

A Notation

Table 1 provides an overview of the notation used in this paper. We distinguish between variables used in the original model described in Section 2 [7] and quantities necessary for our inference method described in Section 3.

Table 1: Notation

$\mathbf{x}_t \in \mathbb{R}^m$	state at time t
$\mathbf{u}_t \in \mathbb{R}^p$	action
$\mathbf{y}_t \in \mathbb{R}^k$	agent's observation of the state
T	number of time steps
A, B, H	system dynamics and agent's observation matrices
$\xi_t, \omega_t, \eta_t, \epsilon_t, \varepsilon_t$	standard Gaussian noise terms
V, W, E	scaling matrices for system, observation and estimation noise
C_1, \dots, C_d	scaling matrices for control-dependent system noise
D_1, \dots, D_c	scaling matrices for state-dependent observation noise
Q_t	state-dependent costs
\tilde{R}_t	control-dependent costs
$\tilde{\mathbf{x}}_t$	agent's state estimate
K_t	filter gain matrices
L_t	control gain matrices
S	researcher's observation matrix
U	scaling matrix for researcher's observation noise
θ	model parameters
N	number of trials
F, \tilde{F}	dynamics of the joint dynamical system
ζ_t, ϑ_t	standard Gaussian noise terms
M, \tilde{M}	scaling of the signal-dependent noises in the joint dynamical system
G	scaling of the signal-independent noise in the joint dynamical system
$\mathbf{o}_t \in \mathbb{R}^s$	experimenter's observation of the state
$\hat{\boldsymbol{\mu}}_t, \hat{\Sigma}_t$	mean and covariance of the Gaussian approximation of state and agent's estimate
$\hat{\boldsymbol{\mu}}_{\tilde{\mathbf{x}} x}, \hat{\Sigma}_{\tilde{\mathbf{x}} x}$	mean and covariance of the experimenter's belief about the agent's estimate

B Approximate optimal control of LQG systems with sensorimotor noise characteristics

For approximately solving a system as described in Section 2, the optimal filters K_t and controllers L_t can be iteratively determined in an alternating fashion, leaving the respective other one constant Todorov [7].

Given filter matrices K_t , the optimal control matrices L_t are computed in form of a backward pass as

$$\begin{aligned}
 L_t &= \left(R_t + B^T P_{t+1}^x B + \sum_i C_i^T (P_{t+1}^x + P_{t+1}^e) C_i \right)^{-1} B^T P_{t+1}^x A \\
 P_t^x &= Q_t + A^T P_{t+1}^x (A - B L_t) + \sum_i D_i^T K_t^T P_{t+1}^e K_t D_i \\
 P_t^e &= A^T P_{t+1}^x B L_t + (A - K_t H)^T P_{t+1}^e (A - K_t H) \\
 s_t &= \text{tr} (P_{t+1}^x V V^T + P_{t+1}^e (V V^T + E E^T + K_t W W^T K_t^T)) + s_{t+1},
 \end{aligned}$$

where we initialize $P_T^x = Q_T, P_T^e = 0, s_T = 0$.

Given optimal control matrices L_t , the optimal filter matrices K_t are computed in form of a forward pass as

$$\begin{aligned} K_t &= A \Sigma_t^e H^T \left(H \Sigma_t^e H^T + W W^T + \sum_i D_i \left(\Sigma_t^e + \Sigma_t^{\tilde{x}} + \Sigma_t^{\tilde{x}e} + \Sigma_t^{\tilde{e}e} \right) D_i^T \right)^{-1} \\ \Sigma_{t+1}^e &= V V^T + E E^T + (A - K_t H) \Sigma_t^e A^T + \sum_i C_i L_t \Sigma_t^{\tilde{x}} L_t^T C_i^T \\ \Sigma_{t+1}^{\tilde{x}} &= E E^T + K_t H \Sigma_t^e A^T + (A - B L_t) \Sigma_t^{\tilde{x}} (A - B L_t)^T \\ &\quad + (A - B L_t) \Sigma_t^{\tilde{x}e} H^T K_t^T + K_t H \Sigma_t^{\tilde{x}e} (A - B L_t)^T \\ \Sigma_{t+1}^{\tilde{x}e} &= (A - B L_t) \Sigma_t^{\tilde{x}e} (A - K_t H)^T - E E^T, \end{aligned}$$

with $\Sigma_t^{\tilde{e}e} = (\Sigma_t^{\tilde{x}e})^T$ and we initialize

$$\begin{aligned} \Sigma_1^e &= \Sigma_1 \\ \Sigma_1^{\tilde{x}} &= \tilde{\mathbf{x}}_1 \tilde{\mathbf{x}}_1^T \\ \Sigma_1^{\tilde{x}e} &= 0. \end{aligned}$$

C Joint update equation derivation

C.1 Fully-observable state

Stacking Eq. (7) and Eq. (8) into a vector, gives

$$\begin{aligned} \begin{bmatrix} \mathbf{x}_{t+1} \\ \tilde{\mathbf{x}}_{t+1} \end{bmatrix} &= \begin{bmatrix} A \mathbf{x}_t - B L \tilde{\mathbf{x}}_t + V \boldsymbol{\xi}_t - \sum_{i=1}^d \epsilon_t^i C_i L_t \tilde{\mathbf{x}}_t, \\ (A - B L_t - K_t H) \tilde{\mathbf{x}}_t + K_t H \mathbf{x}_t + K_t W \boldsymbol{\omega}_t + K_t \sum_{i=1}^c \epsilon_t^i D_i \mathbf{x}_t + E \boldsymbol{\eta}_t \end{bmatrix} \\ &= \begin{bmatrix} A \\ K_t H \end{bmatrix} \mathbf{x}_t + \begin{bmatrix} -B L_t \\ A - B L_t - K_t H \end{bmatrix} \tilde{\mathbf{x}} + \begin{bmatrix} 0 \\ K_t \end{bmatrix} \sum_{i=1}^c \epsilon_t^i D_i \mathbf{x}_t + \sum_{i=1}^d \epsilon_t^i C_i \begin{bmatrix} -L_t \\ 0 \end{bmatrix} \tilde{\mathbf{x}} \\ &\quad + \begin{bmatrix} V & 0 & 0 \\ 0 & K_t W & E \end{bmatrix} \begin{bmatrix} \boldsymbol{\xi}_t \\ \boldsymbol{\omega}_t \\ \boldsymbol{\eta}_t \end{bmatrix} \\ &= \left(\begin{bmatrix} A \\ K_t H \end{bmatrix} + \begin{bmatrix} 0 \\ K_t \end{bmatrix} \sum_{i=1}^c \epsilon_t^i D_i \right) \mathbf{x}_t + \left(\begin{bmatrix} -B L_t \\ A - B L_t - K_t H \end{bmatrix} + \sum_{i=1}^d \epsilon_t^i C_i \begin{bmatrix} -L_t \\ 0 \end{bmatrix} \right) \tilde{\mathbf{x}} \\ &\quad + \begin{bmatrix} V & 0 & 0 \\ 0 & K_t W & E \end{bmatrix} \begin{bmatrix} \boldsymbol{\xi}_t \\ \boldsymbol{\omega}_t \\ \boldsymbol{\eta}_t \end{bmatrix} \\ &=: (F_t + M_t \sum_{i=1}^c \epsilon_t^i D_i) \mathbf{x}_t + (\tilde{F}_t + \sum_{i=1}^d \epsilon_t^i C_i \tilde{M}_t) \tilde{\mathbf{x}} + G_t \boldsymbol{\zeta}_t. \end{aligned}$$

C.2 Partially-observable state

Stacking Eq. (7), Eq. (8), and Eq. (12) into a vector, gives

$$\begin{bmatrix} \mathbf{x}_{t+1} \\ \tilde{\mathbf{x}}_{t+1} \end{bmatrix} = \begin{bmatrix} A \mathbf{x}_t - B L \tilde{\mathbf{x}}_t + V \boldsymbol{\xi}_t - \sum_{i=1}^d \epsilon_t^i C_i L_t \tilde{\mathbf{x}}_t \\ (A - B L_t - K_t H) \tilde{\mathbf{x}}_t + K_t H \mathbf{x}_t + K_t W \boldsymbol{\omega}_t + K_t \sum_{i=1}^c \epsilon_t^i D_i \mathbf{x}_t + E \boldsymbol{\eta}_t \\ S \left(A \mathbf{x}_t - B L \tilde{\mathbf{x}}_t + V \boldsymbol{\xi}_t - \sum_{i=1}^d \epsilon_t^i C_i L_t \tilde{\mathbf{x}}_t \right) + U \boldsymbol{\vartheta}_t \end{bmatrix}$$

$$\begin{aligned}
&= \begin{bmatrix} A \\ K_t H \\ SA \end{bmatrix} \mathbf{x}_t + \begin{bmatrix} -BL_t \\ A - BL_t - K_t H \\ -SBL_t \end{bmatrix} \tilde{\mathbf{x}} + \begin{bmatrix} 0 \\ K_t \\ 0 \end{bmatrix} \sum_{i=1}^c \epsilon_t^i D_i \mathbf{x}_t + \sum_{i=1}^d \epsilon_t^i C_i \begin{bmatrix} -L_t \\ 0 \\ -SL_t \end{bmatrix} \tilde{\mathbf{x}} \\
&\quad + \begin{bmatrix} V & 0 & 0 & 0 \\ 0 & K_t W & E & 0 \\ SV & 0 & 0 & U \end{bmatrix} \begin{bmatrix} \boldsymbol{\xi}_t \\ \boldsymbol{\omega}_t \\ \boldsymbol{\eta}_t \\ \boldsymbol{\vartheta} \end{bmatrix} \\
&= \left(\begin{bmatrix} A \\ K_t H \\ SA \end{bmatrix} + \begin{bmatrix} 0 \\ K_t \\ 0 \end{bmatrix} \sum_{i=1}^c \epsilon_t^i D_i \right) \mathbf{x}_t + \left(\begin{bmatrix} -BL_t \\ A - BL_t - K_t H \\ -SBL_t \end{bmatrix} + \sum_{i=1}^d \epsilon_t^i C_i \begin{bmatrix} -L_t \\ 0 \\ -SL_t \end{bmatrix} \right) \tilde{\mathbf{x}} \\
&\quad + \begin{bmatrix} V & 0 & 0 & 0 \\ 0 & K_t W & E & 0 \\ SV & 0 & 0 & U \end{bmatrix} \begin{bmatrix} \boldsymbol{\xi}_t \\ \boldsymbol{\omega}_t \\ \boldsymbol{\eta}_t \\ \boldsymbol{\vartheta} \end{bmatrix} \\
&=: (F_t + M_t \sum_{i=1}^c \epsilon_t^i D_i) \mathbf{x}_t + (\tilde{F}_t + \sum_{i=1}^d \epsilon_t^i C_i \tilde{M}_t) \tilde{\mathbf{x}} + G_t \boldsymbol{\zeta}_t.
\end{aligned}$$

D Approximate likelihood in case of partially-observable state

In the following, we define $\check{\mathbf{x}}_t$ as the true state \mathbf{x}_t and the agent's belief $\tilde{\mathbf{x}}_t$ stacked together, i.e., $\check{\mathbf{x}}_t = [\mathbf{x}_t, \tilde{\mathbf{x}}_t]$. First, we assume that the belief of the agent's belief at time step t is given by a Gaussian distribution

$$p(\mathbf{x}_t, \tilde{\mathbf{x}}_t \mid \mathbf{o}_{1:t}) = \mathcal{N} \left(\begin{bmatrix} \mathbf{x}_t \\ \tilde{\mathbf{x}}_t \end{bmatrix} \mid \boldsymbol{\mu}_{\check{\mathbf{x}}|o} = \begin{bmatrix} \boldsymbol{\mu}_{x|o} \\ \boldsymbol{\mu}_{\tilde{x}|o} \end{bmatrix}, \Sigma_{\check{\mathbf{x}}|o} = \begin{bmatrix} \Sigma_{x|o} & \Sigma_{x\tilde{x}|o} \\ \Sigma_{\tilde{x}x|o} & \Sigma_{\tilde{x}|o} \end{bmatrix} \right).$$

To approximately propagate $p(\mathbf{x}_t, \tilde{\mathbf{x}}_t \mid \mathbf{o}_{1:t})$ through Eq. (9), we compute the mean and variance of the resulting distribution via moment matching (see Appendix E) and obtain the approximation

$$p(\check{\mathbf{x}}_{t+1}, \mathbf{o}_{t+1} \mid \mathbf{o}_{1:t}) \approx \mathcal{N} \left(\begin{bmatrix} \check{\mathbf{x}}_{t+1} \\ \mathbf{o}_{t+1} \end{bmatrix} \mid \hat{\boldsymbol{\mu}}_{t+1} = \begin{bmatrix} \boldsymbol{\mu}_{\check{\mathbf{x}}} \\ \hat{\boldsymbol{\mu}}_o \end{bmatrix}, \hat{\Sigma}_{t+1} = \begin{bmatrix} \hat{\Sigma}_{\check{\mathbf{x}}} & \hat{\Sigma}_{\check{\mathbf{x}}o} \\ \hat{\Sigma}_{o\check{\mathbf{x}}} & \hat{\Sigma}_o \end{bmatrix} \right), \quad (14)$$

with

$$\begin{aligned}
\hat{\boldsymbol{\mu}}_{t+1} &= F_t \boldsymbol{\mu}_{x|o} + \tilde{F}_t \boldsymbol{\mu}_{\tilde{x}|o} = \check{F} \boldsymbol{\mu}_{\check{\mathbf{x}}|o}, \\
\hat{\Sigma}_{t+1} &= M_t \left(\sum_{i=1}^c D_i (\Sigma_{x|o} + \boldsymbol{\mu}_{x|o} \boldsymbol{\mu}_{x|o}^T) D_i^T \right) M^T + \sum_{i=1}^d C_i \tilde{M} (\Sigma_{\tilde{x}|o} + \boldsymbol{\mu}_{\tilde{x}|o} \boldsymbol{\mu}_{\tilde{x}|o}^T) \tilde{M}^T C_i^T \\
&\quad + \check{F} \Sigma_{\check{\mathbf{x}}|o} \check{F}^T + GG^T,
\end{aligned}$$

where \check{F}_t consists of F_t and \tilde{F}_t vertically stacked, i.e., $\check{F}_t := [F_t^T \quad \tilde{F}_t^T]^T$.

Marginalizing over \mathbf{x}_{t+1} and $\tilde{\mathbf{x}}_{t+1}$ gives an approximation of the likelihood factor of time step $t+1$, $p(\mathbf{o}_{t+1} \mid \mathbf{o}_{1:t}) \approx \mathcal{N}(\hat{\boldsymbol{\mu}}_o, \hat{\Sigma}_{oo})$. On the other hand, conditioning on observation \mathbf{o}_{t+1} gives the belief of the state and the agent's estimate for the following time step, $p(\mathbf{x}_{t+1}, \tilde{\mathbf{x}}_{t+1} \mid \mathbf{o}_{1:t+1}) = \mathcal{N}(\hat{\boldsymbol{\mu}}_{\check{\mathbf{x}}|o}, \hat{\Sigma}_{\check{\mathbf{x}}|o})$, with

$$\hat{\boldsymbol{\mu}}_{\check{\mathbf{x}}|o} = \hat{\boldsymbol{\mu}}_{\check{\mathbf{x}}} + \hat{\Sigma}_{\check{\mathbf{x}}o} \hat{\Sigma}_{oo}^{-1} (\mathbf{o}_{t+1} - \hat{\boldsymbol{\mu}}_o), \quad \hat{\Sigma}_{\check{\mathbf{x}}|o} = \hat{\Sigma}_{\check{\mathbf{x}}\check{\mathbf{x}}} - \hat{\Sigma}_{\check{\mathbf{x}}o} \hat{\Sigma}_{oo}^{-1} \hat{\Sigma}_{o\check{\mathbf{x}}}. \quad (15)$$

We initialize $p(\mathbf{x}_0, \tilde{\mathbf{x}}_0 \mid \mathbf{o}_0)$ with our initial belief of the state and of the agent's belief.

E Derivation of approximate propagation

For a fully-observable state, the goal is to derive a closed-form approximation for $p(\mathbf{x}_{t+1}, \tilde{\mathbf{x}}_{t+1} \mid \mathbf{x}_{1:t})$ when propagating the (approximate) belief $p(\tilde{\mathbf{x}}_t \mid \mathbf{x}_{1:t})$ and the state \mathbf{x}_t through the extended dynamics

model $p(\mathbf{x}_{t+1}, \tilde{\mathbf{x}}_{t+1} \mid \mathbf{x}_t, \tilde{\mathbf{x}}_t)$ by matching the result with a Gaussian distribution. In this section we will consider first the more general case where the state is partially observable, and derive the equations for a fully-observable state as a special case afterwards. As the fully- and partially-observable cases differ in the number of random variables (for partial-observability we have random variables \mathbf{o}_t in addition to the true state \mathbf{x}_t), we will consider a general joint dynamics model $p(\mathbf{w}_{t+1} \mid \mathbf{x}_t, \tilde{\mathbf{x}}_t)$ and general observations \mathbf{o}_t . Then, the goal becomes to derive the approximation for $p(\mathbf{w}_{t+1} \mid \mathbf{o}_{1:t})$ when propagating the belief $p(\mathbf{x}_t, \tilde{\mathbf{x}}_t \mid \mathbf{o}_{1:t})$ through the model $p(\mathbf{w}_{t+1} \mid \mathbf{x}_t, \tilde{\mathbf{x}}_t)$.

We restate the update equation for \mathbf{w}_{t+1} , coinciding with both Eq. (9) and Eq. (13),

$$\mathbf{w}_{t+1} = (F_t + M_t \sum_{i=1}^c \epsilon^i D_i) \mathbf{x}_t + (\tilde{F}_t + \sum_{i=1}^d \epsilon^i C_i \tilde{M}_t) \tilde{\mathbf{x}} + G_t \zeta_t, \quad (16)$$

where the matrices $F_t, M_t, \tilde{F}_t, C_t$, and G_t are partitioned as described in Appendix C.1.

We assume that $p(\mathbf{x}_t, \tilde{\mathbf{x}}_t \mid \mathbf{o}_{1:t})$ follows a Gaussian distribution with

$$p(\mathbf{x}_t, \tilde{\mathbf{x}}_t \mid \mathbf{o}_{1:t}) \sim \mathcal{N} \left(\begin{bmatrix} \mathbf{x}_t \\ \tilde{\mathbf{x}}_t \end{bmatrix} \mid \boldsymbol{\mu}_t = \begin{bmatrix} \boldsymbol{\mu}_x \\ \boldsymbol{\mu}_{\tilde{x}} \end{bmatrix}, \Sigma_t = \begin{bmatrix} \Sigma_{xx} & \Sigma_{x\tilde{x}} \\ \Sigma_{\tilde{x}x} & \Sigma_{\tilde{x}\tilde{x}} \end{bmatrix} \right). \quad (17)$$

To match $p(\mathbf{w}_{t+1} \mid \mathbf{o}_{1:t})$ with a Gaussian distribution, we compute the mean $\hat{\boldsymbol{\mu}}$ and variance $\hat{\Sigma}$ of Eq. (16) where \mathbf{x} and $\tilde{\mathbf{x}}$ are distributed according to Eq. (17). In the following, we will drop time indices to enhance readability. For the mean, we obtain

$$\begin{aligned} \hat{\boldsymbol{\mu}} &= \mathbb{E}_{\mathbf{x}, \epsilon, \tilde{\mathbf{x}}, \epsilon, \zeta} \left[(F + M \sum_{i=1}^c D_i \epsilon^i) \mathbf{x} + (\tilde{F} + \sum_{i=1}^d C_i \epsilon^i \tilde{M}) \tilde{\mathbf{x}} + G \zeta \right] \\ &= \mathbb{E}_{\mathbf{x}, \epsilon} \left[(F + M \sum_{i=1}^c D_i \epsilon^i) \mathbf{x} \right] + \mathbb{E}_{\tilde{\mathbf{x}}, \epsilon} \left[(\tilde{F} + \sum_{i=1}^d C_i \epsilon^i \tilde{M}) \tilde{\mathbf{x}} \right] + \mathbb{E}_{\zeta} [G \zeta] \\ &= \mathbb{E}_{\mathbf{x}} \left[(F + M \sum_{i=1}^c D_i \mathbb{E}_{\epsilon^i} [\epsilon^i]) \mathbf{x} \right] + \mathbb{E}_{\tilde{\mathbf{x}}} \left[(\tilde{F} + \sum_{i=1}^d C_i \mathbb{E}_{\epsilon^i} [\epsilon^i] \tilde{M}) \tilde{\mathbf{x}} \right] + G \mathbb{E}_{\zeta} [\zeta] \\ &= \mathbb{E}_{\mathbf{x}} [F \mathbf{x}] + \mathbb{E}_{\tilde{\mathbf{x}}} [\tilde{F} \tilde{\mathbf{x}}] + G \mathbb{E}_{\zeta} [\zeta] \\ &= F \boldsymbol{\mu}_x + \tilde{F} \boldsymbol{\mu}_{\tilde{x}} \\ &=: \hat{\boldsymbol{\mu}}_x + \hat{\boldsymbol{\mu}}_{\tilde{x}}. \end{aligned}$$

For the variance, we use that ζ is independent of $\check{\mathbf{x}} := [\mathbf{x}^T, \tilde{\mathbf{x}}^T]^T$, therefore

$$\hat{\Sigma} = \hat{\Sigma}_{\check{\mathbf{x}}} + \hat{\Sigma}_{\zeta},$$

with $\hat{\Sigma}_{\zeta} = GG^T$ and we define

$$T_x + T_{\tilde{x}} := (F + M \sum_{i=1}^c D_i \epsilon^i) \mathbf{x} + (\tilde{F} + \sum_{i=1}^d C_i \epsilon^i \tilde{M}) \tilde{\mathbf{x}}.$$

To derive $\hat{\Sigma}_{\check{\mathbf{x}}}$, we first regard the terms $\mathbb{E}_{\mathbf{x}, \epsilon} [T_x T_x^T]$:

$$\begin{aligned} \mathbb{E}_{\mathbf{x}, \epsilon} [T_x T_x^T] &= \mathbb{E}_{\mathbf{x}, \epsilon} \left[(F + M \sum_{i=1}^c D_i \epsilon^i) \mathbf{x} \mathbf{x}^T (F^T + \sum_{i=1}^c \epsilon^i D_i^T M^T) \right] \\ &= \mathbb{E}_{\epsilon} \left[(F + M \sum_{i=1}^c D_i \epsilon^i) \mathbb{E}_{\mathbf{x}} [\mathbf{x} \mathbf{x}^T] (F^T + \sum_{i=1}^c \epsilon^i D_i^T M^T) \right] \\ &= \mathbb{E}_{\epsilon} \left[(F + M \sum_{i=1}^c D_i \epsilon^i) \Upsilon_{xx} (F^T + \sum_{i=1}^c \epsilon^i D_i^T M^T) \right] \end{aligned}$$

$$\begin{aligned}
&= F\Upsilon_{xx}F^T + \mathbb{E}_\epsilon \left[F\Upsilon_{xx} \sum_{i=1}^c \epsilon^i D_i^T M^T \right] + \mathbb{E}_\epsilon \left[M \left(\sum_{i=1}^c D_i \epsilon^i \right) \Upsilon_{xx} F^T \right] \\
&\quad + \mathbb{E}_\epsilon \left[M \left(\sum_{i=1}^c D_i \epsilon^i \right) \Upsilon_{xx} \left(\sum_{i=1}^c \epsilon^i D_i^T M^T \right) \right] \\
&= F\Upsilon_{xx}F^T + F\Upsilon_{xx} \sum_{i=1}^c \mathbb{E}_\epsilon [\epsilon^i] D_i^T M^T + M \left(\sum_{i=1}^c D_i \mathbb{E}_\epsilon [\epsilon^i] \right) \Upsilon_{xx} F^T \\
&\quad + \mathbb{E}_\epsilon \left[M \left(\sum_{i=1}^c D_i \epsilon^i \right) \Upsilon_{xx} \left(\sum_{i=1}^c \epsilon^i D_i^T M^T \right) \right] \\
&= F\Upsilon_{xx}F^T + \mathbb{E}_\epsilon \left[M \left(\sum_{i=1}^c D_i \epsilon^i \right) \Upsilon_{xx} \left(\sum_{i=1}^c \epsilon^i D_i^T M^T \right) \right] \\
&= F\Upsilon_{xx}F^T + \mathbb{E}_\epsilon \left[M \left(\sum_{i=1}^c \sum_{j=1}^c D_i \epsilon^i \Upsilon_{xx} \epsilon^j D_j^T \right) M^T \right] \\
&= F\Upsilon_{xx}F^T + \mathbb{E}_\epsilon \left[M \left(\sum_{i=1}^c D_i \epsilon^i \Upsilon_{xx} \epsilon^i D_i^T \right) M^T \right] \\
&= F\Upsilon_{xx}F^T + M \left(\sum_{i=1}^c D_i \mathbb{E}_{\epsilon^i} [\epsilon^i \epsilon^i] \Upsilon_{xx} D_i^T \right) M^T \\
&= F\Upsilon_{xx}F^T + M \left(\sum_{i=1}^c D_i \Upsilon_{xx} D_i^T \right) M^T,
\end{aligned}$$

where we defined the (raw) second moments of \mathbf{x} as

$$\begin{aligned}
\Upsilon_{\mathbf{x}} &:= \mathbb{E}_{\mathbf{x}} [\mathbf{x}\mathbf{x}^T] \\
&= \Sigma_{\mathbf{x}} + \boldsymbol{\mu}_{\mathbf{x}}\boldsymbol{\mu}_{\mathbf{x}}^T.
\end{aligned}$$

With that, we obtain for $\mathbb{E}_{\mathbf{x},\epsilon} [T_x T_x^T] - \boldsymbol{\mu}_{\mathbf{x}}\boldsymbol{\mu}_{\mathbf{x}}^T$:

$$\begin{aligned}
\mathbb{E}_{\mathbf{x},\epsilon} [T_x T_x^T] - \boldsymbol{\mu}_{\mathbf{x}}\boldsymbol{\mu}_{\mathbf{x}}^T &= F\Upsilon_{xx}F^T + M \left(\sum_{i=1}^c D_i \Upsilon_{xx} D_i^T \right) M^T - F\boldsymbol{\mu}_{\mathbf{x}}\boldsymbol{\mu}_{\mathbf{x}}^T F^T \\
&= F(\Upsilon_{xx} - \boldsymbol{\mu}_{\mathbf{x}}\boldsymbol{\mu}_{\mathbf{x}}^T)F^T + M \left(\sum_{i=1}^c D_i \Upsilon_{xx} D_i^T \right) M^T \\
&= F\Sigma_{xx}F^T + M \left(\sum_{i=1}^c D_i \Upsilon_{xx} D_i^T \right) M^T
\end{aligned}$$

A similar derivation follows for $\mathbb{E}_{\tilde{\mathbf{x}},\epsilon} [T_{\tilde{x}} T_{\tilde{x}}^T] - \boldsymbol{\mu}_{\tilde{\mathbf{x}}}\boldsymbol{\mu}_{\tilde{\mathbf{x}}}^T$, giving

$$\mathbb{E}_{\tilde{\mathbf{x}},\epsilon} [T_{\tilde{x}} T_{\tilde{x}}^T] - \boldsymbol{\mu}_{\tilde{\mathbf{x}}}\boldsymbol{\mu}_{\tilde{\mathbf{x}}}^T = \tilde{F}\tilde{\Sigma}_{\tilde{x}\tilde{x}}\tilde{F}^T + \sum_{i=1}^d C_i \tilde{M} \Upsilon_{\tilde{x}\tilde{x}} \tilde{M}^T C_i^T.$$

Using these intermediate results, we can now compute $\hat{\Sigma}_{\tilde{\mathbf{x}}}$:

$$\begin{aligned}
\hat{\Sigma}_{\tilde{\mathbf{x}}} &= \mathbb{E}_{\mathbf{x},\epsilon,\tilde{\mathbf{x}},\epsilon} [(T_x + T_{\tilde{x}})(T_x + T_{\tilde{x}})^T] - \mathbb{E}_{\mathbf{x},\epsilon,\tilde{\mathbf{x}},\epsilon} [(T_x + T_{\tilde{x}})] \mathbb{E}_{\mathbf{x},\epsilon,\tilde{\mathbf{x}},\epsilon} [(T_x + T_{\tilde{x}})]^T \\
&= \mathbb{E}_{\mathbf{x},\epsilon,\tilde{\mathbf{x}},\epsilon} [(T_x T_x^T + T_x T_{\tilde{x}}^T + T_{\tilde{x}} T_x^T + T_{\tilde{x}} T_{\tilde{x}}^T)] - (\boldsymbol{\mu}_{\mathbf{x}} + \boldsymbol{\mu}_{\tilde{\mathbf{x}}})(\boldsymbol{\mu}_{\mathbf{x}} + \boldsymbol{\mu}_{\tilde{\mathbf{x}}})^T \\
&= \mathbb{E}_{\mathbf{x},\epsilon} [T_x T_x^T] - \boldsymbol{\mu}_{\mathbf{x}}\boldsymbol{\mu}_{\mathbf{x}}^T + \mathbb{E}_{\mathbf{x},\epsilon,\tilde{\mathbf{x}},\epsilon} [T_x T_{\tilde{x}}^T] - \boldsymbol{\mu}_{\mathbf{x}}\boldsymbol{\mu}_{\tilde{\mathbf{x}}}^T + \mathbb{E}_{\mathbf{x},\epsilon,\tilde{\mathbf{x}},\epsilon} [T_{\tilde{x}} T_x^T] \\
&\quad - \boldsymbol{\mu}_{\tilde{\mathbf{x}}}\boldsymbol{\mu}_{\mathbf{x}}^T + \mathbb{E}_{\tilde{\mathbf{x}},\epsilon} [T_{\tilde{x}} T_{\tilde{x}}^T] - \boldsymbol{\mu}_{\tilde{\mathbf{x}}}\boldsymbol{\mu}_{\tilde{\mathbf{x}}}^T
\end{aligned}$$

$$\begin{aligned}
&\stackrel{(*)}{=} F\Sigma_{xx}F^T + M\left(\sum_{i=1}^c D_i\Upsilon_{xx}D_i^T\right)M^T + \tilde{F}\Sigma_{\tilde{x}\tilde{x}}\tilde{F}^T + \sum_{i=1}^d C_i\tilde{M}\Upsilon_{\tilde{x}\tilde{x}}\tilde{M}^TC_i^T \\
&\quad + \mathbb{E}_{\mathbf{x},\boldsymbol{\epsilon},\tilde{\mathbf{x}},\boldsymbol{\epsilon}} [T_xT_x^T] - \boldsymbol{\mu}_x\boldsymbol{\mu}_x^T + \mathbb{E}_{\mathbf{x},\boldsymbol{\epsilon},\tilde{\mathbf{x}},\boldsymbol{\epsilon}} [T_{\tilde{x}}T_{\tilde{x}}^T] - \boldsymbol{\mu}_{\tilde{x}}\boldsymbol{\mu}_{\tilde{x}}^T \\
&= M\left(\sum_{i=1}^c D_i\Upsilon_{xx}D_i^T\right)M^T + \sum_{i=1}^d C_i\tilde{M}\Upsilon_{\tilde{x}\tilde{x}}\tilde{M}^TC_i^T + F\Sigma_{xx}F^T + \tilde{F}\Sigma_{\tilde{x}\tilde{x}}\tilde{F}^T \\
&\quad + F\Sigma_{x\tilde{x}}\tilde{F}^T + \tilde{F}\Sigma_{\tilde{x}x}F^T \\
&= M\left(\sum_{i=1}^c D_i\Upsilon_{xx}D_i^T\right)M^T + \sum_{i=1}^d C_i\tilde{M}\Upsilon_{\tilde{x}\tilde{x}}\tilde{M}^TC_i^T + \check{F}\Sigma\check{F}^T,
\end{aligned}$$

where in (*) we used the previously derived results and defined \check{F}_t as F_t and \tilde{F}_t vertically stacked, i.e., $\check{F}_t := [F_t^T \quad \tilde{F}_t^T]^T$.

By putting everything together, we get $\mathbf{w}_{t+1} \sim \mathcal{N}(\hat{\boldsymbol{\mu}}_w, \hat{\Sigma}_w)$ with

$$\begin{aligned}
\hat{\boldsymbol{\mu}}_w &= F\boldsymbol{\mu}_x + \tilde{F}\boldsymbol{\mu}_{\tilde{x}} = \check{F}\boldsymbol{\mu}_{\check{x}}, \\
\hat{\Sigma}_w &= M\left(\sum_{i=1}^c D_i\Upsilon_{xx}D_i^T\right)M^T + \sum_{i=1}^d C_i\tilde{M}\Upsilon_{\tilde{x}\tilde{x}}\tilde{M}^TC_i^T + \check{F}\Sigma\check{F}^T + GG^T \\
&= M\left(\sum_{i=1}^c D_i(\Sigma_{xx} + \boldsymbol{\mu}_x\boldsymbol{\mu}_x^T)D_i^T\right)M^T + \sum_{i=1}^d C_i\tilde{M}(\Sigma_{\tilde{x}\tilde{x}} + \boldsymbol{\mu}_{\tilde{x}}\boldsymbol{\mu}_{\tilde{x}}^T)\tilde{M}^TC_i^T + \check{F}\Sigma\check{F}^T + GG^T,
\end{aligned} \tag{18}$$

where \check{F}_t consists of F_t and \tilde{F}_t vertically stacked, i.e., $\check{F}_t := [F_t^T, \tilde{F}_t^T]^T$.

E.1 Partially-observable state

In case of partial observability, we have $\mathbf{w}_{t+1} = [\mathbf{x}_{t+1}^T, \tilde{\mathbf{x}}_{t+1}^T, \mathbf{o}_{t+1}^T]^T$ with observations \mathbf{o}_t , so the goal becomes to approximate $p(\mathbf{x}_{t+1}, \tilde{\mathbf{x}}_{t+1}, \mathbf{o}_{t+1} | \mathbf{o}_{1:t})$. If $p(\mathbf{x}_t, \tilde{\mathbf{x}}_t | \mathbf{o}_{1:t})$ is distributed as in Eq. (17), Eq. (18) gives directly the formula for the approximation of $p(\mathbf{x}_{t+1}^T, \tilde{\mathbf{x}}_{t+1}^T, \mathbf{o}_{t+1}^T | \mathbf{o}_{1:t})$.

E.2 Fully-observable state

In case of a fully-observable state, we have $\mathbf{w}_{t+1} = [\mathbf{x}_{t+1}^T, \tilde{\mathbf{x}}_{t+1}^T]^T$, with observation $\mathbf{o}_t = \mathbf{x}_t$, so we are interested in approximating $p(\mathbf{x}_{t+1}, \tilde{\mathbf{x}}_{t+1} | \mathbf{x}_{1:t})$. We assume $p(\tilde{\mathbf{x}}_t | \mathbf{x}_{1:t}) = \mathcal{N}(\boldsymbol{\mu}_{\tilde{x}|\mathbf{x}}, \Sigma_{\tilde{x}|\mathbf{x}})$ and \mathbf{x}_t to be observed and therefore deterministic. We can then informally write

$$p(\mathbf{x}_t, \tilde{\mathbf{x}}_t | \mathbf{x}_{1:t}) \sim \mathcal{N}\left(\begin{bmatrix} \mathbf{x}_t \\ \tilde{\mathbf{x}}_t \end{bmatrix} \middle| \boldsymbol{\mu}_t = \begin{bmatrix} \mathbf{x}_t \\ \boldsymbol{\mu}_{\tilde{x}|\mathbf{x}} \end{bmatrix}, \Sigma_t = \begin{bmatrix} \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \Sigma_{\tilde{x}|\mathbf{x}} \end{bmatrix}\right).$$

Plugging this into Eq. (18), we obtain $\mathbf{w}_{t+1} \sim \mathcal{N}(\hat{\boldsymbol{\mu}}_w, \hat{\Sigma}_w)$, with

$$\begin{aligned}
\hat{\boldsymbol{\mu}}_w &= F\boldsymbol{\mu}_x + \tilde{F}\boldsymbol{\mu}_{\tilde{x}} \\
&= F\mathbf{x}_t + \tilde{F}\boldsymbol{\mu}_{\tilde{x}|\mathbf{x}}, \\
\hat{\Sigma}_w &= M\left(\sum_{i=1}^c D_i\Upsilon_{xx}D_i^T\right)M^T + \sum_{i=1}^d C_i\tilde{M}\Upsilon_{\tilde{x}\tilde{x}}\tilde{M}^TC_i^T + F\Sigma_{xx}F^T + \tilde{F}\Sigma_{\tilde{x}\tilde{x}}\tilde{F}^T \\
&\quad + F\Sigma_{x\tilde{x}}\tilde{F}^T + \tilde{F}\Sigma_{\tilde{x}x}F^T + GG^T \\
&= M\left(\sum_{i=1}^c D_i\mathbf{x}_t\mathbf{x}_t^TD_i^T\right)M^T + \sum_{i=1}^d C_i\tilde{M}\Upsilon_{\tilde{x}|\mathbf{x}}\tilde{M}^TC_i^T + \tilde{F}\Sigma_{\tilde{x}|\mathbf{x}}\tilde{F}^T + GG^T \\
&= M\left(\sum_{i=1}^c D_i\mathbf{x}_t\mathbf{x}_t^TD_i^T\right)M^T + \sum_{i=1}^d C_i\tilde{M}(\Sigma_{\tilde{x}|\mathbf{x}} + \boldsymbol{\mu}_{\tilde{x}|\mathbf{x}}\boldsymbol{\mu}_{\tilde{x}|\mathbf{x}}^T)\tilde{M}^TC_i^T + \tilde{F}\Sigma_{\tilde{x}|\mathbf{x}}\tilde{F}^T + GG^T,
\end{aligned}$$

where the time-dependency of the matrices was omitted to enhance readability.

F Further information on applications

If not stated otherwise, we used 100 trajectories to determine the MLEs. For optimization, we ran the optimizer 10 times with random initial points and took the optimal value to avoid local optima.

F.1 Reaching model

The reaching model was the same one used in the original publication [7], where all details can be found. The cost function which was minimized, is given by

$$(x_p(T) - x_p^*)^2 + (v \cdot x_v(T))^2 + (f \cdot x_f(T))^2 + \frac{r}{M-1} \sum_{k=0}^{M-1} u(k\Delta)^2,$$

where x_p is the position, x_v the velocity, x_f the force, x_p^* the target position, and Δ the time discretization, discretizing T into M time steps, i.e., $\Delta M = T$. Note that there is no explicit parameter for the cost of the end-point position needed because the other parameters are relative to this quantity.

We used the same model for the real reaching data, which were taken from the *Database for Reaching Experiments and Models*⁴ and were previously published [55]. We took the horizontal component of reaching movements towards targets at 0 degrees (right of center) and truncated each trial so that it contained the movement only.

F.2 Saccadic eye movement problem

We used the model by Crevecoeur and Kording [56] with a time discretization of 1.25 ms. The initial angle was set to -10 and the target angle to 10 as shown in Figure 1b of the referenced paper.

F.3 Random models

The random models were inspired by the work of Todorov [7]. For these models, the state space was four-dimensional with an additional dimension for modelling the target that the first dimension of the state should be controlled to. The action space was $n = 2$ -dimensional. The matrices A , B and H of the dynamical system were randomly sampled with

$$A_{ij}, B_{ij}, H_{ij} \sim \mathcal{N}(0, 1).$$

A and B were normalized to 1 using the Frobenius norm. The additive noises V , W , and E were sampled from LKJ-Cholesky distributions

$$V, W \sim \text{LKJ-Cholesky}(1)$$

and the multiplicative noises were sampled with

$$C_{ij}, D_{ij} \sim \text{Uniform}(0, 0.5).$$

The state cost matrices Q_t were set to $Q_{1:T-1} = 0$ and $Q_T = \mathbf{d}\mathbf{d}^T$ with $\mathbf{d} = [1 \ 0 \ \dots \ -1]$ yielding a state cost at the last time step of $(\mathbf{x}_T - 1)(\mathbf{x}_T - 1)^T$. The control cost was parametrized with $R = \text{diag}[r_1, \dots, r_n]$. We used our maximum likelihood method to infer the parameters r_1, \dots, r_n .

G Additional results

In this section we will provide additional results for the reaching and random problems.

G.1 Synthetic reaching data

G.1.1 Comparison to a baseline

Fig. 6 shows the RMSEs of our method in comparison to the two baselines described in Section 4.1.

⁴<https://crcns.org/data-sets/movements/dream/overview>

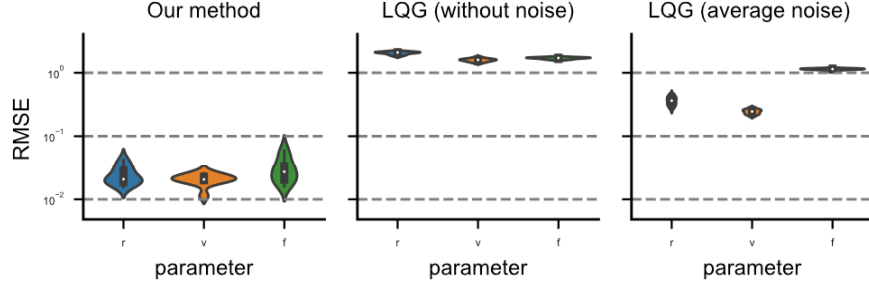


Figure 6: **Evaluation of the proposed inverse optimal control method in terms of RMSE of the maximum-likelihood parameter estimates.** Left: Our method, in which the likelihood is approximated via moment-matching. Middle: Standard LQG (without signal-dependent noise), for which the likelihood of the parameters given the data can be calculated in closed-form. Right: Standard LQG where the additive noise level was set to the average noise magnitude in the trajectories.

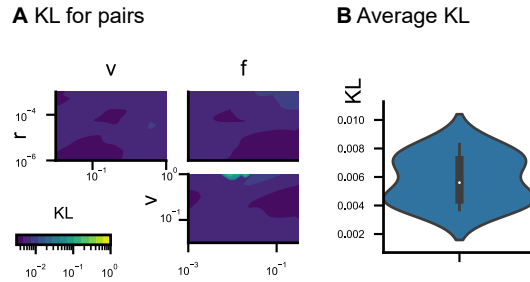


Figure 7: **KL divergences.** As a secondary evaluation metric, we show the KL divergence between trajectories simulated using the true parameters and trajectories simulated using the MLEs. The results are qualitatively similar to the RMSEs and show that the inferred parameters generate trajectories very similar to the observed data.

G.1.2 Kullback–Leibler divergence as evaluation measure

As an alternative evaluation measure, we propose to compare the distributions induced by the true parameters and the maximum likelihood parameters. For this, we estimate empirical distributions of the trajectories by generating 10,000 trajectories and computing the mean and variance for each time step to approximate the distribution for each time step by a Gaussian. The Kullback–Leibler (KL) divergence between two Gaussian distributions $p = \mathcal{N}(\boldsymbol{\mu}_p, \Sigma_p)$ and $q = \mathcal{N}(\boldsymbol{\mu}_q, \Sigma_q)$ can be calculated in closed form as

$$D_{KL}(p \parallel q) = \frac{1}{2} \left[\log \frac{|\Sigma_q|}{|\Sigma_p|} - k + (\boldsymbol{\mu}_p - \boldsymbol{\mu}_q)^T \Sigma_q^{-1} (\boldsymbol{\mu}_p - \boldsymbol{\mu}_q) + \text{tr} \{ \Sigma_q^{-1} \Sigma_p \} \right].$$

Instead of using the KL divergence directly, which is not symmetric, we consider instead a symmetric version,

$$\frac{1}{2} D_{KL}(p \parallel q) + \frac{1}{2} D_{KL}(q \parallel p),$$

and compute the mean over time to aggregate the values over time.

The results like in Fig. 2 E and F with KL divergence instead of RMSE as a metric are shown in Fig. 7.

G.1.3 Partially observable state

We evaluate a version of the reaching model in which the experimenter observes only the position and treats velocity and acceleration as latent variables. The results are qualitatively similar to the fully observed case. However, there are regions in the parameter space where estimates are worse (Fig. 8 A). Additionally, estimates of the parameters representing the penalty on velocity (v) and acceleration (f) are worse by an order of magnitude compared to the fully observed case (Fig. 8 B),

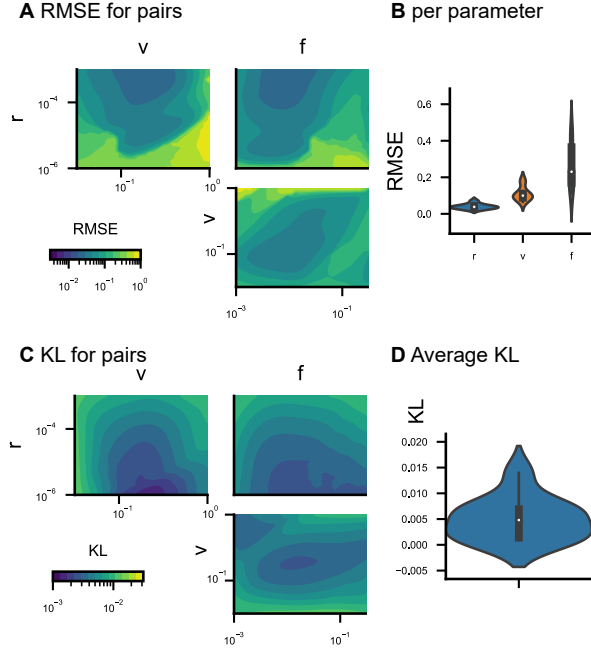


Figure 8: **Evaluation of partially observed model.** RMSEs and KLs for partially observed reaching model.

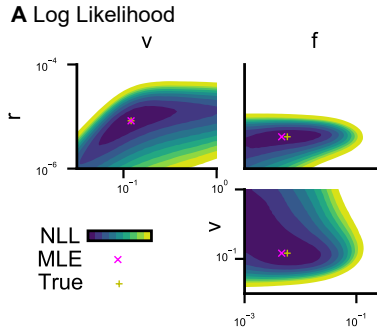


Figure 9: **Negative log likelihood of the partially observed model.**

which is to be expected when only the position is observed. Specifically, the average RMSEs are 4.0×10^{-2} (r), 1.1×10^{-1} (v), and 2.6×10^{-1} (f). However, the parameter errors do not result in large differences in the KL divergence of the simulated trajectories (position only) w.r.t. the observed trajectories, so we suspect that the higher RMSEs in estimated parameters are due to ambiguities in the trajectories for this particular model.

G.1.4 Evaluation of moment matching approximation

In Section 3.2 we introduced a moment matching approximation to make computation of the likelihood tractable. An experimentalist comparing an optimal control model to experimental data might be interested in the influence of this approximation on trajectories. For the reaching model from Section 4.1, we therefore compare the empirical distributions over trajectories estimated using Monte Carlo rollouts (using 10,000 trajectory samples) to the approximate distribution over trajectories determined using our method (given the true parameters). The difference in symmetrized KL (see Appendix G.1.2) between the empirically estimated distribution and our approximation is found to be 1.60×10^{-3} . Additionally, we compare this result to a baseline by replacing the signal-dependent noise by additive noise, for which the trajectory distribution can be calculated in closed form. The additive noise magnitude is chosen as the average of the signal-dependent noise magnitudes for the whole trajectories. Note that this quantity is not directly available and therefore has to be also

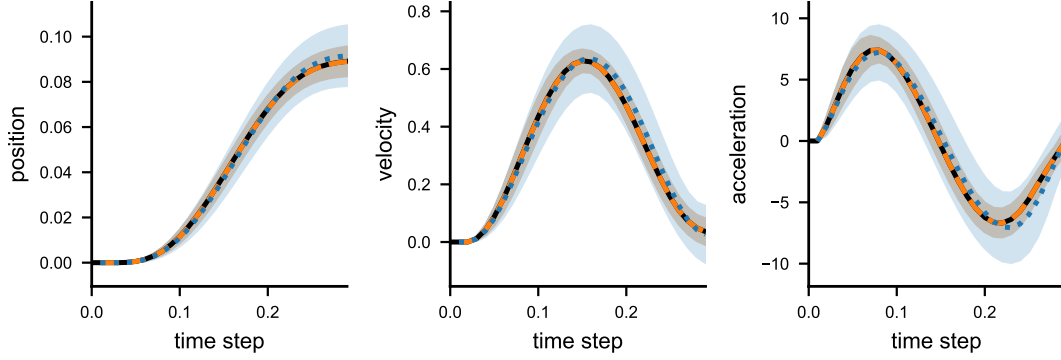


Figure 10: **Empirical evaluation of the moment matching approximation.** Trajectory distributions of the reaching task using a Monte-Carlo approximation (mean: black solid line, $2 \times \text{STD}$ in gray), our moment-matching approximation (orange, dashed), and a baseline where the signal-dependent noise is replaced by additive noise (blue, dotted). The gray and orange areas are hardly distinguishable visually, showing that the moment matching approximation estimates the trajectory distribution very precisely. The baseline overestimates the variance through the signal-dependent noise in early time steps, leading to an overall too high variance of the trajectories.

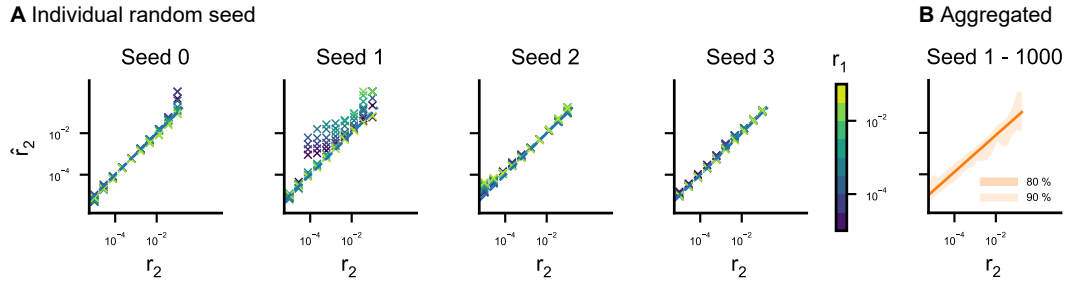


Figure 11: **Random problems (second parameter).** **A** MLEs for a range of parameter values of r_2 for different random problems and different values of r_1 . **B** Aggregated results (median, percentiles) for 1000 random problems.

estimated, e.g., via Monte-Carlo rollouts. The difference in symmetrized KL between the empirical estimate and the baseline is 6.05. A plot of the resulting distributions is shown in Fig. 10. As expected, the moment matching approximation estimates the trajectory distribution very precisely in comparison to the baseline.

G.2 Random problems

In Fig. 11 A we show the errors for a range of parameter values r_2 for different random models. The median and quantiles for the results of 1000 random problems are shown in Fig. 11 B. One can observe that the estimates are generally very close to the true parameters.