

Unified Local-Cloud Decision-Making via Reinforcement Learning

Kathakoli Sengupta, Zhongkai Shangguan, Sandesh Bharadwaj, Sanjay Arora[†],
Eshed Ohn-Bar, and Renato Mancuso

Boston University [†]Red Hat

Abstract. Embodied vision-based real-world systems, such as mobile robots, require a careful balance between energy consumption, compute latency, and safety constraints to optimize operation across dynamic tasks and contexts. As local computation tends to be restricted, offloading the computation, i.e., to a remote server, can save local resources while providing access to high-quality predictions from powerful and large models. However, the resulting communication and latency overhead has led to limited usability of cloud models in dynamic, safety-critical, real-time settings. To effectively address this trade-off, we introduce UniLCD, a novel hybrid inference framework for enabling flexible local-cloud collaboration. By efficiently optimizing a flexible routing module via reinforcement learning and a suitable multi-task objective, UniLCD is specifically designed to support the multiple constraints of safety-critical end-to-end mobile systems. We validate the proposed approach using a challenging, crowded navigation task requiring frequent and timely switching between local and cloud operations. UniLCD demonstrates improved overall performance and efficiency, by over 35% compared to state-of-the-art baselines based on various split computing and early exit strategies. Our code is available at <https://unilcd.github.io/>.

Keywords: Task Offloading · Efficiency · Reinforcement Learning · Navigation

1 Introduction

We are currently undergoing a transformative societal phase as vision-based systems, such as mobile robots and personalized devices, transition from their controlled lab environments and into the real world. However, computational and energy constraints currently hinder the performance of deployed real-world systems, e.g., high-cost systems today may only be able to support lightweight neural network models with limited accuracy and runtime of about an hour before the on-board battery is depleted [23, 26, 40, 64, 69, 90, 123]. Current efficiency bottlenecks may only become worse over time given ever-increasing sensor resolutions and larger models [2, 6, 58, 62, 87, 95].

To address the need for reliable and accurate inference, the currently prevailing approach for processing with increasingly powerful models involves sending data, e.g., an image, to remote cloud servers in order to offload computation [2, 6, 40, 55, 59, 117]. While this practice can benefit real-world systems through delivering high-quality model predictions, the transmission and cloud processing overhead incurs a high latency cost and is therefore not suitable for dynamic real-world agents, such as mobile

systems, which must continuously anticipate and respond to dynamic settings. Safety-critical systems with hard real-time constraints and responsiveness may solely rely on local processing, while leveraging strategies such as model pruning [24, 36, 47, 107] and quantization [48, 80, 99] to support embodied constraints, e.g., limited hardware and battery life. However, on-device lightweight models often suffer from significantly degraded accuracy and in turn can hinder safe decision-making, i.e., missing or falsely detecting nearby pedestrian. Thus, we seek to address a fundamental research question: How to realize robust vision-based systems that can be flexibly optimized for both *safety and real-time efficiency* while operating in dynamic real-world settings?

While shared local-cloud computation frameworks have been extensively studied by prior works, e.g., for various Internet of Things applications [25, 31, 32, 39, 39, 40, 55, 77, 91, 117], the aforementioned methods rarely analyze safety-critical, real-time operation. Instead, standard off-loading frameworks rely on various ad-hoc and inadequate heuristics, e.g., based on a task-agnostic layer-wise split or simplistic statistics [39, 40], and thus cannot accommodate dynamic inference, i.e., based on the difficulty or safety-constraints of the current scenario. For instance, some scenarios may require split-second decisions, such as in the case of dense and dynamic scenarios with crowded surrounding pedestrians. In this work, we develop a more generalized paradigm enabling situation-specific collaboration between cloud computing and local inference to balance the energy cost while maintaining high system accuracy. Our method can be used to flexibly optimize for multiple constraints, i.e., latency, accuracy, efficiency, and safety, while also enabling systems with a highly sustainable operation (a critical need as compute-intensive systems become more widely adopted [92]).

Contributions: Towards facilitating real-time, cloud-enabled systems, we make three key contributions. First, we introduce a novel cloud-local hybrid collaboration framework, UniLCD. Our proposed reinforcement learning-based framework is inspired by the observation that small, low-power consumption models on edge devices are sufficient in some scenarios, while more complex conditions may call for powerful cloud resources. Our approach employs a conditional routing module that learns to dynamically route computation between local and cloud resources. Second, we demonstrate the importance of carefully designing the reward function in order to effectively balance over multiple constraints, i.e., energy consumption, latency, and safe performance in a context-dependent manner. Third, we validate our method and several state-of-the-art baselines by introducing a novel benchmark comprising challenging vision-based crowded navigation tasks that require seamless decision-making with frequent cloud-local switching. We demonstrate our approach to significantly outperform prior methods [39, 40] by 35% across multiple metrics, as well as a proposed ecological score. Our code and benchmark are available at <https://unilcd.github.io/> for future researchers tackling multifaceted challenges in practical, real-world, vision-based systems.

2 Related Work

Cloud-Edge Collaborative Systems: The integration of cloud computing and edge devices has attracted increasing attention in recent years as it can potentially combine the

Table 1: Comparison with Prior Work. By holistically tackling aspects of system safety, efficiency, and latency, our framework is suitable for real-time decision-making in dynamic scenes.

Method	Low Latency Deployment	Edge Deployment	End-to-End Training	Situational Embedding	High Accuracy	Real Time
On-Device	✓	✓	-	✓	✗	✗
On-Cloud	✗	✗	-	✓	✗	✓
Neurosurgeon [40]	✗	✓	✓	✓	✓	✗
Dynamic [27]	✓	✓	✗	✗	✓	✗
Selective Query [39]	✓	✓	✓	✗	✗	✓
Compressive Offloading [111]	✗	✗	✓	✓	✓	✗
Adaptive Offloading [98]	✓	✓	✓	✓	✗	✗
Deep Sequential RL [100]	✓	✓	✓	✓	✓	✗
UniLCD (ours)	✓	✓	✓	✓	✓	✓

advantages from both [17, 25, 79, 84]. Leveraging the cloud’s capacity to employ advanced hardware and deploy large models allows quick execution of computationally intensive tasks with accurate results, while edge devices enable real-time processing and diminish latency by accessing data close to the source. Thus, cloud-edge collaboration systems have been developed across various domains, such as autonomous driving [44] and smart cities [105]. Recently, Kag et al. [39] proposed a hybrid approach that learns to efficiently select queries for cloud processing. However, the method focuses on simplistic image classification domains without integrating sensitivity to computational delays. In contrast, our framework emphasizes the importance of *both prompt and accurate* cloud-edge collaboration and can thus more flexibly support broad applications, e.g., with handling of dynamic and complex safety-critical scenarios.

Real-Time Decision-Making: Real-time decision-making is critical for scenarios that require fast and accurate responses (e.g., autonomous driving [49, 50, 61, 113], energy grid management [88, 106], assistive technologies [16, 33, 96]). State-of-the-art models usually have billions of parameters, resulting in long inference times. Notably, on-going work is primarily concentrated on enhancing data and model processing speed [22, 102], as well as designing dedicated processing chips [41, 89]. In this work, we tackle a complementary direction of cloud-edge collaboration, where we explore the integration of the cloud’s accelerated processing speeds with the deployment of edge devices. This integration aims to effectively mitigate latency concerns while ensuring the reliability of prompt decision-making, which generally results in efficiency and latency costs [4, 66, 78].

Early-Exit Models: Early exit techniques are closely related to our work as they incorporate internal classifiers at various shallow layers, allowing the model to exit earlier during inference while still predicting the correct label, i.e., to save time and energy costs [35, 40, 51, 54, 81, 93, 108]. Yet, despite the fast inference, models can still suffer in latency and accuracy [40]. Li et al. introduced AppealNet [54], an architecture that efficiently processes deep learning tasks by predicting whether inputs can be managed by a resource-constrained edge device or need to be offloaded to a cloud-based model. However, such early exit strategies only addresses on-device efficiency without con-

sidering overall task and safety, leading to potentially unsafe and frequent navigation errors in our robot navigation task, as will be shown in our analysis in Sec 4.

Energy-Efficient Models: Our proposed method primarily aims to present timely results while minimizing the energy cost from both cloud servers and edge devices, which can not only help reduce carbon emissions [3, 12, 45, 121] but also increase battery life and reduce total overhead costs [21, 26, 59]. The promise of energy-efficient models has been substantially developed through various research fields, focusing on optimizing computational efficiency and resource utilization. One aspect of the field of energy-efficient models emphasizes hardware design [38, 73]. However, hardware development is often time-consuming and expensive. Another aspect provides insight into the efficacy of model pruning and quantization techniques that can minimize model size thus helping save energy [37, 110, 122] consumption. In our work, we delve into resource allocation and optimization strategies, seeking to further improve energy efficiency.

Decision-Making Models for Mobile Systems: Robotic companions and assistants designed to operate in human-inhabited environments has made significant progress [9, 42, 114]. Imitation learning [34, 112, 115, 120] involves the system learning from observed behavior, typically demonstrated by a human or professional operator, to guide its decision-making process in mobile systems, such as autonomous vehicles [75, 94] or drones [103]. Reinforcement learning (RL) [60, 70, 104, 109, 118] has also been applied to optimize decision-making processes, allowing robots to navigate efficiently while avoiding obstacles and optimizing path planning [7, 100, 119]. However, the wide-scale deployment of real-time vision-based navigation models on robots today is hindered by ever-increasing computational and hardware constraints [20, 57], making it difficult to deploy accurate but large models on systems. In our work, our RL-based framework can be used to query the cloud efficiently, i.e., to reduce overall energy cost without sacrificing local resources [28, 63, 101, 111]. Our flexible framework is model and platform-agnostic, while automatically adapting across situations and communication settings.

3 Method

Our objective is to learn a *policy* that determines the dynamic allocation of local and cloud resources, optimizing energy efficiency and real-time performance. In the following, we discuss the main components of our method. First, we formulate our real-time task of navigating to a specific goal in Sec. 3.1. Next, we employ imitation learning to train two navigation policies tailored for local and cloud settings, respectively (Sec. 3.2). Third, we discuss how we train our sample-efficient *routing policy* via Proximal Policy Optimization (PPO) [86] (Sec. 3.3). In particular, we introduce our multi-objective reward, designed to optimize energy efficiency while maintaining navigation performance. Fig. 1 depicts an overview of our method.

3.1 Problem Formulation

We formulate our robot navigation task as a real-time sequential decision-making problem from a set of observations $\mathbf{o} = \{\mathbf{I}, \mathbf{p}\} \in \mathcal{O}$, comprising a front-view camera image

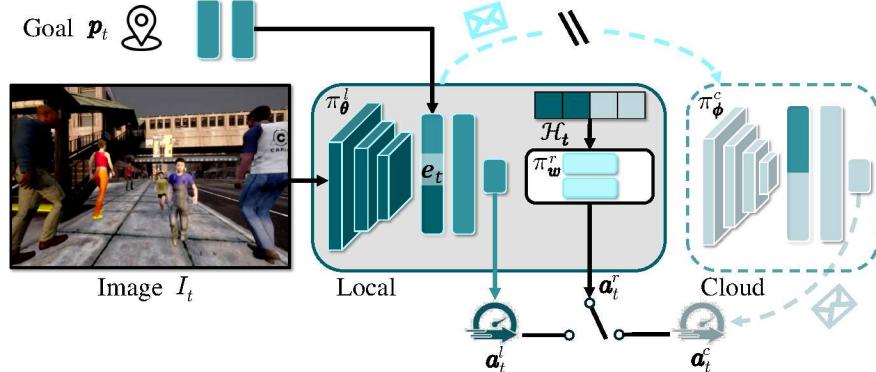


Fig. 1: Overview of UniLCD. Our system comprises a situational routing module, which takes the current embedding and a history of previous actions. Local actions are predicted by a pre-trained lightweight model that can be deployed efficiently on a the mobile system. The sample-efficient routing module, trained via RL, determines whether to implement the local action or transmit the scene embedding to the cloud server model, which is more accurate but computationally expensive and induces latency.

$\mathbf{I} \in \mathbb{R}^{W \times H \times 3}$, the next waypoint $\mathbf{p} \in \mathbb{R}^2$, to a set of actions $\mathbf{a} = \{d, v\} \in \mathcal{A}$, where $d \in \mathbb{R}$ represents the orientation and $v \in \mathbb{R}$ is the speed [14, 19]. Our objective is to obtain a mapping function $f_\Theta: \mathcal{O} \rightarrow \mathcal{A}$ parameterized by $\Theta = [\theta, \phi, \omega]$ for generating actions at each time step to minimize the overall energy cost from both the local device and cloud server [10]. Our approach consists of three policies: a local and a cloud navigation policy π_θ^l and π_ϕ^c with weights $[\theta, \phi]$ that map the observations to actions, producing \mathbf{a}_t^l and \mathbf{a}_t^c , respectively. Additionally, we introduce a routing policy $\pi_\omega^r(\mathbf{o}|\pi_\theta^l, \pi_\phi^c)$ parameterized by ω , to dynamically determine the optimal utilization of local resources versus offloading to the cloud server.

While our method can be optimized end-to-end, we initialize the two navigation policies in an imitation learning manner [14, 29, 76]. We also note that, given the constraints of local device hardware in our setting, the number of parameters for the local navigation policy is generally smaller than that of the cloud navigation policy.

For comparative analysis and sample efficiency, after training the navigation policies, we freeze the parameters θ, ϕ and train the routing policy $\pi_\omega^r(\mathbf{o}|\pi_\theta^l, \pi_\phi^c)$ based on the observations and the two navigation policies using residual RL. It's noteworthy that both the local navigation policy and the routing policy are deployed on a local device, sharing parameters for the initial layers of the neural network, as shown in Fig. 1.

Cloud-Aware Robot Navigation Task: Due to the challenges of obtaining sufficient data in the dynamic real-world scenario [13, 67, 97], we develop a simulation environment tailored for robot navigation through a crowded outdoor setting. Our environment is built upon CARLA (version 0.9.13) [19], an open-source simulator typically employed for testing autonomous driving algorithms. The information transmission between the local and the cloud server potentially introduces a stochastic delay [8, 30, 40, 116].

Algorithm 1 UniLCD’s Routing Policy Training with Reinforcement Learning

```

1: Input: Image  $\mathbf{I}$ , next waypoint  $\mathbf{p}$ , local policy  $\pi_\theta^l$ , cloud policy  $\pi_\phi^c$ 
2: Initialize: Number of iterations  $T$ , history  $\mathcal{H}$ , routing policy  $\pi_\omega^r$ , reply buffer  $\mathcal{S}$ 
3: Collect on policy samples:
4: for  $t = 1$  to  $T$  do
5:   Obtain local action  $\mathbf{a}_t^l$  and embeddings  $\mathbf{e}_t$  using local policy  $\pi_\theta^l(\mathbf{I}_t, \mathbf{p}_t)$ 
6:   Append  $(\mathbf{a}_t^l, 0)$  to history  $\mathcal{H}_t$ 
7:   if  $\pi_\omega^r(\mathcal{H}_t, \mathbf{e}_t) = 0$  then  $\mathbf{a}_t = \mathbf{a}_t^l$ 
8:   else
9:     Send  $\mathbf{e}_t$  to cloud,  $\mathbf{a}_t = \pi_\phi^c(\mathbf{I}_t, \mathbf{p}_t)$ 
10:    Update last value of  $\mathcal{H}_t$  to  $(\mathbf{a}_t, 1)$ 
11:   end if
12:   Compute instant reward using Eq. (2)
13:   if Arrived destination then break
14:   end if
15:   Update replay buffer  $\mathcal{S} = \mathcal{S} \cup \{\mathbf{I}_t, \mathbf{p}_t, \mathcal{H}_t, r_t\}$ 
16:   Update routing policy parameters with PPO
17: end for

```

Toward Energy-Efficiency: The primary objective of our research is to enhance energy efficiency while preserving task performance [31, 40, 77]. In our study, the navigation task is offloaded to the cloud for processing only when local computation is insufficient for achieving optimal task performance. However, communication bottlenecks necessitate prioritizing local processing to conserve energy and minimize latency. Offloading to the cloud is reserved for situations where only the cloud can meet the task performance requirements, such as challenging scenarios like navigating in bad weather [72] or crowded environments. Consequently, we leverage a routing network trained with PPO and suitable for running on local devices. This modular architecture is then optimized for energy efficiency, as discussed below.

3.2 Learning Local and Cloud Policies

In our work, we follow the standard imitation learning approach to train our navigation policies in an offline manner. Specifically, we first collect a dataset $\mathcal{D} = \{\mathbf{I}_i, \mathbf{p}_i, \mathbf{a}_i\}_{i=1}^N$ on diverse and complex routes and weathers using CARLA [19, 74] to simulate real-world scenarios to provide the basis for our navigation policies.

As shown in Fig. 1, both the local and cloud navigation policies comprise (i) a visual semantics feature extractor module for obtaining embeddings from input images, (ii) a multilayer perceptron to extract the feature associated with the robot’s imminent goal point, and (iii) a goal-conditional module that takes concatenation of image embeddings and goal features to predict robot actions, encompassing both direction d and speed v . To share common features and accommodate local computing resource constraints, we first train a robust cloud policy with a large neural network. Subsequently, we freeze the parameters of the first few layers of the pre-trained cloud policy to serve as a shared feature extractor for the local policy. Additional fully connected layers are then added to the extracted features to form the local policy. Therefore, the local policy is significantly

smaller compared to the cloud policy. Only the parameters of the fully connected layer in the local policy are optimized, leading to reduced computational consumption and improved efficiency in both training and inference time.

The learning objective for both local and cloud policies is achieved by minimizing the \mathcal{L}_1 distance:

$$\text{minimize } \mathbb{E}_{(\mathbf{I}, \mathbf{p}, \mathbf{a}) \sim D} [\mathcal{L}_1(\mathbf{a}, \pi(\mathbf{I}, \mathbf{p}))] \quad (1)$$

where π represents navigation policy $\{\pi_\theta^l, \pi_\phi^c\}$.

3.3 Learning a Routing Policy

To effectively balance cloud and local computing resources while achieving good performance at the same time, our study proposes a *routing policy* that seamlessly switches between local and cloud. We follow the standard PPO training process as shown in Algorithm 1 and the formulation of our routing policy is introduced below.

State Space: We denote the current state as $\mathbf{s}_t = \{\mathbf{e}_t, \mathcal{H}_t\}$, where \mathbf{e}_t represents the image embeddings and goal features extracted from the shared feature extractor employed by both local and cloud policies, as discussed in Sec. 3.2, and $\mathcal{H}_t = \{(\mathbf{a}_i, \mathbf{1}_i)\}_{t-k}^t$ is a sequence of the last k history actions, where \mathbf{a}_i represents previous navigation actions and $\mathbf{1}_i$ is an indicator function indicating whether the action is obtained locally or from the cloud.

Action Space: The action produced by our router policy is a binary discrete value, indicating whether to accept the local navigation policy π_θ^l or transmit the embeddings \mathbf{e}_t to the cloud navigation policy π_ϕ^c .

Task Reward: It is challenging to optimize task performance and energy efficiency while considering the sub-optimal actions generated from a small local model and latency-induced cloud model.

We design our reward function encompassing five key components: geodesic reward r_{geo} , speed reward r_{speed} , energy disadvantage bonus r_{energy} , extreme action clip r_{action} , and collision penalty $r_{collision}$. The overall reward function is defined as

$$r = (r_{geo} \cdot r_{speed} \cdot r_{energy} \cdot r_{action})^\alpha - r_{collision} \quad (2)$$

where $\alpha = 1/4$ is a scaling factor that scales the overall reward between $[0, 1]$. By using a multiplicative objective, we ease optimization across different instantaneous components of the reward, i.e., if one of the terms is low, the entire reward is affected. Hence, the multiplicative overall reward objective can also be scale invariant to some extent (however, clipping one of the reward terms can change its importance). In our analysis, we find it to improve optimization and reduce the need for careful hyper-parameter tuning in multi-objective RL. For collision reward, we treat the term separately as it is sparsely observed over selected frames as a large negative reward, as defined below.

Geodesic Reward: We introduce r_{geo} to align the robot's trajectory with the pre-defined path, penalizing the robot for movements that deviate from the path.

$$r_{geo} = (1 - \tanh(d_{geo})) \quad (3)$$

where d_{geo} represents the Euclidean distance from the robot's current position to the nearest waypoint on the pre-defined path.

Speed Reward: The speed reward, denoted as r_{speed} , is designed to motivate robots to reach their destination in the shortest time. Considering the impact of GPU resource utilization on task execution duration, speed reward encourages quick task completion thus minimizing the overall computing resource usage. The maximum speed of a robot in our environment is set to $m_v = 1.5m/s$ according to [53, 71], and we define the speed reward as

$$r_{speed} = v/m_v \quad (4)$$

where v represents the current speed.

Energy Disadvantage: In our pursuit of minimizing energy consumption as the primary objective, and considering the significantly lower energy cost of local computation compared to cloud transmission and computation, we introduce penalties when the robot seeks cloud assistance. Similar to the speed component, we formulate the normalized energy disadvantage reward as

$$r_{energy} = 1 - e/m_e \quad (5)$$

where m_e is the maximum energy cost at one step.

Extreme Action Clip: Due to the property that \tanh converges at extreme action values, leading to a potential risk of converging to local maxima, we introduce an action clipping mechanism to address this concern. Specifically, actions exceeding the maximum allowed value will be directly assigned a reward of 0. This precautionary measure is extended to both standardized directional and speed actions, as shown in Eq. (6)

$$r_{action} = \mathbb{1}(|r_{speed}| < \epsilon) \cdot \mathbb{1}(|\frac{d}{d_m}| < \epsilon) \quad (6)$$

where $\epsilon = 0.97$ is the threshold for clipping the actions, d is the next rotation angle suggested by the navigation model chosen and d_m is the maximum possible rotation, ensuring that extreme actions do not unduly influence the reward function.

Collision Disadvantage: Robot navigation needs to emphasize safety and endeavor to avoid collisions with obstacles, ensuring the successful accomplishment to the destination. To underscore the significance of collision avoidance, we include an episodic termination mechanism to emphasize the severity of collisions. Specifically, a substantially higher negative penalty $r_{collision}$ is assigned and the episode is terminated whenever a collision happens.

4 Experiments

4.1 Implementation Details

Data Collection: To learn the local and cloud navigation policies using imitation learning, we collect crowd navigation data from our CARLA environment. We simulate

environments with varying densities by introducing five, 15, 30, and 70 pedestrians into the ego robot’s path, corresponding to low, medium, dense, and crowd-density settings, respectively. This data is then aggregated to form a comprehensive training dataset for both local and cloud navigation policies. Our data collection approach employs a heuristic policy: we establish 10 different paths with predefined waypoints, which the robot follows from the starting position to the destination. The robot avoids static obstacles and halts when other pedestrians block its path, resuming moving once the path is clear.

Local and Cloud Policies: As discussed in Sec. 3.2, our local and cloud policies share a common feature extractor that processes a 480×480 image and a 2D imminent position vector. Specifically, we utilize the initial layers of RegNet [85] to extract the image features, and two fully connected (FC) layers to capture the position features. For the local policy, the image feature is flattened and concatenated with the position feature, followed by additional FC layers to generate the local actions. For the cloud policy, both the image feature and position feature are transmitted from the local to the cloud. The image feature is then processed by the remaining layers of RegNet [85] and concatenated with the position feature. An MLP is built on top of the combined features to generate the cloud actions.

Routing Policy: Our routing policy follows the standard PPO training process, with a policy network that decides whether to route locally or to the cloud, and a value network that evaluates the chosen actions by estimating the value function. Both the policy and value networks are MLPs with two hidden layers, with the hidden size being 16 for the policy network and 256 for the value network.

Training Protocol: We use AdamW [65] optimizer and train for 200 epochs with a learning rate of 0.0001 using our robot navigation dataset for both local and cloud navigation policies. The routing policy is trained for 1,000 episodes, each consisting at most 1,500 steps, with a discount factor γ set to 0.99. Episodes are truncated if the agent collides with other objects, completes the designated number of steps, reaches the predefined destination, or deviates more than three meters from the predefined route.

4.2 Evaluation Metrics

Task Performance Evaluation: We adapt CARLA’s evaluation metrics [1] to assess the navigation performance of our ego robot. In addition to commonly used metrics (e.g., success rate, route completion, etc.), we propose *Navigation Score*, denoted as NS , similar to the driving score [1], by normalizing the infraction counts per meter. NS is defined as

$$NS = RC \cdot P_I^{IC} \cdot P_{RD} \quad (7)$$

where RC is route completion, P_I represents the infraction penalty for collisions, IC represents the number of robot collisions per meter, and P_{RD} represents the penalty for route deviation. Following the CARLA leaderboard settings, P_I is set to 0.5, and P_{RD} is defined as

$$P_{RD} = \begin{cases} 0.8, & \text{if } RD > \epsilon_{RD} \\ 1.0, & \text{otherwise} \end{cases} \quad (8)$$

where $\epsilon_{RD} = 1.5\text{m}$ is the route deviation threshold.

Energy Evaluation: In addition to the navigation-specific evaluations, we introduce metrics to assess the computational and communication overhead incurred during the execution of our algorithm.

Specifically, the total energy consumption for each episode is defined as

$$\text{Energy} = E_{local} \cdot N_{local} + E_{cloud} \cdot N_{cloud} \quad (9)$$

where N_{local} and N_{cloud} represent the number of steps the local and cloud models are chosen by the routing policy in one episode, respectively. E_{local} and E_{cloud} are the energy costs for processing one observation locally or on the cloud server. We refer to the energy values suggested by [40] and scale them according to our model and image size. In our study, $E_{local} = 0.15J$ is the computation energy cost of the local policy, and $E_{cloud} = 1.5J$ comprises both the computation and communication energy costs of the cloud policy. Further details regarding the calculation can be found in the supplementary material.

Ecological Navigation Score: As our primary objective is to optimize the overall energy consumption while ensuring robot navigation performance, we introduce an *Ecological Navigation Score* (ENS) that balances navigation performance and efficiency considerations.

The ENS is defined as

$$\text{ENS} = P_E \cdot \text{NS} \quad (10)$$

where P_E is the penalty term for energy consumption, and defined as

$$P_E = 1 - \frac{\text{Energy}}{N_E} \quad (11)$$

where $N_E = (E_{local} + E_{cloud}) \cdot (N_{local} + N_{cloud})$ is a normalization factor.

Run-Time: Real-time decision-making is crucial for robot navigation. In our study, the robots need to respond instantly to environmental changes and swiftly switch between local and cloud-based navigation policies to adjust path planning. To demonstrate this real-time response, we report our run-time Frames Per Second (FPS), which includes both the model processing time and the communication time between a local device and the cloud server.

4.3 Results

We use the CARLA simulator (version 0.9.13) [19] to validate the effectiveness of our proposed method. We tested our algorithm on five different routes in CARLA’s Town 10, each with 30 episodes. In testing time, we explored diverse weather conditions (including hard rain, sunny, wet, and sunset) and varying traffic density to simulate realistic conditions [15]. The maximum length of the routes is 40 meters. In this section, we first compare UniLCD with state-of-the-art local-cloud collaboration solutions [39, 40, 98, 100, 111]. Subsequently, we evaluate the performance of UniLCD with

Table 2: Comparing UniLCD with Baselines. We categorize our comparative analysis between individual models (Local and Cloud only), baselines (Deep Learning and RL-based computational offloading research), and our proposed UniLCD variations, which show progressive improvement. \dagger denotes methods transmitting to the cloud raw input data, i.e., instead of an embedding. ENS is Ecological Navigation Score (%), NS is Navigation Score (%), SR is Success Rate (%), RC is Route Completion (%), Infract. is Infraction Rate (/m), Energy is measured in Joules per meter (J/m) and FPS is Frames Per Second.

Method	ENS \uparrow	NS \uparrow	SR \uparrow	RC \uparrow	Infract. \downarrow	Energy \downarrow	FPS \uparrow
\dagger Cloud-Only [85]	0.00	96.47	93.33	98.50	0.03	36.49	7.11
Local-Only [83]	63.43	67.33	0.00	75.23	0.16	4.33	65.40
<i>Baseline Methods:</i>							
Compressive Offloading [111]	13.98	80.16	0.00	80.16	0.00	90.66	1.82
\dagger Selective Query [39]	24.14	61.28	0.00	82.68	0.11	45.35	18.14
\dagger Adaptive Offloading [98]	37.42	40.37	70.00	94.05	1.22	4.80	30.14
Neurosurgeon [40]	39.85	63.10	0.00	80.54	0.03	28.31	12.53
SPINN [52]	36.31	72.75	60.00	92.73	0.35	18.94	20.37
Deep Sequential RL [100]	58.84	61.83	0.00	79.36	0.36	3.77	77.94
<i>UniLCD Module Ablations:</i>							
\dagger Standard Reward	48.35	54.99	0.00	75.23	0.13	3.57	50.20
\dagger Standard Reward w/ History	50.04	57.21	10.00	77.71	0.12	8.38	49.07
\dagger Our Reward (Eq. (2))	48.30	79.90	56.66	91.15	0.19	21.72	16.05
\dagger Our Reward w/ History	71.70	87.71	83.33	94.66	0.11	7.83	33.98
Our Reward (Eq. (2))	57.20	87.39	60.00	91.10	0.06	6.60	12.49
Our Reward w/ History	85.97	94.58	93.33	95.90	0.02	2.90	26.49

various sizes of local models and data-transmission settings. Third, we conduct ablation studies across different crowd-density settings. Finally, we systematically assess the significance of each reward term by omitting one at a time and analyzing their respective impacts, thereby reinforcing the foundational rationale of our reward design. The last two ablation studies can be found in our supplementary.

Comparing UniLCD with Baselines: Tab. 2 depicts our main results, comparing against multiple baselines, including local and cloud-only models. We emphasize that prior work does not usually consider safety-critical, real-time tasks such as social navigation. As shown in Tab. 2, several baselines can achieve a viable route completion (RC) score, however, this comes at a cost, e.g., high infraction rates. This includes early-exit baselines, which provide a limited offloading mechanism. We show the ENS metric to drop to zero when assessing the high energy-consuming cloud model exclusively, demonstrating that conducting all computations in the cloud with the computation-heavy model is energy-consuming. Moreover, the local-only model demonstrates a poor navigation score of 67.33% and a high collision rate of 0.16, indicating that low-accuracy, device-only models are more susceptible to navigation errors. The baseline models show comparable ENS scores to the local-only method but do not improve much in other metrics. Early-exit strategies such as SPINN [52] and Neurosurgeon [40] address on-device efficiency with improved ENS scores up to 39.85%, but perform worse than offloading-based methods in other metrics. Furthermore, we compare our method with DNN-based offloading (Selective Query [39]), which fails to improve due to the

dynamic nature of the task and over-reliance on the cloud. This particular baseline also requires transmitting raw data to the cloud, which is inefficient. The best performing baseline, of Deep Sequential RL [100], achieves an ENS of 58.84%.

UniLCD achieves a state-of-the-art results, with an ENS of 85.97%, achieving state-of-the-art performance. Our analysis also highlights the need for cloud-edge collaborative methods that can combine both the strengths of fast local inference and high-capacity cloud computation.

Impact of Reward Design: Tab. 2 also shows ablation of UniLCD using various rewards. We evaluate UniLCD using the standard reward design (e.g., [123]), where individual reward terms are added. We initially find the additive reward to result in comparable performance to our reward design (ENS of 48.35% compared to 48.30%). However, when added history, the additive reward does not improve, while ours is shown to outperform significantly (ENS of 50.04% for standard reward compared to 71.70% with ours). We attribute this to the normalization effect of the multiplicative function during training. This eases optimization, whereas, in additive overall reward, the contribution of each term may need to be carefully optimized. While it is possible that careful hyperparameter optimization can match these results, this highlights the flexibility and generalization of our objective, which performs well across different model inputs. This finding also uncovers the sensitivity of current methods for the reward design. Finally, we utilize an embedding from the local model as input to UniLCD to provide the necessary context for collision handling and transmit it to the cloud, i.e., instead of raw data to reduce energy costs. This improves overall performance, achieving an ENS of 85.97%, mostly due to reduced energy consumption, while maintaining high performance. The embedding-based communication architecture also reduces collisions due to a higher FPS, resulting in a higher NS.

Local Policy Backbone Ablations: Tab. 3 shows ablations of UniLCD for local backbones of various sizes. We run experiments with UniLCD that transmit raw image data to the cloud using local backbones such as MobileNetV2 (medium), MobileNetV3small (small), and MobileVIT (tiny), coupled with a RegNet (cloud) model. As an imitation model, the medium backbone demonstrates task performance closest to the cloud, with minimal instances of infractions followed by the small and tiny backbones. The UniLCD with history model consistently favors local at instances where both local and cloud models have similar capacity, i.e., the medium backbone. The algorithm successfully sustains an impressive average route completion rate of 94.66% with occasional cloud support and significantly reduces energy consumption to 7.83 J/m using a comparatively weaker local model with a smaller backbone. Further, testing local policy incapable of path-following or collision avoidance with a tiny backbone reveals a complete reliance on cloud assistance but results in the successful completion of all episodes.

We also run ablations of UniLCD that transmit embeddings to the cloud. The local models for these experiments are trained by freezing pre-trained cloud weights where stage 1 is fine-tuned with the feature extractor containing the early layers of the cloud model and stage 2 is fine-tuned with the feature extractor containing the final layers of the cloud model. Observations indicate that UniLCD with history model achieves an ENS of 85.97% solely with the stage 1 backbone. It is noteworthy that both models ex-

Table 3: Local Policy Backbone Ablations. We test the task performance in the dense crowd setting for UniLCD and UniLCD w/ History. \dagger denotes methods transmitting to the cloud raw input data, i.e., instead of an embedding. For UniLCD transmitting data, we use three dimensions of local models: Tiny, Small, and Medium. For UniLCD transmitting embedding, we use local models trained up to two stages: Stage 1 trained with fewer parameters and Stage 2 trained with more parameters. Params is Number of Parameters in each local model (M), ENS is Ecological Navigation Score (%), NS is Navigation Score (%), SR is Success Rate (%), RC is Route Completion (%), Infract. is Infraction Rate (/m), Energy is measured in Joules per meter (J/m), and FPS is Frames Per Second.

Local Model Size	Params	ENS \uparrow	NS \uparrow	SR \uparrow	RC \uparrow	Infract. \downarrow	Energy \downarrow	FPS \uparrow
<i>UniLCD:</i>								
\dagger Tiny	1.37	12.80	93.29	90.00	97.93	0.07	34.01	7.35
\dagger Small	2.54	48.30	79.90	56.66	91.15	0.19	21.72	16.05
\dagger Medium	3.50	50.92	87.87	80.00	95.50	0.12	21.19	15.10
<i>UniLCD w/ History:</i>								
\dagger Tiny	1.37	0.00	91.27	93.33	98.50	0.11	36.52	6.42
\dagger Small	2.54	71.70	87.71	83.33	94.66	0.11	7.83	33.98
\dagger Medium	3.50	73.46	83.86	90.00	96.22	0.15	5.22	11.53
<i>UniLCD:</i>								
Stage 1	0.53	57.20	87.39	60.00	91.10	0.06	6.60	12.49
Stage 2	0.95	74.12	81.54	93.33	91.10	0.16	1.80	65.40
<i>UniLCD w/ History:</i>								
Stage 1	0.53	85.97	94.58	93.33	95.99	0.02	2.90	26.49
Stage 2	0.95	86.78	95.47	93.33	98.15	0.04	1.77	36.50

hibit a bias towards local processing when a backbone comparable to cloud performance (stage 2) is employed, resulting in the highest ENS of 86.78%, and a route completion of 98.15% with minimal infractions, averaging 0.04 per meter. Overall, these findings underscore UniLCD’s proficiency in leveraging cloud resources only when local processing is inadequate.

Performance over Training Progress: We also investigate the evolution of the ENS task metric over the training progress as shown in Fig. 2. Remarkably, our model and reward outperform the standard reward even at the onset of training despite not requiring careful hyper-parameter tuning. Nonetheless, as training progresses, the difference in the baselines becomes even more apparent. The ENS performance is presented over training iterations and averaged over 10 tests, i.e., using different environmental seeds. Results for performance in different environments and settings (e.g., crowd density) can be found in the supplementary. In our comparative analysis, we note that as the complexity of the scenario increases, our baseline approaches show a decline in task performance. This decline is evidenced by an increase in route deviation and infraction rates. While UniLCD can easily navigate comparatively simpler scenarios, its infraction rate increases only slightly. This observation suggests the algorithm’s robust adaptability for effective navigation in diverse and intricate environments.

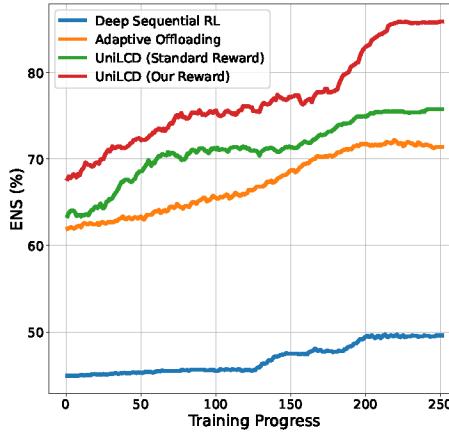


Fig. 2: Training Progress Results. We evaluate performance for different models, including UniLCD trained with a standard, carefully tuned, additive reward vs. UniLCD with the proposed reward function. We find consistently improved performance throughout the entire model training process. Results are shown for averaging across 10 evaluation seeds.

5 Conclusion and Future Work

We envision large-scale robot agents that collaborate and actively execute tasks in complex real-world scenarios. To this end, we design a PPO-based routing policy that learns to seamlessly switch between local and cloud policies depending on the situation and task objective of the observations. Specifically, we address the challenging task of real-time social robot navigation. We demonstrate that the proposed approach results in reduced overall energy cost while maintaining robust navigation performance. Our simulation environment can be used to facilitate more research into both *effective and efficient* navigation in the future. Thus, we aim to facilitate the development of more sustainable and robust societal-scale AI-based systems. Moreover, our approach can be extended to various platforms and real-time decision-making tasks that require both seamless communication and high performance. While we take a step towards quantifying trade-offs in real-time, safety-critical systems in challenging simulation settings, the next step would be to analyze UniLCD within diverse real-world environments and ambient settings (e.g., different communication settings, larger cloud models). Finally, although we consider the overarching objective of minimizing processing computation across both local and cloud devices, other aspects of efficiency can also be considered. For instance, maintaining local infrastructure often requires additional battery and hardware resources, which can be better optimized with cloud-based solutions especially when leveraging large, high-capacity models [11, 18].

Acknowledgments

We thank the Red Hat Collaboratory (award #2024-01-RH07) for supporting this research, and Bassel Mabsout for helpful discussions regarding the problem formulation.

References

1. Carla autonomous driving leaderboard. <https://leaderboard.carla.org/> (2022)
2. Achiam, J., Adler, S., Agarwal, S., Ahmad, L., Akkaya, I., Aleman, F.L., Almeida, D., Altenschmidt, J., Altman, S., Anadkat, S., et al.: GPT-4 technical report. arXiv (2023)
3. Alharbi, H.A., Aldossary, M.: Energy-efficient edge-fog-cloud architecture for iot-based smart agriculture environment. Access (2021)
4. Aujla, G.S., Kumar, N., Zomaya, A.Y., Ranjan, R.: Optimal decision making for big data processing at edge-cloud environment: An sdn perspective. Trans. Ind. Inform. (2017)
5. Booth, S., Knox, W.B., Shah, J., Niekum, S., Stone, P., Allievi, A.: The perils of trial-and-error reward design: misdesign through overfitting and invalid task specifications. In: AAAI (2023)
6. Brohan, A., Brown, N., Carbajal, J., Chebotar, Y., Chen, X., Choromanski, K., Ding, T., Driess, D., Dubey, A., Finn, C., et al.: RT-2: Vision-language-action models transfer web knowledge to robotic control. CoRL (2023)
7. Caruso, M., Regolin, E., Camerota Verdù, F.J., Russo, S.A., Bortolussi, L., Seriani, S.: Robot navigation in crowded environments: A reinforcement learning approach. Machines (2023)
8. Charneski, A.: Modeling network latency. Online (2015)
9. Chen, C., Liu, Y., Kreiss, S., Alahi, A.: Crowd-robot interaction: Crowd-aware robot navigation with attention-based deep reinforcement learning. In: ICRA (2019)
10. Chen, X., Zhang, J., Lin, B., Chen, Z., Wolter, K., Min, G.: Energy-efficient offloading for dnn-based smart iot systems in cloud-edge environments. TPDS (2021)
11. Cheng, Y., Wang, D., Zhou, P., Zhang, T.: Model compression and acceleration for deep neural networks: The principles, progress, and challenges. IEEE Signal Process. Mag. (2018)
12. Chin, R., Morcos, A.S., Marculescu, D.: Pareco: Pareto-aware channel optimization for slimmable neural networks. In: ICML (2020)
13. Chung, M.J., Wang, M.J.: The change of gait parameters during walking at different percentage of preferred walking speed for healthy adults aged 20–60 years. Gait & posture (2010)
14. Codevilla, F., Müller, M., López, A., Koltun, V., Dosovitskiy, A.: End-to-end driving via conditional imitation learning. In: ICRA (2018)
15. Codevilla, F., Santana, E., López, A.M., Gaidon, A.: Exploring the limitations of behavior cloning for autonomous driving. In: ICCV (2019)
16. Dang, B., Zhao, W., Li, Y., Ma, D., Yu, Q., Zhu, E.Y.: Real-time pill identification for the visually impaired using deep learning. arXiv (2024)
17. Ding, S., Lin, D.: Dynamic task allocation for cost-efficient edge cloud computing. In: SCC (2020)
18. Dong, Q., Chen, X., Satyanarayanan, M.: Creating edge ai from cloud-based llms. In: IWM-CSA (2024)
19. Dosovitskiy, A., Ros, G., Codevilla, F., Lopez, A., Koltun, V.: Carla: An open urban driving simulator. In: CoRL (2017)
20. Dudek, G., Jenkin, M.: Computational principles of mobile robotics. CUP (2024)
21. El Haber, E., Nguyen, T.M., Ebrahimi, D., Assi, C.: Computational cost and energy efficient task offloading in hierarchical edge-clouds. In: PIMRC (2018)
22. Frantar, E., Alistarh, D.: Spdy: Accurate pruning with speedup guarantees. In: ICML (2022)
23. Fu, Z., Kumar, A., Malik, J., Pathak, D.: Minimizing energy consumption leads to the emergence of gaits in legged robots. In: CoRL (2021)

24. Gamanayake, C., Jayasinghe, L., Ng, B.K.K., Yuen, C.: Cluster pruning: An efficient filter pruning method for edge ai vision applications. *JSTSP* (2020)
25. Gan, Y., Pan, M., Zhang, R., Ling, Z., Zhao, L., Liu, J., Zhang, S.: Cloud-device collaborative adaptation to continual changing environments in the real-world. In: *CVPR* (2023)
26. Hadidi, R., Cao, J., Xie, Y., Asgari, B., Krishna, T., Kim, H.: Characterizing the deployment of deep neural networks on commercial edge devices. In: *IISWC* (2019)
27. Han, Y., Huang, G., Song, S., Yang, L., Wang, H., Wang, Y.: Dynamic neural networks: A survey. *PAMI* (2021)
28. Hanyao, M., Jin, Y., Qian, Z., Zhang, S., Lu, S.: Edge-assisted online on-device object detection for real-time video analytics. In: *INFOCOM* (2021)
29. Hawke, J., Shen, R., Gurau, C., Sharma, S., Reda, D., Nikolov, N., Mazur, P., Micklethwaite, S., Griffiths, N., Shah, A., et al.: Urban driving with conditional imitation learning. In: *ICRA* (2020)
30. Hooghiemstra, G., Van Mieghem, P.: Delay distributions on fixed internet paths. Delft University of Technology, report (2001)
31. Hu, L., Sun, G., Ren, Y.: Coedge: Exploiting the edge-cloud collaboration for faster deep learning. *Access* (2020)
32. Hu, X., Wang, L., Wong, K.K., Tao, M., Zhang, Y., Zheng, Z.: Edge and central cloud computing: A perfect pairing for high energy efficiency and low-latency. *Trans. Wirel.* (2019)
33. Huang, Z., Shangguan, Z., Zhang, J., Bar, G., Boyd, M., Ohn-Bar, E.: ASSISTER: Assistive navigation via conditional instruction generation. In: *ECCV* (2022)
34. Hussein, A., Gaber, M.M., Elyan, E., Jayne, C.: Imitation learning: A survey of learning methods. *CSUR* (2017)
35. Jazbec, M., Timans, A., Veljković, T.H., Sakmann, K., Zhang, D., Naesseth, C.A., Nalisnick, E.: Fast yet safe: Early-exiting with risk control. *arXiv* (2024)
36. Jiang, Y., Wang, S., Valls, V., Ko, B.J., Lee, W.H., Leung, K.K., Tassiulas, L.: Model pruning enables efficient federated learning on edge devices. *Trans. Neural Netw. Learn. Syst.* (2022)
37. Jin, C., Che, T., Peng, H., Li, Y., Pavone, M.: Learning from teaching regularization: Generalizable correlations should be easy to imitate. *arXiv* (2024)
38. Judd, P., Albericio, J., Hetherington, T., Aamodt, T.M., Moshovos, A.: Stripes: Bit-serial deep neural network computing. In: *MICRO* (2016)
39. Kag, A., Fedorov, I., Gangrade, A., Whatmough, P., Saligrama, V.: Efficient edge inference by selective query. In: *ICLR* (2022)
40. Kang, Y., Hauswald, J., Gao, C., Rovinski, A., Mudge, T., Mars, J., Tang, L.: Neurosurgeon: Collaborative intelligence between the cloud and mobile edge. *SIGARCH* (2017)
41. Karras, K., Pallis, E., Mastorakis, G., Nikoloudakis, Y., Batalla, J.M., Mavromoustakis, C.X., Markakis, E.: A hardware acceleration platform for AI-based inference at the edge. *CSSP* (2020)
42. Katyal, K.D., Hager, G.D., Huang, C.M.: Intent-aware pedestrian prediction for adaptive crowd navigation. In: *ICRA* (2020)
43. Ke, Z., Zou, Q., Liu, J., Qian, S.: Real-time system optimal traffic routing under uncertainties—can physics models boost reinforcement learning? *arXiv* (2024)
44. Kim, S.W., Ko, K., Ko, H., Leung, V.C.: Edge-network-assisted real-time object detection framework for autonomous driving. *Network* (2021)
45. Kimovski, D., Mathá, R., Hammer, J., Mehran, N., Hellwagner, H., Prodan, R.: Cloud, fog, or edge: Where to compute? *Internet Comput.* (2021)
46. Knox, W.B., Allievi, A., Banzhaf, H., Schmitt, F., Stone, P.: Reward (mis) design for autonomous driving. *J. Artif. Intell.* (2023)
47. Kong, Z., Dong, P., Ma, X., Meng, X., Sun, M., Niu, W., Shen, X., Yuan, G., Ren, B., Qin, M., et al.: SPViT: Enabling faster vision transformers via soft token pruning. *arXiv* (2021)

48. Kryzhanovskiy, V., Balitskiy, G., Kozyrskiy, N., Zuruev, A.: Qpp: Real-time quantization parameter prediction for deep neural networks. In: CVPR (2021)
49. Lai, L., Ohn-Bar, E., Arora, S., Yi, J.S.K.: Uncertainty-guided never-ending learning to drive. In: CVPR (2024)
50. Lai, L., Shangguan, Z., Zhang, J., Ohn-Bar, E.: XVO: Generalized visual odometry via cross-modal self-training. In: ICCV (2023)
51. Laskaridis, S., Kouris, A., Lane, N.D.: Adaptive inference through early-exit networks: Design, challenges and directions. In: IWEMDL (2021)
52. Laskaridis, S., Venieris, S.I., Almeida, M., Leontiadis, I., Lane, N.D.: Spinn: synergistic progressive inference of neural networks over device and cloud. In: MobiCom (2020)
53. Levine, R.V., Norenzayan, A.: The pace of life in 31 countries. *J. Cross-Cult. Psychol.* (1999)
54. Li, M., Li, Y., Tian, Y., Jiang, L., Xu, Q.: Appealnet: An efficient and highly-accurate edge/cloud collaborative architecture for dnn inference. In: DAC (2021)
55. Li, X., Dang, Y., Aazam, M., Peng, X., Chen, T., Chen, C.: Energy-efficient computation offloading in vehicular edge cloud computing. *Access* (2020)
56. Li, X., Yang, Y., Yuan, Y., Ma, Y., Huang, Y., Ni, H.: Intelligent vehicle classification system based on deep learning and multi-sensor fusion. *Preprints* (2024)
57. Li, Z., Ren, T., He, X., Liu, C.: Red: A systematic real-time scheduling approach for robotic environmental dynamics. In: RTSS (2023)
58. Lin, B., Zhu, B., Ye, Y., Ning, M., Jin, P., Yuan, L.: Video-LLaVA: Learning united visual representation by alignment before projection. *arXiv* (2023)
59. Lin, B., Huang, Y., Zhang, J., Hu, J., Chen, X., Li, J.: Cost-driven off-loading for dnn-based applications over cloud, edge, and end devices. *Trans. Ind. Inform.* (2019)
60. Lin, L., Wu, W., Shangguan, Z., Wshah, S., Elmoudi, R., Xu, B.: Hpt-rl: Calibrating power system models based on hierarchical parameter tuning and reinforcement learning. In: ICMLA (2020)
61. Lin, Y., Zhang, J.w., Liu, H.: Deep learning based short-term air traffic flow prediction considering temporal-spatial correlation. *Aerosp. Sci. Technol.* (2019)
62. Liu, H., Li, C., Wu, Q., Lee, Y.J.: Visual instruction tuning. *NeurIPS* (2023)
63. Liu, L., Li, H., Gruteser, M.: Edge assisted real-time object detection for mobile augmented reality. In: ICMCN (2019)
64. Liu, R., Xu, X., Shen, Y., Zhu, A., Yu, C., Chen, T., Zhang, Y.: Enhanced detection classification via clustering svm for various robot collaboration task. *arXiv* (2024)
65. Loshchilov, I., Hutter, F.: Decoupled weight decay regularization. *arXiv* (2017)
66. Masoudi, M., Cavdar, C.: Device vs edge computing for mobile services: Delay-aware decision making to minimize power consumption. *Trans. Mob. Comput.* (2020)
67. Mavrogiannis, C., Baldini, F., Wang, A., Zhao, D., Trautman, P., Steinfeld, A., Oh, J.: Core challenges of social robot navigation: A survey. *T-HRI* (2023)
68. Mehta, S., Rastegari, M.: Mobilevit: light-weight, general-purpose, and mobile-friendly vision transformer. *arXiv* (2021)
69. Miki, T., Lee, J., Hwangbo, J., Wellhausen, L., Koltun, V., Hutter, M.: Learning robust perceptive locomotion for quadrupedal robots in the wild. *Sci. Robot.* (2022)
70. Moerland, T.M., Broekens, J., Plaat, A., Jonker, C.M., et al.: Model-based reinforcement learning: A survey. *Found. Trends Mach. Learn.* (2023)
71. Mohler, B.J., Thompson, W.B., Creem-Regehr, S.H., Pick, H.L., Warren, W.H.: Visual flow influences gait transition speed and preferred walking speed. *Exp. Brain Res.* (2007)
72. Muşat, V., Fursa, I., Newman, P., Cuzzolin, F., Bradley, A.: Multi-weather city: Adverse weather stacking for autonomous driving. In: CVPR (2021)
73. Nguyen, D.T., Nguyen, T.N., Kim, H., Lee, H.J.: A high-throughput and power-efficient fpga implementation of yolo cnn for object detection. *VLSI* (2019)

74. Ohn-Bar, E., Prakash, A., Behl, A., Chitta, K., Geiger, A.: Learning situational driving. In: CVPR (2020)
75. Pan, Y., Cheng, C.A., Saigol, K., Lee, K., Yan, X., Theodorou, E., Boots, B.: Agile autonomous driving using end-to-end deep imitation learning. arXiv (2017)
76. Pan, Y., Cheng, C.A., Saigol, K., Lee, K., Yan, X., Theodorou, E.A., Boots, B.: Imitation learning for agile autonomous driving. Int. J. Robot. Res. (2020)
77. Park, S.H., Jeong, S., Na, J., Simeone, O., Shamai, S.: Collaborative cloud and edge mobile computing in c-ran systems with minimal end-to-end latency. Trans. Signal Inf. Process. (2021)
78. Park, S., Kwon, D., Kim, J., Lee, Y.K., Cho, S.: Adaptive real-time offloading decision-making for mobile edges: deep reinforcement learning framework and simulation results. Appl. Sci. (2020)
79. Penmetcha, M., Min, B.C.: A deep reinforcement learning-based dynamic computational offloading method for cloud robotics. Access (2021)
80. Polino, A., Pascanu, R., Alistarh, D.: Model compression via distillation and quantization. In: ICLR (2018)
81. Qendro, L., Campbell, A., Lio, P., Mascolo, C.: Early exit ensembles for uncertainty quantification. In: ML4H (2021)
82. Radosavovic, I., Kosaraju, R.P., Girshick, R., He, K., Dollár, P.: Designing network design spaces. In: CVPR (2020)
83. Sandler, M., Howard, A., Zhu, M., Zhmoginov, A., Chen, L.C.: Mobilenetv2: Inverted residuals and linear bottlenecks. In: CVPR (2018)
84. Satyanarayanan, M., Beckmann, N., Lewis, G.A., Lucia, B.: The role of edge offload for hardware-accelerated mobile devices. In: IWMCSA (2021)
85. Schneider, N., Piewak, F., Stiller, C., Franke, U.: Regnet: Multimodal sensor registration using deep neural networks. In: IV (2017)
86. Schulman, J., Wolski, F., Dhariwal, P., Radford, A., Klimov, O.: Proximal policy optimization algorithms. arXiv (2017)
87. Shah, D., Osiński, B., Levine, S., et al.: LM-Nav: Robotic navigation with large pre-trained models of language, vision, and action. In: CoRL (2023)
88. Shangguan, Z., Lin, L., Wu, W., Xu, B.: Neural process for black-box model optimization under bayesian framework. arXiv (2021)
89. Shao, Y.S., Clemons, J., Venkatesan, R., Zimmer, B., Fojtik, M., Jiang, N., Keller, B., Klinefelter, A., Pinckney, N., Raina, P., et al.: Simba: Scaling deep-learning inference with multi-chip-module-based architecture. In: MICRO (2019)
90. Song, Y., Arora, P., Singh, R., Varadharajan, S.T., Haynes, M., Starner, T.: Going blank comfortably: Positioning monocular head-worn displays when they are inactive. In: ISWC (2023)
91. Song, Y., Arora, P., Varadharajan, S.T., Singh, R., Haynes, M., Starner, T.: Looking from a different angle: Placing head-worn displays near the nose. In: AHs (2024)
92. Sudhakar, S., Sze, V., Karaman, S.: Data centers on wheels: emissions from computing onboard autonomous vehicles. Micro (2022)
93. Tang, S., Wang, Y., Kong, Z., Zhang, T., Li, Y., Ding, C., Wang, Y., Liang, Y., Xu, D.: You need multiple exiting: Dynamic early exiting for accelerating unified vision language model. In: CVPR (2023)
94. Teng, S., Chen, L., Ai, Y., Zhou, Y., Xuanyuan, Z., Hu, X.: Hierarchical interpretable imitation learning for end-to-end autonomous driving. T-IV (2022)
95. Touvron, H., Lavril, T., Izacard, G., Martinet, X., Lachaux, M.A., Lacroix, T., Rozière, B., Goyal, N., Hambro, E., Azhar, F., et al.: LLaMA: Open and efficient foundation language models. arXiv (2023)

96. Treuillet, S., Royer, E.: Outdoor/indoor vision-based localization for blind pedestrian navigation assistance. *Int. J. Image Graph.* (2010)
97. Tsai, C.E., Oh, J.: A generative approach for socially compliant navigation. In: ICRA (2020)
98. Van Tam, N., Hieu, N.Q., Van, N.T.T., Luong, N.C., Niyato, D., Kim, D.I.: Adaptive task offloading in coded edge computing: A deep reinforcement learning approach. *COMML* (2021)
99. Wang, J.: Lightweight and real-time object detection model on edge devices with model quantization. In: JPCS (2021)
100. Wang, J., Hu, J., Min, G., Zhan, W., Ni, Q., Georgalas, N.: Computation offloading in multi-access edge computing using a deep sequential model based on reinforcement learning. *Commun. Mag.* (2019)
101. Wang, J., Hu, J., Min, G., Zomaya, A.Y., Georgalas, N.: Fast adaptive task offloading in edge computing based on meta reinforcement learning. *TPDS* (2020)
102. Wang, K., Liu, Z., Lin, Y., Lin, J., Han, S.: Haq: Hardware-aware automated quantization with mixed precision. In: CVPR (2019)
103. Wang, T., Chang, D.E.: Robust navigation for racing drones based on imitation learning and modularization. In: ICRA (2021)
104. Wiering, M.A., Van Otterlo, M.: Reinforcement learning. *Adapt. Learn. Optim.* (2012)
105. Wu, H., Zhang, Z., Guan, C., Wolter, K., Xu, M.: Collaborate edge and cloud computing with distributed deep learning for smart city internet of things. *IoT-J* (2020)
106. Wu, N., Wang, H.: Deep learning adaptive dynamic programming for real time energy management and control strategy of micro-grid. *J. Clean. Prod.* (2018)
107. Wu, T., Song, C., Zeng, P.: Model pruning based on filter similarity for edge device deployment. *Front. Neurorobot.* (2023)
108. Xia, G., Bouganis, C.S.: Window-based early-exit cascades for uncertainty estimation: When deep ensembles are more efficient than single models. In: ICCV (2023)
109. Yang, M., Li, Z., Ke, Z., Li, M.: A deep reinforcement learning-based ramp metering control framework for improving traffic operation at freeway weaving sections. In: TRB (2019)
110. Yang, T.J., Chen, Y.H., Sze, V.: Designing energy-efficient convolutional neural networks using energy-aware pruning. In: CVPR (2017)
111. Yao, S., Li, J., Liu, D., Wang, T., Liu, S., Shao, H., Abdelzaher, T.: Deep compressive offloading: Speeding up neural network inference by trading edge computation for network latency. In: SenSys (2020)
112. Zhang, J., Huang, Z., Ohn-Bar, E.: Coaching a teachable student. In: CVPR (2023)
113. Zhang, J., Huang, Z., Ray, A., Ohn-Bar, E.: Feedback-guided autonomous driving. In: CVPR (2024)
114. Zhang, J., Zheng, M., Boyd, M., Ohn-Bar, E.: X-World: Accessibility, vision, and autonomy meet. In: ICCV (2021)
115. Zhang, J., Zhu, R., Ohn-Bar, E.: SelfD: Self-learning large-scale driving policies from the web. In: CVPR (2022)
116. Zhang, W., He, J.: Modeling end-to-end delay using pareto distribution. In: ICIMP (2007)
117. Zhang, X., Zhang, H., Zhou, X., Yuan, D.: Energy minimization task offloading mechanism with edge-cloud collaboration in iot networks. In: VTC (2021)
118. Zheng, S., Li, M., Ke, Z., Li, Z.: Coordinated variable speed limit control for consecutive bottlenecks on freeways using multiagent reinforcement learning. *J. Adv. Transp.* (2023)
119. Zhu, K., Zhang, T.: Deep reinforcement learning based mobile robot navigation: A review. *TST* (2021)
120. Zhu, R., Huang, P., Ohn-Bar, E., Saligrama, V.: Learning to drive anywhere. *CoRL* (2023)
121. Zhu, W., Rosendo, A.: Psto: Learning energy-efficient locomotion for quadruped robots. *Machines* (2022)

122. Zhu, Y., Mottaghi, R., Kolve, E., Lim, J.J., Gupta, A., Fei-Fei, L., Farhadi, A.: Target-driven visual navigation in indoor scenes using deep reinforcement learning. In: ICRA (2017)
123. Zhuang, Z., Fu, Z., Wang, J., Atkeson, C., Schwertfeger, S., Finn, C., Zhao, H.: Robot parkour learning. In: CoRL (2023)

Supplementary for Unified Local-Cloud Decision-Making via Residual Reinforcement Learning

Abstract. This supplementary provides additional implementation and ablative details. Specifically, we discuss additional details regarding the introduced cloud and energy-aware environment as well as additional ablative analysis. Our supplementary video depicts qualitative results, comparing roll-outs generated by the proposed model to baseline models.

1 Environment Implementation

Energy Cost Model: To quantify energy consumption, we assume a standard cost per floating point operation model [26, 40, 56]. For instance, a 675 Kb $480 \times 480 \times 3$ image and a 1.37 M parameters MobileVIT [68] model would result in energy consumption on a local GPU that is 0.15 J (based on a factor of 0.095 J per flop [40]). We further add an energy cost for every communication, i.e., transmitting raw data or an embedding to the cloud [39]. Here, energy consumed for local-cloud communication is computed as 6.94×10^{-5} J per byte [40]. When transmitting only an embedding (a 24×24 array) from the local policy as an intermediate input to the cloud policy, the communication energy consumption decreases, e.g., from the original image which requires 1.55 J to 25.18 mJ.

Cloud Communication Model: Cloud transmission involves a latency, which the model should optimize over, e.g., in safety-critical scenarios based on the task. Our communication model resembles the one used by Kang et al [40], where latency is modeled as a Gaussian with an average time of 0.5 s and variance of 0.1 s. To extend our research, we also experimented with modeling latency using a Pareto distribution. Under the Pareto distribution, we observed a 7% reduction in ENS, with an ENS value of 80.45, infractions at 0.03 /m, and energy consumption at 2.76 J/m. While this is a conservative estimate, we will release our simulation, code, and models to facilitate the analysis of RL agents across diverse computational and latency configurations.

1.1 Reward Design

Reward Progression: Fig. 3 shows our reward analysis for RL baselines and UniLCD models. The reward progression indicates that the baselines fail to learn any meaningful information throughout the training and continue to oscillate within a very narrow range. Our UniLCD demonstrates a slight improvement in the initial episodes but exhibits delayed convergence due to its slower learning rate. Remarkably, the UniLCD with history model displays a distinct pattern of rapid advancement, characterized by a steep rise in the initial episodes. Consequently, it is clear that this configuration learns and adapts to the desired behavior more efficiently than the comparative approaches,

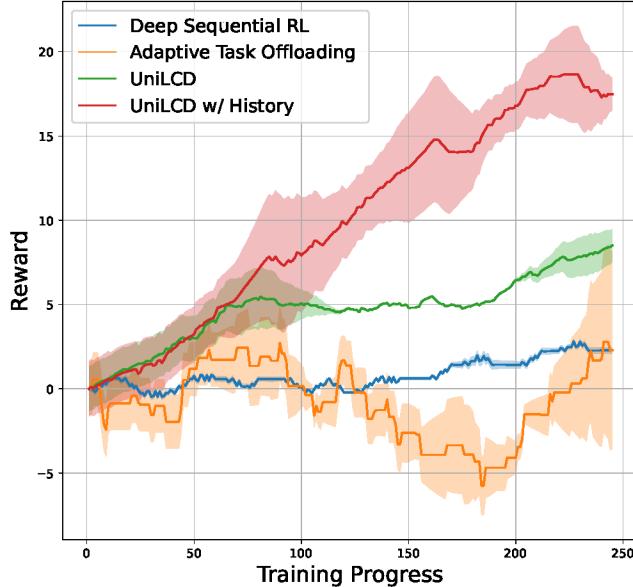


Fig. 3: Training Progression. We evaluate the rewards for our overall framework throughout the training process against baseline models. We demonstrate high sample-efficiency compared to prior methods, particularly when the routing module is given a history input, i.e., prior actions and their source (cloud or local decision). Despite common instabilities in training reinforcement learning models, UniLCD is shown to achieve significantly higher reward early in the training process.

even within a few iterations. The primary benefits in sample efficiency stem from the definition of the reward function definition and the careful design of the routing module and its inputs.

Comparison with Standard Additive Reward: In the main paper, we find it crucial to carefully shape the reward function in order to effectively optimize for several different types of rewards. Our reward is contrasted with current standard approaches that linearly combine the rewards [5, 43, 46, 123], i.e.,

$$r_{\text{standard}} = \alpha_g r_{\text{geo}} + \alpha_s r_{\text{speed}} + \alpha_e r_{\text{energy}} + \alpha_c r_{\text{collision}} \quad (12)$$

adding the geodesic, speed, energy, and collision rewards and $\alpha_g, \alpha_s, \alpha_e, \alpha_c$, are tuned scalar hyperparameters (we perform careful tuning of these with a grid search).

Impact of Reward Components: Table 4 analyzes the impact of various reward terms on task performance. We show each term to contribute to the overall task performance. We note that we do not multiply the collision reward with the other terms in order to emphasize it during training. This provides a **major negative disadvantage to avoid collisions**, as it dominates the overall reward once it occurs (and results in episode termination). Moreover, we show that without penalizing energy consumption of communication to and processing in the cloud, the model effectively learns to rely on the

Table 4: Reward Component Ablation. We investigate the impact of each reward component in our reward design for UniLCD. ENS is Ecological Navigation Score (%), NS is Navigation Score (%), SR is Success Rate (%), RC is Route Completion (%), Infract. is Infraction Rate (/m), Energy is measured in Joules per meter (J/m), and FPS is Frames Per Second (empirically measured and averaged throughout the trials).

Reward	ENS↑	NS↑	SR↑	RC↑	Infract.↓	Energy↓	FPS↑
All Terms	85.97	94.58	93.33	95.90	0.02	2.90	26.49
w/o r_{geo}	67.04	74.65	0.00	76.22	0.03	7.42	58.82
w/o r_{speed}	66.05	72.66	0.00	77.34	0.09	6.61	65.40
w/o r_{energy}	0.00	93.53	90.00	95.50	0.03	43.75	6.42

more powerful model. Finally, while the imitation-learned models take as input the goal, adding a geodesic reward term is found to help ensure the routing module and final action prediction do not result in a policy that deviates from the planned path.

Real-world Validation: Our simulation experiments are based on real-world network latency calculations, as detailed in Sec 1. To extend our research into real-world validations, we conducted a small-scale experiment using an RC vehicle platform (Traxxas X-Maxx) equipped with a Jetson Nano. In this setup, reward terms were computed in a self-supervised manner using an EAI XL2 LiDAR with a minimum range of 0.1 meters. Given the need for high sample efficiency in real-world reinforcement learning (RL), we performed an ablation study comparing Proximal Policy Optimization (PPO) and CrossQ. In Line 16 of Algorithm 1, CrossQ learns a policy in approximately five minutes, demonstrating a 21% improvement over [31].

Further Work: While our work takes a step towards developing vision-based systems and effectively address multifaceted objectives, several simplifications were necessary. Due to the difficulty and sample inefficiency associated with the optimization of RL agents for energy, latency, and safety constraints – all are crucial in real-world settings – we adopted a step-by-step training approach. Specifically, the local and cloud policies were trained using imitation learning by assuming expert trajectories. Our UniLCD pipeline can readily support additional model fine-tuning and incorporation of consumption-based reward terms, e.g., both the cloud and local models can be fine-tuned using RL, yet this requires additional study in the future.

1.2 Performance Over Task Difficulty

The results in the main paper are shown for the high (30 pedestrians) density settings of our environment, spawned along the route (Fig. 4). For completeness, we report in Table 5 the model performance across additional pedestrian density settings along the path, starting from a low-density scenario with 5 pedestrians to a crowded scenario with 70 pedestrians. As shown in the table, several models degrade in performance, particularly in crowded settings. Remarkably, our final UniLCD model exhibits only slight degradation (1%), in terms of ENS, in the most challenging and crowded settings.



Fig. 4: Examples of Different Environmental Settings in Our Robot Navigation Environment. We vary the pedestrian density in order to stress-test the proposed UniLCD method. Pedestrian count along the path ranges from 5 (Low), 15 (Medium), 30 (High), and 70 (Crowd).

Table 5: Pedestrian Density Ablation. Model performance is shown for different pedestrian densities along the path, from Low (mostly empty, 5 pedestrians), Medium (15 pedestrians), High (30 pedestrians), and up to Crowded (70 pedestrians). \dagger denotes methods transmitting to the cloud raw input data instead of an embedding. ENS is Ecological Navigation Score (%), NS is Navigation Score (%), SR is Success Rate (%), RC is Route Completion (%), Infract is Infraction Rate (/m), Energy is measured in Joules per meter (J/m), and FPS is Frames Per Second.

Method	Setting	ENS \uparrow	NS \uparrow	SR \uparrow	RC \uparrow	Infract. \downarrow	Energy \downarrow	FPS \uparrow
\dagger Cloud-Only [82]	Low	0.00	98.41	95.00	99.10	0.01	33.93	8.91
	Medium	0.00	97.14	93.33	98.50	0.02	35.11	8.74
	High	0.00	96.47	93.33	98.50	0.03	36.49	7.11
	Crowd	0.00	88.24	92.00	95.24	0.11	48.93	6.52
Local-Only [83]	Low	66.18	69.70	0.00	75.23	0.11	3.78	78.19
	Medium	65.99	69.70	0.00	75.23	0.11	3.98	73.67
	High	63.43	67.33	0.00	75.23	0.16	4.33	68.40
	Crowd	35.47	37.55	0.00	63.16	0.75	4.94	65.40
<i>Baseline Methods:</i>								
SPINN [52]	Low	43.71	85.13	70.00	95.12	0.16	16.88	23.84
	Medium	39.94	80.60	66.66	93.24	0.21	16.74	24.21
	High	36.31	72.75	60.00	92.73	0.35	18.94	20.37
	Crowd	32.07	53.22	60.00	90.76	0.77	17.16	21.84
Compressive Offloading [111]	Low	11.81	80.43	0.00	80.43	0.00	88.23	2.46
	Medium	12.38	80.16	0.00	80.16	0.00	89.12	2.17
	High	13.98	80.16	0.00	80.16	0.00	90.66	1.82
	Crowd	10.06	75.66	0.00	75.66	0.00	92.66	1.19
\dagger Selective Query [39]	Low	27.87	64.78	0.00	82.68	0.03	42.64	21.57
	Medium	25.40	60.63	0.00	80.12	0.08	44.87	19.39
	High	24.14	61.28	0.00	82.68	0.11	45.35	18.14
	Crowd	14.90	44.27	0.00	79.36	0.52	51.72	14.38
Neurosurgeon [40]	Low	43.79	67.62	0.00	85.12	0.01	25.62	12.53
	Medium	42.12	65.12	0.00	82.54	0.02	26.48	13.21
	High	39.85	63.10	0.00	80.54	0.03	28.31	12.53
	Crowd	32.94	51.19	0.00	75.58	0.24	29.18	9.61
\dagger Adaptive Offloading [98]	Low	89.38	97.15	90.0	98.51	0.02	5.02	34.91
	Medium	51.87	55.56	80.00	97.42	0.81	4.22	33.39
	High	37.42	40.37	70.00	94.05	1.22	4.80	30.14
	Crowd	33.96	37.67	80.00	95.38	1.34	6.39	28.71
Deep Sequential RL [100]	Low	70.15	73.02	0.00	79.36	0.12	3.07	78.10
	Medium	66.66	69.56	0.00	79.36	0.19	3.25	78.03
	High	58.84	61.83	0.00	79.36	0.36	3.77	77.94
	Crowd	26.47	27.93	0.00	75.28	1.43	4.30	76.12
<i>UniLCD Variations:</i>								
\dagger Standard Reward	Low	49.37	56.15	0.00	75.23	0.10	2.47	52.63
	Medium	49.37	56.15	0.00	75.23	0.10	3.12	51.18
	High	48.35	54.99	0.00	75.23	0.13	3.57	50.20
	Crowd	30.60	34.80	0.00	75.23	0.79	4.86	47.74
\dagger Standard Reward w/ History	Low	51.44	58.81	10.00	77.71	0.08	7.02	51.42
	Medium	50.38	57.60	10.00	77.71	0.11	7.57	50.20
	High	50.04	57.20	10.00	77.71	0.12	8.38	49.07
	Crowd	35.22	40.26	10.00	75.23	0.58	10.57	43.16
\dagger Our Reward	Low	51.82	85.73	60.00	90.00	0.07	18.49	20.82
	Medium	51.20	84.71	60.00	92.70	0.13	20.57	17.14
	High	48.30	79.90	56.66	91.15	0.19	21.72	16.05
	Crowd	38.79	64.16	60.00	86.45	0.43	24.57	11.20
\dagger Our Reward w/ History	Low	73.72	90.17	83.33	94.66	0.07	5.78	36.61
	Medium	72.20	88.32	83.33	94.66	0.10	6.51	34.73
	High	71.70	87.71	83.33	94.66	0.11	7.83	33.98
	Crowd	59.04	72.22	83.33	93.33	0.37	10.83	30.62
Our Reward	Low	61.88	94.54	60.00	95.20	0.01	5.51	14.18
	Medium	60.68	92.71	60.00	94.66	0.03	5.88	13.72
	High	57.20	87.39	60.00	91.10	0.06	6.60	12.49
	Crowd	50.22	76.74	60.00	90.00	0.23	8.43	10.73
Our Reward w/ History	Low	88.10	96.92	93.33	97.60	0.01	2.58	30.02
	Medium	87.06	95.77	90.00	96.44	0.01	2.63	29.12
	High	85.97	94.58	93.33	95.90	0.02	2.90	26.49
	Crowd	83.04	91.36	93.33	95.90	0.07	3.12	22.84