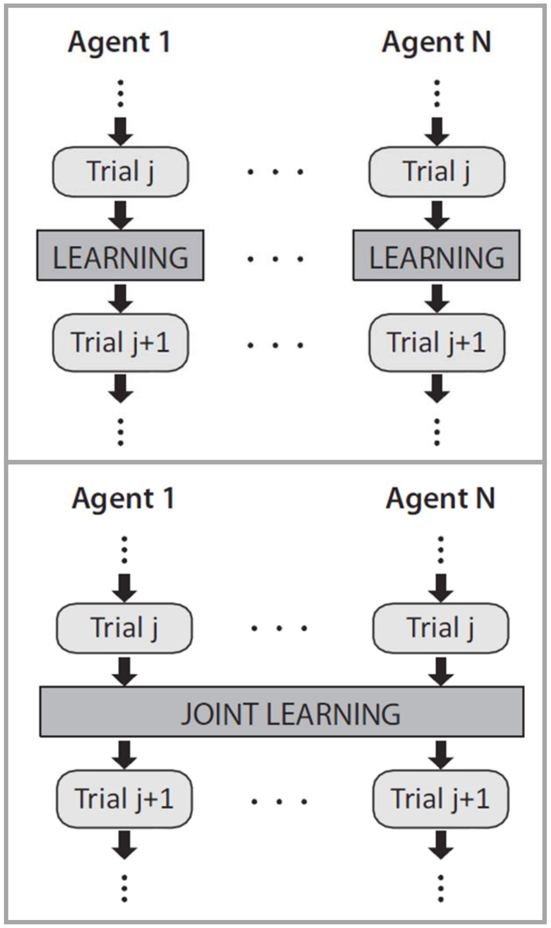

Q2: How critical is the underlying similarity assumption?

SIMILARITY ASSUMPTION

$$d^{\text{common}} \sim \mathcal{N}(0,\ \sigma^{\text{common}})$$
$$d^{i,\text{ind}} \sim \mathcal{N}(0,\ \sigma^{\text{ind}})$$

True values. For $d^i = d^{\text{common}} + d^{i,\text{ind}}$

$$d^i \sim \mathcal{N}(0, \sigma) \qquad \sigma^{\text{common}} = \epsilon\,\sigma$$
$$\sigma^{\text{ind}} = (1-\epsilon)\sigma$$

Defines degree of similarity.

ASSUME THAT DEGREE OF SIMILARITY IS UNKNOWN.

Nominal values ("our best guess").

$$d^i \sim \mathcal{N}(0, \sigma) \qquad \bar{\sigma}^{\text{common}} = \bar{\epsilon}\,\sigma$$
$$\bar{\sigma}^{\text{ind}} = (1-\bar{\epsilon})\sigma$$

⇨ **SOLVE KALMAN FILTER EQUATIONS UNDER NEW ASSUMPTIONS**

SENSITIVITY ANALYSIS – RESULTS

JOINT ESTIMATION PERFORMANCE IS DEGRADED.

LEMMA: Sufficient condition

$$\epsilon \geq \bar{\epsilon} \quad \Rightarrow \quad R \geq 1$$

Underestimate similarity → Joint estimation remains beneficial.

Worst case. Assume agents are identical and they are not, then joint estimation does NOT converge.

CONCLUSION

In the proposed framework, where we learn open-loop input corrections...

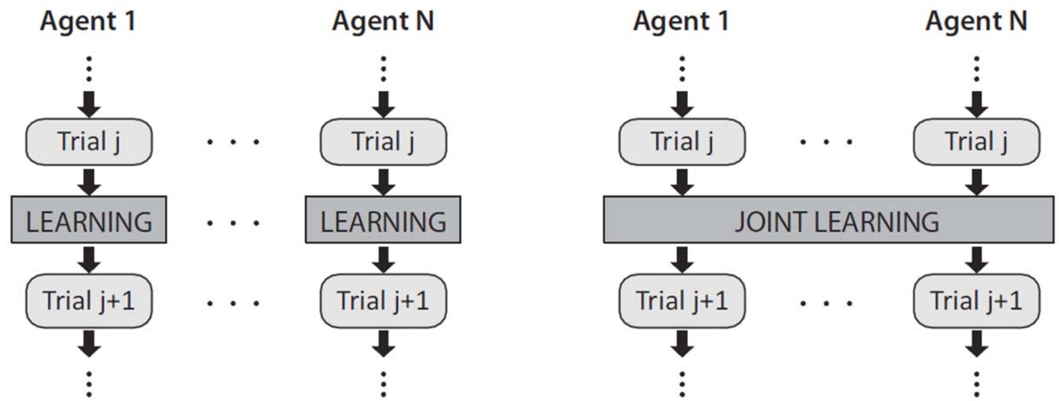

TAKE HOME MESSAGE:

(1) Joint learning good only if high similarity of unknown disturbance can be guaranteed

(2) For joint learning, it's always safer to underestimate similarity.

Choose independent learning as default since benefit of joint learning is minor for most cases.

Sensitivity of Joint Estimation in Multi-Agent Iterative Learning Control

Angela Schoellig and Raffaello D‘Andrea

Institute for Dynamic Systems and Control

ETH Zurich, Switzerland