

HIERARCHICAL DECISION MAKING WITH STRUCTURED POLICIES: A PRINCIPLED DESIGN VIA INVERSE OPTIMIZATION

Anonymous authors

Paper under double-blind review

ABSTRACT

Hierarchical decision-making frameworks are pivotal for addressing complex control tasks, enabling agents to decompose intricate problems into manageable sub-goals. Despite their promise, existing hierarchical policies face critical limitations: (i) reinforcement learning (RL)-based methods struggle to guarantee strict constraint satisfaction, and (ii) optimization-based approaches often rely on myopic and computationally prohibitive formulations. To reconcile these trade-offs, hierarchical RL-OC architectures have emerged as a promising paradigm. However, the formulation of the lower-level optimization within these frameworks remains underexplored, often relying on heuristic or myopic objectives. In this work, we propose a principled framework that systematically integrates upper-level goal abstraction with structured lower-level decision making. We adopt an inverse optimization approach to inform the structure of the lower-level problem from expert demonstrations, ensuring that the objective of the lower-level policy remains aligned with the overall long-term task goal. To validate the approach, our framework is evaluated on distinct decision making tasks: network-based resource allocation and continuous collision avoidance. Empirical results demonstrate that our method outperforms Learning-based OC, End-to-end RL and Representative Hierarchical RL in both efficiency and decision quality.

1 INTRODUCTION

Real-time decision-making in cyber-physical systems, such as robotics, autonomous driving, power grids, and transportation (Jendoubi & Bouffard, 2023; Zhou et al., 2024; Liang et al., 2025), is inherently challenging due to high-dimensional state and action spaces, nonlinear dynamics, and complex physical constraints. Existing solutions largely stem from optimal control (OC) and deep reinforcement learning (RL). OC-based methods aim to optimize system performance over infinite or long horizons while ensuring stability and feasibility. These methods are well-suited for safety-critical systems due to their theoretical guarantees, but may scale poorly in high-dimensional or nonlinear settings and often require accurate models. In contrast, RL-based approaches directly learn a policy from interactions with the environment, which scale well to complex tasks and do not require explicit modeling of complex system dynamics. Nevertheless, these approaches require extensive training and lack safety or constraint satisfaction guarantees due to their black-box nature. These trade-offs have motivated growing interest in combining OC-based and RL-based methods to exploit the strengths of both paradigms.

A promising approach for combining OC- and RL-based methods is through a hierarchical architecture that decomposes decision-making into two sequential subproblems (Lew et al., 2023; Karnchanachari et al., 2020). The upper-level employs an RL policy for strategic planning, such as generating subgoals, while the lower level uses OC to ensure safe and feasible execution. This hierarchical architecture not only enhances scalability and feasibility but also aligns with human cognition, as humans tend to perform abstract planning guided by intrinsic motivation, grounded by fast, lower-level execution.

Despite the promise of hierarchical RL-OC frameworks, the formulation of the lower-level optimization problem remains underexplored. The lower-level controller must be both computationally efficient and aligned with the upper-level goals, since a poorly designed formulation may inad-

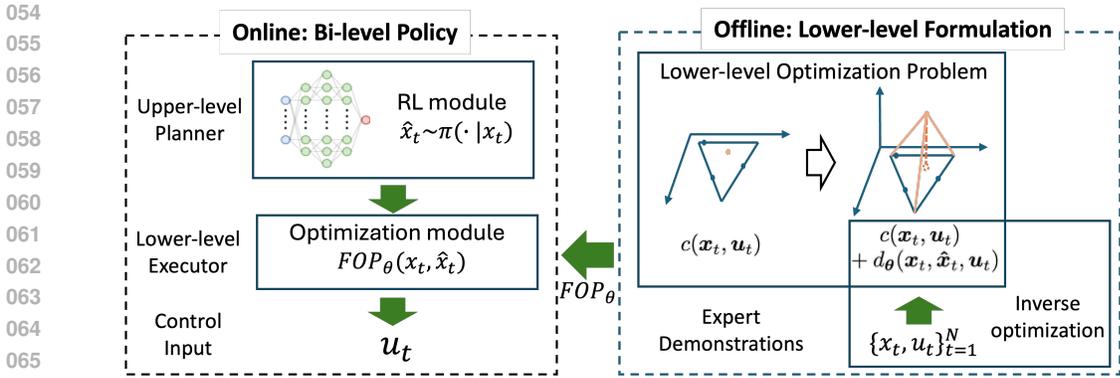


Figure 1: We propose an RL-OC hierarchical decision-making framework with lower-level policy informed by inverse optimization.

vertently exclude high-quality solutions. Existing approaches have several limitations. First, most hierarchical methods adopt long-horizon OC formulations at the lower level to preserve stability and feasibility guarantees (Song & Scaramuzza; Cheng et al., 2024; Landgraf et al., 2022). However, such long-horizon MPC formulations are well known to incur prohibitive computational complexity for real-time applications (Karamanakos et al., 2014; Krishnamoorthy et al., 2020). Second, recent efforts leverage single-step OC to enhance computation efficiency by directly generating the value or constraints on the desired next state (Gammelli et al., 2023; Schmidt et al., 2024). Nevertheless, these formulations typically rely on myopic objectives without an appropriately designed formulation. As established in classical optimal control theory, short horizons without accurate terminal value approximations cause solutions to overlook longer-term impacts, potentially leading to sub-optimal trajectories (Rawlings et al., 2020; Lowrey et al., 2018). The above challenges highlight the importance of formulating the lower-level optimization problem in a way that reduces sub-optimality while simultaneously ensuring tractable computational complexity.

To address these limitations, we propose an inverse optimization framework to guide the design of the lower-level policy from a handful of demonstrations, which provide valuable insights (Figure 1). We cast the design of the lower-level cost function as an inverse optimization problem. For a special class of lower-optimization problems with linear cost functions, we provide a theoretical characterization of the conditions under which the expert demonstrations are optimal. Once the formulation is established, we employ efficient methods to solve the inverse problem, determining the exact mathematical formulation that best fits the observed data. To validate the effectiveness of our approach, we demonstrate our method on the problem of autonomous vehicle rebalancing, supply chain inventory management and mobile robot navigation. The improvements in learning the formulation are validated from multiple perspectives. In light of the above discussion, we summarize the main contributions of this work as follows:

- We propose an inverse optimization-based approach to systematically design the lower-level control policy within a hierarchical RL-OC framework.
- We provide theoretical analysis for a special class of problems with broad applications, proposing a tractable cost structure and efficient inverse optimization formulation which ensures inverse-feasibility, forward-stability, and computational tractability.
- We demonstrate the effectiveness of the proposed framework on several scenarios from different fields, showcasing its practical relevance and potential impact in real-world applications.

2 RELATED WORK

This work is related to the literature on hierarchical structured control policies. Depending on whether learning-based or model-based approaches are used at each level, existing works can be broadly classified into two categories: (i) hierarchical reinforcement learning (HRL) that employs RL at each level, and (ii) learning-based optimal control, e.g., frameworks integrating RL and MPC,

whereby a upper-level policy learns desired states or goals, and a simplified lower-level MPC ensures safe and feasible execution.

Hierarchical Reinforcement Learning HRL decomposes a complex, difficult-to-solve problem into multiple simpler, smaller problems by setting subgoals (Kulkarni et al., 2016; Vezhnevets et al., 2017; Ma et al., 2021; Xie et al., 2021; Eppe et al., 2022; Qi et al., 2022; Huang et al., 2022; Gu et al., 2023; Mao et al., 2024; Luo et al., 2024; Zhang et al., 2024b; Hirt et al., 2024). We focus on how hierarchical policies utilize various forms of intrinsic motivation by setting subgoals. Naveed et al. (2021) develop a hierarchical reinforcement learning framework for autonomous vehicle trajectory planning, where a upper-level policy selects maneuver options, and a lower-level planner generates waypoints accordingly. Vezhnevets et al. (2017) applies Feudal Networks for hierarchical reinforcement learning. The manager module sets abstract goals that are conveyed to and enacted by the Worker module. Another common way to set goals is to consider the desired states. In Nachum et al. (2018), the author sets the upper-level actions to be goal states and reward the lower-level policy for performing actions that yield an observation close to matching the desired goal. In recent years, a growing body of work has investigated how to define subgoals and how to search efficiently within the subgoal space (Liu et al., 2021; Ma et al., 2023). Nevertheless, these studies have paid limited attention to the relationship between subgoals and the agent’s final actions. Moreover, as previously discussed, incorporating constraints to ensure safety strictly is challenging in reinforcement learning framework. In next section, we focus on how learning, especially reinforcement learning, interplays with optimization in the previous literature.

Learning-based OC. In the control community, various real-world control problems are solved by using learning-based OC. In many existing works, learning-based methods are often applied to learn cost functions or system dynamics (Lenz et al., 2015; Coulson et al., 2019; Hewing et al., 2020; Dogan et al., 2023; Zhang et al., 2024a; Lu et al., 2024; Zhang et al., 2024a; Dinkla et al., 2026). However, solving optimal control problems in real time still poses challenges when long control horizons are used, due to the high dimensionality of variables and the complexity of constraints. Prior work has attempted to reduce the computational burden by approximating the long-horizon MPC problem with a single-step formulation and learning terminal cost to alleviate myopic behavior (Abdufattokhov et al., 2021; Alsmeyer et al., 2024). However, the learning paradigm of these methods typically optimizes a surrogate objective (e.g., immediate fit to reference) rather than the final cumulative rewards directly (Zhang et al., 2024a). Instead, Gammelli et al. (2023) propose to leverage reinforcement learning to learn upper-level actions to shorten the control horizon of network flow control problem. In this framework, reinforcement learning is applied to provide reference trajectories that guide the lower-level executor toward maximizing cumulative rewards. However, the lower-level optimization problem still hinges on myopic objectives and does not rely on an explicitly and appropriately designed formulation. In addition, Schmidt et al. (2024) study hierarchical RL-OC in an offline setting and focus on generating upper-level subgoals. Their framework assumes that the offline data are optimal for the lower-level control problem, which can be sometimes unrealistic and leaves the design of the lower-level OC formulation unaddressed.

Related to this line of research, our work investigates how to formulate the lower-level optimization problem in hierarchical RL-OC framework to alleviate sub-optimality issues stemming from structural myopia.

3 METHODOLOGY

3.1 PROBLEM SETTING AND PRELIMINARY

Let us consider a general multi-step decision-making problem formulated in (1).

$$\begin{aligned}
 & \min_{\{\mathbf{u}_t\}_{t=0}^{\infty}} \limsup_{T \rightarrow \infty} \frac{1}{T} \sum_{t=0}^{T-1} c(\mathbf{x}_t, \mathbf{u}_t) \\
 & \text{s.t. } \quad \mathbf{x}_{t+1} = f(\mathbf{x}_t) + g(\mathbf{x}_t)\mathbf{u}_t, \quad \forall t \geq 0 \\
 & \quad \quad \mathbf{x}_t \in \mathcal{X}_t, \quad \forall t \geq 0 \\
 & \quad \quad \mathbf{u}_t \in \mathcal{U}_t, \quad \forall t \geq 0
 \end{aligned} \tag{1}$$

where $\mathbf{x}_t \in \mathbb{R}^n$ is the system state at time step t , $\mathbf{u}_t \in \mathbb{R}^m$ is the control input, and $c(\mathbf{x}_t, \mathbf{u}_t)$ is the stage cost function assumed to be convex. The system dynamics are assumed to be control-affine (i.e., linear in \mathbf{u}_t) with functions $f(\cdot)$ and $g(\cdot)$, which cover a wide range of applications in cyber-physical systems. The feasible sets \mathcal{X}_t and \mathcal{U}_t encode admissible states and inputs, respectively, which are assumed to be convex. The initial state x_0 is known.

Due to the infinite horizon and the presence of complex constraints, solving Problem (1) is often computationally intractable. A standard workaround is a finite-horizon approximation. However, OC problems with long time horizons can still be computationally challenging for large-scale systems, which do not satisfy the real-time requirements of practical applications. To address this issue, we generalize the bi-level decision-making framework proposed by Gammelli et al. (2023) to a more general problem setting.

3.2 BI-LEVEL FRAMEWORK

The hierarchical RL-OC framework is implemented as an end-to-end system shown in Problem (2).

$$\begin{aligned} \pi^* \in \arg \max_{\pi \in \Pi} \mathbb{E}_\tau \left[\sum_{t=0}^{\infty} \gamma^t c(\mathbf{u}_t, \mathbf{x}_t) \right] \\ \text{s.t. } h_t \sim \pi(h_t | \mathbf{x}_t) \\ \mathbf{u}_t = \text{FOP}(h_t, \mathbf{x}_t) \end{aligned} \quad (2)$$

The overall policy π^* composes an upper-level RL policy π and the solution to a lower-level optimization problem FOP. The upper-level policy encodes task-relevant abstract information or goals to produce the intrinsic subgoal h_t . The lower-level optimization module receives both the intrinsic subgoal h_t and the current state \mathbf{x}_t to compute the control input \mathbf{u}_t that is feasible.

Remark 1 (Practical requirements for RL-OC frameworks). There are two requirements for Problem (2). First, for real-time deployment, FOP should capture operational constraints and be computationally efficient to solve. Second, to enable stable training, the intrinsic subgoal h_t is preferably low-dimensional.

Following Remark 1, a commonly used form of subgoals is a linear transformation of the desired next state $\mathbf{x}_t^{\text{des}}$ using a known matrix C , denoted by $\hat{\mathbf{x}}_t = C\mathbf{x}_t^{\text{des}}$ (Gammelli et al., 2023; Schmidt et al., 2024). This subgoal guides the system toward the desired state by inducing appropriate lower-level actions. Here, in scenarios with high state dimensions, the transformation matrix C serves to compress the action space of the upper-level RL by mapping the high-dimensional system state to a lower-dimensional planning space. The transformation matrix C is typically set as an identity matrix in scenarios with low state dimensions. Overall, the bi-level decision-making framework is given in (3) and (4).

$$\begin{aligned} \pi^* \in \arg \max_{\pi \in \Pi} \mathbb{E}_\tau \left[\sum_{t=0}^{\infty} \gamma^t c(\mathbf{u}_t, \mathbf{x}_t) \right] \\ \text{s.t. } \hat{\mathbf{x}}_t \sim \pi(\hat{\mathbf{x}}_t | \mathbf{x}_t) \\ \mathbf{u}_t \in \text{FOP}(\mathbf{x}_t, \hat{\mathbf{x}}_t) \end{aligned} \quad (3)$$

$$\begin{aligned} \text{FOP}(\mathbf{x}_t, \hat{\mathbf{x}}_t) := \arg \min_{\mathbf{u}_t} c(\mathbf{x}_t, \mathbf{u}_t) \\ \text{s.t. } Cf(\mathbf{x}_t) + Cg(\mathbf{x}_t)\mathbf{u}_t = \hat{\mathbf{x}}_t, \\ f(\mathbf{x}_t) + g(\mathbf{x}_t)\mathbf{u}_t \in \mathcal{X}_{t+1}, \mathbf{u}_t \in \mathcal{U}_t \end{aligned} \quad (4)$$

By leveraging intrinsic subgoals learned by the upper-level policy, the lower-level executor is able to make near-instantaneous, subgoal-conditioned decisions, which is particularly advantageous in time-sensitive or high-dimensional environments. However, the formulation of FOP still remains ambiguous in the bi-level framework, as discussed in Remark 2.

Remark 2 (Importance of proper lower-level formulations). We highlight two considerations for Problem (4). First, the transformed desired next state $\hat{\mathbf{x}}_t$ produced by the upper-level RL may be infeasible in practice, e.g., planned trajectories in robotics tasks, which require the design of a cost function to penalize violations. Such a cost function must be aligned with the overarching objective

of the decision-making problem to avoid sub-optimality. Second, even when \hat{x} is feasible, the formulation of the lower-level optimization still requires careful design to mitigate myopic decisions, particularly when the transformation matrix C reduces the state dimension. Specifically, there may exist multiple actions \mathbf{u}_t that satisfy the constraint $Cf(\mathbf{x}_t) + Cg(\mathbf{x}_t)\mathbf{u}_t = \hat{\mathbf{x}}_t$ but yield different next states under the true dynamics, which will influence future rewards. Therefore, it is important to properly design the lower-level problem to align the action selection with the overarching objective.

In the next subsection, we will present an inverse optimization approach to inform the design of the lower-level optimization problem.

3.3 LEARNING PROBLEM FORMULATION: INVERSE OPTIMIZATION

3.3.1 GENERAL FRAMEWORK

We assume access to a set of expert data $\{\mathbf{x}_t, \mathbf{u}_t\}_{t \in \mathcal{T}_e}$ generated by an existing decision-making policy, such as Model Predictive Control (MPC) or other high-quality heuristics ensuring strict safety guarantees, where \mathcal{T}_e represents the set of time steps at which the dataset is collected. We make three remarks regarding the expert dataset. First, as high-quality solution methods may not be suitable for making real-time decisions in practice, the expert data can be derived offline rather than from real-world operations. Second, for the same reason, the dataset is assumed to be small, which makes supervised learning methods such as imitation learning less suitable. Third, \mathcal{T}_e is chosen to cover diverse operating conditions and does not have to consist of consecutive time steps.

The core of our approach is to recover latent structure encoded in the expert data. Specifically, we parameterize the lower-level optimization problem as

$$\text{FOP}_{\theta}(\mathbf{x}_t, \hat{\mathbf{x}}_t) := \arg \min_{\mathbf{u}_t} \{c(\mathbf{x}_t, \mathbf{u}_t) + d_{\theta}(\mathbf{x}_t, \hat{\mathbf{x}}_t, \mathbf{u}_t) \mid \mathbf{a}_t + B_t \mathbf{u}_t = \hat{\mathbf{x}}_t, \mathbf{u}_t \in \hat{\mathcal{U}}_t\} \quad (5)$$

where $\mathbf{a}_t := Cf(\mathbf{x}_t)$, $B_t := Cg(\mathbf{x}_t)$, and $\hat{\mathcal{U}}_t := \mathcal{U}_t \cap \{\mathbf{u}_t \mid f(\mathbf{x}_t) + g(\mathbf{x}_t)\mathbf{u}_t \in \mathcal{X}_{t+1}\}$. The term $d_{\theta}(\cdot)$ parameterized by parameter θ is assumed to be convex and will be properly designed to align the lower-level optimization problem to the objective of the overarching decision-making problem.

We then convert the design of the lower-level optimization problem into estimating θ from the expert dataset. Specifically, we aim to find θ via inverse optimization such that for each expert pair $(\mathbf{x}_t, \mathbf{u}_t)$, the action \mathbf{u}_t is (approximately) optimal for $\text{FOP}_{\theta}(\mathbf{x}_t, \hat{\mathbf{x}}_t)$ under some subgoal $\hat{\mathbf{x}}_t$. Specifically, let $\mathcal{U}^{\text{opt}}(\theta, \mathbf{x}_t, \hat{\mathbf{x}}_t)$ denote the optimal solution set corresponding to parameter θ , state \mathbf{x}_t , and subgoal $\hat{\mathbf{x}}_t$. In practice, since the lower-level optimization is convex, $\mathcal{U}^{\text{opt}}(\theta, \mathbf{x}_t, \hat{\mathbf{x}}_t)$ can be characterized by the Karush–Kuhn–Tucker (KKT) conditions. Then, the inverse optimization problem can be formulated as

$$\min_{\theta} \left\{ \kappa h(\theta) + \frac{1}{|\mathcal{T}_e|} \sum_{t \in \mathcal{T}_e} \ell(\mathbf{u}_t, \mathcal{U}^{\text{opt}}(\theta, \mathbf{x}_t, \hat{\mathbf{x}}_t)) \mid \theta \in \Theta \right\}. \quad (6)$$

where the objective function is a weighted combination of two parts: (i) the sum of losses with each loss $\ell(\mathbf{u}_t, \mathcal{U}_t^{\text{opt}}(\theta))$ indicating the deviation of expert action \mathbf{u}_t from the optimal solution set $\mathcal{U}_t^{\text{opt}}(\theta)$, i.e., the sub-optimality of the expert action, and (ii) a user-defined, application-specific regularization term $h(\theta)$ representing prior information or user preference regarding θ .

Remark 3 (Inverse optimization for lower-level formulations). We make three remarks regarding the inverse optimization framework. First, we focus on learning the objective function of FOP_{θ} rather than its constraints, as constraints are typically dictated by physical requirements. Second, without loss of generality, we assume the constraint $\mathbf{a}_t + B_t \mathbf{u}_t = \hat{\mathbf{x}}_t$ is feasible. If this is not the case, we can relax it as $\mathbf{a}_t + B_t \mathbf{u}_t = \hat{\mathbf{x}}_t + \epsilon_t$, and augment the action, corresponding matrix, and cost function as $\tilde{\mathbf{u}}_t = [\mathbf{u}_t; \epsilon_t]$, $\tilde{B}_t = [B_t, I]$, and $d_{\theta}(\mathbf{x}_t, \hat{\mathbf{x}}_t, \tilde{\mathbf{u}}_t)$, respectively, which yields a lower-level optimizer of the same form. Third, we focus on common cases where the subgoal $\hat{\mathbf{x}}_t$ can be retrieved from expert data, such as the desired next state or state representations as mentioned above. For other cases where subgoals cannot be directly obtained, they can be introduced as additional latent variables $\{\hat{\mathbf{x}}_t\}_{t \in \mathcal{T}_e}$ and estimated jointly with θ .

Without loss of generality, we assume $\tilde{\mathcal{U}}_t$ is a polytope represented by $\tilde{\mathcal{U}}_t = \{u \mid H_t u \leq \mathbf{b}\}$. Under this assumption, the learning problem for θ can be formulated as

$$\min_{\theta, \{\mathbf{w}_t, \boldsymbol{\lambda}_t, \boldsymbol{\epsilon}_t\}_{t \in \mathcal{T}_e}} \sum_{t \in \mathcal{T}_e} \|\boldsymbol{\epsilon}_t\|_2^2 + \kappa h(\boldsymbol{\theta}) \quad (7a)$$

$$\text{s.t. } \mathbf{0} \in \partial_{\mathbf{u}} (c(\mathbf{x}_t, \mathbf{u}_t) + d_{\theta}(\mathbf{x}_t, \hat{\mathbf{x}}_t, \mathbf{u}_t)) + \boldsymbol{\lambda}_t^T H_t + \mathbf{w}_t^T B_t, \forall t \quad (7b)$$

$$\mathbf{a}_t + B_t \mathbf{u}_t = \hat{\mathbf{x}}_t, \quad H_t \mathbf{u}_t \leq \mathbf{b}, \forall t \quad (7c)$$

$$\boldsymbol{\lambda}_t \geq 0, \forall t \quad (7d)$$

$$\varphi(\mathbf{u}_t, \boldsymbol{\lambda}_t, \mathbf{w}_t) + \boldsymbol{\epsilon}_t = c(\mathbf{x}_t, \mathbf{u}_t) + d_{\theta}(\mathbf{x}_t, \hat{\mathbf{x}}_t, \mathbf{u}_t), \forall t \quad (7e)$$

where $\partial_{\mathbf{u}}$ denotes the subdifferential of the cost function of FOP_{θ} , which generalizes the gradient to allow for non-smooth objectives. The vectors \mathbf{w}_t and $\boldsymbol{\lambda}_t$ are dual variables. The objective function in (7a) follows the structure of Problem (6), penalizing the sum of squared duality gaps while regularizing θ via $h(\cdot)$. Constraints (7b)-(7e) are the KKT conditions, where (7e) relaxes the strong duality condition with a duality gap $\boldsymbol{\epsilon}_t$.

Concretely, we minimize the duality gap in Problem (7) instead of the distance between given expert decisions and optimal solutions to the forward optimization problem. This aims to reduce the computational time when Problem (7) is non-convex, which will be analyzed in more detail in the following sections. Once the preliminary formulation of FOP_{θ} is established and the corresponding inverse optimization problem is formulated, we leverage the expert data to infer the unknown parameters within the model, with the objective of recovering a parameterization under which the observed decisions are (approximately) optimal.

3.3.2 SPECIAL CASE WITH THEORETICAL ANALYSIS

We next consider a special class of decision-making problems with a linear stage cost function $\mathbf{c}^T \mathbf{u}_t$ and state-independent function $g(\cdot)$ (i.e., $B_t = B, \forall t$). We show that under this setting, our framework can achieve desirable properties such as inverse feasibility, forward stability, and computation efficiency. At the end of this section, we generalize our conclusions to a broader class of problems with quadratic stage cost functions.

We make the following remarks on this special setting. First, this setting captures a broad range of applications in resource allocation, logistics, and energy systems, where strict constraint satisfaction and fast computation are critical. Second, despite the linear cost structure, the original multi-stage decision-making problem (1) can still be challenging to solve in real-world operations due to (i) the potentially nonlinear structures of $f(\cdot)$ and $g(\cdot)$ and (ii) potentially high state and action dimensions. Therefore, a hierarchical decision-making problem is still required for computational efficiency.

In this special setting, we aim to derive an FOP formulation such that the inverse optimization is always feasible. This is important, as it can allow us to utilize any type of expert data. In this case, we formulate the estimation of θ as

$$\min_{\theta} \{ \varphi(\theta) \mid \theta \in \boldsymbol{\theta}_t^{\text{inv}}(\mathbf{u}_t) \forall t \in \mathcal{T}_e, \theta \in \Theta \}. \quad (8)$$

where $\boldsymbol{\theta}_t^{\text{inv}}(\mathbf{u}_t) := \{ \theta \mid \mathbf{u}_t \in \mathcal{U}_t^{\text{opt}}(\theta) \}$ is the inverse feasible set, i.e., the set of parameter values under which the expert action \mathbf{u}_t belongs to the optimal solution set $\mathcal{U}_t^{\text{opt}}(\theta)$ of the lower-level problem $FOP_{\theta}(\mathbf{x}_t, \hat{\mathbf{x}}_t)$. The function $\varphi(\theta)$ denotes a user-defined, application-specific objective that resolves indeterminacy by selecting among multiple feasible parameters.

As stated in Proposition 1, Problem (8) can be infeasible without a properly designed cost structure $d_{\theta}(\cdot)$. Counterexamples are constructed in Appendix A.1.

Proposition 1. *The inverse feasible set for Problem (5)*

$$\boldsymbol{\theta}_t^{\text{inv}}(\mathbf{u}_t) := \{ \theta \mid \mathbf{u}_t \in \mathcal{U}_t^{\text{opt}}(\theta) \}.$$

is not guaranteed to be non-empty, where $\mathcal{U}_t^{\text{opt}}(\theta)$ is defined by the KKT conditions of Problem (5) and θ denotes unknown parameters in the problem.

To address the sub-optimality, we propose an inverse-optimization-guided design procedure that leverages criteria such as inverse feasibility and forward stability to inform the lower-level formulation.

Stage 1: Validating optimality of expert decisions. We first construct a preliminary formulation of FOP_θ such that the historical decisions $\{\mathbf{u}_t\}_{t=1}^T$ are possible to be optimal under the observed states $\{\mathbf{x}_t\}_{t=1}^T$. We achieve this goal by validating the feasibility of the inverse problem.

We select the terminal cost $d_\theta(\cdot)$ represented by the summation of ReLU-based regularization terms, which demonstrates high efficiency despite its simplicity. Propositions 2 and 3 show the motivation to design such cost structure from both geometric and algebraic perspectives respectively.

$$\text{FOP}(\mathbf{x}_t, \hat{\mathbf{x}}_t) := \arg \min_{\mathbf{u}_t} \left\{ \mathbf{c}^\top \mathbf{u}_t + \sum_{k=1}^K \max\{\boldsymbol{\theta}_k^\top \mathbf{u}_t - \nu_k, 0\} \mid \mathbf{a}_t + B\mathbf{u}_t = \hat{\mathbf{x}}_t, \mathbf{u}_t \in \tilde{\mathcal{U}}_t \right\} \quad (9)$$

Proposition 2. *The inverse feasible set for Problem (9), $\boldsymbol{\theta}_t^{\text{inv}}(\mathbf{u}_t) := \{\boldsymbol{\theta} \mid \mathbf{u}_t \in \mathcal{U}_t^{\text{opt}}(\boldsymbol{\theta})\}$, is always non-empty, where $\mathcal{U}_t^{\text{opt}}(\boldsymbol{\theta})$ is defined by KKT conditions of Problem (9).*

Proposition 3. *By appropriately enlarging the decision space with auxiliary variables in Problem (9), any given set of expert decisions can be made optimal in the lifted space.*

Stage 2: Ensuring forward stability. Although Problem (9) has proven to make expert data optimal, i.e., the inverse feasibility is guaranteed, another critical issue in inverse optimization, named forward stability, proposed by Shahmoradi & Lee (2022), is not guaranteed in the case of a linear program. The definition of forward stability is defined as follows.

Definition 1 (Forward Instability (Shahmoradi & Lee, 2022)). *Given a set of expert observations $\hat{\mathcal{U}}$, the forward instability of an inverse solution $\hat{\boldsymbol{\theta}} \in \boldsymbol{\theta}^{\text{inv}}(\hat{\mathcal{U}})$ is defined as*

$$\max_{u \in \mathcal{U}^{\text{opt}}(\hat{\boldsymbol{\theta}})} \left\{ d(\hat{\mathcal{U}}, u) \right\}, \quad (10)$$

where $\mathcal{U}^{\text{opt}}(\hat{\boldsymbol{\theta}})$ denotes the set of forward optimal solutions corresponding to $\hat{\boldsymbol{\theta}}$. This value quantifies the worst-case distance between a forward solution u induced by $\hat{\boldsymbol{\theta}}$ and the expert data $\hat{\mathcal{U}}$, measuring how unstable the inverse solution $\hat{\boldsymbol{\theta}}$ can be.

To solve this issue, we include small quadratic terms into objective function to ensure strong convexity of objective function, such that forward stability is improved and the expert decisions are approximately optimal at the same time. Overall, we formulate the forward optimization problem as follows, which improves forward stability without sacrificing computational tractability at the same time.

$$\begin{aligned} \text{FOP}(\hat{\mathbf{x}}_t, \mathbf{x}_t) := & \arg \min_{\mathbf{u}_t} \mathbf{c}^\top \mathbf{u}_t + \sum_{k=1}^K \max\{\boldsymbol{\theta}_k^\top \mathbf{u}_t - \nu_k, 0\} + l * \|\mathbf{u}_t\|_2^2 \\ \text{s.t.} \quad & \mathbf{a}_t + B\mathbf{u}_t = \hat{\mathbf{x}}_t, t \in \mathcal{T}_e \\ & \mathbf{u}_t \in \mathcal{U}_t, t \in \mathcal{T}_e \end{aligned} \quad (11)$$

We propose the following inverse optimization formulation, where dual variables are denoted by $\mathbf{z}_{kt}, \mathbf{y}_{kt}, \mathbf{w}_t, \boldsymbol{\lambda}_t$ and τ_{kt} is the auxiliary variable. $g(\boldsymbol{\lambda}_t, \mathbf{z}_t, \mathbf{y}_t, \mathbf{w}_t)$ is the dual objective function with explicit form in this problem. To keep the formulation concise, we omit the explicit form of this function.

$$\min_{\mathbf{z}_{kt}, \mathbf{y}_{kt}, \boldsymbol{\lambda}_t, \mathbf{w}_t, \boldsymbol{\epsilon}_t, \boldsymbol{\theta}^{(k)}} \sum_{t \in \mathcal{T}_e} \|\boldsymbol{\epsilon}_t\|_2^2 + \rho \sum_k \|\boldsymbol{\theta}^{(k)}\|_2^2 \quad (12a)$$

$$\text{s.t.} \quad 1 - \mathbf{z}_{kt} - \mathbf{y}_{kt} = 0, k = 1, \dots, K, t \in \mathcal{T}_e \quad (12b)$$

$$\mathbf{c}^\top + \boldsymbol{\lambda}_t^\top H + \mathbf{w}_t^\top B + \sum_k \mathbf{z}_{kt} \boldsymbol{\theta}_k^\top + 2l\mathbf{u}_t = 0, t \in \mathcal{T}_e \quad (12c)$$

$$g(\boldsymbol{\lambda}_t, \mathbf{z}_t, \mathbf{y}_t, \mathbf{w}_t) + \boldsymbol{\epsilon}_t = \mathbf{c}^\top \mathbf{u}_t + \sum_{k=1}^K \tau_{kt} + l\mathbf{u}_t^\top \mathbf{u}_t, t \in \mathcal{T}_e \quad (12d)$$

$$\tau_{kt} \geq \boldsymbol{\theta}_k^\top \mathbf{u}_t - \nu_k, k = 1, \dots, K, t \in \mathcal{T}_e \quad (12e)$$

$$\mathbf{z}_{kt}, \mathbf{y}_{kt}, \boldsymbol{\lambda}_t, \tau_{kt} \geq 0, k = 1, \dots, K, t \in \mathcal{T}_e \quad (12f)$$

We minimize duality gap together with regularizing norm of $\boldsymbol{\theta}_k$ for robustness. This reformulation aims to reduce the computational time required to solve the non-convex problem which poses

378 significant challenges for computational tractability. Instead of minimizing distance between opti-
 379 mal solutions and given expert decisions, this formulation reduces the number of bilinear terms by
 380 avoiding introducing new variables corresponding to the optimal solutions of the forward problem
 381 11. What’s more, quadratic term is introduced to ensure that forward stability is improved. Propo-
 382 sition 4 demonstrates that the distance between the input and optimal solutions remains bounded.
 383 By introducing a sufficient number of ReLU-based terms, we can constrain the values of ϵ_0 within
 384 a small range, thereby controlling the upper bound of the distance.

385 **Proposition 4.** *Problem (12) yields optimal solutions \mathbf{u}_t^* satisfying $|f(\mathbf{u}_t^*) - f(\mathbf{u}_t)| \leq \epsilon_0$, s.t.*

$$386 \quad \|\mathbf{u}_t^* - \mathbf{u}_t\| \leq \sqrt{\frac{\epsilon_0}{l}} \quad (13)$$

387 where $\epsilon_0 = \max_t \epsilon_t^*$ and \mathbf{u}_t is the given expert decision in time step t .

389 For solving Problem (12), it can be reformulated as a Mixed Integer Program and solved using the
 390 spatial branch-and-bound algorithm within a reasonable time (Smith & Pantelides, 1999). By mini-
 391 mizing the duality gap instead of the distance between optimal solutions and given expert solutions,
 392 the number of bilinear terms is reduced from $\mathcal{O}(|\mathcal{A}|K + |\mathcal{A}|^2)$ to $\mathcal{O}(|\mathcal{A}|K)$. Thus, the number of
 393 constraints introduced by the construction of the McCormick envelope can be significantly reduced.

394 For the broader class of control problems with quadratic objective functions, we can still infer pa-
 395 rameters as shown in Problem (12). Proposition 4 can be extended as shown below.

396 **Corollary 1.** *When the objective function is given by $\mathbf{u}^T R \mathbf{u} + \mathbf{x}^T Q \mathbf{x}$ with the assumption $R \succ 0$,
 397 the upper bound will be achieved by $\sqrt{\frac{\lambda_{\min}(R)}{l}}$, where $\lambda_{\min}(R)$ denotes the minimum eigenvalue
 398 of R .*

402 4 CASE STUDY

403 In this section, we evaluate the proposed framework through three case studies from different fields:
 404 (1) autonomous vehicle rebalancing; (2) supply chain inventory management problems, and (3)
 405 mobile robot navigation. The results highlight its ability to enhance decision quality and provide
 406 interpretable solutions in dynamic environments. The details of the three environments are given in
 407 Appendix B.

408 4.1 EXPERIMENTAL SETUP

409 **Data Collection:** We generate expert demonstrations by simulating a full-horizon MPC controller
 410 over a finite horizon T . The dataset consists of state-action pairs collected offline, designed to cover
 411 diverse operating conditions for each distinct task. Crucially, we utilize only a handful of these
 412 demonstrations to infer the formulation of the lower-level optimization policy via inverse optimiza-
 413 tion

414 **RL Implementation:** We employ A2C with Graph Neural Networks for the AV and Supply Chain
 415 tasks, and Soft Actor-Critic (SAC) for the Mobile Robot task. Detailed network architectures and
 416 hyperparameters are provided in Appendix C.1.

417 **Baselines:** We compare our results to:

- 418 1. MPC: MPC that receives perfect information of the system dynamics and future evolution
 419 of task-specific environmental variables.
- 420 2. Bi-level-unchanged: Bi-level framework where lower-level optimization policy is un-
 421 changed with original single-step cost (Gammelli et al., 2023).
- 422 3. Bi-level-cvxpy: Embed optimization problem as a differential layer and solve for the opti-
 423 mal parameters using gradient-based method (Agrawal et al., 2019).
- 424 4. End-to-end RL: End-to-end reinforcement learning that trains a single model to directly
 425 map inputs to the final control actions.
- 426 5. Learning-Based Terminal Cost Approximation: Assume quadratic structure of terminal
 427 cost and estimate the value function (Abdulfattokhov et al., 2021).

4.2 MODEL EVALUATION

We compare the bi-level framework equipped with a learned lower-level policy against variant baselines. Table 1 and 2 present the system performance. Our Bi-level-learned approach improves upon the Bi-level-unchanged and Bi-level-cvxpy baseline by employing a refined lower-level formulation. And our bi-level framework significantly outperforms end-to-end architectures, revealing a substantial performance gap.

Moreover, compared with MPC, the bi-level framework shortens the planning horizon from T steps to one step, which significantly reduces decision-making time. We compare the time spent to make decisions for one step in Table 3. It shows a significant reduction in runtime which indicates that our method can significantly improve computational efficiency without substantially compromising solution quality, thereby enabling faster real-time decision-making.

These results highlight the effectiveness of optimizing the lower-level problem structure in enhancing overall system performance. Besides, scalability analysis, sensitivity analysis, interpretations of the results and discussions of methods are provided and discussed in the Appendix C.

Table 1: Performance Comparison on Autonomous Vehicles Rebalancing and Supply Chain Inventory Management.

Method	Autonomous Vehicle Rebalancing		Inventory Management	
	Reward	Served Demand	Reward	Served Demand
MPC	35725 (± 41.6)	3203 (± 3.07)	11335 (± 34.3)	1337 (± 5.61)
End-to-end RL	20989 (± 213.5)	2334.8 (± 14.9)	2764 (± 90.9)	476 (± 18.4)
Bi-level-unchanged	33406 (± 128.1)	2993 (± 8.06)	7491 (± 26.5)	778 (± 2.62)
Bi-level-cvxpy	34403 (± 108.7)	3086 (± 7.06)	9268 (± 32.1)	912 (± 2.20)
Bi-level-learned (Ours)	35019 (± 53.0)	3141 (± 6.14)	9442 (± 42.3)	930 (± 1.17)

Table 2: Performance Comparison on Mobile Robot Navigation.

Method	Travel Time (s)	Path Length (m)	Energy (J)
MPC	3.50	3.50	4.81
End-to-end RL	7.60 \pm 0.11	4.29 \pm 0.02	6.78 \pm 0.23
Bi-level-unchanged	9.14 \pm 0.08	4.35 \pm 0.02	6.20 \pm 0.13
Bi-level-learned (Ours)	4.82 \pm 0.30	4.16 \pm 0.19	2.92 \pm 1.04
Bi-level-cvxpy	*	*	*

An asterisk (*) indicates that no valid result was obtained for this method.

Table 3: Comparison of Computational Time (in seconds) Across Various Scenarios

Scenarios	Our Method	MPC
AV Rebalancing	0.025	1.44
SCIM	0.034	0.082
Mobile Robot	0.0187	0.0275

5 FUTURE WORK

There are several promising directions for future research. First, to enhance the generalizability of the proposed framework, it’s critical to make the learning procedure more generalizable. We aim to extend it to settings involving more complex system dynamics beyond the current control-affine formulation to figure out how the current framework can be adapted to different problems more efficiently. Second, as the inverse optimization component can become computationally intensive in large-scale scenarios, developing more scalable and efficient algorithms for inverse problem solving

486 remains an important avenue. Besides, the subgoal in the hierarchical framework may not be re-
487 trived from expert demonstrations which requires further analysis to extend the applicability of the
488 proposed method.

490 6 CONCLUSION

491
492 In this work, we propose a hierarchical reinforcement learning and optimization framework that
493 addresses key challenges in real-time, safety-critical decision-making scenarios. By leveraging in-
494 trinsic motivation, the upper-level policy generates subgoals that abstract planning objectives, while
495 the lower-level controller executes these subgoals through a structured optimization problem. A
496 central contribution of this work lies in the data-driven design of the lower-level optimization for-
497 mulation based on expert demonstrations. Our method is evaluated on several real-world scenarios
498 from different fields, where it demonstrates strong empirical performance. Furthermore, our model
499 exhibits improved alignment with expert decisions and offers interpretable, structured control poli-
500 cies. This work highlights promising directions for combining learning-based goal abstraction with
501 structured optimization. Our results suggest that such structured, bi-level approaches are promising
502 for scaling decision-making in dynamic and safety-critical domains.

504 REFERENCES

- 505 Shokhjaton Abdufattokhov, Mario Zanon, and Alberto Bemporad. Learning convex terminal costs
506 for complexity reduction in mpc. In *2021 60th IEEE Conference on Decision and Control (CDC)*,
507 pp. 2163–2168. IEEE, 2021.
- 508 Akshay Agrawal, Brandon Amos, Shane Barratt, Stephen Boyd, Steven Diamond, and J Zico Kolter.
509 Differentiable convex optimization layers. *Advances in neural information processing systems*,
510 32, 2019.
- 511 Hendrik Alsmeyer, Anton Savchenko, and Rolf Findeisen. Neural horizon model predictive control-
512 increasing computational efficiency with neural networks. In *2024 American Control Conference*
513 *(ACC)*, pp. 1646–1651. IEEE, 2024.
- 514 Zijun Cheng, Xianlin Zeng, Hao Fang, Gang Wang, and Lihua Dou. Hierarchical mpc-based mo-
515 tion planning for automated vehicles in parallel autonomy. *Unmanned Systems*, 12(05):927–938,
516 2024.
- 517 Jeremy Coulson, John Lygeros, and Florian Dörfler. Data-enabled predictive control: In the shallows
518 of the deepc. In *2019 18th European control conference (ECC)*, pp. 307–312. IEEE, 2019.
- 519 Rogier Dinkla, Tom Oomen, Sebastiaan Paul Mulders, and Jan-Willem van Wingerden. Closed-loop
520 data-enabled predictive control and its equivalence with closed-loop subspace predictive control.
521 *Automatica*, 183:112556, 2026.
- 522 Ilgin Dogan, Zuo-Jun Max Shen, and Anil Aswani. Regret analysis of learning-based mpc with
523 partially unknown cost function. *IEEE Transactions on Automatic Control*, 69(5):3246–3253,
524 2023.
- 525 Manfred Eppe, Christian Gumbsch, Matthias Kerzel, Phuong DH Nguyen, Martin V Butz, and
526 Stefan Wermter. Intelligent problem-solving as integrated hierarchical reinforcement learning.
527 *Nature Machine Intelligence*, 4(1):11–20, 2022.
- 528 Daniele Gammelli, James Harrison, Kaidi Yang, Marco Pavone, Filipe Rodrigues, and Francisco C
529 Pereira. Graph reinforcement learning for network control via bi-level optimization. *arXiv*
530 *preprint arXiv:2305.09129*, 2023.
- 531 Ziqing Gu, Lingping Gao, Haitong Ma, Shengbo Eben Li, Sifa Zheng, Wei Jing, and Junbo Chen.
532 Safe-state enhancement method for autonomous driving via direct hierarchical reinforcement
533 learning. *IEEE Transactions on Intelligent Transportation Systems*, 24(9):9966–9983, 2023.
- 534 Lukas Hewing, Kim P Wabersich, Marcel Menner, and Melanie N Zeilinger. Learning-based model
535 predictive control: Toward safe learning in control. *Annual Review of Control, Robotics, and*
536 *Autonomous Systems*, 3(1):269–296, 2020.

- 540 Sebastian Hirt, Maik Pfefferkorn, Ali Mesbah, and Rolf Findeisen. Stability-informed bayesian
541 optimization for mpc cost function learning. *IFAC-PapersOnLine*, 58(18):208–213, 2024.
542
- 543 Keke Huang, Ke Wei, Fanbiao Li, Chunhua Yang, and Weihua Gui. Lstm-mpc: A deep learning
544 based predictive control method for multimode process control. *IEEE Transactions on Industrial
545 Electronics*, 70(11):11544–11554, 2022.
- 546 Imen Jendoubi and François Bouffard. Multi-agent hierarchical reinforcement learning for energy
547 management. *Applied Energy*, 332:120500, 2023.
548
- 549 Petros Karamanakos, Tobias Geyer, Nikolaos Oikonomou, Frederick D Kieferndorf, and Stefanos
550 Manias. Direct model predictive control: A review of strategies that achieve long prediction
551 intervals for power electronics. *IEEE Industrial Electronics Magazine*, 8(1):32–43, 2014.
552
- 553 Napat Karnchanachari, Miguel Iglesia Valls, David Hoeller, and Marco Hutter. Practical reinforce-
554 ment learning for mpc: Learning from sparse objectives in under an hour on a real robot. In
555 *Learning for Dynamics and Control*, pp. 211–224. PMLR, 2020.
- 556 Dinesh Krishnamoorthy, Lorenz T Biegler, and Johannes Jäschke. Adaptive horizon economic non-
557 linear model predictive control. *Journal of Process Control*, 92:108–118, 2020.
558
- 559 Tejas D Kulkarni, Karthik Narasimhan, Ardavan Saeedi, and Josh Tenenbaum. Hierarchical deep
560 reinforcement learning: Integrating temporal abstraction and intrinsic motivation. *Advances in
561 neural information processing systems*, 29, 2016.
- 562 Daniel Landgraf, Andreas Völz, Georgios Kontes, Christopher Mutschler, and Knut Graichen. Hier-
563 archical learning for model predictive collision avoidance. *IFAC-PapersOnLine*, 55(20):355–360,
564 2022.
565
- 566 Ian Lenz, Ross A Knepper, and Ashutosh Saxena. Deepmpc: Learning deep latent features for
567 model predictive control. In *Robotics: Science and Systems*, volume 10, pp. 25. Rome, Italy,
568 2015.
- 569 Thomas Lew, Sumeet Singh, Mario Prats, Jeffrey Bingham, Jonathan Weisz, Benjie Holson, Xi-
570 aohan Zhang, Vikas Sindhwani, Yao Lu, Fei Xia, et al. Robotic table wiping via reinforcement
571 learning and whole-body trajectory optimization. In *2023 IEEE International Conference on
572 Robotics and Automation (ICRA)*, pp. 7184–7190. IEEE, 2023.
573
- 574 Jinhao Liang, Chaopeng Tan, Longhao Yan, Jingyuan Zhou, Guodong Yin, and Kaidi Yang.
575 Interaction-aware trajectory prediction for safe motion planning in autonomous driving: A
576 transformer-transfer learning approach. *IEEE Transactions on Intelligent Transportation Systems*,
577 2025.
- 578 Chenghao Liu, Fei Zhu, Quan Liu, and Yuchen Fu. Hierarchical reinforcement learning with auto-
579 matic sub-goal identification. *IEEE/CAA journal of automatica sinica*, 8(10):1686–1696, 2021.
580
- 581 Kendall Lowrey, Aravind Rajeswaran, Sham Kakade, Emanuel Todorov, and Igor Mordatch. Plan
582 online, learn offline: Efficient learning and exploration via model-based control. *arXiv preprint
583 arXiv:1811.01848*, 2018.
584
- 585 Yiwen Lu, Zishuo Li, Yihan Zhou, Na Li, and Yilin Mo. Mpc-inspired reinforcement learning for
586 verifiable model-free control. In *6th Annual Learning for Dynamics & Control Conference*, pp.
587 399–413. PMLR, 2024.
- 588 Yu Luo, Tianying Ji, Fuchun Sun, Huaping Liu, Jianwei Zhang, Mingxuan Jing, and Wenbing
589 Huang. Goal-conditioned hierarchical reinforcement learning with high-level model approxi-
590 mation. *IEEE Transactions on Neural Networks and Learning Systems*, 2024.
591
- 592 Haozhe Ma, Thanh Vinh Vo, and Tze-Yun Leong. Human-ai collaborative sub-goal optimization in
593 hierarchical reinforcement learning. In *Proceedings of the AAAI symposium series*, volume 1, pp.
86–89, 2023.

- 594 Yi Ma, Xiaotian Hao, Jianye Hao, Jiawen Lu, Xing Liu, Tong Xialiang, Mingxuan Yuan, Zhigang
595 Li, Jie Tang, and Zhaopeng Meng. A hierarchical reinforcement learning based optimization
596 framework for large-scale dynamic pickup and delivery problems. *Advances in neural information*
597 *processing systems*, 34:23609–23620, 2021.
- 598
- 599 Zhiqi Mao, Yang Liu, and Xiaobo Qu. Integrating big data analytics in autonomous driving: An
600 unsupervised hierarchical reinforcement learning approach. *Transportation Research Part C:*
601 *Emerging Technologies*, 162:104606, 2024.
- 602
- 603 Ofir Nachum, Shixiang Shane Gu, Honglak Lee, and Sergey Levine. Data-efficient hierarchical
604 reinforcement learning. *Advances in neural information processing systems*, 31, 2018.
- 605
- 606 Kaleb Ben Naveed, Zhiqian Qiao, and John M Dolan. Trajectory planning for autonomous vehicles
607 using hierarchical reinforcement learning. In *2021 IEEE International Intelligent Transportation*
608 *Systems Conference (ITSC)*, pp. 601–606. IEEE, 2021.
- 609
- 610 Chunyang Qi, Yiwen Zhu, Chuanxue Song, Guangfu Yan, Feng Xiao, Xu Zhang, Jingwei Cao,
611 Shixin Song, et al. Hierarchical reinforcement learning based energy management strategy for
612 hybrid electric vehicle. *Energy*, 238:121703, 2022.
- 613
- 614 James Blake Rawlings, David Q Mayne, Moritz Diehl, et al. *Model predictive control: theory,*
615 *computation, and design*, volume 2. Nob Hill Publishing Madison, WI, 2020.
- 616
- 617 Carolin Schmidt, Daniele Gammelli, James Harrison, Marco Pavone, and Filipe Rodrigues. Offline
618 hierarchical reinforcement learning via inverse optimization. *arXiv preprint arXiv:2410.07933*,
619 2024.
- 620
- 621 Zahed Shahmoradi and Taewoo Lee. Quantile inverse optimization: Improving stability in inverse
622 linear programming. *Operations research*, 70(4):2538–2562, 2022.
- 623
- 624 Edward MB Smith and Constantinos C Pantelides. A symbolic reformulation/spatial branch-and-
625 bound algorithm for the global optimisation of nonconvex minlps. *Computers & chemical engi-*
626 *neering*, 23(4-5):457–478, 1999.
- 627
- 628 Yunlong Song and Davide Scaramuzza. Learning high-level policies for model predictive control.
629 in 2020 iee. In *RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp.
630 7629–7636.
- 631
- 632 Alexander Sasha Vezhnevets, Simon Osindero, Tom Schaul, Nicolas Heess, Max Jaderberg, David
633 Silver, and Koray Kavukcuoglu. Feudal networks for hierarchical reinforcement learning. In
634 *International conference on machine learning*, pp. 3540–3549. PMLR, 2017.
- 635
- 636 Ruobing Xie, Shaoliang Zhang, Rui Wang, Feng Xia, and Leyu Lin. Hierarchical reinforcement
637 learning for integrated recommendation. In *Proceedings of the AAAI conference on artificial*
638 *intelligence*, volume 35, pp. 4521–4528, 2021.
- 639
- 640 Fawang Zhang, Jingliang Duan, Haoyuan Xu, Hao Chen, Hui Liu, Shida Nie, and Shengbo Eben
641 Li. Inverse model predictive control: Learning optimal control cost functions for mpc. *IEEE*
642 *Transactions on Industrial Informatics*, 2024a.
- 643
- 644 Zhaofan Zhang, Yanan Xiao, Lu Jiang, Dingqi Yang, Minghao Yin, and Pengyang Wang. Spatial-
645 temporal interplay in human mobility: A hierarchical reinforcement learning approach with hy-
646 pergraph representation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, vol-
647 ume 38, pp. 9396–9404, 2024b.
- 646
- 647 Jingyuan Zhou, Longhao Yan, and Kaidi Yang. Enhancing system-level safety in mixed-autonomy
platoon via safe reinforcement learning. *IEEE Transactions on Intelligent Vehicles*, 2024.

648 A PROOF OF PROPOSITIONS

649 A.1 PROPOSITION 1

650 The inverse feasible set for Problem (4)

$$651 \boldsymbol{\theta}_t^{\text{inv}}(\mathbf{u}_t) := \{\boldsymbol{\theta} \mid \mathbf{u}_t \in \mathcal{U}_t^{\text{opt}}(\boldsymbol{\theta})\}.$$

652 is not guaranteed to be non-empty, where $\mathcal{U}_t^{\text{opt}}(\boldsymbol{\theta})$ is defined by KKT optimality conditions of Problem (4).

653 *Proof.* We explore whether there exists $\boldsymbol{\theta}$ belonging to the inverse feasible set for Problem (4), which is making \mathbf{u}_t to satisfy all KKT conditions.

- 654 • KKT conditions for Problem (4):

- 655 1. Stationary Conditions:

$$656 c^T + \boldsymbol{\lambda}_t^T H + \mathbf{w}_t^T B = 0$$

- 657 2. Complementary Slackness:

$$658 \boldsymbol{\lambda}_t^T (H\mathbf{u}_t - h) = 0$$

- 659 3. Dual Feasibility and Primal Feasibility:

$$660 \boldsymbol{\lambda}_t \geq 0$$

$$661 H\mathbf{u}_t \leq h$$

662 We assume $H\mathbf{u}_t \leq h$ is not active, which means $\boldsymbol{\lambda}_t = 0$. The system $B^T \mathbf{w}_t = -c$ has no feasible solution if the vector $-c$ does not lie in the column space of B^T . For example, when the system is overdetermined and $\text{rank}(B^T)$ is larger than dimension of \mathbf{w}_t . \square

663 A.2 PROPOSITION 2

664 The inverse feasible set for Problem (9)

$$665 \boldsymbol{\theta}_t^{\text{inv}}(\mathbf{u}_t) := \{\boldsymbol{\theta} \mid \mathbf{u}_t \in \mathcal{U}_t^{\text{opt}}(\boldsymbol{\theta})\}.$$

666 is always non-empty, where $\mathcal{U}_t^{\text{opt}}(\boldsymbol{\theta})$ is defined by KKT optimality conditions of Problem (9).

667 *Proof.* We explore whether there exists $\boldsymbol{\theta}$ belonging to the inverse feasible set for Problem (9), which is making \mathbf{u}_t to satisfy all KKT conditions.

- 668 • KKT conditions for Problem 9:

- 669 1. Stationary Conditions:

$$670 c^T + \sum_{k=1}^K \mathbf{z}_{kt} \boldsymbol{\theta}_k^T + \boldsymbol{\lambda}_t^T H + \mathbf{w}_t^T B = 0, \quad \text{for } t \in \mathcal{T}_e$$

$$671 1 - \mathbf{z}_{kt} - \mathbf{y}_{kt} = 0, \quad \text{for } k = 1, \dots, K \text{ and } t \in \mathcal{T}_e$$

- 672 2. Complementary Slackness:

$$673 \mathbf{z}_{kt}(\tau_{kt} - \boldsymbol{\theta}_k^T \mathbf{u}_t + \nu_k) = 0, \quad \text{for } k = 1, \dots, K \text{ and } t \in \mathcal{T}_e$$

$$674 \mathbf{y}_{kt} \tau_{kt} = 0, \quad \text{for } k = 1, \dots, K \text{ and } t \in \mathcal{T}_e$$

$$675 \boldsymbol{\lambda}_t^T (H\mathbf{u}_t - h) = 0, \quad \text{for } t \in \mathcal{T}_e$$

- 676 3. Dual Feasibility and Primal Feasibility:

$$677 \mathbf{z}_{kt}, \mathbf{y}_{kt}, \boldsymbol{\lambda}_t \geq 0, \quad \text{for } k = 1, \dots, K \text{ and } t \in \mathcal{T}_e$$

$$678 H\mathbf{u}_t \leq h, \quad \text{for } t \in \mathcal{T}_e$$

$$679 \tau_{kt} \geq \boldsymbol{\theta}_k^T \mathbf{u}_t - \nu_k, \tau_{kt} \geq 0, \quad \text{for } k = 1, \dots, K \text{ and } t \in \mathcal{T}_e$$

702 There exist trivial feasible solutions to satisfy all KKT conditions:

703 We assume $z_{kt} = 1$, $\mathbf{y}_{kt} = 0$ and $\boldsymbol{\lambda}_t = 0$. Then the set of KKT conditions are transformed into

$$704 \quad c^T + \sum_{k=1}^K \boldsymbol{\theta}_k^T + \mathbf{w}_t^T B = 0, \quad \text{for } t \in \mathcal{T}_e \quad (14)$$

$$705 \quad \tau_{kt} = \boldsymbol{\theta}_k^T \mathbf{u}_t - \nu_k, \quad \tau_{kt} \geq 0, \quad \text{for } k = 1, \dots, K \text{ and } t \in \mathcal{T}_e$$

706 We must find $\{\boldsymbol{\theta}_k\}_{k=1}^K$ to satisfy $c^T + \sum_{k=1}^K \boldsymbol{\theta}_k^T + \mathbf{w}_t^T B = 0$ and let $\nu_k = \min_{t=1}^T \{\boldsymbol{\theta}_k^T \mathbf{u}_t\}$ \square

707 A.3 PROPOSITION 3

708 By appropriately enlarging the decision space with auxiliary variables in Problem (9), any given set of expert decisions can be made optimal in the lifted space.

709 *Proof.* Consider a collection of expert decisions $\{u_i\}_{i=1}^N$ with $u_i \in \mathbb{R}^n$. Problem (9) can be reformulated as a linear program by introducing auxiliary variables $\{t_k\}_{k=1}^K$, which effectively enlarges the decision space from \mathbb{R}^n to \mathbb{R}^{n+K} . For clarity, let us focus on the case $K = 1$. By adding the constraints

$$710 \quad t_k \geq \boldsymbol{\theta}_k^T u_i - \nu, \quad (15)$$

711 We can at least guarantee that all expert decisions $\{u_i\}_{i=1}^N$ lie on the same n -dimensional face of the feasible polytope in the augmented space by making the new added constraint active. This face can be made parallel to the hyperplane defined by the objective function by adjusting the new added constraint.

712 This situation, however, represents an extreme case. In practice, the set $\{u_i\}_{i=1}^N$ may lie within a face of dimension strictly smaller than n . To illustrate, suppose $u_i \in \mathbb{R}^2$. If all expert decisions lie on a one-dimensional face (an edge) of the feasible polytope, then they remain on a corresponding one-dimensional face in the lifted \mathbb{R}^3 space. Assume that this edge arises as the intersection of two adjacent facets, whose supporting hyperplanes have normal vectors \vec{n}_1 and \vec{n}_2 , respectively. The direction vector of the edge is then given by

$$713 \quad \vec{d}_{\text{edge}} = \vec{n}_1 \times \vec{n}_2, \quad (16)$$

714 since it is simultaneously orthogonal to both \vec{n}_1 and \vec{n}_2 .

715 For the expert decisions to be optimal, it suffices that the objective hyperplane be parallel to this edge. Equivalently, the edge direction \vec{d}_{edge} must be orthogonal to the objective normal \vec{n}_{obj} , i.e.,

$$716 \quad \vec{d}_{\text{edge}} \cdot \vec{n}_{\text{obj}} = 0. \quad (17)$$

717 This condition establishes the required relationship between the optimal geometry and the objective orientation. Thus, we can solve the equation for feasible unknown parameters in the terminal cost. \square

718 A.4 PROPOSITION 4

719 Problem (12) yields optimal solutions \mathbf{u}_t^* satisfying $|f(\mathbf{u}_t^*) - f(\mathbf{u}_t)| \leq \epsilon_0$, s.t.

$$720 \quad \|\mathbf{u}_t^* - \mathbf{u}_t\| \leq \sqrt{\frac{\epsilon_0}{l}} \quad (18)$$

721 *Proof.* First, we show that the feasible region of Problem (12) is non-empty. Since the strong duality constraint has been relaxed, the main task reduces to verifying the stationarity condition

$$722 \quad c^T + \boldsymbol{\lambda}_t^T H + \mathbf{w}_t^T B + \sum_k z_{kt} \boldsymbol{\theta}_k^T + 2l\mathbf{u}_t = 0, \quad t = 1, \dots, T. \quad (19)$$

For the i -th expert decision, suppose there are n_i active inequality constraints. Each such active constraint introduces one additional degree of freedom in the stationarity condition. Together with the N equality constraints, the total number of effective constraints across T time periods is

$$\sum_{i=1}^T n_i + NT. \quad (20)$$

Since the decision variable c lies in \mathbb{R}^M , the stationarity condition can be satisfied provided that the number of auxiliary terms is large enough to compensate for the remaining degrees of freedom. More precisely, the required number of ReLU terms is bounded above by

$$\max \left\{ 0, T - \frac{\sum_{i=1}^T n_i + NT}{M} \right\}. \quad (21)$$

This establishes the feasibility of the relaxed problem. Next, we derive the upper bound of distance.

Based on strong convexity of $f(u)$, we have

$$f(\mathbf{u}_t) \geq f(\mathbf{u}_t^*) + \nabla f(\mathbf{u}_t^*)^\top (\mathbf{u}_t - \mathbf{u}_t^*) + \frac{\mu}{2} \|\mathbf{u}_t - \mathbf{u}_t^*\|^2 \quad (22)$$

where $\nabla^2 f(\mathbf{u}_t) \geq \mu I > 0$. By stationary condition $\nabla f(\mathbf{u}_t^*) = 0$,

$$f(\mathbf{u}_t) - f(\mathbf{u}_t^*) \geq \frac{\mu}{2} \|\mathbf{u}_t - \mathbf{u}_t^*\|^2 \quad (23)$$

The dual objective function $g(\boldsymbol{\lambda}_t, \mathbf{z}_t, \mathbf{y}_t, \mathbf{w}_t) \leq f(\mathbf{u}_t^*) \leq f(\mathbf{u}_t)$ and we have $f(\mathbf{u}_t) - g(\boldsymbol{\lambda}_t, \mathbf{z}_t, \mathbf{y}_t, \mathbf{w}_t) \leq \epsilon_0$ by constraining strong duality gap, s.t.

$$f(\mathbf{u}_t) - f(\mathbf{u}_t^*) \leq \epsilon_0 \quad (24)$$

Thus, we have

$$\|\mathbf{u}_t - \mathbf{u}_t^*\|^2 \leq \epsilon_0 * \frac{2}{\mu} = \frac{\epsilon_0}{l} \quad (25)$$

□

B DETAILS OF CASE STUDIES

We consider three simulation scenarios described as follows:

Autonomous Vehicle Rebalancing. Autonomous Mobility-on-Demand (AMoD) systems are an evolving mode of transportation in which a centrally coordinated fleet of self-driving vehicles dynamically serves travel requests. In real-world systems, the effectiveness of rebalancing strategies is central to the overall system performance, with sub-optimal strategies potentially exacerbating congestion through unnecessary trips or increased passenger waiting time. The control of these systems is typically formulated as a large network optimization problem. This framework comprises two stages: (1) determining the desired distribution of idle vehicles \hat{q}_t through the use of the learned policy $\pi_\phi(\hat{q}_t | s_t)$ by reinforcement learning, (2) converting this distribution to a passenger flow g_{ij}^t and rebalancing flow f_{ij}^t by solving a linear control problem.

Supply Chain Inventory Management. In the supply chain inventory management task, we aim to determine the optimal ordering and distribution strategies across a network of interconnected warehouses and retail stores to satisfy customer demand while minimizing storage, production and transportation costs. We choose upper-level goals as (i) desired production in warehouse nodes \hat{w}_i^t and (ii) desired inventory in store nodes \hat{q}_i^t . And the lower-level optimization module determines the amount of commodities w_i^t to order in each warehouse, and the shipping flow f_{ij}^t from warehouses to stores.

Mobile Robot Navigation. In the mobile robot navigation task, a ground robot moves from a start location to a target destination while avoiding static obstacles. The control problem is typically formulated in a hierarchical framework. At the high level, a learned policy $\pi_\theta(\hat{p}_t | s_t)$ outputs intermediate waypoints, conditioned on the robot's state, the goal position, and obstacle information. At the low level, these actions are converted into executable control inputs $(v_t, \boldsymbol{\omega}_t)$ for the unicycle dynamics, ensuring feasibility under kinematic and dynamic constraints.

The detailed formulations of the lower-level optimization problems for these three environments are given as follows.

B.1 AUTONOMOUS VEHICLES REBALANCING

The lower-level optimization policy is represented in a matrix form as follows.

$$\begin{aligned}
\min_{f_t, g_t} c^T f_t - p^T g_t + \sum_k \max\{-\boldsymbol{\theta}_f^{(k)T} f_t + \boldsymbol{\theta}_g^{(k)T} g_t - \mu^{(k)}, 0\} \\
s.t. \quad q_t - A(f_t + g_t) \geq \hat{q}_t \\
q_t - G(f_t + g_t) \geq 0 \\
f_t \geq 0, g_t \geq 0 \\
g_t \leq \boldsymbol{\lambda}_t
\end{aligned} \tag{26}$$

where A is the instance matrix, $\boldsymbol{\lambda}_t$ is the travel requests in time slot t and f_t, g_t denote rebalancing flow and passenger flow respectively.

Proposition 5. *The inverse feasible set for Problem 26*

$$\boldsymbol{\theta}_t^{*inv}(\mathbf{f}_t, \mathbf{g}_t) := \{\boldsymbol{\theta} \mid \mathbf{f}_t, \mathbf{g}_t \in \mathcal{X}_t^{*opt}(\boldsymbol{\theta})\}. \tag{27}$$

is always non-empty, where $\mathcal{X}_t^{*opt}(\boldsymbol{\theta})$ is defined by KKT optimality conditions of Problem 26.

Next, we will prove this proposition and also show that the solved parameters have practical significance which has not been discussed further in the general setting.

Proof. The inverse optimization solving for unknown parameters in Problem 26 is formulated as follows

$$\begin{aligned}
\min_{\mathbf{x}_t, \mathbf{y}_t, \mathbf{w}_t, v_{1t}, v_{2t}, z_{1t}, z_{2t}, \boldsymbol{\tau}_t, \boldsymbol{\mu}} \sum_k \|(\boldsymbol{\theta}^{(k)})\|^2 \\
1 - z_{1t}^{(k)} - z_{2t}^{(k)} = 0, \quad k = 1, \dots, K, \quad t = 1, \dots, T \\
c^T + \mathbf{x}_t^T A + \mathbf{y}_t^T G - \mathbf{w}_t^T - \sum_k z_{1t}^{(k)} \boldsymbol{\theta}_f^{(k)T} + 2l_1 f_t = 0, \quad t = 1, \dots, T \\
-p^T + \mathbf{x}_t^T A + \mathbf{y}_t^T G - v_{1t}^T + v_{2t}^T + \sum_k z_{1t}^{(k)} \boldsymbol{\theta}_g^{(k)T} + 2l_2 g_t = 0, \quad t = 1, \dots, T \\
\mathbf{x}_t^T (q_t - A(f_t + g_t) - \hat{q}_t) = 0, \quad t = 1, \dots, T \\
\mathbf{y}_t^T (q_t - G(f_t + g_t)) = 0, \quad t = 1, \dots, T \\
\mathbf{w}_t^T f_t = 0, \quad t = 1, \dots, T \\
v_{1t}^T g_t = 0, \quad t = 1, \dots, T \\
v_{2t}^T (g_t - \boldsymbol{\lambda}_t) = 0, \quad t = 1, \dots, T \\
\text{(other feasibility constraints)}
\end{aligned} \tag{28}$$

Given historical data $\{f_i, g_i, q_i, \lambda_i\}_{i=1}^N$, we consider the worst-case scenario where

$$Gf_i + Gg_i \neq q_i, \quad f_i \neq 0, \quad g_i \neq 0, \quad g_i \neq \lambda_i \quad \forall i = 1, \dots, N.$$

This implies that all corresponding dual variables

$$\mathbf{y}_i = \mathbf{w}_i = v_{1i} = v_{2i} = 0 \quad \forall i = 1, \dots, N.$$

We proceed with the analysis under this assumption. Moreover, We further set

$$z_{1i}^{(k)} = 1, \quad z_{2i}^{(k)} = 0$$

864 which simplifies inverse problem corresponding to Problem 26 into the following form:
865

$$\begin{aligned}
866 & \min_{\boldsymbol{\theta}^{(k)}, \mu^{(k)}, x_i} \sum_k \|\boldsymbol{\theta}^{(k)}\|^2 \\
867 & \text{s.t.} \quad -\boldsymbol{\theta}_f^{(k)\top} f_i + \boldsymbol{\theta}_g^{(k)\top} g_i - \mu^{(k)} = \tau_i^{(k)}, \quad \forall k = 1, \dots, K, \\
868 & c^\top + x_i^\top A - \sum_k \boldsymbol{\theta}_f^{(k)\top} = 0, \\
869 & -p^\top + x_i^\top A + \sum_k \boldsymbol{\theta}_g^{(k)\top} = 0, \\
870 & \text{(other feasibility constraints).}
\end{aligned} \tag{29}$$

871 Now assume $x_i \preceq \epsilon \mathbf{1}$ for some small $\epsilon > 0$. Then we have

$$872 \quad -\epsilon \mathbf{1} \preceq x_i^\top A \preceq \epsilon \mathbf{1}.$$

873 Suppose $\mathbf{c}, \mathbf{p} \in \mathbb{R}_{++}^n$. We can select ϵ such that

$$874 \quad \min_{1 \leq j \leq M} c_j \geq \epsilon, \quad \min_{1 \leq j \leq M} p_j \geq \epsilon.$$

875 Then the constraints in 29 reduce to

$$\begin{aligned}
876 & -\boldsymbol{\theta}_f^{(k)\top} f_i + \boldsymbol{\theta}_g^{(k)\top} g_i - \mu^{(k)} \geq 0, \quad \forall k = 1, \dots, K, \\
877 & c^\top + x_i^\top A - \sum_k \boldsymbol{\theta}_f^{(k)\top} = 0, \\
878 & -p^\top + x_i^\top A + \sum_k \boldsymbol{\theta}_g^{(k)\top} = 0.
\end{aligned} \tag{30}$$

879 Assume further that $K = 1$ (for simplicity). Then the last two equalities yield:

$$880 \quad \boldsymbol{\theta}_f = \mathbf{c} + A^\top x_i, \quad \boldsymbol{\theta}_g = \mathbf{p} - A^\top x_i.$$

881 In the context of autonomous vehicle rebalancing, it is reasonable to assume that the total revenue exceeds the cost, i.e., $p^\top g_i > c^\top f_i$. Then we compute:

$$\begin{aligned}
882 & -\boldsymbol{\theta}_f^\top f_i + \boldsymbol{\theta}_g^\top g_i - \mu = -(c^\top + x_i^\top A) f_i + (p^\top - x_i^\top A) g_i - \mu \\
883 & = (p^\top g_i - c^\top f_i) - x_i^\top A (f_i + g_i) - \mu \\
884 & \geq (p^\top g_i - c^\top f_i) - \epsilon \mathbf{1}^\top (f_i + g_i) - \mu.
\end{aligned} \tag{31}$$

885 Therefore, we can choose ϵ sufficiently small and μ accordingly, such that the right-hand side of
886 31 remains non-negative. This guarantees that the constraints in 30 are satisfied, hence the problem
887 admits a feasible solution. \square

888 Problem 26 is constructed so that the historical expert decisions $\{f_i, g_i\}_{i=1}^N$ are optimal solutions
889 to the forward problem. There exist model parameters $\{\boldsymbol{\theta}_f^{(k)}, \boldsymbol{\theta}_g^{(k)}, \mu^{(k)}\}_{k=1}^K$ and desired next states
890 $\{\hat{g}_i\}_{i=1}^N$, which can be learned via reinforcement learning, such that the expert decisions become
891 approximately optimal.

B.2 SUPPLY CHAIN INVENTORY MANAGEMENT

The lower-level policy is formulated as follows

$$\min_{f_{ij}^t, \mathbf{w}_i^t} \sum_{i \in V_W} m_i^O \cdot \mathbf{w}_i^t + \sum_{(i,j) \in \mathcal{E}} m_{ij}^T \cdot f_{ij}^t + \sum_{i \in V_S} |\epsilon_i^f| + \sum_{i \in V_W} |\epsilon_i^w| + \sum_{k=1}^K \max\{\boldsymbol{\theta}_k^T \mathbf{f} - \mu_k, 0\} \quad (32a)$$

$$\text{s.t.} \quad \sum_{j \in N^-(i)} f_{ji}^t + \epsilon_i^f = \hat{q}_i^{t+1}, i \in V_S \quad (32b)$$

$$q_i^t + \sum_{j \in N^-(i)} f_{ji}^t - d_i^t \leq c_i^t, i \in V_S \quad (32c)$$

$$\sum_{j \in N^+(i)} f_{ij}^t \leq q_i^t, i \in V_W \quad (32d)$$

$$q_i^t + \mathbf{w}_i^t - \sum_{j \in N^+(i)} f_{ij}^t \leq c_i^t, i \in V_W \quad (32e)$$

$$\mathbf{w}_i^t + \epsilon_i^w = \hat{w}_i^t, i \in V_W \quad (32f)$$

$$f_{ij}^t \geq 0, (i, j) \in \mathcal{E}, \mathbf{w}_i^t \geq 0, i \in V_W \quad (32g)$$

where m^O, m^T, d, c are production cost, transportation cost, customer demand and capacity respectively, \mathbf{w}_i denotes order quantity in warehouse i and f_{ij} denotes the departing flow from warehouse i to store j , and q denotes the inventory level.

B.3 MOBILE ROBOT

$$\min_{\mathbf{u}_t = (v_t, \boldsymbol{\omega}_t)} c_p \|p_{t+1} - \hat{p}_t\|^2 + c_h \|\boldsymbol{\theta}_{t+1} - \hat{\boldsymbol{\theta}}_t\|^2 + c_u \|\mathbf{u}_t\|^2 + \sum_{k=1}^K \max\{\alpha_k \mathbf{u}_t - \beta_k, 0\} \quad (33a)$$

$$\text{s.t.} \quad x_{t+1} = x_t + T_s v_t \cos \boldsymbol{\theta}_t \quad (33b)$$

$$\mathbf{y}_{t+1} = \mathbf{y}_t + T_s v_t \sin \boldsymbol{\theta}_t \quad (33c)$$

$$\boldsymbol{\theta}_{t+1} = \boldsymbol{\theta}_t + T_s \boldsymbol{\omega}_t \quad (33d)$$

$$v_{\min} \leq v_t \leq v_{\max} \quad (33e)$$

$$\boldsymbol{\omega}_{\min} \leq \boldsymbol{\omega}_t \leq \boldsymbol{\omega}_{\max} \quad (33f)$$

$$\|p_{t+1} - p_j^{\text{obs}}\| \geq r_j + d_{\text{safe}}, \quad \forall j = 1, \dots, M \quad (33g)$$

where v_t and $\boldsymbol{\omega}_t$ denote speed and angular speed of mobile robot. $p_t = (x_t, \mathbf{y}_t)$ represents its position and \hat{p}_t is the reference waypoint generated by upper-level RL module. The objective is to reach the final goal and avoid obstacles.

C FURTHER EXPERIMENTAL RESULTS AND DISCUSSIONS

C.1 RL METHODOLOGY AND EXPERIMENTAL SETUP

- **Autonomous Vehicle Rebalancing:** For the Autonomous Vehicle Rebalancing experiment, we train an A2C-GNN agent as in (Gammelli et al., 2023) using on-policy updates with a discount factor of 0.99, a maximum of 50 decision steps per episode, and 10000 training episodes. The actor and critic share the same EdgeConv-based GNN encoder and each head is implemented as a two-layer multilayer perceptron with 256 hidden units, and both networks are optimized with the Adam optimizer with a learning rate of 5×10^{-5} . At the end of each episode we normalize the discounted returns, perform a single A2C update of actor and critic parameters with gradient-norm clipping at 5 to improve stability, and we save the checkpoint achieving the highest cumulative training reward for later evaluation.

- **Inventory Management:** We adopt a graph neural network-based Advantage Actor–Critic (A2C) method as in (Gammelli et al., 2023) to learn the control policy. Training is performed with an on-policy A2C update, using a discount factor of 0.99, a maximum episode length of 30 steps, and up to 20,000 training episodes. Both actor and critic are optimized with Adam with learning rate $5e-5$, and gradient norm clipping is applied to improve training stability.
- **Mobile Robot Navigation:** For the mobile robot, we use an off-policy Soft Actor–Critic (SAC) algorithm. SAC is trained with a replay buffer of size 10^6 , a mini-batch size of 256, Adam optimizer with learning rate 3×10^{-4} for both actor and critic, discount factor $\gamma = 0.99$, soft target network updates with $\tau = 0.005$, and automatic entropy tuning with target entropy set to $-\dim(\mathcal{A})$. Gradient updates start once the replay buffer contains more than 400 transitions, and the agent is trained for 500 episodes, periodically saving the actor network during training.

C.2 MOTIVATIONS FOR INVERSE OPTIMIZATION APPROACH

We justify our choice of the proposed inverse optimization framework from two perspectives: the structural design of the cost function and the efficacy of the parameter estimation.

C.2.1 STRUCTURAL DESIGN GUIDED BY THEORETICAL PROPERTIES

The primary motivation for adopting the inverse optimization framework lies in its ability to theoretically inform the structure of the lower-level cost function. Standard approaches often heuristically assume a fixed cost form (e.g., quadratic). We implement the method proposed by Abdulfattokhov et al. (2021) which assumes quadratic structure of terminal cost. We compare the performance on autonomous vehicle rebalancing task. As table 4 suggests, there exists a large optimality gap.

Table 4: Performance Comparison on Autonomous Vehicle (AV) Rebalancing Task.

Scenario	Method	Performance Metrics	
		Reward	Served Demand
Scenario 1	MPC	35725 (± 41.6)	3203 (± 3.07)
	Bi-level-learned (Ours)	35019 (± 53.0)	3141 (± 6.14)
	Value Approximation	24774	2395
Scenario 2	MPC	68637 (± 194)	7540 (± 14.1)
	Bi-level-learned (Ours)	61880 (± 370.3)	6619 (± 22.3)
	Value Approximation	29905	3160

C.2.2 PARAMETER ESTIMATION EFFICACY

Beyond the cost formulation itself, we evaluate Inverse Optimization as a solver against two alternative learning paradigms. Under the assumption that RL enable to offer optimal upper-level goals, we depart the learning of cost structure from the bi-level framework training and explore which method mimics the expert decisions best.

Table 5: Similarity scores between expert decisions and decisions chosen by different methods.

Method	AV Rebalancing		Inventory Management
	Reb. flow	Pax. flow	Commodity. flow
Value Approximation	0.522	0.943	0.461
cvxpylayers	0.541	0.946	0.445
Inverse Optimization	0.703	0.958	0.997

Based on Table 5, inverse optimization outperforms the other two methods to large extent. First, value function approximation method doesn’t perform well in both tasks since it cannot embed the

designed structure of lower-level policy explicitly. Second, cvxpylayers perform well on AV Rebalancing task in certain scenario, but fail in some other tasks. It's due to the following reasons. First, such gradient-based optimization method is sensitive to the differentiation. When the derivative of objective with respect to the unknown parameter is small, the algorithm will struggle to optimize the parameters. Second, the optimal distribution of parameters may be complex, and thus challenging to estimate within the scope of commonly distribution families.

C.3 INTERPRETABILITY

Interpretations of Optimal Parameters. The optimal parameters inferred by inverse optimization are visualized in the following figures, which show that different weights are assigned to different edges. For example, the district "0" of higher value motivates the repositioning of idle vehicles towards this area, particularly along paths with low travel costs. Thus, a significant penalty is applied to insufficient rebalancing flows on the corresponding routes. The yellow line in the figure, which denotes high inferred penalty weights, achieves this motivation. In scenario 2, the situation becomes more complex, but we can still notice some "important" routes and the "not important" ones are not penalized.

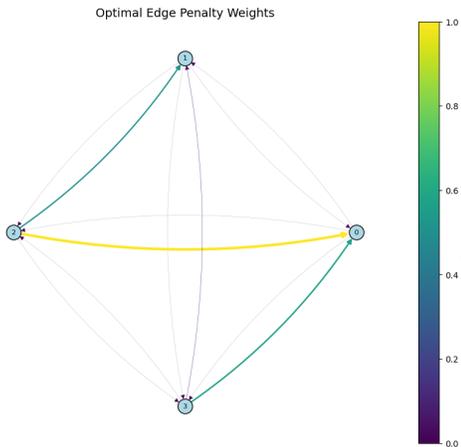


Figure 2: Optimal Weights in Scenario 1

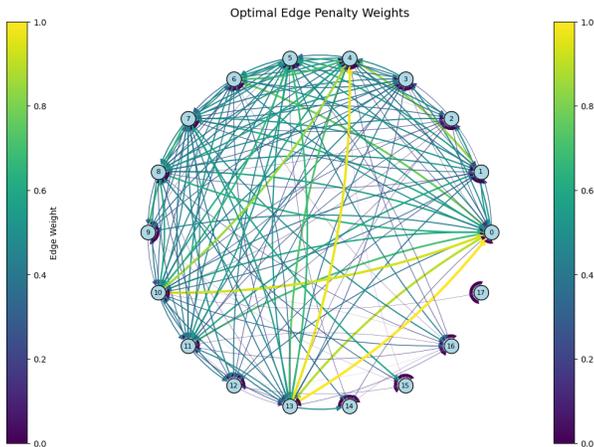


Figure 3: Optimal Weights in Scenario 2

Interpretations of Outputs. To gain deeper insights into the effectiveness of the inverse method and explain model interpretability, we analyze how decisions made by two bi-level frameworks deviate from trajectories made by MPC. Table 6 shows that Bi-level-learned framework make decisions much closer to MPC trajectories, which indicates that the learned lower-level policy based on expert demonstrations extracts information efficiently compared with the original policy with short-sighted cost functions.

Table 6: Comparison of the cosine similarity and Manhattan distance between the rebalancing flow f obtained from different bi-level models and the flow generated by the MPC.

Formulation	Cosine Similarity	Manhattan Distance
Bi-level-unchanged	0.477	44
Bi-level-learned	0.976	7

C.4 SENSITIVITY ANALYSIS

In this subsection, we analyze how the number of ReLU terms included in the cost function and the noise in expert data influence the performance of inverse optimization. We use cosine similarity and Manhattan distance between the noise-free expert decisions and decisions made by solving the lower-level optimization problem with different settings.

- Sensitivity to Number of ReLU Terms

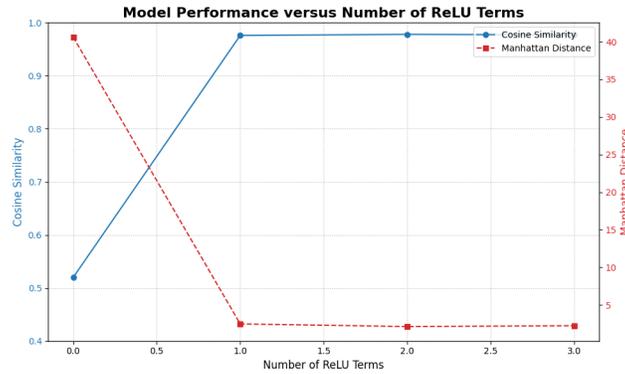


Figure 4: Sensitivity to the Number of ReLU Terms

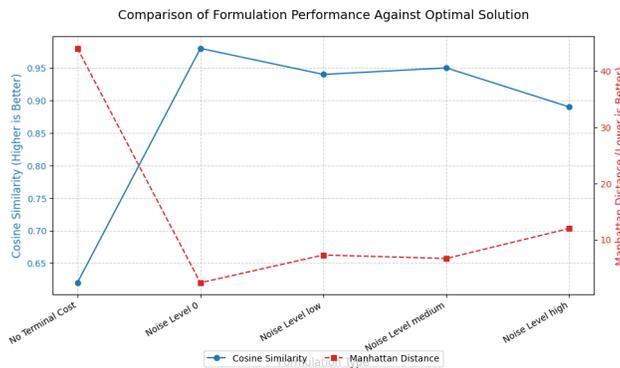


Figure 5: Sensitivity to Noise in Offline Data

Based on Figure 4, we notice that the quality of solutions doesn't vary a lot as number of ReLU terms varies. We can infer well-performed parameters via inverse optimization utilizing only a handful of expert decisions and very few ReLU terms.

- Sensitivity to Data Noise

Sometimes given demonstrations are not optimal, thus we figure out how the noise affect inverse optimization. As figure 5 shows, the solution is much closer to optimal expert demonstrations although some noise exists in offline data. Moreover, when we utilize noisy expert demonstrations to derive lower-level problem formulation and apply it in the bi-level framework, the framework improves the reward compared to the one achieved by original noisy expert demonstrations by 1.7%. Although the improvement is not so significant, it suggests a promising direction that we can employ our inverse optimization-guided bi-level framework to extract insights from sub-optimal offline data and make better decisions.

C.5 SOLVE-TIME SCALING FOR INVERSE OPTIMIZATION

We systematically vary the problem dimension, the number of ReLU terms in the objective, and the number of expert demonstrations, and plot the resulting wall-clock solve times as shown in Figure 6.

We have observed that solve-time increases predictably with problem dimensions and the number of demonstrations, remaining tractable for typical system sizes without imposing prohibitive bottlenecks. The primary driver of complexity is the number of ReLU terms (K). While $K = 1$ is computationally lightweight, increasing K introduces significant combinatorial complexity. To address this, we reformulate the inverse problem as a Mixed-Integer Program (MIP) by treating the auxiliary dual variables associated with ReLU activation as binary. This reformulation is math-

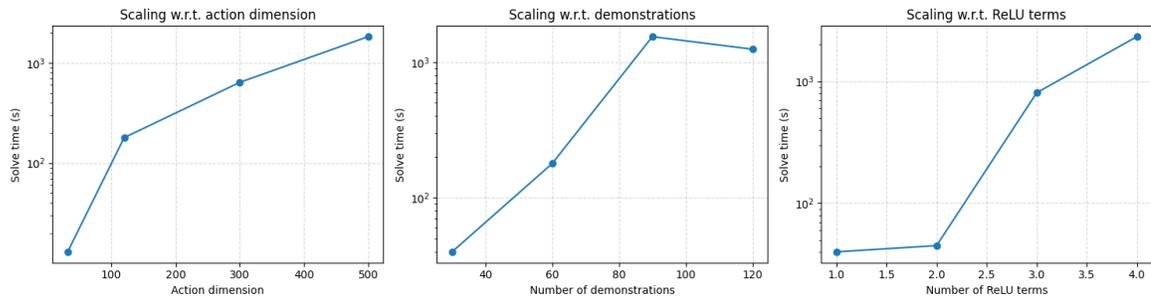


Figure 6: Solve-time Scaling Curve

ematically rigorous—aligning with the inherent discrete logic of ReLU functions—and leverages efficient branch-and-bound solvers to drastically accelerate convergence without compromising solution quality. We also note that solve times can vary even for instances of identical size, driven by the specific geometry of the expert demonstrations and the complexity of the active constraint set required to rationalize the data.

D USE OF LLMs

We utilized a large language model (LLM) to assist with improving the clarity, grammar, and overall readability of the text. The use of the LLM was strictly limited to language editing and refinement. All scientific contributions, including the core ideas, the methodological framework, the experimental design, and the mathematical derivations, are the original work of the authors.