

Trust-Region Neural Moving Horizon Estimation for Robots

Bingheng Wang, Xuyang Chen, and Lin Zhao

Abstract—Accurate disturbance estimation is essential for safe robot operations. The recently proposed neural moving horizon estimation (NeuroMHE), which uses a portable neural network to model the MHE’s weightings, has shown promise in further pushing the accuracy and efficiency boundary. Currently, NeuroMHE is trained through gradient descent, with its gradient computed recursively using a Kalman filter. This paper proposes a trust-region policy optimization method for training NeuroMHE. We achieve this by providing the second-order derivatives of MHE, referred to as the MHE Hessian. Remarkably, we show that many of the intermediate results used to obtain the gradient, especially the Kalman filter, can be efficiently reused to compute the MHE Hessian. This offers linear computational complexity with respect to the MHE horizon. As a case study, we evaluate the proposed trust region NeuroMHE on real quadrotor flight data for disturbance estimation. Our approach demonstrates highly efficient training in under 5 min using only 100 data points. It outperforms a state-of-the-art neural estimator by up to 68.1% in force estimation accuracy, utilizing only 1.4% of its network parameters. Furthermore, our method showcases enhanced robustness to network initialization compared to the gradient descent counterpart.

I. INTRODUCTION

Disturbances are widespread in robot operations, arising from various sources like system uncertainties [1], aerodynamic influences [2], ground friction forces [3], and hydrodynamic disturbances [4]. Addressing these disturbances is vital in robotic control to prevent significant performance degradation [5], [6]. However, developing a versatile model to capture these disturbances across environments is typically impractical due to their complexity. Therefore, online precise disturbance estimation that adapts to environments is crucial for safe and effective robot operations.

Existing works in disturbance estimation mainly utilize either model-based [7]–[11] or model-free [12]–[14] methods. While model-based estimators are data-efficient, they generally have limited dynamic response range, since the models used are typically limited to specific slowly-varying disturbances [7]–[9]. Another challenge lies in their heavy reliance on manually tuning numerous parameters [10], [11], which demands significant experimental efforts and expert knowledge. In contrast, model-free approaches using neural networks (NN) can achieve high performance across various

This work was supported by the Singapore Ministry of Education Tier 1 Academic Research Fund (22-5460-A0001) and Tier 2 AcRF (T2EP20123-0037). (Corresponding author: Lin Zhao.)

These authors are with the Department of Electrical and Computer Engineering, National University of Singapore, 4 Engineering Drive 3, 117583 Singapore, Singapore {wangbingheng, chenxuyang}@u.nus.edu, elezhli@nus.edu.sg

The source code of this work is available at <https://github.com/BinghengNUS/TR-NeuroMHE>




Fig. 1: Learning pipelines of the trust-region NeuroMHE. Currently, NeuroMHE is trained via gradient descent with its gradient computed recursively using a Kalman filter. This paper enhances NeuroMHE training with the second-order trust-region method. Interestingly, we show that the MHE Hessian can be obtained recursively using the same Kalman filter with just minor modifications to its inputs.

disturbance scenarios with minimal expert knowledge. Examples include estimating system uncertainties for robotic joints [12] and aerodynamic effects for aerial robots [13], [14]. In practice, these NN-based estimators typically use large network models for training, along with substantial disturbance data and complicated learning curricula.

To leverage advantages from both model-free and model-based methods, we recently proposed the neural moving horizon estimation (NeuroMHE) [15]. This algorithm fuses a portable neural network into an MHE to realize accurate estimation and fast online adaptation (See the dashed block in Fig. 1). MHE is an optimal state estimator using control theory for dynamic nonlinear optimization online, with a receding horizon strategy. It has consistently demonstrated superior performance in disturbance estimation for both ground [16] and aerial [17], [18] robots. Nevertheless, the parameter tuning remains a challenge for vanilla MHE, especially when dealing with dynamic and nonlinear parameters [19]. In contrast, NeuroMHE tackles this challenge by modeling its parameters using the neural network trained through gradient descent. The gradient includes the derivative of MHE’s optimal solution trajectory with respect to (w.r.t) the parameters, computed recursively using a Kalman filter. This approach results in linear computational complexity relative to the MHE horizon, significantly improving training efficiency over state-of-the-art auto-tuning MHE methods [20], [21].

In this paper, we enhance NeuroMHE training by providing second-order derivative information of MHE (i.e., the MHE Hessian). This allows us to employ the trust-region method (TRM), a second-order optimization technique, for training NeuroMHE. The trust-region method offers adaptive step-size updates, faster convergence, and improved robustness to initialization. It has shown efficacy across machine

learning and robotics, including policy optimization [22] and, more recently, synergizing optimal control with inverse reinforcement learning [23]. Fig. 1 outlines the learning pipelines of the trust-region NeuroMHE. Two key ingredients of our algorithm are the gradient and the Hessian trajectories of the MHE’s optimal solutions w.r.t the tuning parameters in the backward passes. By differentiating the Karush-Kuhn-Tucker (KKT) conditions of the MHE optimization problem in the forward pass, we can efficiently obtain the MHE gradient trajectory using a Kalman filter, as developed in [15]. Interestingly, we further show that the same Kalman filter can be reused to compute the MHE Hessian trajectory with only minor modifications to its inputs, which are acquired through double differentiation of the KKT conditions. This preserves the linear computational complexity w.r.t the MHE horizon, further ensuring the scalability of our method to high-dimension estimation problem.

We evaluate the proposed trust-region NeuroMHE in estimating complex aerodynamic disturbances using open-source real quadrotor flight data as a case study. Compared to the state-of-the-art estimator NeuroBEM [13], our approach achieves: 1) highly efficient training less than 5 min with only 100 data points; 2) using only 1.4% of the NeuroBEM’s network parameters; and 3) reducing the overall force estimation error by up to 68.1%. Further comparisons with the gradient descent counterpart show that, although computing the MHE Hessian trajectory naturally takes longer than the gradient trajectory, our method significantly reduces both the total training episodes and the overall training time. Additionally, our approach exhibits strong robustness to the initialization of the neural network. In summary, our contributions are threefold:

- 1) We propose a trust-region policy optimization method for training the recently developed NeuroMHE.
- 2) We show that much of the computation used to obtain the MHE gradient trajectory, particularly the Kalman filter, can be efficiently reused for computing the MHE Hessian trajectory in a recursive manner.
- 3) We validate the effectiveness of our approach using a real flight dataset, showing highly efficient training and superior performance compared to the state-of-the-art method in aerodynamic force estimation accuracy.

The rest of this paper is organized as follows. Section II briefly reviews the formulation of NeuroMHE. In Section III, we detail the process of reusing the Kalman filter to compute the MHE Hessian. Section IV develops the proposed trust-region NeuroMHE training method. Simulation results on the real dataset are reported in Section V. We conclude this paper and discuss our future work in Section VI.

II. PRELIMINARY OF NEUROMHE

A. Moving Horizon Estimation

MHE is a model-based optimal state estimator. Without loss of generality, we assume that the robotic dynamics model used in MHE takes the following form:

$$\dot{\mathbf{x}} = \mathbf{f}(\mathbf{x}, \mathbf{u}, \mathbf{w}), \quad \mathbf{y} = \mathbf{h}(\mathbf{x}) + \boldsymbol{\nu}, \quad (1)$$

where $\mathbf{x} \in \mathbb{R}^n$ is the system state, $\mathbf{u} \in \mathbb{R}^m$ is the control input, \mathbf{f} is the system model, $\mathbf{w} \in \mathbb{R}^w$ is the process noise, \mathbf{h} is the output function, and $\mathbf{y} \in \mathbb{R}^l$ denotes the measurement affected by the noise $\boldsymbol{\nu} \in \mathbb{R}^l$. At time t , given the most recent measurements $\mathbf{y}_t = \{\mathbf{y}_k\}_{k=t-N}^t$ and control inputs $\mathbf{u}_t = \{\mathbf{u}_k\}_{k=t-N}^{t-1}$ collected in a moving data window of horizon N , the MHE estimator optimizes over $\mathbf{x} = \{\mathbf{x}_k\}_{k=t-N}^t$ and $\mathbf{w} = \{\mathbf{w}_k\}_{k=t-N}^{t-1}$ at each time step $t \geq N$ by solving the following nonlinear optimization problem online.

$$\begin{aligned} \min_{\mathbf{x}, \mathbf{w}} J = & \underbrace{\frac{1}{2} \|\mathbf{x}_{t-N} - \hat{\mathbf{x}}_{t-N}\|_{\mathbf{P}}^2}_{\text{arrival cost}} \\ & + \underbrace{\frac{1}{2} \sum_{k=t-N}^t \|\mathbf{y}_k - \mathbf{h}(\mathbf{x}_k)\|_{\mathbf{R}_k}^2 + \frac{1}{2} \sum_{k=t-N}^{t-1} \|\mathbf{w}_k\|_{\mathbf{Q}_k}^2}_{\text{running cost}} \end{aligned} \quad (2a)$$

$$\text{s.t. } \mathbf{x}_{k+1} = \mathbf{f}(\mathbf{x}_k, \mathbf{u}_k, \mathbf{w}_k, \Delta t), \quad (2b)$$

where $\mathbf{P} \in \mathbb{R}^{n \times n}$, $\mathbf{R}_k \in \mathbb{R}^{l \times l}$, and $\mathbf{Q}_k \in \mathbb{R}^{w \times w}$ are the positive-definite weighting matrices, \mathbf{f} is the discrete-time model of \mathbf{f} with the step-size of Δt for predicting the state, and $\hat{\mathbf{x}}_{t-N}$ is the filter priori, chosen as the MHE’s optimal estimate $\hat{\mathbf{x}}_{t-N|t-1}$ of \mathbf{x}_{t-N} obtained at $t-1$ [24]. We denote by $\hat{\mathbf{x}} = \{\hat{\mathbf{x}}_k|t\}_{k=t-N}^t$ and $\hat{\mathbf{w}} = \{\hat{\mathbf{w}}_k|t\}_{k=t-N}^{t-1}$ the MHE’s optimal solutions to Problem (2) at time step t .

B. Formulation of NeuroMHE

The weighting matrices in the cost function (2a) are the tuning parameters, which determine the MHE performance. For ease of presentation, we collect them in a vector $\boldsymbol{\theta} = [\text{vec}(\mathbf{P}), \{\text{vec}(\mathbf{R}_k)\}_{k=t-N}^t, \{\text{vec}(\mathbf{Q}_k)\}_{k=t-N}^{t-1}] \in \mathbb{R}^p$ where $\text{vec}(\cdot)$ denotes the vectorization of a given matrix. Tuning $\boldsymbol{\theta}$ typically necessitates expert knowledge of the covariances of the noises entering robotic systems. In practice, identifying these noise covariances is difficult due to their large number and strong coupling. It becomes even more demanding when these values exhibit dynamic behavior.

NeuroMHE tackles the above challenge by modeling $\boldsymbol{\theta}$ with a portable neural network:

$$\boldsymbol{\theta} = \mathbf{f}_{\boldsymbol{\varpi}}(\boldsymbol{\chi}). \quad (3)$$

The neural network can accept various signals as its inputs $\boldsymbol{\chi}$, including measurements \mathbf{y} , controls \mathbf{u} , and other external signals \mathbf{o} , depending on specific applications. Its parameters $\boldsymbol{\varpi}$ are trained using powerful machine learning techniques. Here, $\boldsymbol{\varpi}$ represents the tuning parameters for NeuroMHE, and we parameterize Problem (2) as NeuroMHE($\boldsymbol{\varpi}$), with the MHE’s solution denoted as $\hat{\mathbf{x}}(\boldsymbol{\varpi})$. To tune MHE, a high-level optimization problem can be formulated as follows:

$$\min_{\boldsymbol{\varpi}} L(\hat{\mathbf{x}}(\boldsymbol{\varpi})) \quad (4a)$$

$$\text{s.t. } \hat{\mathbf{x}}(\boldsymbol{\varpi}) \text{ generated by NeuroMHE}(\boldsymbol{\varpi}). \quad (4b)$$

where $L(\hat{\mathbf{x}}(\boldsymbol{\varpi}))$ is a loss function assessing NeuroMHE’s performance. It can be built upon estimation errors when

ground truth disturbance data is available. Alternatively, it can be tailored to penalize tracking errors in applications involving robust trajectory tracking control.

III. ANALYTICAL HESSIAN TRAJECTORY

Our previous work [15] solved Problem (4) using gradient descent, with the gradient computed recursively through a Kalman filter. In this section, our objective is to efficiently reuse the same Kalman filter and much of the computation initially employed to obtain the gradient. We intend to apply these resources in computing the MHE Hessian for training NeuroMHE with second-order optimization techniques.

A. MHE Gradient

The gradient of $L(\hat{\mathbf{x}}(\boldsymbol{\varpi}))$ w.r.t $\boldsymbol{\varpi}$ can be computed using the chain rule:

$$\nabla_{\boldsymbol{\varpi}} L(\hat{\mathbf{x}}(\boldsymbol{\varpi})) = \nabla_{\hat{\mathbf{x}}} L(\hat{\mathbf{x}}) \nabla_{\boldsymbol{\theta}} \hat{\mathbf{x}}(\boldsymbol{\theta}) \nabla_{\boldsymbol{\varpi}} \boldsymbol{\theta}(\boldsymbol{\varpi}). \quad (5)$$

The gradients $\nabla_{\hat{\mathbf{x}}} L$ and $\nabla_{\boldsymbol{\varpi}} \boldsymbol{\theta}$ are straightforward to obtain as both $L(\hat{\mathbf{x}})$ and $\boldsymbol{\theta}(\boldsymbol{\varpi})$ are explicit functions. The main challenge lies in the computation of $\nabla_{\boldsymbol{\theta}} \hat{\mathbf{x}}$ (the MHE gradient). Recall that $\hat{\mathbf{x}}(\boldsymbol{\theta})$ is the MHE's solution, with its dependence on $\boldsymbol{\theta}$ implicitly defined via the KKT conditions. This implicit definition enables us to compute $\nabla_{\boldsymbol{\theta}} \hat{\mathbf{x}}$ by differentiating through the KKT conditions of Problem (2). To proceed, we associate the equality constraints (2b) with the dual variables $\boldsymbol{\lambda} = \{\boldsymbol{\lambda}_k\}_{k=t-N}^{t-1}$ and denote $\boldsymbol{\lambda}^* \in \mathbb{R}^n$ as their optimal values. The corresponding Lagrangian can be written as

$$\mathcal{L} = J + \sum_{k=t-N}^{t-1} \boldsymbol{\lambda}_k^T (\mathbf{x}_{k+1} - \mathbf{f}(\mathbf{x}_k, \mathbf{u}_k, \mathbf{w}_k, \Delta t)). \quad (6)$$

Then the KKT conditions of Problem (2) at $\hat{\mathbf{x}}$, $\hat{\mathbf{w}}$, and $\boldsymbol{\lambda}^*$ are given by

$$\nabla_{\hat{\mathbf{x}}_{t-N|t}} \mathcal{L}(\hat{\mathbf{x}}_{t-N|t}, \hat{\mathbf{x}}_{t-N}, \hat{\mathbf{w}}_{t-N|t}, \boldsymbol{\lambda}_{t-N}^*, \boldsymbol{\theta}) = \mathbf{0}, \quad (7a)$$

$$\nabla_{\hat{\mathbf{x}}_{k|t}} \mathcal{L}(\hat{\mathbf{x}}_{k|t}, \hat{\mathbf{w}}_{k|t}, \boldsymbol{\lambda}_k^*, \boldsymbol{\lambda}_{k-1}^*, \boldsymbol{\theta}) = \mathbf{0}, \quad (7b)$$

$$\nabla_{\hat{\mathbf{w}}_{k|t}} \mathcal{L}(\hat{\mathbf{x}}_{k|t}, \hat{\mathbf{w}}_{k|t}, \boldsymbol{\lambda}_k^*, \boldsymbol{\theta}) = \mathbf{0}, \quad (7c)$$

$$\nabla_{\boldsymbol{\lambda}_k^*} \mathcal{L} = \hat{\mathbf{x}}_{k+1|t} - \mathbf{f}(\hat{\mathbf{x}}_{k|t}, \mathbf{u}_k, \hat{\mathbf{w}}_{k|t}, \Delta t) = \mathbf{0}, \quad (7d)$$

Differentiating the KKT conditions (7) w.r.t $\boldsymbol{\theta}$ yields the following differential KKT conditions.

$$\begin{aligned} \frac{d\nabla_{\hat{\mathbf{x}}_{t-N|t}} \mathcal{L}}{d\boldsymbol{\theta}} &= \mathbf{L}_{t-N}^{xx} \nabla_{\boldsymbol{\theta}} \hat{\mathbf{x}}_{t-N|t} - \mathbf{P} \nabla_{\boldsymbol{\theta}} \hat{\mathbf{x}}_{t-N} + \mathbf{L}_{t-N}^{x\boldsymbol{\theta}} \\ &\quad + \mathbf{L}_{t-N}^{xw} \nabla_{\boldsymbol{\theta}} \hat{\mathbf{w}}_{t-N|t} - \mathbf{F}_{t-N}^T \nabla_{\boldsymbol{\theta}} \boldsymbol{\lambda}_{t-N}^* = \mathbf{0}, \end{aligned} \quad (8a)$$

$$\begin{aligned} \frac{d\nabla_{\hat{\mathbf{x}}_{k|t}} \mathcal{L}}{d\boldsymbol{\theta}} &= \mathbf{L}_k^{xx} \nabla_{\boldsymbol{\theta}} \hat{\mathbf{x}}_{k|t} + \mathbf{L}_k^{xw} \nabla_{\boldsymbol{\theta}} \hat{\mathbf{w}}_{k|t} - \mathbf{F}_k^T \nabla_{\boldsymbol{\theta}} \boldsymbol{\lambda}_k^* \\ &\quad + \nabla_{\boldsymbol{\theta}} \boldsymbol{\lambda}_{k-1}^* + \mathbf{L}_k^{x\boldsymbol{\theta}} = \mathbf{0}, \end{aligned} \quad (8b)$$

$$\begin{aligned} \frac{d\nabla_{\hat{\mathbf{w}}_{k|t}} \mathcal{L}}{d\boldsymbol{\theta}} &= \mathbf{L}_k^{wx} \nabla_{\boldsymbol{\theta}} \hat{\mathbf{x}}_{k|t} + \mathbf{L}_k^{ww} \nabla_{\boldsymbol{\theta}} \hat{\mathbf{w}}_{k|t} \\ &\quad - \mathbf{G}_k^T \nabla_{\boldsymbol{\theta}} \boldsymbol{\lambda}_k^* + \mathbf{L}_k^{w\boldsymbol{\theta}} = \mathbf{0}, \end{aligned} \quad (8c)$$

$$\frac{d\nabla_{\boldsymbol{\lambda}_k^*} \mathcal{L}}{d\boldsymbol{\theta}} = \nabla_{\boldsymbol{\theta}} \hat{\mathbf{x}}_{k+1|t} - \mathbf{F}_k \nabla_{\boldsymbol{\theta}} \hat{\mathbf{x}}_{k|t} - \mathbf{G}_k \nabla_{\boldsymbol{\theta}} \hat{\mathbf{w}}_{k|t} = \mathbf{0}, \quad (8d)$$

where $\mathbf{F}_k = \nabla_{\hat{\mathbf{x}}_{k|t}} \mathbf{f}$, $\mathbf{G}_k = \nabla_{\hat{\mathbf{w}}_{k|t}} \mathbf{f}$, $\mathbf{H}_k = \nabla_{\hat{\mathbf{x}}_{k|t}} \mathbf{h}$, $\boldsymbol{\lambda}_t^* = \mathbf{0}$, $\mathbf{L}_k^{xx} = \frac{\partial^2 \mathcal{L}}{\partial \hat{\mathbf{x}}_{k|t}^2}$, $\mathbf{L}_k^{xw} = \frac{\partial^2 \mathcal{L}}{\partial \hat{\mathbf{x}}_{k|t} \partial \hat{\mathbf{w}}_{k|t}}$, $\mathbf{L}_k^{x\boldsymbol{\theta}} = \frac{\partial^2 \mathcal{L}}{\partial \hat{\mathbf{x}}_{k|t} \partial \boldsymbol{\theta}}$, $\mathbf{L}_k^{wx} = \frac{\partial^2 \mathcal{L}}{\partial \hat{\mathbf{w}}_{k|t} \partial \hat{\mathbf{x}}_{k|t}}$, and $\mathbf{L}_k^{w\boldsymbol{\theta}} = \frac{\partial^2 \mathcal{L}}{\partial \hat{\mathbf{w}}_{k|t} \partial \boldsymbol{\theta}}$.

The above differential KKT conditions (8) can be easily shown to be the same as the KKT conditions of the following linear MHE problem:

$$\begin{aligned} \min_{\mathbf{X}, \mathbf{W}} J_2 &= \frac{1}{2} \text{Tr} \left\| \mathbf{X}_{t-N} - \hat{\mathbf{X}}_{t-N} \right\|_{\mathbf{P}}^2 \\ &\quad + \text{Tr} \sum_{k=t-N}^t \left(\frac{1}{2} \mathbf{X}_k^T \bar{\mathbf{L}}_k^{xx} \mathbf{X}_k + \mathbf{W}_k^T \mathbf{L}_k^{wx} \mathbf{X}_k \right) \\ &\quad + \text{Tr} \sum_{k=t-N}^{t-1} \left(\frac{1}{2} \mathbf{W}_k^T \mathbf{L}_k^{ww} \mathbf{W}_k + (\mathbf{L}_k^{w\boldsymbol{\theta}})^T \mathbf{W}_k \right) \\ &\quad + \text{Tr} \sum_{k=t-N}^t \left((\mathbf{L}_k^{x\boldsymbol{\theta}})^T \mathbf{X}_k \right) \end{aligned} \quad (9a)$$

$$\text{s.t. } \mathbf{X}_{k+1} = \mathbf{F}_k \mathbf{X}_k + \mathbf{G}_k \mathbf{W}_k, \quad (9b)$$

where $\mathbf{X} = \{\mathbf{X}_k\}_{k=t-N}^t$, $\mathbf{W} = \{\mathbf{W}_k\}_{k=t-N}^{t-1}$, $\bar{\mathbf{L}}_k^{xx} = \frac{\partial^2 \bar{\mathcal{L}}}{\partial \hat{\mathbf{x}}_{k|t}^2}$ with $\bar{\mathcal{L}}$ obtained by excluding the arrival cost (See (2a)) from \mathcal{L} , and $\text{Tr}(\cdot)$ is the matrix trace. Hence, the optimal solution $\hat{\mathbf{X}}$ to Problem (9) is exactly the gradient $\nabla_{\boldsymbol{\theta}} \hat{\mathbf{x}}$, denoted as $\hat{\mathbf{X}} = \{\hat{\mathbf{X}}_k\}_{k=t-N}^t = \{\nabla_{\boldsymbol{\theta}} \hat{\mathbf{x}}_{k|t}\}_{k=t-N}^t$.

Note that (9) has a quadratic cost function (9a) with linear dynamics constraints (9b). It is possible to obtain $\hat{\mathbf{X}}$ analytically. In [15], we proposed a recursive method for computing $\hat{\mathbf{X}}$ using a Kalman filter. The details of this method will be presented in Algorithm 1 (See III-C). Next, we will focus on an interesting observation: the Hessian of the loss function can also be efficiently computed using a new auxiliary linear MHE system. Remarkably, this approach leverages the same Kalman filter and much of the computation that was already used to obtain the gradient.

B. MHE Hessian

Let $\mathbf{H}_{\hat{\mathbf{x}}} L$ denote the Hessian of L w.r.t $\hat{\mathbf{x}}$, $\mathbf{H}_{\boldsymbol{\theta}} \hat{\mathbf{x}} := \text{dvec}(\nabla_{\boldsymbol{\theta}} \hat{\mathbf{x}}) / \partial \boldsymbol{\theta}$ the MHE Hessian, and \otimes the Kronecker product. The total second-order derivative of $L(\hat{\mathbf{x}}(\boldsymbol{\varpi}))$ w.r.t $\boldsymbol{\varpi}$ can be obtained using the chain rule of matrix calculus:

$$\begin{aligned} \mathbf{H}_{\boldsymbol{\varpi}} L(\hat{\mathbf{x}}(\boldsymbol{\varpi})) &= (\nabla_{\boldsymbol{\theta}} \hat{\mathbf{x}})^T \mathbf{H}_{\hat{\mathbf{x}}} L \nabla_{\boldsymbol{\theta}} \hat{\mathbf{x}} (\nabla_{\boldsymbol{\varpi}} \boldsymbol{\theta})^2 \\ &\quad + (\nabla_{\boldsymbol{\varpi}} \boldsymbol{\theta})^T \mathbf{H}_{\boldsymbol{\theta}} \hat{\mathbf{x}} \left[(\nabla_{\hat{\mathbf{x}}} L)^T \otimes \mathbf{I} \right] \nabla_{\boldsymbol{\varpi}} \boldsymbol{\theta} \\ &\quad + [(\nabla_{\hat{\mathbf{x}}} L \nabla_{\boldsymbol{\theta}} \hat{\mathbf{x}}) \otimes \mathbf{I}] \mathbf{H}_{\boldsymbol{\varpi}} \boldsymbol{\theta}. \end{aligned} \quad (10)$$

Note that all the first-order derivatives have been obtained in computing the gradient of $L(\hat{\mathbf{x}}(\boldsymbol{\varpi}))$ w.r.t $\boldsymbol{\varpi}$. Computing the second-order derivatives $\mathbf{H}_{\hat{\mathbf{x}}} L$ and $\mathbf{H}_{\boldsymbol{\theta}} \hat{\mathbf{x}}$ is straightforward, which is explicit differentiation of the loss function and the NN, respectively. The major challenge lies in the computation of $\mathbf{H}_{\boldsymbol{\theta}} \hat{\mathbf{x}}$.

Given the implicit dependence of $\hat{\mathbf{x}}$ on $\boldsymbol{\theta}$ through KKT conditions, we are motivated to differentiate the KKT condi-

tions (7) w.r.t θ twice. This results in the following second-order differential KKT conditions.

$$\begin{aligned} H_\theta \nabla_{\hat{x}_{t-N|t}} \mathcal{L} &= \dot{\mathbf{L}}_{t-N}^{xx} H_\theta \hat{x}_{t-N|t} - \dot{\mathbf{P}} H_\theta \hat{x}_{t-N} + \dot{\mathbf{L}}_{t-N}^{x\theta} \\ &+ \dot{\mathbf{L}}_{t-N}^{xw} H_\theta \hat{w}_{t-N|t} - \dot{\mathbf{F}}_{t-N}^T H_\theta \lambda_{t-N}^* = \mathbf{0}, \end{aligned} \quad (11a)$$

$$\begin{aligned} H_\theta \nabla_{\hat{x}_{k|t}} \mathcal{L} &= \dot{\mathbf{L}}_k^{xx} H_\theta \hat{x}_{k|t} + \dot{\mathbf{L}}_k^{xw} H_\theta \hat{w}_{k|t} - \dot{\mathbf{F}}_k^T H_\theta \lambda_k^* \\ &+ H_\theta \lambda_{k-1}^* + \dot{\mathbf{L}}_k^{x\theta} = \mathbf{0}, \end{aligned} \quad (11b)$$

$$\begin{aligned} H_\theta \nabla_{\hat{w}_{k|t}} \mathcal{L} &= \dot{\mathbf{L}}_k^{wx} H_\theta \hat{x}_{k|t} + \dot{\mathbf{L}}_k^{ww} H_\theta \hat{w}_{k|t} \\ &- \dot{\mathbf{G}}_k^T H_\theta \lambda_k^* + \dot{\mathbf{L}}_k^{w\theta} = \mathbf{0}, \end{aligned} \quad (11c)$$

$$\begin{aligned} H_\theta \nabla_{\lambda_k^*} \mathcal{L} &= H_\theta \hat{x}_{k+1|t} - \dot{\mathbf{F}}_k H_\theta \hat{x}_{k|t} - \dot{\mathbf{G}}_k H_\theta \hat{w}_{k|t} \\ &+ \dot{\mathbf{L}}_k^{\lambda\theta} = \mathbf{0}, \end{aligned} \quad (11d)$$

where the coefficient matrices are defined below:

$$\begin{aligned} \dot{\mathbf{L}}_k^{xx} &= \mathbf{I} \otimes L_k^{xx}, \quad \dot{\mathbf{F}}_k = \mathbf{I} \otimes F_k, \quad \dot{\mathbf{G}}_k = \mathbf{I} \otimes G_k, \\ \dot{\mathbf{L}}_k^{xw} &= \mathbf{I} \otimes L_k^{xw}, \quad \dot{\mathbf{L}}_k^{wx} = \mathbf{I} \otimes L_k^{wx}, \quad \dot{\mathbf{P}} = \mathbf{I} \otimes P, \\ \dot{\mathbf{L}}_{t-N}^{x\theta} &= (\nabla_{\hat{x}_{t-N|t}} \bar{f}_0) \hat{X}_{t-N|t} + (\nabla_{\lambda_{t-N}^*} \bar{f}_0) \Lambda_{t-N}^* \\ &+ (\nabla_{\hat{w}_{t-N|t}} \bar{f}_0) \hat{W}_{t-N|t} + \nabla_{\theta} \bar{f}_0, \\ \dot{\mathbf{L}}_k^{x\theta} &= (\nabla_{\hat{x}_{k|t}} \bar{f}_k) \hat{X}_{k|t} + (\nabla_{\hat{w}_{k|t}} \bar{f}_k) \hat{W}_{k|t} \\ &+ (\nabla_{\lambda_k^*} \bar{f}_k) \Lambda_k^* + \nabla_{\theta} \bar{f}_k, \\ \dot{\mathbf{L}}_k^{w\theta} &= (\nabla_{\hat{x}_{k|t}} \bar{g}_k) \hat{X}_{k|t} + (\nabla_{\hat{w}_{k|t}} \bar{g}_k) \hat{W}_{k|t} \\ &+ (\nabla_{\lambda_k^*} \bar{g}_k) \Lambda_k^* + \nabla_{\theta} \bar{g}_k, \\ \dot{\mathbf{L}}_k^{\lambda\theta} &= (\nabla_{\hat{x}_{k|t}} \bar{h}_k) \hat{X}_{k|t} + (\nabla_{\hat{w}_{k|t}} \bar{h}_k) \hat{W}_{k|t}, \end{aligned} \quad (12)$$

with \bar{f}_0 , \bar{f}_k , \bar{g}_k , and \bar{h}_k being the left hand sides of (8a)-(8d).

It is interesting to notice that (11) and (8) share exactly the same structure. The only exception lies in (11d) which has an additional term $\dot{\mathbf{L}}_k^{\lambda\theta}$, not present in (8d). This term can be considered as a known input to the linear system. As shown in III-A and [15], (8) can represent the KKT conditions of the auxiliary MHE optimization problem (9), which is subject to the linear system (9b) constructed from (8d). Given this observation and concept, we can interpret (11) as the KKT conditions of another auxiliary MHE optimization problem. It is subject to a new linear system constructed from (11d):

$$\hat{X}_{k+1} = \dot{\mathbf{F}}_k \hat{X}_k + \dot{\mathbf{G}}_k \hat{W}_k - \dot{\mathbf{L}}_k^{\lambda\theta}, \quad (13)$$

where \hat{X}_k and \hat{W}_k denote $H_\theta x_k$ and $H_\theta w_k$, respectively. The related cost function can be established directly from (9a) by updating the latter's coefficient matrices, for example, by replacing P with \dot{P} defined in (12). Consequently, we can efficiently reuse Algorithm 1 to solve for the MHE Hessian $H_\theta \hat{x}$, which is the optimal solution \hat{X} to the new auxiliary linear MHE problem. Only the algorithm's inputs need minor modifications as follows:

- Replace \hat{X}_{t-N} with \hat{X}_{t-N} (which is approximated by $H_\theta \hat{x}_{t-N|t-1}$ obtained at $t-1$);
- Update the matrices in (16) using (12). In particular, add $\dot{\mathbf{L}}_k^{\lambda\theta}$ to (15a) such that the updated A_k becomes $\hat{A}_k = \dot{\mathbf{G}}_k \left(\dot{\mathbf{L}}_k^{ww} \right)^{-1} \dot{\mathbf{L}}_k^{w\theta} + \dot{\mathbf{L}}_k^{\lambda\theta}$. This can be derived based on the example used to obtain (15a) (See III-C).

C. Kalman Filter-based Solver

First, a Kalman filter (KF) is solved to obtain the estimates $\left\{ \hat{X}_{k|k}^{\text{KF}} \right\}_{k=t-N}^t$: the initial condition is given by

$$\hat{X}_{t-N|t-N}^{\text{KF}} = (\mathbf{I} + C_{t-N} S_{t-N}) \hat{X}_{t-N} + C_{t-N} T_{t-N}. \quad (14)$$

where \mathbf{I} is an identity matrix, C_{t-N} is defined in (15c) with $P_{t-N} = P^{-1}$, S_{t-N} and T_{t-N} are given in (16), and \hat{X}_{t-N} is approximated by $\nabla_{\theta} \hat{x}_{t-N|t-1}$ obtained at $t-1$. Then, the remaining estimates are generated through the following equations from $k = t - N + 1$ to t :

$$\hat{X}_{k|k-1} = \bar{F}_{k-1} \hat{X}_{k-1}^{\text{KF}} - A_{k-1}, \quad (15a)$$

$$P_k = \bar{F}_{k-1} C_{k-1} \bar{F}_{k-1}^T + B_{k-1}, \quad (15b)$$

$$C_k = (\mathbf{I} - P_k S_k)^{-1} P_k, \quad (15c)$$

$$\hat{X}_{k|k}^{\text{KF}} = (\mathbf{I} + C_k S_k) \hat{X}_{k|k-1} + C_k T_k. \quad (15d)$$

where A_k , B_k , S_k , T_k , and \bar{F}_k are defined below:

$$\begin{aligned} A_k &= G_k (L_k^{ww})^{-1} L_k^{w\theta}, \\ B_k &= G_k (L_k^{ww})^{-1} G_k^T, \\ S_k &= L_k^{xw} (L_k^{ww})^{-1} L_k^{wx} - \bar{L}_k^{xx}, \quad S_t = -\bar{L}_t^{xx}, \\ T_k &= L_k^{xw} (L_k^{ww})^{-1} L_k^{w\theta} - L_k^{x\theta}, \quad T_t = -L_t^{x\theta}, \\ \bar{F}_k &= F_k - G_k (L_k^{ww})^{-1} L_k^{wx}. \end{aligned} \quad (16)$$

Second, the new dual variables $\Lambda^* = \{\Lambda_k^*\}_{k=t-N}^{t-1}$ are computed using the following equation backward in time from $k = t$ to $t - N + 1$, starting with $\Lambda_t^* = \mathbf{0}$:

$$\Lambda_{k-1}^* = (\mathbf{I} + S_k C_k) \bar{F}_k^T \Lambda_k^* + S_k \hat{X}_{k|k}^{\text{KF}} + T_k. \quad (17)$$

Third, the optimal estimates \hat{X} are obtained using

$$\hat{X}_{k|t} = \hat{X}_{k|k}^{\text{KF}} + C_k \bar{F}_k^T \Lambda_k^*, \quad (18)$$

iteratively from $k = t - N$ to t .

The above recursions are summarized in Algorithm 1.

Algorithm 1: Solving for \hat{X} using a Kalman filter

Input: \hat{X}_{t-N} and the matrices in (16);

```

1  def Kalman_Filter_based_Gradient_Solver:
2  |   Set  $\hat{X}_{t-N|t-N}^{\text{KF}}$  using (14);
3  |   for  $k \leftarrow t - N + 1$  to  $t$  by 1 do
4  |     |   Update  $\hat{X}_{k|k}^{\text{KF}}$  using the Kalman filter (15);
5  |   end for
6  |   for  $k \leftarrow t$  to  $t - N + 1$  by -1 do
7  |     |   Update  $\Lambda_{k-1}^*$  using (17) with  $\Lambda_t^* = \mathbf{0}$ ;
8  |   end for
9  |   for  $k \leftarrow t - N$  to  $t$  by 1 do
10 |     |   Update  $\hat{X}_{k|t}$  using (18);
11 |   end for
12  return  $\left\{ \hat{X}_{k|t} \right\}_{k=t-N}^t$ 
Output:  $\nabla_{\theta} \hat{x} = \hat{X}$ 

```

In the Kalman filter (15), (15a) is the state predictor, (15b) handles the error covariance prediction, (15c) addresses the error covariance correction, and (15d) is the state corrector. These equations are derived from the differential KKT conditions (8) by induction. For example, we obtain (15a) by following steps: 1) solve for $\hat{\mathbf{W}}_{k|t} := \nabla_{\theta} \hat{\omega}_{k|t}$ from (8c); 2) plug $\hat{\mathbf{W}}_{k|t}$ and $\hat{\mathbf{X}}_{k|t}$ (obtained from (18)) into (8d); 3) define the resulting equation, excluding the term related to $\Lambda_k^* := \nabla_{\theta} \lambda_k^*$, as $\hat{\mathbf{X}}_{k+1|k}$ using \mathbf{A}_k and $\bar{\mathbf{F}}_k$ given in (16). The detailed theoretical justification of (15) can be found in Appendix-C of [15]. Note that (15) is not presented in the standard form of a Kalman filter for the purpose of a clearer and more inductive proof, as explained in [15]. The Kalman filter provides analytical gradients in a recursive form, significantly enhancing computational efficiency.

IV. TRUST-REGION LEARNING FRAMEWORK

We aim to train the neural network's parameters ϖ using the trust-region method, which addresses Problem (4) by creating a trust region around the current parameters ϖ_t . Within this region, it minimizes a locally quadratic approximation of $L(\hat{\mathbf{x}}(\varpi_t))$ to determine a candidate update μ for ϖ_t . The corresponding optimization problem is as follows:

$$\min_{\mu} \text{TRM}(\mu) := L + \mu^T \nabla_{\varpi_t} L + \frac{1}{2} \mu^T (H_{\varpi_t} L) \mu \quad (19a)$$

$$\text{s.t. } \|\mu\| \leq \Delta_t. \quad (19b)$$

Given predefined thresholds $0 \leq \xi_1 \leq \xi_2 < \xi_3$, $\kappa_1 \in (0, 1)$, $\kappa_2 > 1$, and the upper bound of the radius $\bar{\Delta}$, the method iteratively explores the range of trust region by updating its radius $\Delta_t \in \mathbb{R}_+$ with the following law [25]:

$$\Delta_{t+1} = \begin{cases} \kappa_1 \Delta_t & \text{if } \rho_t < \xi_2 \\ \min(\kappa_2 \Delta_t, \bar{\Delta}) & \text{if } \rho_t > \xi_3 \& \|\mu_t\| = \Delta_t \\ \Delta_t & \text{otherwise} \end{cases} \quad (20)$$

Here, ρ_t denotes the current ratio between the actual and predicted reductions on the loss function

$$\rho_t = \frac{L(\hat{\mathbf{x}}(\varpi_t)) - L(\hat{\mathbf{x}}(\varpi_t + \mu_t))}{\text{TRM}(\mathbf{0}) - \text{TRM}(\mu_t)}, \quad (21)$$

which is used to decide whether to update ϖ_t : if $\rho_t > \xi_1$, ϖ_t will be updated using μ_t . We summarize the trust-region policy optimization for training NeuroMHE in Algorithm 2.

V. NUMERICAL CASE STUDY

We numerically validate the performance of trust-region NeuroMHE in estimating complex aerodynamic disturbances on quadrotors using a real flight dataset from various agile flights [13]. We compare it with the gradient descent counterpart and the state-of-the-art estimator NeuroBEM [13] to showcase the following advantages: 1) highly-efficient training; 2) improved robustness to network initialization; and 3) superior force estimation accuracy.

In alignment with NeuroBEM, NeuroMHE is trained using estimation errors. In Algorithm 2, we define the desired states \mathcal{S} using the ground truth disturbance forces \mathbf{F} and torques $\boldsymbol{\tau}$ from the real dataset and the quadrotor model. For

Algorithm 2: Trust-region NeuroMHE training

Input: The desired states \mathcal{S} and the measurements \mathcal{Y} .

```

1 while  $L_{\text{mean}}$  not converged do
2   for  $t \leftarrow 0$  to  $T$  do
3     Obtain  $\hat{\mathbf{x}}(\varpi_t)$  by solving NeuroMHE( $\varpi_t$ )
       with the most recent measurements  $\mathbf{y}_t \in \mathcal{Y}$ ;
4     Compute  $L(\hat{\mathbf{x}}(\varpi_t))$  using  $\hat{\mathbf{x}}(\varpi_t)$  and  $\mathbf{s} \in \mathcal{S}$ ;
5     Compute  $\nabla_{\varpi_t} L$  using (5) and Algorithm 1;
6     Compute  $H_{\varpi_t} L$  using (10) and Algorithm 1
       with its inputs updated by (12);
7     Obtain  $\mu_t$  by solving TRM Problem (19);
8     Obtain  $\hat{\mathbf{x}}(\varpi_t + \mu_t)$  by solving
       NeuroMHE( $\varpi_t + \mu_t$ ) with the same  $\mathbf{y}_t$ ;
9     Compute  $L(\hat{\mathbf{x}}(\varpi_t + \mu_t))$  using  $\hat{\mathbf{x}}(\varpi_t + \mu_t)$ 
       and the same  $\mathbf{s}$ ;
10    Update  $\rho_t$  using (21);
11    Update  $\Delta_t$  using (20);
12    if  $\rho_t > \xi_1$  then
13      | Update  $\varpi_{t+1} \leftarrow \varpi_t + \mu_t$ ;
14    else
15      |  $\varpi_{t+1} \leftarrow \varpi_t$ ;
16    end if
17  end for
18  Compute the mean loss  $L_{\text{mean}} = \frac{1}{T} \sum_{t=0}^T L(\hat{\mathbf{x}})$ 
       for the next episode
19 end while

```

NeuroMHE, the measurements \mathcal{Y} are set as the quadrotor's linear \mathbf{v} and angular velocities $\boldsymbol{\omega}$. By comparison, the NeuroBEM's inputs also comprise motor speeds, requiring specific sensors and autopilot firmware. To address the state estimation problem (2), we integrate the velocities with the unmeasurable disturbances to create the augmented state $\mathbf{x} = [\mathbf{v}; \mathbf{F}; \boldsymbol{\omega}; \boldsymbol{\tau}] \in \mathbb{R}^{12}$. Assuming no loss of generality, the disturbance dynamics are driven by the process noises $\mathbf{w}_f \in \mathbb{R}^3$ and $\mathbf{w}_\tau \in \mathbb{R}^3$. The system model for state prediction in MHE is as follows:

$$\dot{\mathbf{v}} = m^{-1}(-mge_3 + \mathbf{F}), \quad \dot{\mathbf{F}} = \mathbf{w}_f, \quad (22a)$$

$$\dot{\boldsymbol{\omega}} = \mathbf{J}^{-1}(-\boldsymbol{\omega}^\times \mathbf{J} \boldsymbol{\omega} + \boldsymbol{\tau}), \quad \dot{\boldsymbol{\tau}} = \mathbf{w}_\tau, \quad (22b)$$

where $g = 9.81$ m/s², $\mathbf{e}_3 = [0; 0; 1]$, $m = 0.772$ kg, and $\mathbf{J} = \text{diag}(0.0025, 0.0021, 0.0043)$ kgm²¹.

The weighting matrices of NeuroMHE have the following dimensions: $\mathbf{P} \in \mathbb{R}^{12 \times 12}$, $\mathbf{R}_k \in \mathbb{R}^{6 \times 6}$, and $\mathbf{Q}_k \in \mathbb{R}^{6 \times 6}$. Instead of training the neural network to generate all \mathbf{R}_k and \mathbf{Q}_k over an MHE horizon, we introduce two forgetting factors $\gamma_{1,2} \in (0, 1)$ to parameterize the time-varying \mathbf{R}_k and \mathbf{Q}_k by $\mathbf{R}_k = \gamma_1^{t-k} \mathbf{R}_t$ and $\mathbf{Q}_k = \gamma_2^{t-1-k} \mathbf{Q}_{t-1}$. To minimize the network's size, we configure \mathbf{P} , \mathbf{R}_t , and \mathbf{Q}_{t-1} as diagonal matrices. Finally, the diagonal elements are specified as $P. = \zeta + p^2$, $R. = \zeta + r^2$, and $Q. = \zeta + q^2$, where $\zeta > 0$, to guarantee the positive definiteness. The gradient (5)

¹These inertial values are reported at the NeuroBEM's website https://rpg.ifi.uzh.ch/neuro_bem/Readme.html.

and Hessian (10) are updated accordingly to incorporate the above parameterization. Our neural network has two hidden layers with the Leaky-ReLU activation function. It takes the measurement $\mathbf{y} = [\mathbf{v}; \boldsymbol{\omega}]$ as inputs, has 8 neurons in each hidden layer, and outputs $\Theta = [p_{1:12}, \gamma_1, r_{1:6}, \gamma_2, q_{1:6}] \in \mathbb{R}^{26}$, resulting in a total of 362 network parameters.

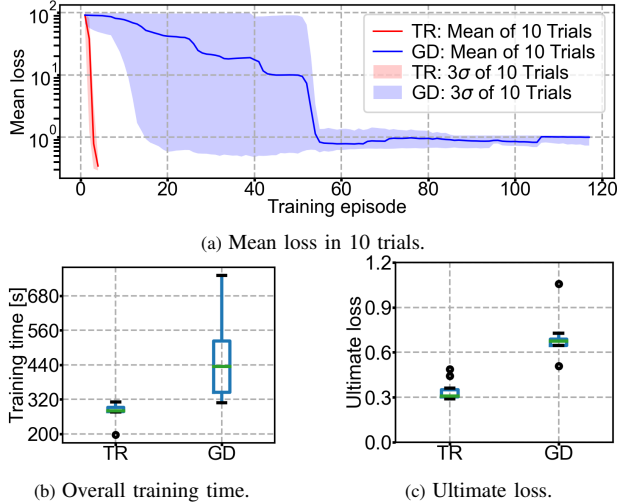


Fig. 2: Comparison of the training performance between the gradient descent (GD) and the trust-region (TR) methods. We randomly initialize the neural network in 10 trials with the Kaiming initialization, a common method for ReLU [26]. For the GD method, we set the learning rate to 1×10^{-4} by balancing between training stability and performance. We carefully select the trials of gradient descent such that the untrained mean loss closely matches that of the TR method, enabling fair comparisons. The ultimate loss shown in Fig. 2c corresponds to the loss value in the last episode.

For training, we select a 0.25-s-long trajectory segment from a Figure-8 flight collected at 400Hz (a total of 100 data points). One training episode amounts to training NeuroMHE over this 0.25-s-long dataset once. For testing, we adopt the same dataset as in [13] to compare NeuroMHE with NeuroBEM across 13 unseen agile trajectories. We implement our algorithm using CasADi [27] and conduct the simulations in a workstation with an Intel Core i7-11700K processor.

TABLE I: Runtime of Algorithm 1 for Different MHE Horizons

Horizon N	10	20	40	60	80
MHE Hessian [ms]	67.5	137.8	253.4	379.3	516.1
MHE Gradient [ms]	1.83	3.74	6.97	12.36	14.93

Before showing the estimation performance of trust-region NeuroMHE, we assess its training efficiency by comparing the CPU runtime of Algorithm 1 and the overall training time with gradient descent. Table I shows that both methods exhibit roughly linear computational complexity w.r.t the horizon, indicating scalability for large MHE problems. Computing the MHE Hessian takes longer than the MHE gradient due to the \otimes operation, which expands the matrix dimension (See (12)). Despite this, the trust-region method significantly reduces the overall training time (under 5 min) by up to 63% (See Fig. 2b) while producing better training results (See Fig. 2c). Given the quadrotor’s fast dynamics in agile flights, we set $N = 10$ in the subsequent tests.




Fig. 3: Comparison of the force estimation performance between NeuroMHE and NeuroBEM on an aggressive Figure-8 flight test dataset as used in [13].

TABLE II: RMSE Comparisons on the Figure-8 Flight Test Dataset

Method	F_{xy} [N]	F_z [N]	τ_{xy} [Nm]	τ_z [Nm]	F [N]	τ [Nm]
NeuroBEM	0.51	1.08	0.03	0.01	1.20	0.03
Ours	0.32	0.31	0.02	0.01	0.45	0.02

Fig. 3 compares the performance of trust-region NeuroMHE with NeuroBEM in estimating force on an unseen flight test dataset. Our method maintains high accuracy, even in challenging force spikes where NeuroBEM experiences substantial performance degradation. Table II presents the related quantitative comparison using Root-Mean-Square errors (RMSEs). Our approach significantly outperforms NeuroBEM in the planar and vertical force estimations, reducing the RMSEs by 37.3% and 71.3%, respectively. Further comparisons on the entire test dataset² show that our method substantially improves the overall force estimation by up to 68.1% across most trajectories. Note that this superior generalizability is achieved by training a lightweight network with 362 parameters on just 0.25-s-long data. By comparison, NeuroBEM requires 3150-s-long flight data to train a high-capacity neural network with 25k parameters.

VI. CONCLUSION AND FUTURE WORK

This paper proposed a trust-region policy optimization method for training NeuroMHE. Our critical insight is that most of the computation initially used to obtain the MHE gradient, especially the Kalman filter, can be efficiently reused for calculating the MHE Hessian in a recursive form. Through extensive simulations using real flight data, we have shown that our method provides highly-efficient training, accurate estimation with fast online adaptation to various challenging scenarios, and improved robustness to network initialization.

Our future work includes conducting a theoretical analysis on the stability of trust-region NeuroMHE, testing its efficacy in solving large scale estimation problems with various activation functions in the neural network, and simultaneously learning optimal controllers and estimators using second-order optimization techniques.

²The comparison results on the entire 13 unseen test flight trajectories are available at <https://github.com/BinghengNUS/TR-NeuroMHE>.

REFERENCES

- [1] M. L. Corradini and G. Orlando, "Robust tracking control of mobile robots in the presence of uncertainties in the dynamical model," *Journal of robotic systems*, vol. 18, no. 6, pp. 317–323, 2001.
- [2] C. Powers, D. Mellinger, A. Kushleyev, B. Kothmann, and V. Kumar, "Influence of aerodynamics and proximity effects in quadrotor flight," in *Experimental Robotics: The 13th International Symposium on Experimental Robotics*. Springer, 2013, pp. 289–302.
- [3] E. Hashemi, M. Pirani, A. Khajepour, B. Fidan, A. Kasaiezadeh, S.-K. Chen, and B. Litkouhi, "Integrated estimation structure for the tire friction forces in ground vehicles," in *2016 IEEE International Conference on Advanced Intelligent Mechatronics (AIM)*. IEEE, 2016, pp. 1657–1662.
- [4] Z. Chu, F. Wang, T. Lei, and C. Luo, "Path planning based on deep reinforcement learning for autonomous underwater vehicles under ocean current disturbance," *IEEE Transactions on Intelligent Vehicles*, vol. 8, no. 1, pp. 108–120, 2022.
- [5] P. Coelho and U. Nunes, "Path-following control of mobile robots in presence of uncertainties," *IEEE Transactions on Robotics*, vol. 21, no. 2, pp. 252–261, 2005.
- [6] B. S. Park, S. J. Yoo, J. B. Park, and Y. H. Choi, "Adaptive neural sliding mode control of nonholonomic wheeled mobile robots with model uncertainty," *IEEE Transactions on Control Systems Technology*, vol. 17, no. 1, pp. 207–214, 2008.
- [7] B. Yüksel, C. Secchi, H. H. Bühlhoff, and A. Franchi, "A nonlinear force observer for quadrotors and application to physical interactive tasks," in *2014 IEEE/ASME international conference on advanced intelligent mechatronics*. IEEE, 2014, pp. 433–440.
- [8] T. Tomić and S. Haddadin, "A unified framework for external wrench estimation, interaction control and collision reflexes for flying robots," in *2014 IEEE/RSJ international conference on intelligent robots and systems*. IEEE, 2014, pp. 4197–4204.
- [9] J. Huang, S. Ri, T. Fukuda, and Y. Wang, "A disturbance observer based sliding mode control for a class of underactuated robotic system with mismatched uncertainties," *IEEE Transactions on Automatic Control*, vol. 64, no. 6, pp. 2480–2487, 2018.
- [10] D. Hentzen, T. Stastny, R. Siegwart, and R. Brockers, "Disturbance estimation and rejection for high-precision multirotor position control," in *2019 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2019, pp. 2797–2804.
- [11] C. D. McKinnon and A. P. Schoellig, "Estimating and reacting to forces and torques resulting from common aerodynamic disturbances acting on quadrotors," *Robotics and Autonomous Systems*, vol. 123, p. 103314, 2020.
- [12] X. Liu, C. Yang, Z. Chen, M. Wang, and C.-Y. Su, "Neuro-adaptive observer based control of flexible joint robot," *Neurocomputing*, vol. 275, pp. 73–82, 2018.
- [13] L. Bauersfeld, E. Kaufmann, P. Foehn, S. Sun, and D. Scaramuzza, "Neurobem: Hybrid aerodynamic quadrotor model," *Proceedings of Robotics: Science and Systems XVII*, p. 42, 2021.
- [14] G. Shi, W. Hönig, X. Shi, Y. Yue, and S.-J. Chung, "Neural-swarm2: Planning and control of heterogeneous multirotor swarms using learned interactions," *IEEE Transactions on Robotics*, vol. 38, no. 2, pp. 1063–1079, 2021.
- [15] B. Wang, Z. Ma, S. Lai, and L. Zhao, "Neural moving horizon estimation for robust flight control," *IEEE Transactions on Robotics*, vol. 40, pp. 639–659, 2024.
- [16] T. Kraus, H. J. Ferreau, E. Kayacan, H. Ramon, J. De Baerdemaeker, M. Diehl, and W. Saeys, "Moving horizon estimation and nonlinear model predictive control for autonomous agricultural vehicles," *Computers and electronics in agriculture*, vol. 98, pp. 25–33, 2013.
- [17] B. Wang, Z. Ma, S. Lai, L. Zhao, and T. H. Lee, "Differentiable moving horizon estimation for robust flight control," in *2021 60th IEEE Conference on Decision and Control (CDC)*. IEEE, 2021, pp. 3563–3568.
- [18] A. Papadimitriou, H. Jafari, S. S. Mansouri, and G. Nikolakopoulos, "External force estimation and disturbance rejection for micro aerial vehicles," *Expert Systems with Applications*, vol. 200, p. 116883, 2022.
- [19] D. G. Robertson, J. H. Lee, and J. B. Rawlings, "A moving horizon-based approach for least-squares estimation," *AIChE Journal*, vol. 42, no. 8, pp. 2209–2224, 1996.
- [20] H. N. Esfahani, A. B. Kordabad, and S. Gros, "Reinforcement learning based on mpc/mhe for unmodeled and partially observable dynamics," in *2021 American Control Conference (ACC)*. IEEE, 2021, pp. 2121–2126.
- [21] S. Muntwiler, K. P. Wabersich, and M. N. Zeilinger, "Learning-based moving horizon estimation through differentiable convex optimization layers," in *Learning for Dynamics and Control Conference*. PMLR, 2022, pp. 153–165.
- [22] J. Schulman, S. Levine, P. Abbeel, M. Jordan, and P. Moritz, "Trust region policy optimization," in *International conference on machine learning*. PMLR, 2015, pp. 1889–1897.
- [23] K. Cao and L. Xie, "Trust-region inverse reinforcement learning," *IEEE Transactions on Automatic Control*, 2023.
- [24] A. Alessandri, M. Baglietto, G. Battistelli, and V. Zavala, "Advances in moving horizon estimation for nonlinear systems," in *49th IEEE Conference on Decision and Control (CDC)*. IEEE, 2010, pp. 5681–5688.
- [25] J. Nocedal and S. J. Wright, *Numerical optimization*. Springer, 1999.
- [26] K. He, X. Zhang, S. Ren, and J. Sun, "Delving deep into rectifiers: Surpassing human-level performance on imagenet classification," in *Proceedings of the IEEE international conference on computer vision*, 2015, pp. 1026–1034.
- [27] J. A. Andersson, J. Gillis, G. Horn, J. B. Rawlings, and M. Diehl, "Casadi: a software framework for nonlinear optimization and optimal control," *Mathematical Programming Computation*, vol. 11, no. 1, pp. 1–36, 2019.