Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

Active Inference for Control and Torque Disturbance Estimation: A Quantitative Single-Degree-of-Freedom Experimental Validation

Hyungseok Ryu
Gwangju Institute of Science and Technology
Pilwon Hur
pilwonhur@gist.ac.kr

Gwangju Institute of Science and Technology

Article

Keywords:

Posted Date: November 13th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-11503630/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: No competing interests reported.

Active Inference for Control and Torque Disturbance Estimation: A Quantitative Single-Degree-of-Freedom Experimental Validation

Hyungseok Ryu[1] **and Pilwon Hur**[1,*]

[1]Gwangju Institute of Science and Technology, Department of Mechanical and Robotics Engineering, Gwangju, 61005, Republic of Korea
[*]pilwonhur@gist.ac.kr

ABSTRACT

Residual torque components not captured by nominal dynamics degrade tracking and complicate disturbance compensation in torque-controlled robots. Active inference (AIF) updates state estimates and control actions jointly through free-energy minimization, but how faithfully its internal quantities represent residual torque disturbances has not been quantified directly. We experimentally evaluated an AIF-based controller for simultaneous position tracking and disturbance estimation on a single-degree-of-freedom torque-controlled testbed, where an unknown payload induced repeatable model mismatch, and compared it with a baseline controller and a nonlinear disturbance observer (NDOB)-based controller under matched sinusoidal- and step-tracking conditions. Estimation fidelity was evaluated against a residual torque reference constructed from torque sensing and nominal-model subtraction. Under nominal conditions, the NDOB-based controller tracked the sinusoid most accurately. Under payload-disturbed conditions, the AIF-based controller showed the smallest position root mean squared error (RMSE) in sinusoidal tracking and the smallest steady-state and integral absolute errors in step tracking, and its disturbance estimate had lower RMSE than that of the NDOB-based controller, by approximately 25% and 28% under sinusoidal and step inputs, respectively. These results suggest that AIF can improve tracking robustness while yielding an internal quantity consistent with the residual torque reference under controlled model mismatch.

Introduction

In wearable and physically interactive robots, joint-level torque demand is shaped not only by the robot dynamics but also by interaction-induced effects arising from the user and the environment. Because such effects are not always fully captured by a nominal model, they can appear as residual torque disturbances that influence tracking performance and assistance behavior. Accurate estimation of such residual torque disturbances is therefore important for reliable torque control in systems that operate under physical interaction[1,2].

From a control perspective, this challenge can be viewed more generally as the problem of estimating residual torque disturbances that are not fully captured by nominal dynamics. Such disturbances may arise from model uncertainty, unmodeled loading, frictional effects, or physical interaction[3]. At the joint level, they are commonly represented as residual torque disturbances acting on the system. A standard approach to handling this problem is to augment the tracking controller with a disturbance-estimation mechanism that reconstructs an equivalent disturbance from measured signals and nominal dynamics[4–6].

Among observer-based approaches, nonlinear disturbance observers (NDOBs) are widely used to estimate such residual disturbances in torque-controlled systems[3,7]. NDOB-based methods reconstruct an equivalent disturbance from measured motion and nominal system dynamics, allowing unmodeled effects to be compensated without direct disturbance measurement. In practical implementations, NDOB-based formulations are particularly useful because they can handle nonlinear system behavior without requiring explicit acceleration sensing[8]. However, their estimation performance remains dependent on modeling accuracy, observer gain selection, and filtering design. This dependence means that the balance between fast disturbance response and robustness to noise is often governed by external tuning, and that the fidelity of the estimated disturbance may vary across operating conditions[9].

Active inference (AIF), derived from variational inference and free-energy minimization, provides an alternative perspective in which estimation and control updates are performed simultaneously by minimizing prediction errors under an internal generative model[10,11]. Within this framework, internal state estimates and control actions are not treated as separate processes, but are updated together so that sensory observations remain consistent with model-based expectations. Under model mismatch, unmodeled effects can influence sensory prediction errors and the associated internal updates, thereby giving rise to internal

quantities that may be related to the residual torque acting on the system. This structural difference motivates examining whether the internal quantities arising in AIF can provide a faithful disturbance-related representation under model mismatch.

AIF has already been explored in a range of robotic control problems, including manipulator control, adaptive behavior, and torque-level control[12–14]. In these studies, the primary emphasis has generally been placed on tracking performance, adaptation, or robustness of the overall control behavior. By contrast, the extent to which AIF can faithfully represent residual torque disturbances has received much less direct quantitative evaluation, particularly against a sensor-based disturbance reference under matched experimental conditions[15]. This distinction is important, because satisfactory tracking performance does not by itself establish whether the internal quantities of the controller provide an accurate and interpretable representation of the disturbance acting on the system.

The present study does not attempt to reproduce the full complexity of human-in-the-loop interaction. Instead, it examines disturbance-estimation fidelity in a controlled setting where a repeatable residual torque is introduced through payload mismatch. A single-degree-of-freedom (1-DoF) rotational testbed is adopted to reduce the confounding effects of multi-joint coupling and interaction geometry, thereby allowing the disturbance source and the resulting estimation error to be interpreted more transparently[16]. In this setting, a residual-torque reference is constructed separately from the controller internals using torque sensing together with nominal-model subtraction, enabling a direct assessment of how closely the disturbance-related internal quantity produced by the AIF-based controller agrees with that reference under controlled mismatch conditions. In addition, we state the exact identities that relate the tracking and disturbance-estimation errors of the AIF-based controller to its prediction and inference errors, and we delimit what the present analysis does and does not establish about closed-loop stability.

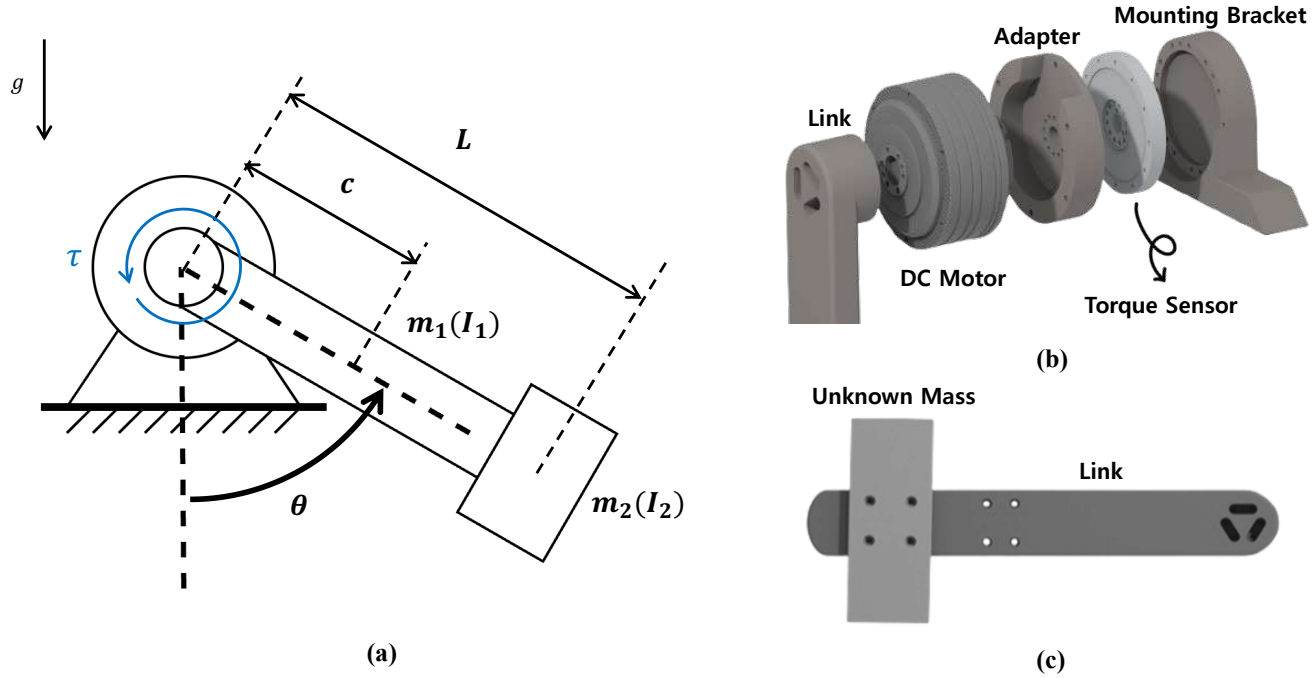


Figure 1. 1-DoF rotational testbed used for validation. (a) Modeling schematic with link and distal payload. (b) Hardware assembly with DC motor and torque sensor. (c) Unknown mass attachment used to induce payload-dependent model mismatch.

Methods

System modeling and disturbance definition

We consider a 1-DoF rotational system driven by a torque-controlled actuator. The system consists of a rigid link and an unknown payload attached at the distal end, as shown in Fig. 1(a). The joint dynamics are described by

$$\tau + \tau_{\text{ext}} = M(\theta)\ddot{\theta} + G(\theta) + F(\dot{\theta}), \tag{1}$$

where τ denotes the applied motor torque, θ is the joint angle, $M(\theta)$ is the inertia term, $G(\theta)$ is the gravitational torque, $F(\dot{\theta})$ represents joint friction, and τ_{ext} denotes an external interaction torque. Because the system has only one rotational joint, Coriolis and centrifugal terms are absent.

Model uncertainty is introduced by attaching an unknown payload to the distal end of the link. Relative to the nominal model used by the controllers, this payload changes the effective inertia and gravity of the system. The corresponding dynamic terms are decomposed as

$$\begin{aligned} M(\theta) &= \hat{M}(\theta) + \Delta M(\theta), \\ G(\theta) &= \hat{G}(\theta) + \Delta G(\theta), \end{aligned} \tag{2}$$

where $\hat{M}(\theta)$ and $\hat{G}(\theta)$ denote the nominal inertia and gravity terms, and $\Delta M(\theta)$ and $\Delta G(\theta)$ represent the mismatch induced by the payload. These mismatch terms are not explicitly identified and are instead treated as unknown functions of the joint configuration. For the single rigid link the inertia terms are constants, and the argument θ is retained only for generality. All controllers use the nominal inertia of the unloaded link, $\hat{M} = m_1 L^2/3 = 0.0408$ kg·m^2 (the uniform-rod value adopted in the implementation), and the nominal gravity torque $\hat{G}(\theta) = m_1 g c \sin\theta$ (Table 1). Here $m_1 = 0.605$ kg is the link mass of the nominal model, $L = 0.45$ m the link length, $c = 0.184$ m the distance from the joint to the center of mass of the link, g the gravitational acceleration, and $\theta = 0$ the downward-hanging stable equilibrium; the payload does not enter the nominal model.

From the nominal-model perspective, the residual torque not captured by the nominal dynamics is defined as a lumped disturbance torque,

$$\tau_d = \Delta M(\theta)\ddot{\theta} + \Delta G(\theta) - \tau_{\text{ext}}. \tag{3}$$

Accordingly, the system can be rewritten in nominal form as

$$\tau = \hat{M}(\theta)\ddot{\theta} + \hat{G}(\theta) + F(\dot{\theta}) + \tau_d. \tag{4}$$

Under this representation, τ_d collects the effects that are not explained by the nominal model. This definition aggregates the effects of inertial and gravitational mismatch together with any unmodeled external interaction torque into a single disturbance term acting relative to the nominal model. Joint friction is not included in τ_d and is instead modeled and compensated separately so that the disturbance definition remains consistent with the nominal-form representation used throughout the controller design and evaluation.

No external torque is applied in the experiments. Accordingly, τ_{ext} is assumed negligible in practice, i.e.,

$$\tau_{\text{ext}} \approx 0, \tag{5}$$

and is retained in the above expressions only for notational completeness. As a result, the dominant source of disturbance in the present experiments is the residual torque induced by payload mismatch relative to the nominal model.

Friction characteristics were identified experimentally under unloaded, constant-velocity conditions so that actuator-side friction could be isolated as much as possible from payload-induced effects. The measured friction profile showed a pronounced increase near zero velocity, followed by a lower friction level at larger velocities together with a viscous component. To describe this behavior, a Stribeck-type friction model was adopted[17,18],

$$F(\dot{\theta}) = \left[\tau_c + (\tau_s - \tau_c)\exp\left(-\left(\frac{|\dot{\theta}|}{v_s}\right)^2\right) + k_v|\dot{\theta}|\right]\operatorname{sgn}(\dot{\theta}), \tag{6}$$

where τ_s and τ_c denote the static and Coulomb friction torques, v_s is the Stribeck velocity parameter, k_v is the viscous friction coefficient, and $\operatorname{sgn}(\cdot)$ denotes the sign function. The model parameters were obtained by nonlinear least-squares fitting to the measured friction data. Fits for the two directions of motion gave $(\tau_c, \tau_s, v_s, k_v) = (0.456, 0.534, 0.233, 0.020)$ and $(0.454, 0.523, 0.260, 0.021)$ (N·m, N·m, rad/s, N·m·s/rad); all controllers use one symmetric parameter set, $\tau_c = 0.462$ N·m, $\tau_s = 0.537$ N·m, $v_s = 0.216$ rad/s and $k_v = 0.0125$ N·m·s/rad. In implementation, a deadband $|\dot{\theta}| \le 0.04$ rad/s was introduced around zero velocity to avoid the discontinuity of the sign term; inside it the friction model returns zero, except in position intervals near the two ends of the motion range, where it returns a constant torque of ∓ 0.802 N·m. This friction model, with the parameters above, is the one used in the friction compensation of the baseline and NDOB-based controllers and in the generative model and prior of the AIF-based controller; the same parameters enter the disturbance reference (27). Where the AIF-based controller needs the friction slope (in the gradients of the free energy with respect to μ_x'), it uses the derivative of the moving Stribeck branch,

$$F'(\dot{\theta}) = k_v - \frac{2(\tau_s - \tau_c)|\dot{\theta}|}{v_s^2}\exp\left(-\left(\frac{|\dot{\theta}|}{v_s}\right)^2\right), \tag{7}$$

which is negative between the deadband edge and $|\dot{\theta}| \approx 0.47$ rad/s (minimum -0.29 N·m·s/rad at 0.15 rad/s).

Baseline and NDOB-based control

For comparison, we implemented a nominal baseline controller and an NDOB-based controller. Both controllers were constructed using the same nominal rigid-body model, the same friction compensation, and the same tracking gains so that the effect of disturbance estimation could be evaluated under matched control conditions.

Baseline controller The baseline controller was constructed using nominal feedback linearization with gravity and friction compensation. The nominal control torque is given by

$$\tau_{\text{base}} = \hat{M}\left[\ddot{\theta}_d + K_p e_\theta + K_d \dot{e}_\theta\right] + \hat{G}(\theta) + F(\dot{\theta}), \tag{8}$$

where $\theta_d(t)$ is the reference position,

$$e_\theta = \theta_d - \theta, \qquad \dot{e}_\theta = \dot{\theta}_d - \dot{\theta}, \tag{9}$$

and K_p and K_d are acceleration-level gains (units s^{-2} and s^{-1}), so that the nominal closed-loop error dynamics are $\ddot{e}_\theta + K_d \dot{e}_\theta + K_p e_\theta = 0$ and the effective torque stiffness and damping are $\hat{M}K_p = 4.49$ N·m/rad and $\hat{M}K_d = 0.82$ N·m·s/rad. This controller does not include any explicit disturbance estimation and serves as the nominal reference for evaluating the effect of disturbance compensation.

NDOB-based controller To compensate for model uncertainty and lumped disturbances, an NDOB-based controller was implemented based on a 1-DoF realization of the acceleration-free NDOB of Mohammadi et al.[8]. The observer reconstructs the lumped disturbance torque without requiring joint acceleration measurements. In the convention of that work the plant is written as $\hat{M}\ddot{\theta} + \hat{N}(\theta, \dot{\theta}) = \tau + d$ with

$$\hat{N}(\theta, \dot{\theta}) = \hat{G}(\theta) + F(\dot{\theta}), \tag{10}$$

so that, by (4), the observed disturbance is $d = -\tau_d$ with τ_d defined in (3). For the present 1-DoF system with constant nominal inertia, the implemented observer, with auxiliary state z and gain ℓ, is

$$\begin{aligned} \dot{z} &= -\ell z + \ell\big(\hat{N}(\theta, \dot{\theta}) - \tau - p\big), \qquad p = X^{-1}\dot{\theta}, \\ \hat{d} &= z + p, \qquad \ell = X^{-1}/\hat{M}, \end{aligned} \tag{11}$$

integrated by the forward-Euler rule with the control period $\Delta t = 2$ ms, $z_{k+1} = z_k + \Delta t\, \dot{z}_k$, where $X > 0$ is the observer design parameter and τ is the applied torque command, including disturbance and friction compensation.

Following the design procedure of Mohammadi et al.[8], the observer parameter has to satisfy $X^{-1} \geq \zeta/2 + \beta\hat{M}$, where $\zeta = \sup_t |\dot{\hat{M}}(t)|$ and β is the required convergence rate of the observer error; because the nominal inertia is constant, $\zeta = 0$. The value $X^{-1} = 2.5$ kg·m^2/s ($\ell = 61.2\ \mathrm{s}^{-1}$), used in all experiments, satisfies the inequality for every $\beta \leq 61\ \mathrm{s}^{-1}$; for $\beta = 50\ \mathrm{s}^{-1}$, which corresponds to a 2% settling time of about 78 ms for an ideal first-order observer response, the margin is 22%. The observer inertia $\hat{M}$ is not adapted to the payload.

The NDOB-based control input is then defined as

$$\tau_{\text{NDOB}} = \tau_{\text{base}} - \hat{d} = \tau_{\text{base}} + \hat{\tau}_d, \tag{12}$$

i.e., the estimate $\hat{\tau}_d = -\hat{d}$ of the lumped disturbance τ_d is added to the baseline torque, and the same estimate (with this sign) is the quantity compared with the disturbance reference in the Results. Differentiating (11) along the plant gives $\dot{\hat{d}} = \ell\,(d - \hat{d})$. Under this formulation, the NDOB-based controller extends the baseline control law by compensating the estimated lumped disturbance torque. The observer parameter was kept fixed throughout all experiments, and the NDOB-based controller was evaluated under the same nominal model, friction compensation, sampling and tracking-gain settings used for the other methods.

AIF-based control

The AIF-based controller was formulated so that state inference and control action were updated simultaneously under an internal generative model[11,19]. In this framework, the hidden state μ_x represents the internal belief over joint position, while the causal state μ_v represents an internally inferred torque-like quantity driving the expected system dynamics.

Generative models The sensory generative model assumes that the measured joint position is generated directly from the hidden state,

$$g(\mu_x, \mu_v) = \mu_x. \tag{13}$$

The state-dynamics generative model describes the expected joint acceleration under the internal model as

$$f(\mu_x, \mu_x', \mu_v) = \frac{-\hat{G}(\mu_x) - F(\mu_x') + \mu_v}{\hat{M}(\mu_x)}. \tag{14}$$

Here, μ_x' denotes the internally inferred joint velocity. The prior generative model encodes the desired trajectory through a feedback-linearization-like structure,

$$\begin{aligned}\eta(\mu_x, \mu_x', r) = \hat{M}(\mu_x)\left[\ddot{\theta}_d + K_d(\dot{\theta}_d - \mu_x') + K_p(\theta_d - \mu_x)\right] \\ + \hat{G}(\mu_x) + F(\mu_x'),\end{aligned} \tag{15}$$

where $r = (\theta_d, \dot{\theta}_d, \ddot{\theta}_d)$ denotes the reference trajectory. This prior specifies the torque-like quantity that would be required, under the nominal model, to drive the inferred state toward the desired trajectory. In this sense, $\eta(\mu_x, \mu_x', r)$ plays the role of a model-based prior expectation over the causal state μ_v.

Prediction errors and free energy The sensory data and the hidden state are represented in generalized coordinates of motion up to second order, $\tilde{s} = (s, s', s'')$ and $\tilde{\mu}_x = (\mu_x, \mu_x', \mu_x'')$, where s and s' are the raw encoder position and velocity. No acceleration is measured: the entry s'' and its precision are both set to zero, so it does not contribute to the free energy or to its gradients. The second-order embedding is retained because the dynamics prediction error $\mu_x'' - f(\mu_x, \mu_x', \mu_v)$ incorporates the rigid-body dynamics through which inertia mismatch and disturbance enter the model. The causal state is likewise represented in generalized coordinates, $\tilde{\mu}_v = (\mu_v, \mu_v', \mu_v'')$. Prediction errors are defined at the sensory, state-dynamics, and prior levels; all three are 3-vectors, $\tilde{\boldsymbol{\varepsilon}}_s$, $\tilde{\boldsymbol{\varepsilon}}_x$ and $\tilde{\boldsymbol{\varepsilon}}_v$, with

$$\begin{aligned}\tilde{\boldsymbol{\varepsilon}}_s &= \tilde{s} - \tilde{\mu}_x = \left(s - \mu_x,\ s' - \mu_x',\ s'' - \mu_x''\right), \\ \tilde{\boldsymbol{\varepsilon}}_x &= \left(0,\ \mu_x'' - f(\mu_x, \mu_x', \mu_v),\ \frac{\hat{G}'(\mu_x)\mu_x' + F'(\mu_x')\mu_x'' - \mu_v'}{\hat{M}}\right), \\ \tilde{\boldsymbol{\varepsilon}}_v &= \left(\eta(\mu_x, \mu_x', r) - \mu_v,\ -\mu_v',\ -\mu_v''\right),\end{aligned} \tag{16}$$

where ε_x denotes the second entry of $\tilde{\boldsymbol{\varepsilon}}_x$ and ε_v the first entry of $\tilde{\boldsymbol{\varepsilon}}_v$, as used in the identities (24) and (25) below. The term F' is the implemented friction slope (7). Using these prediction errors, the variational free energy is defined as

$$\mathscr{F} = \frac{1}{2}\left(\tilde{\boldsymbol{\varepsilon}}_s^\top \Pi_s \tilde{\boldsymbol{\varepsilon}}_s + \tilde{\boldsymbol{\varepsilon}}_x^\top \Pi_x \tilde{\boldsymbol{\varepsilon}}_x + \tilde{\boldsymbol{\varepsilon}}_v^\top \Pi_v \tilde{\boldsymbol{\varepsilon}}_v\right), \tag{17}$$

where $\Pi_s = \mathrm{diag}(\pi_{s,0}, \pi_{s,1}, \pi_{s,2}) = \mathrm{diag}(e^7, e^6, 0)$ weights the position, velocity, and (absent) acceleration sensory errors $\varepsilon_{s,0}$, $\varepsilon_{s,1}$ and $\varepsilon_{s,2}$ (the entries of $\tilde{\boldsymbol{\varepsilon}}_s$ in (16)), and $\Pi_x = I_3$ and $\Pi_v = \mathrm{diag}(e^2, 1, 1)$ weight the state-dynamics and prior prediction errors (Table 1).

State and action updates The hidden and causal states follow generalized filtering[20], i.e., gradient descent on $\mathscr{F}$ combined with the shift operator D that advances each generalized coordinate by the next one, $D(\mu, \mu', \mu'') = (\mu', \mu'', 0)$,

$$\dot{\tilde{\mu}}_x = D\tilde{\mu}_x - \kappa_x \frac{\partial \mathscr{F}}{\partial \tilde{\mu}_x}, \qquad \dot{\tilde{\mu}}_v = D\tilde{\mu}_v - \kappa_v \frac{\partial \mathscr{F}}{\partial \tilde{\mu}_v}, \tag{18}$$

so that, componentwise, $\dot{\mu}_x = \mu_x' - \kappa_x \partial \mathscr{F}/\partial \mu_x$, $\dot{\mu}_x' = \mu_x'' - \kappa_x \partial \mathscr{F}/\partial \mu_x'$ and $\dot{\mu}_x'' = -\kappa_x \partial \mathscr{F}/\partial \mu_x''$, where κ_x and κ_v are the update rates for the hidden and causal states. The control action a is defined as the applied motor torque command, i.e., $\tau = a$, and is updated along the sensory gradient through prescribed action sensitivities $\partial_a s = \partial_a s' = 1$, $\partial_a s'' = 0$, which are dimensionless design constants of the controller rather than mechanical derivatives:

$$\dot{a} = -\kappa_a \sum_{i=0}^{2} \pi_{s,i}\, \varepsilon_{s,i}\, \partial_a s^{(i)} = -\kappa_a\left(\pi_{s,0}\varepsilon_{s,0} + \pi_{s,1}\varepsilon_{s,1}\right), \tag{19}$$

with κ_a denoting the action update rate. Because $\mathscr{F}$ has no explicit action argument, (19) is the definition of the action law, in which the free-energy gradient is propagated to the action through the assumed sensitivities of the sensory channels. In this way, state inference and control are coupled within a single variational optimization process rather than being separated into distinct estimation and control modules. The rates κ_x, κ_v and κ_a are continuous-time rates, and all updates are integrated with the forward-Euler step $\Delta t = 2$ ms of the 500 Hz loop, e.g.

$$a_{k+1} = \mathrm{sat}_{10}\left(a_k - \Delta t\, \kappa_a \frac{\partial \mathscr{F}}{\partial a}\right), \tag{20}$$

where sat_{10} clamps the action to ± 10 N·m directly after each update (the motor command is limited to the same range; no back-calculation anti-windup is implemented). In each 2 ms cycle the controller reads the encoder, updates the beliefs and the action, and sends the new action to the motor in the same cycle, so the controller software introduces no one-sample delay between measurement and command; the transport latency between the controller node, the motor node and the motor driver (see Experimental setup) is not accounted for here.

In the implementation, the precision matrices were chosen to balance tracking responsiveness against sensitivity to measurement noise while maintaining stable closed-loop behavior. The update rates were selected to ensure stable convergence and robust real-time tracking, with final values refined through preliminary trials.

Disturbance estimation in AIF Although the AIF-based controller does not explicitly include a disturbance observer, a disturbance-related residual torque estimate can be defined from the discrepancy between the applied control action a and the inferred causal state μ_v:

$$\hat{\tau}_d^{\mathrm{AIF}} = a - \mu_v. \tag{21}$$

This interpretation follows from the nominal-form dynamics, where the residual disturbance torque is given by

$$\tau_d = \tau - \left(\hat{M}(\theta)\ddot{\theta} + \hat{G}(\theta) + F(\dot{\theta})\right), \tag{22}$$

while rearranging the state-dynamics generative model with ε_x from (16) yields

$$\mu_v = \hat{M}(\mu_x)\left[\mu_x'' - \varepsilon_x\right] + \hat{G}(\mu_x) + F(\mu_x'). \tag{23}$$

Subtracting the two expressions gives the exact estimation-error identity

$$\hat{\tau}_d^{\mathrm{AIF}} - \tau_d = \hat{M}\left[\varepsilon_x - (\mu_x'' - \ddot{\theta})\right] - \left[\hat{G}(\mu_x) + F(\mu_x') - \hat{G}(\theta) - F(\dot{\theta})\right], \tag{24}$$

which holds along every trajectory because it uses only the definitions in (16) and the nominal form (4) with $\tau = a$; it is used in the stability analysis below. With the precisions of Table 1 the sensory channels keep μ_x and μ_x' close to the measured position and velocity, so the bracketed model-mismatch term is small except briefly at the velocity reversals, where μ_x' and $\dot{\theta}$ can lie on different branches of the discontinuous friction law (6) and the bracket can reach the branch jump (0.54 N·m at the deadband edge, or 1.34 N·m in the end intervals). The inferred acceleration μ_x'', by contrast, is not driven by any measurement: in a logged payload-disturbed sinusoidal trial it stayed between -0.003 and 0.071 rad/s^2, whereas the measured acceleration, obtained offline by differentiating the 6 Hz low-pass filtered velocity (see Experimental setup), ranged from -2.15 to 1.48 rad/s^2. Consequently $a - \mu_v$ estimates $\tau_d + \hat{M}\ddot{\theta}$ up to the term $\hat{M}\varepsilon_x$: it includes the nominal inertial torque $\hat{M}\ddot{\theta}$ during motion (at most 0.09 N·m in the logged trial), a term that vanishes at rest. In the same trial the dynamics residual reached $\varepsilon_x = \pm 32.7$ rad/s^2 ($\hat{M}\varepsilon_x \approx \pm 1.3$ N·m), which equals, to within 0.1%, the jump of the friction term of the generative model between the Stribeck branch at the deadband edge (0.54 N·m) and the ∓ 0.802 N·m branch (1.34 N·m in total), divided by $\hat{M}$; we therefore attribute these extremes to the branch switching at the velocity reversals, where the short excursions of the estimate visible in Fig. 5 occur. The raw quantity $a - \mu_v$ is the estimator evaluated in the Results.

Stability analysis

We state what follows exactly from the generative model as printed and from the nominal form of the plant, and we delimit what is not established. Throughout, e_θ and $\dot{e}_\theta$ are the tracking errors (9), and $\delta_0 = \mu_x - \theta$, $\delta_1 = \mu_x' - \dot{\theta}$ and $\delta_2 = \mu_x'' - \ddot{\theta}$ are the inference errors.

Exact identities and conditional tracking bound Eliminating μ_v between ε_x and ε_v in (16) gives, without approximation and without using the plant equation,

$$\ddot{e}_\theta + K_d\dot{e}_\theta + K_p e_\theta = w, \qquad w = \delta_2 + K_d\delta_1 + K_p\delta_0 - \varepsilon_x + \varepsilon_v/\hat{M}; \tag{25}$$

gravity and friction cancel because the same $\hat{G}$ and F enter f and η. Together with the estimation-error identity (24), this relates the tracking and estimation errors of the controller exactly to its prediction and inference errors. If $\limsup_{t\to\infty}|w(t)| \le W$, then

$$\limsup_{t\to\infty}|e_\theta| \le 0.00910\,W, \qquad \limsup_{t\to\infty}|\dot{e}_\theta| \le 0.0724\,W, \tag{26}$$

where the constants (units s^2 and s for $K_p = 110$ s^{-2} and $K_d = 20$ s^{-1}) are the L_1 norms of the impulse response $h(t) = e^{-10t}\sin(\sqrt{10}\,t)/\sqrt{10}$ of (25) and of $\dot{h}$, $\|h\|_1 = \coth(\pi\sqrt{10}/2)/K_p = 0.009092$ and $\|\dot{h}\|_1 = 0.07240$, rounded up to three significant figures; both bounds are attained by a square-wave residual at the damped frequency $\sqrt{10}$ rad/s with a suitably chosen phase. Equation (25) does not establish boundedness of the endogenous residual w, whose size is a property of the coupled plant–inference–action loop.

Scope The two identities hold for the implemented controller because they use only the definitions (16), the generative model as printed and, for (24), the nominal form (4) with $\tau = a$, where a is the action after the clamp in (20). The conditional bound (26) is a property of the prior dynamics alone. A stability certificate for the coupled plant–inference–action loop is not established here: it would have to account for the generalized-coordinate residuals and shift terms, the Stribeck friction with its deadband, the torque limit and the 500 Hz sampled updates. The free energy is not a Lyapunov function of that loop, because the sensory data are time-varying and the shift terms in (18) are not in general descent directions of $\mathscr{F}$. The closed-loop behavior reported in the Results is therefore experimental evidence for the tested conditions and parameters, not a stability guarantee.

Disturbance reference construction

To quantitatively evaluate disturbance-estimation fidelity, a sensor-based residual-torque reference was constructed from torque sensing and nominal-model subtraction. The reference was defined as the residual torque that remained after subtracting the nominal rigid-body contribution, including gravity and modeled friction, from the measured joint torque. Under this definition, the reference disturbance is given by

$$\tau_d^{\mathrm{ref}} = \tau_{\mathrm{meas}} - \left(\hat{M}(\theta)\ddot{\theta} + \hat{G}(\theta) + F(\dot{\theta})\right), \tag{27}$$

where τ_{meas} denotes the torque measured by the torque sensor mounted between the actuator and a fixed bracket (see Experimental setup), a measured signal distinct from the commanded torque τ; (27) applies to it the same nominal-model subtraction by which τ_d is defined in (4).

The nominal-model term in (27) was computed using the same nominal inertia, gravity, and friction model used by the controllers (Table 1). Joint position and velocity were obtained from the encoder, and joint acceleration was estimated by differentiating the zero-phase filtered velocity signal. Zero-phase filtering was used in the offline construction of the disturbance reference so that phase lag introduced by low-pass filtering did not bias the reference quantity. In this way, the disturbance reference represents the portion of the measured torque that cannot be explained by the nominal model under the same sensing and modeling assumptions used throughout the control design.

Because no external contact was applied in the experiments, the dominant component of τ_d^{ref} corresponds to the residual torque induced by payload-related mismatch relative to the nominal model. The resulting quantity is therefore consistent with the lumped disturbance definition in (3), while remaining experimentally accessible through torque sensing and nominal-model subtraction. In this study, τ_d^{ref} was used as the reference quantity for quantitatively comparing the disturbance estimates produced by the NDOB-based and AIF-based controllers.

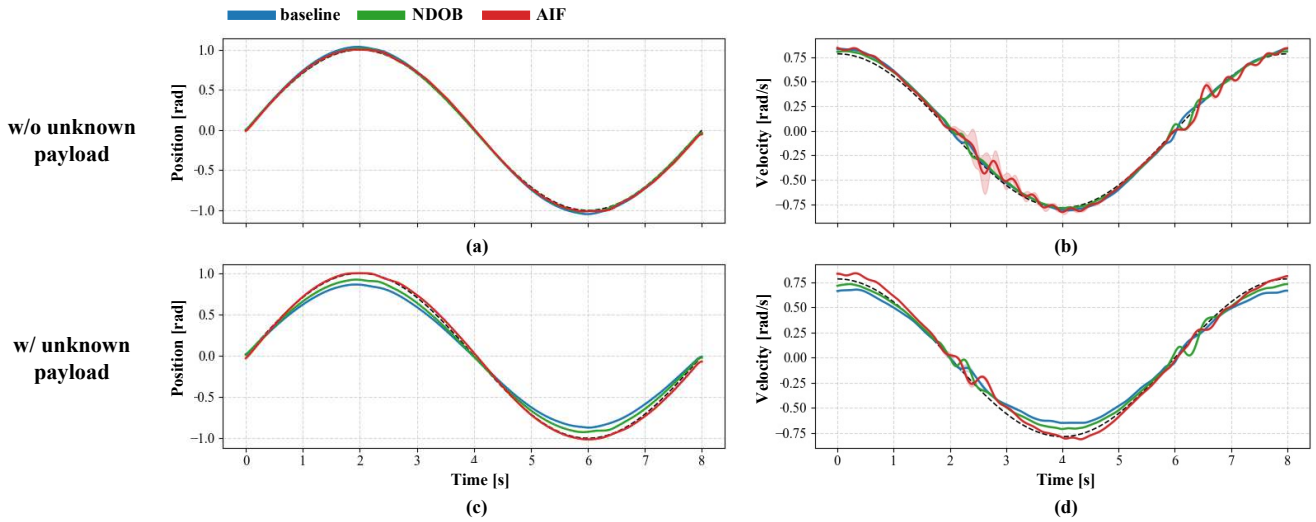


Figure 2. Sinusoidal trajectory tracking results under nominal (w/o unknown payload) (a, b) and payload-disturbed (w/ unknown payload) (c, d) conditions. Position (a, c) and velocity (b, d) tracking are compared across the baseline, NDOB-based, and AIF-based controllers. Traces show trial means; shaded bands denote the standard deviation across trials. Under the payload-disturbed condition, the AIF-based controller shows closer agreement with the sinusoidal reference waveform than the baseline and NDOB-based controllers.

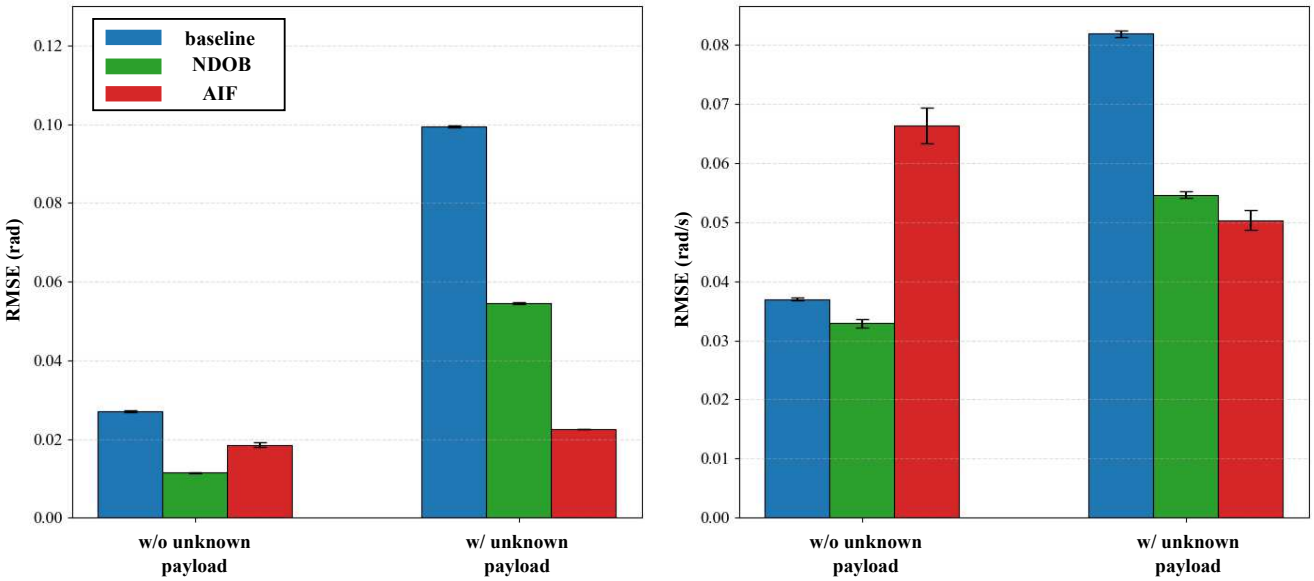


Figure 3. RMSE of position and velocity tracking under nominal (w/o unknown payload) and payload-disturbed (w/ unknown payload) conditions. Bars indicate mean RMSE over repeated trials, with error bars denoting standard deviation. The AIF-based controller achieves the lowest position RMSE under the payload-disturbed condition, while the NDOB-based controller improves velocity tracking relative to the baseline.

Table 1. Parameters of the testbed, the controllers and the AIF inference, used consistently across all experiments. m_1, c and L are the link mass of the nominal model, the distance from the joint to the center of mass of the link and the link length; $\hat{M} = m_1 L^2/3$ is the nominal inertia used by all controllers and by the disturbance reference; m_2 and I_{m_2} are the payload mass and its inertia about the joint; K_p and K_d are the tracking gains; Π_s, Π_x and Π_v are the precisions; $\partial_a s^{(i)}$ are the action sensitivities; κ_x, κ_v and κ_a are the update rates in the sinusoidal (SIN) and step (STEP) experiments; τ_c, τ_s, v_s and k_v are the friction parameters; X^{-1} is the NDOB design parameter; Δt is the control period.

Parameter	Value	Parameter	Value
m_1	0.605 kg	K_p	110 s^{-2}
$\hat{M}$	0.0408 kg·m^2	K_d	20 s^{-1}
c	0.184 m	Π_x	I_3
m_2	0.342 kg	Π_s	diag(e^7, e^6, 0)
I_{m_2}	0.069 kg·m^2	Π_v	diag(e^2, 1, 1)
L	0.45 m	$\partial_a s$, $\partial_a s'$, $\partial_a s''$	1.0, 1.0, 0.0
$\kappa_x, \kappa_v, \kappa_a$ (SIN)	0.2, 0.4, 0.4	$\kappa_x, \kappa_v, \kappa_a$ (STEP)	0.25, 0.4, 0.4
τ_c, τ_s	0.462, 0.537 N·m	v_s, k_v	0.216 rad/s, 0.0125 N·m·s/rad
X^{-1} (NDOB)	2.5 kg·m^2/s	Δt; torque limit	2 ms; $\pm$10 N·m

Experimental setup and evaluation metrics

The experimental platform consisted of a torque-controlled quasi-direct-drive brushless DC actuator (AK10-9, CubeMars), a torque sensor (ATSB50-D84, Aidin Robotics) mounted between the actuator and a fixed bracket, and a custom mechanical linkage, as shown in Fig. 1(b) and (c). Joint torque was measured directly through the torque sensor, while joint position and velocity were obtained from the motor-integrated encoder; the applied torque τ used in the analysis is the torque command sent to the motor driver. These signals were used for control and for constructing the disturbance reference.

The control loop and data acquisition were executed at 500 Hz. The controllers were implemented as Robot Operating System 2 (ROS 2) nodes on an embedded computer (Jetson Orin NX, NVIDIA; Ubuntu 22.04, ROS 2 Humble): in each 2 ms timer callback the controller computed the torque command from the current raw encoder position and velocity and published it, and the motor node forwarded it to the motor driver on receipt. None of the three controllers filtered its in-loop signals. For analysis and post-processing, measured signals were low-pass filtered using a fourth-order Butterworth filter with a cutoff frequency of 6 Hz. For offline disturbance-reference construction and step-response timing metrics, zero-phase filtering was

used to avoid phase lag, and joint acceleration was obtained by differentiating the filtered velocity signal.

Two reference trajectories were used: a sinusoidal trajectory (amplitude: 1 rad, period: 8 s) and a unit step trajectory, realized as a square reference alternating between $\theta_d = 0$ and 1 rad with $\dot{\theta}_d = \ddot{\theta}_d = 0$ set directly (the position jumps were not differentiated). Each trajectory was tested under two conditions: a nominal condition without additional load and a disturbed condition with an unknown payload attached to the distal end of the link. For each controller and condition, 10 repeated trials were conducted. In the sinusoidal experiments, the initial cycle was excluded from analysis to reduce startup transients.

Table 1 summarizes the physical, control, and inference parameters used in the experiments. Nominal model parameters were obtained from direct measurements and estimates from a computer-aided design (CAD) model. Tracking gains were tuned under nominal conditions, and the AIF precision parameters and update rates were selected to maintain stable real-time behavior while balancing responsiveness and sensitivity to measurement noise.

Performance was evaluated in terms of tracking accuracy and disturbance-estimation fidelity. For sinusoidal tracking, the root mean squared error (RMSE) of the joint position error was used. For step tracking, rise time, settling time, steady-state error, and integral absolute error (IAE) were computed. Disturbance-estimation fidelity was quantified as the RMSE between each controller's disturbance estimate and the constructed disturbance reference. Unless otherwise stated, metrics were computed for each trial and then summarized across repeated trials by their mean and standard deviation.

Code availability

The controller code (the baseline, NDOB-based and AIF-based controllers implemented as ROS 2 nodes) is available from the corresponding author on reasonable request.

Use of large language models

Large language models were used to assist with the analytical derivations, numerical checks and text editing during the revision of this manuscript; the authors verified all derivations, numbers and statements and take full responsibility for the content.

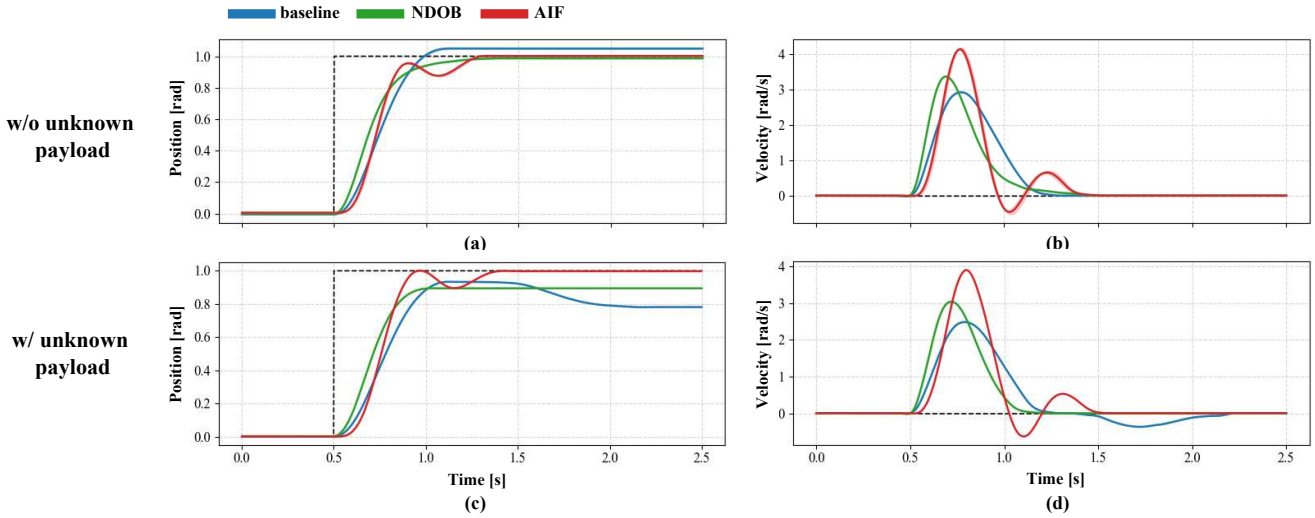


Figure 4. Step position tracking responses under nominal (w/o unknown payload) (a, b) and payload-disturbed (w/ unknown payload) (c, d) conditions. Joint position (a, c) and velocity (b, d) are shown for the baseline, NDOB-based, and AIF-based controllers. Traces show trial means; shaded bands denote the standard deviation across trials. Under the nominal condition the NDOB-based and AIF-based controllers settle close to the commanded position, whereas the baseline controller retains a visible residual offset (Table 2); noticeable steady-state offset and velocity deviations are observed for the baseline and NDOB-based controllers under the payload-disturbed condition.

Results

Quantitative accuracy of position tracking

Steady-state tracking under sinusoidal reference

Steady-state tracking under a sinusoidal position reference is depicted in Fig. 2 for nominal and payload-disturbed conditions. In the nominal condition, the position trajectories closely follow the sinusoidal reference for all controllers, and the corresponding

Table 2. Step-response metrics for position tracking under nominal (w/o unknown payload) and payload-disturbed (w/ unknown payload) conditions. Metrics are computed from the position error $e_\theta(t) = \theta_d(t) - \theta(t)$ over a fixed evaluation window following the step command.

Metric	w/o unknown payload			w/ unknown payload		
	Baseline	NDOB	AIF	Baseline	NDOB	AIF
Rise time t_r [s]	0.3640	0.4820	**0.2720**	N/A†	N/A†	**0.2820**
Settling time t_s [s]	N/A†	0.7520	**0.7500**	N/A†	N/A†	**0.8280**
Steady-state error e_{ss} [rad]	0.0719	0.0091	**0.0035**	0.2176	0.1055	**0.0023**
IAE [rad·s]	0.3528	**0.2329**	0.2809	0.5108	0.3952	**0.2838**

Definitions. t_r is the time for $\theta(t)$ to rise from 5% to 95% of the step amplitude (first crossings of the position signal after the offline 6 Hz low-pass filtering described in Experimental setup). t_s is the time after which $|e_\theta(t)|$ remains within $\pm 2\%$ of the step amplitude for the remainder of the evaluation window. e_{ss} is the mean absolute position error over the final 0.5 s of the evaluation window, i.e., over $t \in [T-0.5, T]$. IAE is the integral of absolute error over the post-step evaluation window, $\int_0^T |e_\theta(t)|\,dt$, where $t = 0$ denotes the step command time and $T = 2$ s. Bold marks the best value in each condition.
N/A†. Not applicable: within the $T = 2$ s evaluation window the response did not reach 95% of the step amplitude (t_r) or did not remain within the $\pm 2\%$ band (t_s).

velocity responses remain consistent with the reference waveform, apart from a short oscillation in the velocity of the AIF-based controller after each reversal, which is reflected in its larger velocity RMSE under the nominal condition (Fig. 3). With the unknown payload attached, the baseline controller shows visible waveform distortion in position together with a mismatch in velocity, indicating increased sensitivity to payload-induced model mismatch. The NDOB-based controller reduces this distortion relative to the baseline, whereas the AIF-based controller preserves the sinusoidal reference profile most closely among the compared methods in both position and velocity.

The corresponding trial-averaged RMSE results are depicted in Fig. 3. Under the nominal condition, the AIF-based controller achieves a position RMSE of 0.0186 ± 0.0006 rad, which is lower than that of the baseline controller (about 0.027 rad) but higher than that of the NDOB-based controller (about 0.011 rad). Under the payload-disturbed condition, tracking errors increase for all controllers; however, the AIF-based controller exhibits the smallest position RMSE among the compared methods (0.023 ± 0.002 rad, against about 0.100 and 0.055 rad for the baseline and NDOB-based controllers). Relative to the baseline and NDOB-based controllers, the AIF-based controller shows a substantial reduction in tracking error under payload-induced model mismatch. These results show that the AIF-based controller achieved the lowest steady-state tracking error under the disturbed sinusoidal condition.

Transient tracking under step reference

Transient tracking under a step position reference is depicted in Fig. 4 for nominal and payload-disturbed conditions. In the nominal condition, all controllers rise rapidly toward the commanded position, and the velocity traces show a single dominant peak associated with the step transition. Under the disturbed condition, the baseline controller fails to eliminate the steady-state offset, resulting in a persistent deviation from the target position. The NDOB-based controller reduces this offset relative to the baseline but retains a visible residual error in the final position. In contrast, the AIF-based controller converges closest to the commanded position among the compared methods after the transient phase.

Step-response metrics are summarized in Table 2. Under the nominal condition, both the NDOB-based and AIF-based controllers substantially reduce the steady-state error from 0.0719 rad (baseline) to below 0.01 rad, with the AIF-based controller achieving the smallest steady-state error of 0.0035 rad, whereas the NDOB-based controller has the lowest IAE under this condition (0.2329 against 0.2809 rad·s for the AIF-based controller). Under the payload-disturbed condition, the baseline and NDOB-based responses fail to satisfy the settling criterion within the evaluation window, resulting in large residual offsets ($e_{ss} = 0.2176$ rad and 0.1055 rad, respectively). In contrast, the AIF-based controller converges to the smallest steady-state error ($e_{ss} = 0.0023$ rad) and, under this condition, achieves the lowest IAE among the compared methods (0.2838 rad·s). The AIF-based controller also has the shortest rise time in both conditions and is, under the payload-disturbed condition, the only controller that settles within the evaluation window (Table 2). These results show that the AIF-based controller achieved the smallest final-position error under the disturbed step condition.

Quantitative accuracy of torque disturbance estimation

Steady-state disturbance estimation under sinusoidal reference

Steady-state torque disturbance estimation under a sinusoidal reference is depicted in Fig. 5. The dashed curve denotes the constructed residual torque reference corresponding to the lumped disturbance defined in the system dynamics, and the solid

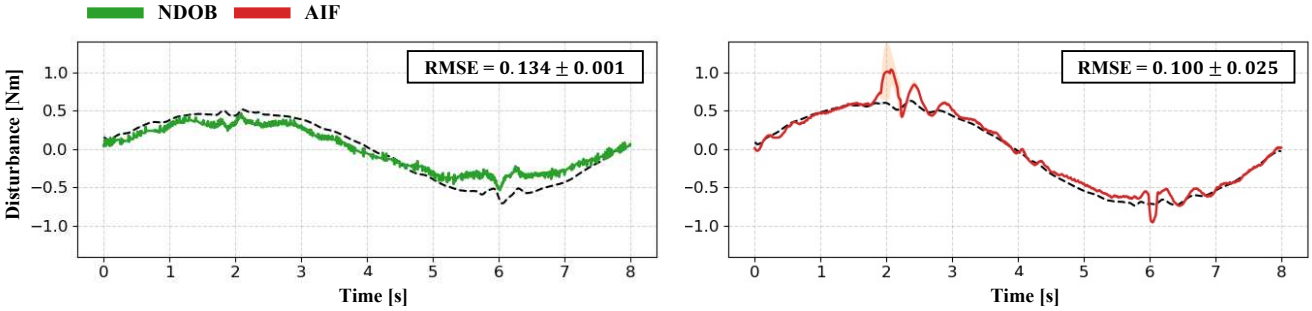


Figure 5. Disturbance estimation results during sinusoidal motion under the payload-disturbed condition. Disturbance estimates from the NDOB-based and AIF-based controllers are compared against the disturbance reference constructed from torque sensing and nominal-model subtraction (dashed line), which is constructed from the trials of the respective controller because it depends on the measured motion; curves show trial means, with shaded bands denoting the across-trial standard deviation; the inset values are the disturbance-estimation RMSE (mean ± standard deviation across trials). Closer agreement with the reference and lower RMSE are observed for the AIF-based estimate over the motion cycle.

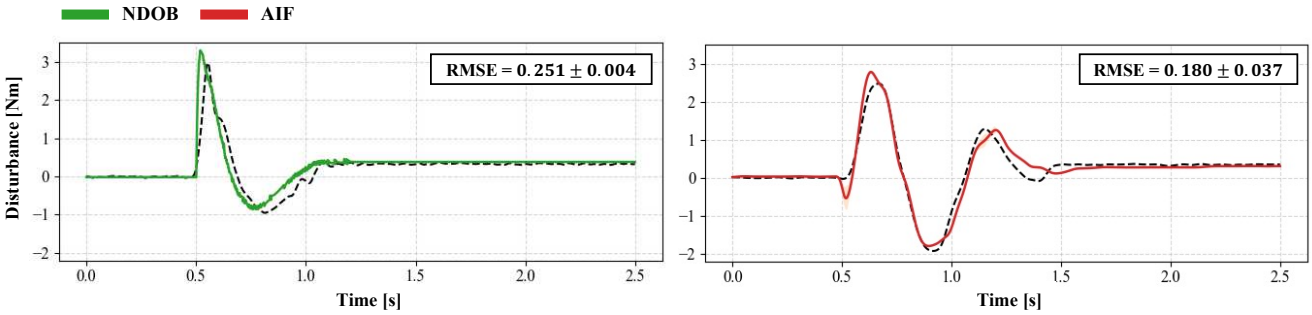


Figure 6. Disturbance estimation results during step motion under the payload-disturbed condition. The dashed black line denotes the disturbance reference constructed from torque sensing and nominal-model subtraction, constructed from the trials of the respective controller because it depends on the measured motion; curves show trial means, with shaded bands denoting the across-trial standard deviation; the inset values are the disturbance-estimation RMSE (mean ± standard deviation across trials). The AIF-based estimate shows smaller deviation from the reference than the NDOB-based estimate during the transient response, as reflected in the RMSE values.

curves show the disturbance estimates produced by the NDOB-based and AIF-based controllers. Both methods capture the overall periodic trend of the disturbance profile; however, the AIF-based estimate shows closer agreement with the residual torque reference over the motion cycle, particularly in the negative phase where the NDOB-based estimate exhibits a larger deviation. The AIF-based estimate departs from the reference briefly at the velocity reversals (near $t = 2$ and 6 s), where the friction term of the generative model switches branch (Methods); these excursions are also the main source of the across-trial spread visible in the shaded band. Quantitatively, the AIF-based controller reduces the disturbance-estimation RMSE from 0.134 ± 0.001 N·m (NDOB) to 0.100 ± 0.025 N·m, i.e., by approximately 25% relative to the NDOB-based controller under the sinusoidal steady-state condition, indicating improved agreement with the constructed residual torque reference.

Transient disturbance estimation under step reference

Transient torque disturbance estimation under a step reference is depicted in Fig. 6. Following the step transition, the residual torque reference shows a pronounced transient excursion before settling to a non-zero steady level, consistent with the payload-induced mismatch under the commanded motion. Both the NDOB-based and AIF-based controllers reproduce the main transient structure, but differences are observed during the peak and recovery phases: the NDOB-based estimate leads its reference by a few tens of milliseconds and overshoots the peak, which produces large deviations during the steep rise and fall of the excursion, whereas the AIF-based estimate follows the peak, the negative excursion and the rebound of its reference more closely, apart from a brief initial dip and a slight lag in the rebound. In terms of quantitative accuracy, the AIF-based controller achieves an approximately 28% reduction in disturbance-estimation RMSE compared with the NDOB-based controller during the step response (0.180 ± 0.037 against 0.251 ± 0.004 N·m), indicating improved disturbance-estimation fidelity under transient

conditions. The AIF-based estimate includes the nominal inertial torque $\hat{M}\ddot{\theta}$ during the transient, by the identity (24).

Discussion

This study examined active inference as a practical framework for simultaneous position tracking and torque-disturbance estimation on a torque-controlled rotary joint under nominal and payload-disturbed conditions. Across the presented experiments, the advantage of the AIF-based controller became most evident under payload-induced model mismatch, where it showed smaller tracking errors than the baseline and NDOB-based controllers, and closer agreement with the residual torque reference than the NDOB-based controller. In the sinusoidal experiments, the AIF-based controller preserved the reference waveform most closely under disturbance (Figs. 2–3), while in the step experiments it achieved the smallest steady-state error in both conditions and, under the payload-disturbed condition, the lowest IAE among the compared methods (Table 2). In the disturbance-estimation results, the AIF-based controller produced lower RMSE than the NDOB-based controller under both sinusoidal and step references (Figs. 5–6), indicating closer agreement with the residual torque reference.

An important feature of the results is that the relative benefit of AIF was limited under nominal conditions but became substantially clearer once payload-induced mismatch was introduced. This behavior is consistent with the structure of the problem itself. When the nominal model remains sufficiently accurate, the need for explicit compensation of residual torque is inherently reduced, and observer-based compensation with appropriate bandwidth can provide comparable tracking performance[3,21]. Under disturbed conditions, however, the nominal model no longer explains the torque demand accurately, and the residual torque becomes a larger component of the closed-loop behavior. In this regime, the AIF-based controller showed the smallest tracking errors among the compared methods (Figs. 2–4), suggesting that its advantage lies not in improving nominal tracking performance itself, but in handling the residual torque demand arising from model mismatch[22]. It should be noted that, although the three controllers share the numerical gains K_p and K_d, these gains define the nominal error dynamics of the baseline and NDOB-based controllers but enter the AIF-based controller only through its prior. The AIF action is the time integral of the precision-weighted sensory prediction errors (19), so it keeps changing as long as these errors persist. The baseline control law contains no integral action. The correspondence between active inference and integral control has been established for linear generative models[23]. The comparison is therefore between controller structures rather than between matched loop gains.

The disturbance-estimation results are consistent with the interpretation of the AIF-derived residual term as a disturbance-related quantity. In the present formulation, the discrepancy between the applied control action and the inferred causal state was used as an estimate of the lumped residual torque. This interpretation is supported by the exact estimation-error identity (24): the discrepancy $a - \mu_v$ differs from the lumped disturbance by the nominal inertial term $\hat{M}(\ddot{\theta} - \mu_x'')$, by $\hat{M}\varepsilon_x$ and by the model-mismatch bracket, and since μ_x'' remained near zero in the logged sinusoidal trial (Methods), it includes the nominal inertial torque $\hat{M}\ddot{\theta}$ during motion, a term that vanishes at rest. From this perspective, the closer agreement of the AIF-based estimate with the residual torque reference (Figs. 5–6) is consistent with this identity: with μ_x'' near zero and ε_x small away from the reversals, $a - \mu_v$ reduces to the applied torque minus the nominal gravity and friction torques evaluated at the beliefs, so the agreement reflects how closely the beliefs track the measurements and how accurately the commanded torque is realized, together with the small inertial term. The difference between the two estimates can be understood in part from the structural distinction between the two formulations. The NDOB-based controller reconstructs the disturbance through an auxiliary observer appended to the nominal control law, whereas the AIF-based controller updates hidden states, causal states, and action within a shared prediction-error objective. This point is important because good tracking alone does not guarantee that the controller-derived residual term corresponds to the disturbance acting on the system.

The significance of the present experiments lies in the use of a controlled payload mismatch setting together with a residual torque reference that could be constructed and evaluated transparently. Although the broader motivation of this work is related to physically interactive and wearable robotic systems, direct validation in human-interactive settings would make disturbance interpretation substantially more difficult because interaction torques vary in complex, time-varying, and task-dependent ways. By introducing an unknown payload instead, the experiments produced a repeatable and state-dependent residual torque while preserving a disturbance-reference construction that remained experimentally accessible through torque sensing and nominal-model subtraction. The present results should therefore be interpreted not as a full validation of human-interaction torque estimation, but as a quantitative validation of disturbance-estimation fidelity in the simplest controlled setting where such comparison can be carried out transparently.

The transient responses also reveal a practically relevant trade-off. Under payload-disturbed step conditions, the AIF-based controller achieved the most accurate final-position regulation (Table 2), but this improvement was accompanied by a more aggressive transient response than that of the other methods (Fig. 4): a larger peak velocity and a small overshoot of the raw position signal (mean 0.018 rad, at most 0.035 rad across the rising payload steps). This behavior is consistent with an action-update mechanism that integrates prediction errors (19) during abrupt transitions, and it is reflected in the shortest rise times of Table 2, suggesting that improved steady-state regulation and improved disturbance-estimation fidelity may come at the

cost of a more aggressive transient response depending on the choice of precision parameters and action sensitivities[12,24]. For this reason, the parameters used in the present study were selected to ensure stable real-time operation and reliable closed-loop behavior across the tested conditions. The comparison was not intended to reflect trajectory-specific over-tuning, but rather the relative behavior of the formulations under a common and practically relevant parameter setting.

Although the controlled 1-DoF setting enabled a transparent comparison between estimated and reference residual torques, it does not capture the full complexity of human-generated interaction torque or multi-joint robot dynamics. In addition, the disturbance reference was obtained from torque sensing and nominal-model subtraction; although this provided an experimentally accessible residual-torque quantity constructed separately from the controller internals, it should not be interpreted as an exact decomposition of all physical disturbance sources[25]. The disturbance considered here is therefore best understood as a lumped residual torque dominated by payload-induced mismatch. The analysis in the Methods establishes exact identities between the tracking and estimation errors and the prediction and inference errors of the controller, together with a bound that is conditional on the residual w of (25), which combines the prediction and inference errors, remaining bounded; a stability certificate for the coupled plant–inference–action loop under the implemented friction law with its deadband, the torque limit and the 500 Hz sampled updates is not established and is left for future work. Even so, the formulation itself is not inherently restricted to single-joint systems, and extending the present evaluation to multi-DoF systems remains an important direction for future work.

References

1. Lenzi, T. *et al.* Measuring human–robot interaction on wearable robots: A distributed approach. *Mechatronics* **21**, 1123–1131, DOI: 10.1016/j.mechatronics.2011.04.003 (2011).
2. Liang, X. *et al.* Adaptive human–robot interaction torque estimation with high accuracy and strong tracking ability for a lower limb rehabilitation robot. *IEEE/ASME Trans. Mechatron.* **29**, 4814–4825, DOI: 10.1109/tmech.2024.3394491 (2024).
3. Chen, W.-H., Yang, J., Guo, L. & Li, S. Disturbance-observer-based control and related methods—an overview. *IEEE Trans. Ind. Electron.* **63**, 1083–1095, DOI: 10.1109/tie.2015.2478397 (2016).
4. Chen, W.-H. Disturbance observer based control for nonlinear systems. *IEEE/ASME Trans. Mechatron.* **9**, 706–710, DOI: 10.1109/tmech.2004.839034 (2004).
5. Solsona, J., Valla, M. I. & Muravchik, C. Nonlinear control of a permanent magnet synchronous motor with disturbance torque estimation. *IEEE Trans. Energy Convers.* **15**, 163–168, DOI: 10.1109/60.866994 (2000).
6. Haddadin, S., De Luca, A. & Albu-Schäffer, A. Robot collisions: A survey on detection, isolation, and identification. *IEEE Trans. Robot.* **33**, 1292–1312, DOI: 10.1109/TRO.2017.2723903 (2017).
7. Chen, W.-H., Ballance, D. J., Gawthrop, P. J. & O'Reilly, J. A nonlinear disturbance observer for robotic manipulators. *IEEE Trans. Ind. Electron.* **47**, 932–938, DOI: 10.1109/41.857974 (2000).
8. Mohammadi, A., Tavakoli, M., Marquez, H. J. & Hashemzadeh, F. Nonlinear disturbance observer design for robotic manipulators. *Control Eng. Pract.* **21**, 253–267, DOI: 10.1016/j.conengprac.2012.10.008 (2013).
9. Sariyildiz, E. & Ohnishi, K. Stability and robustness of disturbance-observer-based motion control systems. *IEEE Trans. Ind. Electron.* **62**, 414–422, DOI: 10.1109/tie.2014.2327009 (2015).
10. Friston, K. The free-energy principle: a unified brain theory? *Nat. Rev. Neurosci.* **11**, 127–138, DOI: 10.1038/nrn2787 (2010).
11. Friston, K., FitzGerald, T., Rigoli, F., Schwartenbeck, P. & Pezzulo, G. Active inference: A process theory. *Neural Comput.* **29**, 1–49, DOI: 10.1162/neco_a_00912 (2017).
12. Pio-Lopez, L., Nizard, A., Friston, K. & Pezzulo, G. Active inference and robot control: a case study. *J. R. Soc. Interface* **13**, 20160616, DOI: 10.1098/rsif.2016.0616 (2016).
13. Pezzato, C., Ferrari, R. & Corbato, C. H. A novel adaptive controller for robot manipulators based on active inference. *IEEE Robot. Autom. Lett.* **5**, 2973–2980, DOI: 10.1109/lra.2020.2974451 (2020).
14. Baioumy, M., Duckworth, P., Lacerda, B. & Hawes, N. Active inference for integrated state-estimation, control, and learning. In *2021 IEEE International Conference on Robotics and Automation (ICRA)*, 4665–4671, DOI: 10.1109/ICRA48506.2021.9562009 (IEEE, 2021).
15. Meo, C., Franzese, G., Pezzato, C., Spahn, M. & Lanillos, P. Adaptation through prediction: Multisensory active inference torque control. *IEEE Trans. Cogn. Dev. Syst.* **15**, 32–41, DOI: 10.1109/tcds.2022.3156664 (2023).

16. Thomas, G. C. *et al.* Formulating and deploying strength amplification controllers for lower-body walking exoskeletons. *Front. Robot. AI* **8**, 720231, DOI: 10.3389/frobt.2021.720231 (2021).

17. Armstrong-Hélouvry, B., Dupont, P. & Canudas de Wit, C. A survey of models, analysis tools and compensation methods for the control of machines with friction. *Automatica* **30**, 1083–1138, DOI: 10.1016/0005-1098(94)90209-7 (1994).

18. Olsson, H., Åström, K. J., Canudas de Wit, C., Gäfvert, M. & Lischinsky, P. Friction models and friction compensation. *Eur. J. Control.* **4**, 176–195, DOI: 10.1016/S0947-3580(98)70113-X (1998).

19. Friston, K. J., Daunizeau, J., Kilner, J. & Kiebel, S. J. Action and behavior: a free-energy formulation. *Biol. Cybern.* **102**, 227–260, DOI: 10.1007/s00422-010-0364-z (2010).

20. Friston, K., Stephan, K., Li, B. & Daunizeau, J. Generalised filtering. *Math. Probl. Eng.* **2010**, 621670, DOI: 10.1155/2010/621670 (2010).

21. Back, J. & Shim, H. Adding robustness to nominal output-feedback controllers for uncertain nonlinear systems: A nonlinear version of disturbance observer. *Automatica* **44**, 2528–2537, DOI: 10.1016/j.automatica.2008.02.024 (2008).

22. Ha, W. & Back, J. A disturbance observer-based robust tracking controller for uncertain robot manipulators. *Int. J. Control Autom. Syst.* **16**, 417–425, DOI: 10.1007/s12555-017-0188-x (2018).

23. Baltieri, M. & Buckley, C. L. PID control as a process of active inference with linear generative models. *Entropy* **21**, 257, DOI: 10.3390/e21030257 (2019).

24. Baioumy, M., Mattamala, M., Duckworth, P., Lacerda, B. & Hawes, N. Adaptive manipulator control using active inference with precision learning. In *UKRAS20 Conference: "Robots into the real world" Proceedings*, vol. 3 of *UK-RAS Conference for PhD and Early Career Researchers Proceedings*, 105–107, DOI: 10.31256/gz8eu7c (EPSRC UK-RAS Network, 2020).

25. Mamedov, S. & Mikhel, S. Practical aspects of model-based collision detection. *Front. Robot. AI* **7**, 571574, DOI: 10.3389/frobt.2020.571574 (2020).

Data availability

The experimental datasets generated and analyzed during this study, together with the scripts used to process the data and generate the figures and tables, are available from the corresponding author on reasonable request.

Author contributions statement

H.R. designed the methodology, performed the experiments, analyzed the data, and wrote the original manuscript. P.H. conceived the study, supervised the study, contributed to the interpretation of the results, and reviewed and edited the manuscript. Both authors derived the error identities of the revised Methods and revised the manuscript.

Funding

This work was supported by the National Rehabilitation Center (NRC), Republic of Korea under Grant NRCTR-EX23003. This work was also supported by the National Research Foundation of Korea (NRF) grant funded by the Korean government (Ministry of Science and ICT, MSIT) (RS-2026-25492395).

Additional information

Competing interests The authors declare no competing interests.