

Tabular Reinforcement Learning Methods for Artificial Intelligence Tasks Offloading in Smart Eye-Wears

ABEDNEGO WAMUHINDO KAMBALE, HAMTA SEDGHANI, FEDERICA FILIPPINI, GIACOMO VERTICALE, and DANILO ARDAGNA, Politecnico di Milano, Italy

Virtual and Extended Reality technologies are increasingly adopted in fields such as healthcare, entertainment, and education. These applications heavily rely on Smart Eye-Wears (SEWs) and Artificial Intelligence (AI) to provide users with new ways to perceive their environment. However, SEWs face limitations in computational power, memory, and battery life. Offloading computations to external servers is a prominent example of edge computation. However, this also presents considerable challenges due to delays caused by varying network conditions and server workloads. This paper proposes self-adaptive techniques based on tabular reinforcement learning (RL) to optimize the offloading of Deep Neural Network tasks between the SEW, the user's smartphone, and cloud servers. The goal is to maintain a high-quality user experience while minimizing energy consumption and 5G connection costs. We evaluated our framework under varying 5G and WiFi bandwidths and cloud latency. The results show that Q-learning, SARSA, and Expected SARSA achieve near-optimal policies, with Q-learning demonstrating superior performance in reducing execution time violations (approximately at 10%) and improving agent stability. Additionally, our approach offers a more favorable trade-off between energy efficiency and execution time violations compared to two baseline methods. Real-system experiments reveal that the proposed solution can double SEW battery life with respect to local computation while maintaining a good quality of service, with only 11% execution time violations. These findings highlight the effectiveness of our approach in managing resources and enhancing the overall user experience in SEW AI applications.

CCS Concepts: • **Computing methodologies** → **Intelligent agents**; • **Information systems** → *Computing platforms*.

Additional Key Words and Phrases: Smart glasses, Smart Eye-Wear, Reinforcement Learning, Edge Computing, Task offloading

ACM Reference Format:

Abednego Wamuhindo Kambale, Hamta Sedghani, Federica Filippini, Giacomo Verticale, and Danilo Ardagna. 2018. Tabular Reinforcement Learning Methods for Artificial Intelligence Tasks Offloading in Smart Eye-Wears. In *Proceedings of Make sure to enter the correct conference title from your rights confirmation email (Conference acronym 'XX)*. ACM, New York, NY, USA, 36 pages. <https://doi.org/XXXXXXXX.XXXXXXX>

1 INTRODUCTION

Artificial Intelligence (AI) recently became pervasive, finding application in various domains such as healthcare, agriculture, and entertainment [10, 18, 33, 40]. In all these fields, AI has yielded promising outcomes: e.g., in healthcare, AI algorithms play a pivotal role in expediting diagnoses and improving

Authors' Contact Information: [Abednego Wamuhindo Kambale](mailto:abednegowamuhindo.kambale@polimi.it), abednegowamuhindo.kambale@polimi.it; [Hamta Sedghani](mailto:hamta.sedghani@polimi.it), hamta.sedghani@polimi.it; [Federica Filippini](mailto:federica.filippini@polimi.it), federica.filippini@polimi.it; [Giacomo Verticale](mailto:giacomo.verticale@polimi.it), giacomo.verticale@polimi.it; [Danilo Ardagna](mailto:daniilo.ardagna@polimi.it), daniilo.ardagna@polimi.it, Politecnico di Milano, Milan, Milan, Italy.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

Conference acronym 'XX, June 03–05, 2018, Woodstock, NY

© 2018 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 978-1-4503-XXXX-X/18/06

<https://doi.org/XXXXXXXX.XXXXXXX>

50 patient well-being [10]. Meanwhile, in agriculture, AI aids farmers in making informed decisions to
51 optimize crop performance while minimizing resource wastage [33]. In the realm of entertainment,
52 AI applications have elevated user satisfaction levels and delivered more efficient and captivating
53 experiences [18, 40]. In an effort to offer advanced features, many smart device manufacturers are
54 integrating AI models into various user devices (e.g., smartphones) as well as into smaller, resource-
55 constrained devices (e.g., smart glasses) [50]. These AI applications predominantly rely on complex
56 image-based processes and frequently involve the utilization of Deep Neural Networks (DNNs). Due
57 to their large number of parameters, these DNNs necessitate considerable computational power
58 and memory, particularly in applications where real-time processing is required.

59 Although significant advancements have been made in enhancing the memory and computational
60 capabilities of IoT and edge devices, challenges related to limited battery capacity and computational
61 power persist when executing complex AI models [4]. Depending on the circumstances, these
62 constraints may make it impractical to execute all the DNN layers on a single device. Moreover,
63 the limitation posed by battery capacity detrimentally affects the user experience.

64 Several studies [23, 24, 35] proposed DNN partitioning as a solution to address these challenges.
65 DNN partitioning involves splitting the DNN into different parts, each suitable for running on
66 different devices across a computing continuum. This partitioning offers various configurations,
67 allowing the system to switch between them to optimize performance. An emerging application for
68 this technology is found in smart glasses (smart eye-wear) and virtual/augmented reality scenarios.
69 Smart Eye-Wear (SEW) devices collect data from their surroundings, process a portion of the DNN
70 locally, and delegate the remaining computation to a mobile phone or the cloud exploiting all the
71 resources available in the computing continuum. However, communication between the SEW and
72 the mobile phone, as well as between the mobile phone and the cloud, may experience fluctuations,
73 potentially impacting application performance. Therefore, the most appropriate configuration to
74 use in each time interval needs to be carefully determined to ensure a satisfying user experience
75 while adapting to fluctuations in network bandwidth and cloud back-end latency.

76 This paper introduces a self-adaptive runtime optimization framework, leveraging Reinforcement
77 Learning (RL) to distribute the computation of AI applications over various computing units. The
78 goal is to reduce energy and 5G networking costs from the user's viewpoint while maintaining
79 end-to-end latency at levels that preserve a satisfactory user experience.

80 The main contributions of this paper are the following:

- 81 - We design a three-layer computing continuum architecture to model the different compo-
82 nents when managing the execution of SEW AI applications. We formulate the general
83 problem as a Markov Decision Process (MDP) that can be solved using different approaches.
84 We consider image-processing inference applications where frames are provided as input
85 to the DNN at a given rate.
- 86 - We propose an RL-based runtime management solution in which an agent selects the optimal
87 configuration to minimize energy consumption and 5G network costs from the user's
88 perspective, while ensuring that end-to-end latency requirements are met. We implement
89 our proposed solution by comparing different tabular methods including Q-learning [49],
90 Double Q-learning [17], weighted Q-learning [9], SARSA [42] and Expected SARSA [43] to
91 evaluate their performance and identify the ones yielding the most promising results.
- 92 - We evaluate the proposed solution by accounting for network bandwidth and cloud latency
93 variability, as well as by adjusting execution time constraints and the distribution of im-
94 portance weights for the agent's objectives. The results identify Q-learning as the most
95 effective method, with execution time violations around 10%, suggesting that the strategy
96 is promising for efficient resource allocation and user satisfaction.

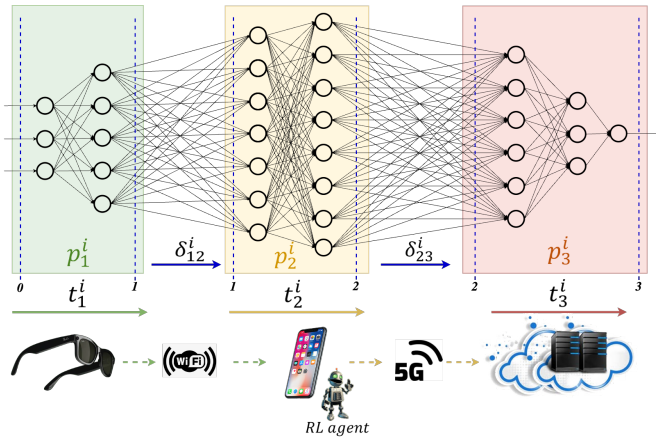


Fig. 1. Reference scenario for performing AI tasks.

- We compared our solution with two baselines: *Neurosurgeon*, a seminal work presented in [23], and the static solution of the underlying optimization problem, referred to as Opt^* , which considers the average WiFi and 5G bandwidths. Our approach demonstrated superior performance compared to both baselines, reducing execution time violations by 42% compared to Opt^* and achieving up to 41% and 19% energy savings compared to *Neurosurgeon* and Opt^* , respectively, while maintaining a comparable level of execution time violations to *Neurosurgeon*.
- Finally, we validate the performance of the trained agent on a real system prototype comprising the Microsoft HoloLens 2 [32], focusing on energy consumption and execution time constraint violations. Our agent achieves a 47% reduction in HoloLens energy consumption compared to running the entire DNN locally, effectively doubling the device’s battery life while incurring only 11% of execution time violations. These results demonstrate the practical utility of our approach.

The remainder of this paper is organized as follows. Section 2 details the scenario considered in this work. Section 3 presents the system model and formally defines the optimization problem, while Section 4 presents the RL-based formulation. In Section 5 we present the experimental setup and discuss the results of different analyses, including the validation on the real system prototype. Section 6 provides an overview the related literature proposals. Finally, Section 7 concludes the paper and presents future perspectives.

2 REFERENCE SCENARIO

This section presents the reference scenario foundational to our research. The SEW device runs an AI application encompassing a DNN capable of executing various AI functionalities related to image processing, e.g., object classification, detection or tracking.

The primary objective is to ensure a positive user experience by setting time thresholds on the AI application execution. For example, in the case of an object-tracking application, it is crucial to maintain a minimum application frame rate of 30 frames per second (fps), which translates to an end-to-end execution time of less than 33 ms [28]. For low-frame rate videos, a rate of 1–5 fps (i.e., 200–1,000 ms) is acceptable [2]. For Simultaneous Localization And Mapping (SLAM) applications, an acceptable frame rate falls within the range of 15–30 fps, implying that the end-to-end execution time should range in about 30–66 ms [15, 25].

148 **Figure 1** depicts the three primary elements in our system. In the remainder of this paper we
149 assume a forward-looking scenario where the reference hardware utilizes next-generation System-
150 on-Chips (SoCs) as projected in the Snapdragon series [38]. The SEW device (based on, e.g., a
151 Snapdragon AR2 chip [38]) will establish a connection with the mobile phone through WiFi 7.
152 Additionally, we assume that it incorporates components, including CPU and an AI neural engine,
153 to execute specific computations. Analogously, the mobile phone (integrating, e.g., a Snapdragon 8
154 Gen 2 chip [39]) features high-end computing components, such as a GPU.

155 The SEW collects data from its surroundings and transmits them to the mobile phone. Depending
156 on the current system state (e.g., battery level or network quality), a RL agent running on the
157 phone determines whether the acquired data are sent directly to the phone itself or undergo local
158 DNN processing before transmission. As already mentioned, a DNN-based AI application can be
159 characterized by different partitioning points (see **Figure 2**), which enable the dynamic computation
160 offloading [7, 23]. The phone will run the DNN from the layer where the SEW terminated the
161 execution, either completing the processing, if it has sufficient resources, or offloading part of it
162 to the cloud according to the RL agent decision. In such cases, the cloud handles the remaining
163 computations and returns the final results. In other terms, the RL agent is responsible for determining
164 the appropriate DNN configuration to run. Note that, as in shown in [23], only specific partitioning
165 points are relevant for optimizing energy consumption and data transfers among different system
166 components, and therefore are considered as viable candidates. When the RL agent selects a
167 configuration, it directly informs which partition should run on the SEW, on the mobile phone,
168 or in the cloud. The RL agent takes one action every few seconds, or whenever a given number
169 of consecutive violations (e.g., five consecutive frames processed at the application level) are
170 observed prior to the conclusion of the periodic control interval. The main goal is to choose a setup
171 that reduces energy usage and communication costs while ensuring adherence to execution time
172 constraints, thus maintaining a satisfactory user experience.

173 The various components of the system communicate through distinct network domains. The
174 phone uses a 5G network to transmit data to the cloud server. This results in a significant reduction
175 in transfer delay between the connected entities [3, 6] since 5G mmWave systems offer substantial
176 bandwidth, albeit with a notable degree of variability [34]. At the same time, WiFi interference can
177 become a concern when multiple devices operate at the same frequency, resulting in disruptions and
178 fluctuations in network throughput. This interference, which occurs, e.g., when two SEW end-users
179 come into close proximity, can result from overlapping signals and transmissions, reducing the
180 quality and reliability of the wireless connection. Consequently, the network throughput may
181 exhibit fluctuations, impacting the overall performance of WiFi-based systems and applications
182 [36]. Furthermore, cloud servers, which function as the computational backbone of the computing
183 continuum, exhibit inherent performance variability due to fluctuating workloads. This variability
184 impacts server queuing times, which in turn influence the overall end-to-end latency of AI appli-
185 cations, potentially affecting their performance and responsiveness [1]. While our management
186 system is designed to make local decisions to manage such variability, it is important to note that
187 cloud resource allocation remains beyond our control.

188 The total application processing time includes several components: the local execution time on
189 the SEW, the transmission latency from the SEW to the smartphone, the processing time on the
190 smartphone, the data-transfer latency from the smartphone to the cloud, and, finally, the cloud
191 processing time. In this study, we will neglect the time required to return results, as it is negligible
192 when compared to the overall execution time and data transfer delays. For example, in the context
193 of a classification task, the returned results include only the labels of the classified objects. On the
194 other hand, the results of an object-detection task encompass objects labels, their bounding boxes,
195
196

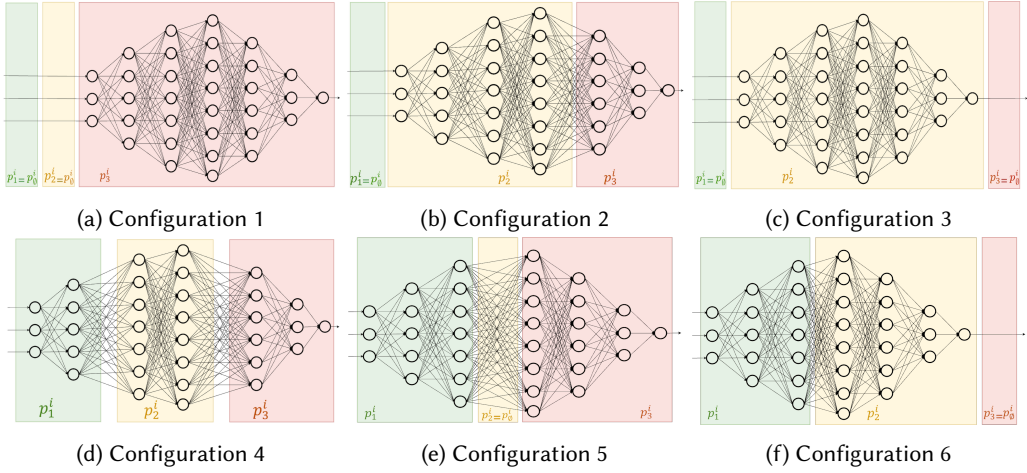


Fig. 2. Candidate configurations of the DNN.

and their class probabilities; while the size is surely larger than in a classification task, it is still significantly smaller than the intermediate tensors transmitted during the processing.

3 SYSTEM MODEL AND PROBLEM DEFINITION

In this section, we present the system model and formulate the runtime optimization problem of AI applications partitioning in the SEW computing continuum. In particular, Section 3.1 presents the DNN partition model. Section 3.2 describes the communication model we consider to determine the data transfer rates between system components for different network domains and the characterization of the 5G connection cost. Section 3.3 introduces how the latency is computed at the different system levels, Section 3.4 presents the energy consumption model, and Section 3.5 provides the complete problem formulation. For the reader's convenience, all the problem parameters and variables detailed in the next sections are summarized in Appendix A, Table 4.

3.1 DNN partitions model

In this work, we consider image-based AI applications that process requests supplied in the form of image frames with a frequency Λ (expressed in frames per second).

The decision made by the agent dictates whether the application should be run entirely on the SEW itself or if some DNN layers should be offloaded to either the smartphone or the cloud server [23] (in this case, Λ is also the rate at which the SEW interacts with the phone). This choice corresponds to the selection of a specific DNN configuration, and it is motivated by the fact that the SEW might have limited computational power, not enough to ensure that the whole AI application will be executed within the required time frame. Moreover, its battery may have limited residual capacity. Depending on the system state, the choice of a specific configuration is guided by the minimization of incurred costs. At runtime, in case the offloading is needed, the first partition of the selected DNN configuration will be executed on the SEW, the second one on the phone, and the last one on the cloud server. We denote by $\mathcal{D} = \{1, 2, 3\}$ the set of the computing component indices, so that $d = 1$ denotes the SEW, $d = 2$ the mobile phone and $d = 3$ the cloud server. Moreover, we define the set \mathcal{K} of all candidate configurations. Each element $\kappa^i \in \mathcal{K}$ represents a system configuration including three DNN partitions, defined as $\kappa^i = \{p_d^i\}_{d \in \mathcal{D}}$.

Figure 2 illustrates representative candidate configurations for a DNN. To simplify the modeling and the identification of the partition that is executed on a specific device, we introduce a partition p_0^i that does not include any DNN layer, meaning that no computation is executed. This allows us to establish a one-to-one correspondence between partitions and devices, and therefore to characterize a priori where each partition is deployed. Indeed, for every configuration κ^i , the first partition p_1^i is always executed on the SEW, the second partition p_2^i on the mobile phone and the third partition p_3^i on the cloud server. Figures 2a and 2c show cases where the whole DNN is executed on the cloud server, or the mobile phone, respectively; accordingly, the remaining partitions correspond to p_0^i . Figures 2b, 2e and 2f represent cases where a single DNN partitioning point is chosen: two partitions run on the phone and the cloud, the SEW and the phone, or the SEW and the cloud, respectively, while the third computational component remains idle (i.e., it is assigned p_0^i). Finally, Figure 2d illustrates the case of three partitions, executed by all the three computational components.

Throughout the application execution, the agent might switch from one configuration to another; for instance, Figures 2a and 2b might represent the solutions adopted when the SEW battery level is low. We assume that all the partitions fit in the devices memory, since those violating this constraint can be filtered a priori.

We denote the control time-period as τ . All the problem parameters we introduce in the following are measured in each of these time windows; unless differently specified, their value, which may vary from one interval to the other, will be considered as constant within τ . This allows us to drop the explicit time-dependency from the parameters definition, and, instead, add an apex τ to the parameter name only when we want to mark its relationship with a specific time slot.

We define x^i as a binary variable that equals 1 if configuration κ^i is selected, and we denote by $|\mathcal{K}|$ the total number of configurations. Note that one single configuration must be chosen in each time slot τ , which can be expressed by the following constraint:

$$\sum_{i=1}^{|\mathcal{K}|} x^i = 1. \quad (1)$$

When a computing component does not run the entire DNN or its last layer, it sends the intermediate tensor to the next component. Specifically, the SEW sends the data to the smartphone and this to the cloud. We denote by δ_{dh}^i the amount of per-request data related to configuration κ^i that is sent by partition p_d^i to partition p_h^i . Thus, δ_{12}^i is the amount of per-request data that are sent from the SEW to the smartphone, and δ_{23}^i from the smartphone to the cloud. When everything is executed on the SEW, we have $\delta_{12}^i = \delta_{23}^i = 0$, while in case everything is offloaded to the cloud $\delta_{12}^i = \delta_{23}^i = \delta_0$, where δ_0 is the size of the input tensor.

The parameter μ_d^i , measured in floating point operations, denotes the per-request computational workload of running the partition p_d^i . This information can be easily gathered by development frameworks such as PyTorch, TensorFlow, etc. This quantity, as first approximation [52], determines the energy consumption of the computing resources of the device d when running a partition p_d^i . Moreover, to each partition p_d^i we associate a partition latency t_d^i , which defines the processing time of partition p_d^i when running on the associated device. This latency can be computed from a partition latency model that predicts the DNN latency on the specific device [30] or by performing a runtime profiling of the DNN [11, 26] (which is the approach we follow in this work).

3.2 Network latency and cost model

As in [45] we estimate networks latency through the access time (T), bandwidth (ϕ) and `data_size` to transfer:

$$l_{network} = \left(T + \frac{data_size}{\phi} \right), \quad (2)$$

where T and ϕ depend on the network domain we consider (WiFi7 or 5G). In particular, T_{5G} depends on the geographic location of the network node that the user wants to access. For simplicity and as discussed in [34], we consider it as invariant with respect to the time, even though it may sometimes fluctuate due to the solicitation of the network (for instance, if many users are concurrently accessing the network, they may experience an access time greater than the one they would observe when the network is not saturated).

Since information about the network access time and bandwidth are not easily accessible in practical applications, we focus on the network throughput that characterizes each domain, denoted as r_{WiFi7} and r_{5G} , respectively. The network throughput can be influenced by many factors, among which the network access time, network congestion, and packet loss, and it intervenes in the computation of both the data-transfer time and the energy consumption associated with it. Therefore, we can use it to express Equation (2) as:

$$l_{WiFi7} = \frac{data_size}{r_{WiFi7}}, \quad l_{5G} = \frac{data_size}{r_{5G}}. \quad (3)$$

Note that the WiFi7 throughput may be reduced when two users approach each other, due to the interference between their two WiFi hotspots. As already mentioned, the network throughput is measured in each time interval τ , thus we should denote it by r_{WiFi7}^τ and r_{5G}^τ for the two network domains. Since we consider it to be constant within the time interval τ , we drop the apex in the rest of the formulation.

We assume that users enable the WiFi7 hotspot to connect the SEW to the phone, while they leverage 5G to connect from the phone to the cloud. Therefore, the data-transfer latencies l_{sp} between the SEW device and the smartphone, and l_{pc} between the smartphone and the cloud are:

$$l_{sp} = \frac{\sum_{i=1}^{|\mathcal{K}|} \delta_{12}^i x^i}{r_{WiFi7}}, \quad l_{pc} = \frac{\sum_{i=1}^{|\mathcal{K}|} \delta_{23}^i x^i}{r_{5G}}. \quad (4)$$

Note that, by Constraint (1), only one x^i is equal to 1.

The user is interested in minimizing the cost c_{5G} associated with the 5G connection used to offload the computation from the smartphone to the cloud. This depends on the size δ_{23}^i of the data sent from the smartphone to the cloud, and on the cost g paid to transfer one byte of data, i.e.:

$$c_{5G} = \sum_{i=1}^{|\mathcal{K}|} \Lambda g \delta_{23}^i x^i. \quad (5)$$

On the other hand, the WiFi7 hotspot between the SEW and the phone does not require any internet connection, therefore we do not consider WiFi costs in our model.

3.3 Application latency model

We recall that the end-to-end latency comprises different entity execution times, and the data-transfer latency between those entities. Let l_{SEW} and l_{phone} be the SEW and phone execution times, respectively, which can be expressed as:

$$l_{SEW} = \sum_{i=1}^{|\mathcal{K}|} t_1^i x^i, \quad l_{phone} = \sum_{i=1}^{|\mathcal{K}|} t_2^i x^i. \quad (6)$$

Similarly, we consider a simple cloud model in which the latency and its variance are measured in the control time slots τ . We denote the total latency on the cloud system by l_{cloud}^τ and define it as $l_{cloud}^\tau = f(\lambda^\tau, \sum_i t_3^i x^i)$, where λ^τ is the (unknown) workload injected in the cloud in the time slot τ .

Exploiting the information on the data-transfer latencies l_{sp} and l_{pc} introduced in the previous section, the end-to-end latency is given by:

$$l_{total} = l_{SEW} + l_{phone} + l_{sp} + l_{pc} + l_{cloud}. \quad (7)$$

3.4 Energy consumption model

The power consumption of the devices includes four components, namely:

- The power consumed by the SEW and phone when running a partition locally, denoted as P_{exeSEW} and $P_{exePhone}$ respectively, which consider the CPU and AI neural engine/GPU energy consumption models;
- The power consumption of the wireless network interface (WNI) and 5G interface (5GI) when uploading the intermediate tensor to the smartphone and cloud, denoted as P_{wnu} and P_{5gu} respectively;
- The power consumed by the WNI and 5GI when retrieving the remote execution results, denoted as P_{wnd} and P_{5gd} respectively;
- The power consumed by the WNI and 5GI when the SEW and phone are in *idle* state, denoted as P_{wnidle} and P_{5gidle} respectively. The idle state is enabled when the SEW and phone are waiting for the results from the execution of the remote partitions.

The power terms P_{exeSEW} and $P_{exePhone}$ can be computed as in [52] and are given by:

$$P_{exeSEW} = \sum_{i=1}^{|\mathcal{K}|} \Lambda z_{SEW} \mu_1^i x^i \quad (8)$$

$$P_{exePhone} = \sum_{i=1}^{|\mathcal{K}|} \Lambda z_{phone} \mu_2^i x^i,$$

where z_{SEW} and z_{phone} are the energy consumption for performing one floating point operation on SEW and phone, respectively.

When the devices offload all or some parts of the computation, they consume the power P_{wnu} and P_{5gu} to send the intermediate tensor. Let θ_{SEW} and θ_{phone} denote the power consumption of the SEW WNI and phone 5G interface when transmitting data, respectively. The value of P_{wnu} and P_{5gu} are computed as proposed in [41]:

$$P_{wnu} = \Lambda \theta_{SEW} l_{sp} = \frac{\sum_{i=1}^{|\mathcal{K}|} \Lambda \theta_{SEW} \delta_{12}^i x^i}{r_{WiFi7}} \quad (9)$$

$$P_{5gu} = \Lambda \theta_{phone} l_{pc} = \frac{\sum_{i=1}^{|\mathcal{K}|} \Lambda \theta_{phone} \delta_{23}^i x^i}{r_{5G}}.$$

In this work, we neglect the power consumption of downloading back the result (P_{wnd} and P_{5gd}) and while in the idle state (P_{wndle} and P_{5gidle}). Therefore, we define the total power consumption P_{SEW} and P_{phone} as:

$$P_{SEW} = P_{exeSEW} + P_{wnu}, \quad P_{phone} = P_{exePhone} + P_{5gu}. \quad (10)$$

Note that, since all the terms in Equations (8) and (9) are powers, the energy consumption in a specific time slot with duration τ is computed as follows:

$$E_{SEW}^\tau = P_{SEW}\tau, \quad E_{phone}^\tau = P_{phone}\tau. \quad (11)$$

3.5 Problem definition

As we mentioned in Section 3.1, the agent must choose an appropriate DNN configuration depending on the system state. We formulate this decision problem as the minimization of the total processing cost as follows:

$$\min_{x^i} \left(\alpha(P_{SEW} + P_{phone}) + c_{5G} \right) \tau \quad (P1a)$$

subject to Constraint (1) and:

$$l_{total} < L_{max} \quad (P1b)$$

$$x^i \in \{0, 1\}, \forall i \in \{1, 2, \dots, |\mathcal{K}|\}. \quad (P1c)$$

The goal is to minimize the energy consumption defined in the previous section and the 5G connection cost. In the objective function above, α denotes the energy unit price (measured in \$/J); moreover, L_{max} is the maximum time latency that guarantees a satisfactory user experience. This work does not deal with the cost of the cloud service since we are trying to optimize the system configuration from the user's perspective (while cloud costs are incurred by the application provider). The optimization of the overall system (which includes the decision, e.g., on the number of virtual machines to run in the cloud back-end) will be considered as part of our future work. Note that, in a specific time-slot τ , the optimization problem can be solved by inspection, by considering all the possible system configurations, with a complexity of $O(|\mathcal{K}|)$ (as done by Neurosurgeon [23]). However, due to the variability in system performance between time slots, leveraging RL offers the advantage of learning system dynamics and making proactive decisions, enabling the achievement of better energy-performance trade-offs (as it will be shown in Section 5).

4 REINFORCEMENT LEARNING APPROACH

In this section we model the runtime management of AI applications as an MDP (see Section 4.1). The easiest way to solve MDPs is to use dynamic programming, but in some cases, as the one we consider, this is not possible due to the state space size or the unknown system dynamic. Therefore, we present an RL-based approach to tackle the MDP in Section 4.2.

4.1 Problem formulation as an MDP

We formulate the SEW offloading problem as a discrete-time infinite-horizon MDP [44]. This can be defined by a 5-tuple $\langle \mathcal{S}, \mathcal{A}, P, c, \gamma \rangle$, where \mathcal{S} represents the infinite set of all the possible states; $\mathcal{A}(s)$ denotes the finite set of all possible actions in state s ; $P(s'|s, a)$ is the transition probability from a given state s to a state s' given an action $a \in \mathcal{A}(s)$; $c(s, a, s')$ is the immediate cost incurred when

an action a is executed in state s and the system transits to state s' ; and $\gamma \in [0, 1]$ is a discount factor that adjusts the importance of future costs.

We define the agent state as:

$$s = (r_{\text{WiFi7}}, r_{5G}, l_{\text{SEW}}, l_{\text{phone}}, l_{\text{cloud}}). \quad (12)$$

Among these, the data rates r_{WiFi7} and r_{5G} , and the cloud latency l_{cloud} are exogenous parameters that are not under the agent control. Therefore, their variability is independent on the chosen actions and is simply observed from the environment. Moreover, we assume that the phone works in light-load conditions, i.e., there are no other running tasks that compete for its resources. On the other hand, l_{SEW} and l_{phone} change according to the actions, but their value is known a priori: indeed, we can estimate the execution times t_1^i and t_2^i of the partitions executed on the SEW and the smartphone for all the configurations κ^i (see Section 6.1). If we discretize all the continuous variables with 10 values each, the dimension of \mathbf{S} is in the order of 10^5 .

An action consists in selecting a configuration κ^i , thus the number of actions equals the number of configurations $|\mathcal{K}|$. Note that, in some states, the agent may choose not to change the configuration selected in the previous time window; we model this scenario by introducing an action η that represents the *do nothing* choice. The set of all the possible actions $\mathbf{A}(s)$ is hence given by:

$$\mathbf{A}(s) = \{a^1, a^2, \dots, a^{|\mathcal{K}|}\} \cup \{\eta\}, \quad (13)$$

where each a^i represents the action of raising to 1 the variable x^i , thus selecting configuration κ^i .

We associate a cost $c(s, a, s')$ to each triplet state-action-next state; $c(s, a, s')$ embeds the energy costs $c_{e_{\text{SEW}}}(s, a) = \alpha e_{\text{SEW}}$ and $c_{e_{\text{phone}}}(s, a) = \alpha e_{\text{phone}}$, the 5G connection cost $c_{5G}(s, a)$, and the penalty $c_{\text{lat}}(s, a, s')$ we pay when violating the maximum latency requirement (Constraint (P1b)). Moreover, if the agent chooses an action different from η in a given time slot, the system must be reconfigured, which implies some overhead; therefore, we introduce a reconfiguration penalty c_{rcfg} , defined as $c_{\text{rcfg}} = \mathbb{1}_{\{a \neq \eta\}}$, meaning that we pay a penalty equal to 1 every time the chosen action is different from η , 0 otherwise. Similarly, the cost $c_{\text{lat}}(s, a, s')$ incurred when the system violates the latency constraint can be simply defined as $c_{\text{lat}}(s, a, s') = \mathbb{1}_{\{l_{\text{total}} > L_{\text{max}}\}}$. We combine the different costs using a simple additive weighting approach [51], writing:

$$c(s, a, s') = \omega_{e_{\text{SEW}}} \frac{c_{e_{\text{SEW}}}}{c_{e_{\text{SEW}}, \text{max}}} + \omega_{e_{\text{phone}}} \frac{c_{e_{\text{phone}}}}{c_{e_{\text{phone}}, \text{max}}} + \omega_{5G} \frac{c_{5G}}{c_{5G, \text{max}}} + \omega_{\text{lat}} c_{\text{lat}} + \omega_{\text{rcfg}} c_{\text{rcfg}}, \quad (14)$$

where $\omega_{e_{\text{SEW}}}$, $\omega_{e_{\text{phone}}}$, ω_{5G} , ω_{lat} and ω_{rcfg} are non-negative weights that sum up to 1. $c_{e_{\text{SEW}}, \text{max}}$, $c_{e_{\text{phone}}, \text{max}}$ and $c_{5G, \text{max}}$ are the normalization parameters for the energy cost on SEW and phone and the cost of the 5G connection, respectively.

When we have a high battery level, we may want to promote the local computation on the SEW, while, when the battery level drops, we may foster the computation offloading to the mobile phone or the cloud. Therefore, we might vary the weights according to the actual battery level.

The system dynamic is stochastic; hence, the transition probability matrix cannot be defined. As a consequence, it is impossible to solve this MDP using traditional dynamic programming. The next section will present the model-free RL-based approaches we propose to tackle this problem.

4.2 Reinforcement Learning approaches

RL approaches aim to learn an optimal policy by interacting directly with the system and leveraging the collected experience [44]. A basic RL scenario can be described as follows: considering an

Algorithm 1 ϵ -greedy tabular methods.

```

491 1: Input:  $S, A, C, T, lr \in [0, 1], \gamma \in [0, 1], \epsilon \in [0, 1]$ , set of weights  $\mathbf{W}$ , set of  $L_{max}$ , method
492 2: Initialization:
493 3:  $Q : S \times A \rightarrow \mathbb{R}$  initialize arbitrarily
494 4: Start in state  $s \in S$ 
495 5: for  $1, \dots, T$  do
496 6:   Set  $\mathbf{W}$  or  $L_{max}$  according to the step
497 7:    $a \leftarrow \epsilon$ -greedyPolicy( $s, A, \epsilon$ )
498 8:   Take action  $a$ , observe new state  $s'$  and get reward  $R$  considering  $\mathbf{W}$ 
499 9:   Update  $Q$  table(s) based on method
500 10:   $s \leftarrow s'$ 
501 11: end for
502 12: return  $Q$ 

```

infinite horizon of discrete time steps, an agent takes an action at each time step to minimize the long-term cost. When the agent takes an action, it considers the current state of the environment; its choice results in paying an immediate cost and moving to a new state. This transition, as well as the immediate cost, is often stochastic, since it usually depends on exogenous factors that are unknown to the agent. The agent stores $Q(s, a)$, which is an estimate (called Q-function) of the expected long-term costs that result from the execution of an action a in a state s . The agent uses these estimates to decide which action to take at the next step, and then updates $Q(s, a)$ based on the observed cost, thus improving its policy. RL approaches assume that the system complies with the Markov property and is stationary. While these properties might not hold in real systems, some works have successfully used RL techniques in non-Markovian scenarios [14, 29].

In this work, we consider several Q-tabular methods, i.e., algorithms that leverage a matrix (called Q-table) to store the estimated $Q(s, a)$ values. All of them are model-free and do not require knowledge of the system dynamics. We consider a simple ϵ -greedy action selection approach that, at each step τ , chooses the greedy action (i.e. $a^\tau = \operatorname{argmin}_{a \in A(s^\tau)} Q(s^\tau, a)$) with probability $1 - \epsilon$, *exploiting* its knowledge of the system, whereas it picks a random action with probability ϵ to *explore* the space, enhancing its knowledge of the application. The ϵ -greedy policy guarantees the exploration of sub-optimal actions, albeit with low probability, while choosing the optimal action for a given state most of the times.

Since most RL methods are highly problem-dependent, we conduct a comprehensive comparison among different tabular methods, namely Q-learning [49], SARSA [42], ExpectedSARSA [43], Double Q-learning [17] and Weighted Q-learning [9], to investigate the effects on our problem of different levels of exploration and of the trade-off between sample efficiency and computational time. Moreover, we first train our agents in a simulated environment, and later move to a real system where they can further update the developed policies.

Algorithm 1 summarizes the approach we followed for offline training while Algorithm 2 describes the RL agent behavior when deployed on a real system prototype. Algorithm 1 receives as inputs the state and action spaces, the cost function, the number of steps T , and some RL hyper-parameters as the learning rate lr , γ and ϵ (line 1). Then, the Q-table of state-action values is randomly initialized (lines 3), and a state is randomly sampled from the environment (line 4). The general outline of all tabular methods is presented in lines 5–11: in each step, the importance weights \mathbf{W} or the threshold L_{max} are set according to the user preference to account for energy efficiency (more details in Section 5); then, the agent chooses an action following the ϵ -greedy policy, observes the new state s' and receives the reward (lines 6–8). The Q table is updated depending on the selected method, and the environment moves to the next state (lines 9–10). If we denote

Algorithm 2 Algorithm used to update the RL agent on real system prototype.

```

540 1: Input:  $S, A, C, T, lr \in [0, 1], \gamma \in [0, 1], \varepsilon \in [0, 1]$ , set of weights  $\mathbf{W}$ , set of  $L_{max}$ , control time-period  $\tau$ ,
541    $Max\_consecutive\_violations$ , method
542 2: Initialization:
543 3:  $Q : S \times A \rightarrow \mathbb{R}$  initialize with learned Q table
544 4: Initialize violations and elapsed_time
545 5: Start in state  $s \in S$ 
546 6: for  $1, \dots, T$  do
547 7:   Set  $\mathbf{W}$  or  $L_{max}$  according to battery level or user preference
548 8:    $a \leftarrow \varepsilon$ -greedyPolicy( $s, A, \varepsilon$ )
549 9:   while elapsed_time <  $\tau$  and violations <  $Max\_consecutive\_violations$  do
550 10:     Take action  $a$ , observe state  $s'$  and get reward  $R$  considering  $\mathbf{W}$ 
551 11:     Store  $(s, a, R, s')$  in transition buffer  $T_b$ 
552 12:     if latency >  $L_{max}$  then
553 13:       violations  $\leftarrow$  violations + 1
554 14:     else
555 15:       violations  $\leftarrow$  0
556 16:     end if
557 17:   end while
558 18:   Update Q table(s) based on  $T_b$  and method
559 19:   Reset violations, flush the buffer  $T_b$  and reset elapsed_time
560 20: end for
561 21: return  $Q$ 

```

by R the immediate reward, i.e., the negative of the cost defined in the previous section, line 9 in Algorithm 1 can be detailed for the different methods as follows.

- **Q-learning** learns the optimal policy directly by maximizing rewards through a greedy action-selection strategy, eliminating the possibility of the agent taking an exploration step from the second step in the update function; accordingly, it uses the action a' that maximizes the Q value (i.e., minimizes the cost) to update the Q table:

$$Q(s, a) \leftarrow (1 - lr)Q(s, a) + lr(R + \gamma \max_{a'} Q(s', a')).$$

- **Double Q-learning** exploits two Q-tables (i.e., two value estimates, which makes the algorithm more stable) to reduce bias and avoid the overestimation of action values often incurred by Q-Learning. To perform the update, it randomly chooses one of the Q tables (Q^A and Q^B), selects the action a' that maximizes its value, and updates it based on the other Q value. For instance, if Q^A is selected, we set:

$$a' = \operatorname{argmax}_a Q^A(s', a)$$

$$Q^A(s, a) \leftarrow (1 - lr)Q^A(s, a) + lr(R + \gamma Q^B(s', a')).$$

- **Weighted Q-learning** aims to address a Double Q-Learning drawback, i.e., the fact that the relative performance of the two estimators is highly problem-dependent. To increase robustness, it estimates the maximum expected value relying on a weighted average of sample means; weights are computed using Gaussian approximations for the sample means distributions. For each action $a_i \in A$, it defines w_i as the number of times a_i minimizes the samples, divided by the total number of samples. Then, the weighted Q is computed as $W = w^T Q(s', a)$, where \mathbf{a} is the vector of all actions. Finally, Q is updated as:

$$Q(s, a) \leftarrow (1 - lr)Q(s, a) + lr(R + \gamma W).$$

- **SARSA** incorporates an exploration step by consistently following the ε -greedy strategy to choose action a' , and it updates the Q table as follows:

$$Q(s, a) \leftarrow (1 - lr)Q(s, a) + lr(R + \gamma Q(s', a')).$$

Directly learning from the exploration policy (on-policy) makes SARSA more conservative and stable than Q-Learning, even though the latter tends to converge more quickly to an optimal policy.

- **Expected SARSA** updates the Q table based on a weighted average of the Q-values for all the possible actions in the next state, taking into account the likelihood of each action under the current policy, instead of stochastically sampling state-action values. Accordingly:

$$Q(s, a) \leftarrow (1 - lr)Q(s, a) + lr(R + \gamma \sum_{a'} \pi(a'|s')Q(s', a')),$$

where $\pi(a'|s')$ is the probability of selecting action a' in state s' , according to the current policy, that is updated as follows:

$$\pi(a'|s') = \begin{cases} \frac{1-\varepsilon}{|\arg\max_{a''} Q(s', a'')|} + \frac{\varepsilon}{|A|} & \text{if } Q(s', a') = \max_{a''} Q(s', a'') \\ \frac{\varepsilon}{|A|} & \text{otherwise.} \end{cases}$$

When running [Algorithm 2](#) in a real system the Q-table is initialized with the Q-table learned during offline training (line 3). Also, all the other variables needed by the algorithm are initialized in line 4. The importance weights \mathbf{W} and L_{max} remain constant or are updated based on user preferences or battery level (line 7). The agent follows an ε -greedy policy, executing actions and storing transitions (s, a, R, s') in a buffer T_b during the control time (10 seconds). No Q-table updates occur during this period, but constraints violations are recorded (lines 8–14). Once the control time elapses, the agent updates its Q-table using transitions from T_b and requests a new action (lines 18–20). If the number of consecutive violations exceeds the acceptable threshold, the agent requests a new action even before the control time ends (line 9). After each cycle, we reset the violation count and elapsed time, and clear the buffer T_b (line 19).

5 EXPERIMENTAL RESULTS

This section evaluates the performance of the RL agent under different scenarios. [Section 5.1](#) presents the experimental setup, with a focus on the network bandwidth and cloud performance variability. [Section 5.2](#) introduces the alternative methods we considered as baseline for the evaluation. [Section 5.3](#) and [Section 5.4](#) present the results of the analyses varying the importance weights and the maximum time latency, respectively. In [Section 5.5](#), we compare our approach against the baselines, focusing on energy efficiency and execution-time violations. Finally, in [Section 5.6](#) we validate the best-performing agent (based on Q-learning) by considering a prototype application run in a real edge system.

5.1 Experimental setup

We considered a reference AI application for object tracking. These can be used in an indoor scenario, where the user wearing the SEW may want to identify various objects in her/his environment, or outdoor for, e.g., people tracking. A rate of 1–5 fps is considered acceptable for low-frame rate videos [2].

According to [31], DNN-based object-tracking frameworks incorporate CNNs for object detection. These frameworks include post-processing operations such as bounding boxes filtering and drawing,

638 as well as object assignment in different frames. In our experiments, we focused on the object
 639 detection phase, and we adopted the widely-used YOLOv5 model. YOLOv5 utilizes a CNN as its
 640 backbone and incorporates additional operations in its neck and head to enable object detection. A
 641 profiling campaign was carried out to measure the execution time of different partitions, considering
 642 HoloLens 2 as a representative of the computing power that will be available in the next generation
 643 SEW, a Samsung Galaxy S23¹ as the second device on the second layer, and a Dell Precision 5480
 644 PC as the cloud server. The HoloLens 2 comprises a Qualcomm Snapdragon 850 System-On-Chip
 645 (SoC) with an Octa-core, 2840 MHz, Kryo 385, 64-bit ARM processor, 4 GB of RAM and Adreno
 646 630 GPU. The Samsung S23 includes a Snapdragon 8 Gen 2 SoC, 8GB RAM, 128 GB storage, an
 647 Octa-core processor with a maximum frequency of 3.36 GHz, and an Adreno 740 GPU which is
 648 compliant with the announced Qualcomm SEW-phone AR ecosystem. The PC is characterized by
 649 64GB of RAM, a 13th Gen Intel Core i7-13800Hx20 processor, and an Nvidia 8GB GPU. All three
 650 devices were connected to the same WiFi hotspot provided by the Samsung S23. The HoloLens 2
 651 and the Samsung S23 support WiFi 5 and WiFi 6E, respectively. The DNN time and energy profiling
 652 were performed using a C# application on HoloLens 2 (implemented using Microsoft Visual Studio
 653 and Unity), an Android server on the mobile phone, and the PC running a docker container for the
 654 third partition of the DNN model. The details of YOLOv5 model partitioning and configuration
 655 parameters are presented in Appendix B.2.

656 We set the parameter ε of the agents ε -greedy policy to 0.05 in all the experiments. The learning
 657 rate is initialized to 1 and exponentially decays with the number of training steps; it is reset to 1
 658 at the beginning of every interval when the weights or the latency threshold are changed, to let
 659 the agent learn the new scenario. It is important to note that, since the environment changes are
 660 triggered by the users (e.g., depending on the battery level of users' devices), the learning rate can
 661 be updated accordingly. The discount factor γ in all experiments is 0.99

662 As explained in Appendix B.2, YOLOv5 offers 104 feasible configurations. To enhance the experi-
 663 ment robustness, we generated 10 random applications with a random number of configurations
 664 $|\mathcal{K}| \in [80, 104]$. Moreover, for each application (with a fixed configuration) we performed 10 runs
 665 of the RL agent considering 10 different random seeds. In the following, we will report the average
 666 results among all 10 applications and their 10 random runs.

667 For each configuration κ^i , the execution times t_1^i , t_2^i and t_3^i (see Section 3.1) are randomly sampled
 668 from a range having as mean the profiled value, and bounded by $\pm 20\%$. This work considers both
 669 WiFi and 5G base traces as the foundation for generating additional traces across all application
 670 instances. The WiFi base trace, consisting of 3000 rows, was obtained from profiling conducted as
 671 described above. On the other hand, the 5G base trace [34], derived from real-time measurements,
 672 comprised 11024 data samples, with the maximum uplink throughput reaching 230.75 Mbps. The
 673 base traces were replayed to reach the total runtime of each experiment. Various operations were
 674 applied during replay to eliminate periodicity, including random shifts, insertion of random noise
 675 ($\pm 10\%$ of the value at time step t), and inversion of the trace in its middle. Different 5G and WiFi
 676 traces for each application are considered across the 10 runs. Additionally, for each partition p_3^i
 677 executed on the cloud, a random cloud latency is sampled using an exponential distribution with
 678 $\lambda = 1/t_3^i$, as in [46]. Examples of 5G, WiFi and cloud latency traces are presented in Appendix B.1.

679 The agent control time period τ is set to 10 s but the next action is also triggered if 5 consecutive
 680 frames (sampled every L_{max}) incur in a latency threshold violations.

681 The action space of the agent for each application includes the number of corresponding config-
 682 urations plus η (*do nothing*). We considered an MDP with an infinite horizon; however, for each
 683 experiment, we run the agent for a total of 1.5×10^6 steps, which we collectively refer to as an
 684

685 ¹https://www.phonearena.com/phones/Samsung-Galaxy-S23_id11999

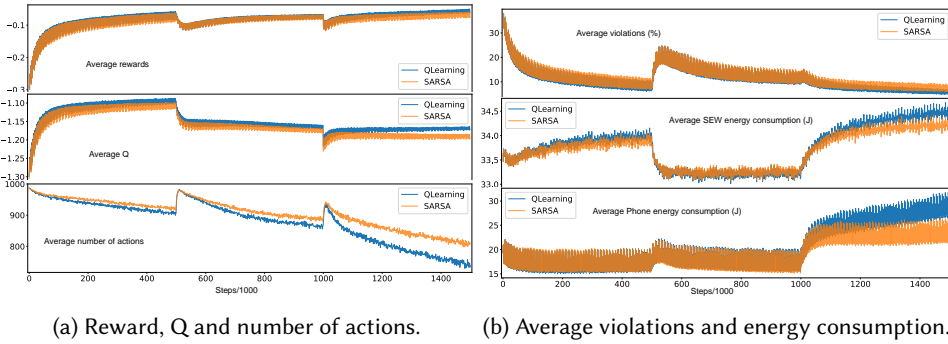


Fig. 3. Experimental results when varying the importance weights (average across 10 applications and 10 instances each).

episode. We will call the periods $[0, 5] \times 10^5$, $[5, 10] \times 10^5$, and $[1, 1.5] \times 10^6$ the first, second, and third interval, respectively.

In all experiments, we compute the reward as the negative of the cost.

5.2 Alternative Methods for Evaluation

To evaluate the effectiveness of our approach we consider two baselines: the first one, denoted by Opt^* , is a static setup that represents the optimal solution of Problem (P1) when the average WiFi and 5G throughput values derived from the traces are considered. This setup enables us to evaluate whether dynamic runtime offloading provides any benefits over a static pre-computed solution.

The second baseline we consider is *Neurosurgeon* [23], a state-of-the-art algorithm that selects the optimal partition point based on the previous sample of the networks throughput, aiming to minimize either total latency or energy consumption, but not both simultaneously. In contrast, our approach incorporates both latency and energy consumption into the agent's decision-making process, weighting each factor according to its importance in the objective function.

Since *Neurosurgeon* uses a single partition point, executing part of the DNN on the mobile device and offloading the rest to either an edge or a cloud server, we extended it to support two partition points. This change enables handling three computational layers, facilitating a fair comparison with our approach.

5.3 Varying importance weights scenario

This section presents the results of a scenario where the latency constraint threshold L_{max} is set to 500 ms for all steps, while the cost-function importance weights (see Equation (14)) change in the three episode intervals. This change in the weights allows to simulate what happens when the SEW battery capacity decreases and then the device is recharged. In particular, we set $\omega_{lat} = 0.93$, $\omega_{eSEW} = 0.03$, $\omega_{e_{phone}} = 0.02$, $\omega_{5G} = 0$ and $\omega_{rcfg} = 0.02$ in the first interval to put more emphasis on the execution time when the SEW is fully charged; in the second interval, since the battery capacity has decreased, we consider $\omega_{lat} = 0.4$, $\omega_{eSEW} = 0.5$, $\omega_{e_{phone}} = 0.05$, $\omega_{5G} = 0.03$ and $\omega_{rcfg} = 0.02$ to slightly relax the latency constraint and shift the focus on SEW energy consumption. Finally, in the third interval we assume the SEW has been partially recharged, therefore we have $\omega_{lat} = 0.86$, $\omega_{eSEW} = 0.07$, $\omega_{e_{phone}} = 0.02$, $\omega_{5G} = 0.03$ and $\omega_{rcfg} = 0.02$.

Figure 3 illustrates the comparison between Q-learning and SARSA obtained after training the agent over 1.5×10^6 steps, considering several significant metrics including the agent reward over time, the Q function value, the percentage number of violations, the number of reconfigurations,

736 and the energy consumption on SEW and phone. All the metrics are computed as a moving average
 737 over a window of 1000 steps (recall that average values are computed across 10 applications and 10
 738 random runs for each application as well). Note that Q-learning and SARSA provided better results
 739 compared with other methods in terms of violation; therefore, we only reported their results to
 740 enhance the figure clarity. A full comparison among the methods is shown in Appendix B.3.

741 In the initial steps of the first interval, the lack of knowledge about system dynamics leads to
 742 increased exploration and a higher rate of violations (see Figure 3b, first line), resulting in greater
 743 energy consumption (see Figure 3b, second line) and more frequent reconfigurations (see Figure 3a,
 744 third line). This phase is also characterized by lower Q-values and rewards, as observed in Figure 3a
 745 (first and second line). However, after approximately 10^5 steps, the agent progressively learns to
 746 adapt, leading to a reduction in the violation rate to around 7%. This adaptation results in lower
 747 energy consumption and fewer reconfigurations, ultimately yielding higher rewards.

748 At the beginning of the second interval, as the weight associated with latency constraint violations
 749 decreases and the weight for SEW energy consumption increases, the agent adapts by reducing
 750 energy consumption on the SEW while increasing energy consumption on the phone, which is
 751 assigned a lower weight. Consequently, the reward experiences a slight decline, and the percentage
 752 of violations increases, reaching 27% in the initial steps ($[5, 5.5] \times 10^5$ in all plots of Figure 3).
 753 However, the agent gradually learns to select configurations that reduce both the number of
 754 violations and energy consumption also on the phone (see steps after 5.5×10^5 in all plots of
 755 Figure 3). Finally, when the weight associated with violations increases again in the last interval,
 756 the agent compensates by increasing energy consumption on both the SEW and the phone to
 757 optimize the reward.

758 Figure 4 provides a more in-depth analysis of the agents' behavior on the first instance of the
 759 first application, which is characterized by 84 candidate configurations. This specific instance was
 760 selected for its representative behavior, closely aligning with the average. Figure 4 highlights the five
 761 most frequently chosen configurations per interval. Overall, configurations 68 to 83 dominate the
 762 first and third intervals, while configurations 36 and 44 also appear frequently in the second interval.
 763 The rationale behind this lies in the fact that, even though the importance weights change, the
 764 violation weight remains dominant in the first and last intervals; consequently, the agent prioritizes
 765 addressing violations, opting for configurations that execute the DNN as locally as possible to
 766 minimize transmission delays. In the second interval, however, the SEW energy consumption weight
 767 is dominant; therefore, the agent favors the configurations with remote execution to minimize
 768 the SEW energy consumption. As an example, at the end of the training, configuration 71 is the
 769 most popular in the third interval for both Q-learning and SARSA: it runs the DNN totally on the
 770 SEW and phone to avoid the 5G delay associated with transmitting the tensors to the cloud. In the
 771 second interval, configuration 36 is the most popular: it runs a small part of the DNN on SEW and
 772 phone, while transferring the substantial portion of the DNN computation, thereby preventing
 773 unnecessary energy consumption on the SEW device.

774 Figures 5a and 5b show the percentage number of violations in the whole episode and in the last
 775 5×10^4 steps of each interval, respectively, for all methods. They also report the average over the
 776 three intervals, referred to as *total* in the legend.

777 According to Figure 3 and Figure 5b, we conclude that, as it can be expected due to the lack of
 778 any prior experience, the agent needs more steps to converge in the first interval, and it reaches a
 779 lower number of violations by consuming more energy on the SEW and phone. Indeed, despite the
 780 weight for violations is lower in the third interval than in the first interval, we experience more
 781 violations in the last 5×10^4 steps of the first interval than in the third one, while the constraint
 782 threshold is the same. Figure 5c shows that the methods slightly reduce energy consumption
 783 in the second interval. This occurs because the violation and energy consumption weights are
 784





Fig. 4. Distribution of configurations while changing importance weights (single instance).



(a) Total violations per interval. (b) Violations in the last 50000 steps. (c) Total EC per interval. (d) EC rate with respect to Q-learning.

Fig. 5. Violation rates and energy consumption (EC) at different intervals when changing the importance weights (average across 10 applications and 10 instances each).

lower and higher, respectively, compared to the first interval, and the agent prioritizes minimizing energy consumption to maximizing the overall reward. However, in the third interval, the methods significantly increase energy consumption due to the updated weights. This last aspect is particularly true for Q-learning, due to its nature of determining the optimal action in each state.

Figure 5a shows that all methods experience the lowest percentage number of violations in the last interval, and that Q-learning is the method achieving the lowest values in all intervals. Figure 5b highlights that the agent reduces the violations between 21% and 25% in Q-learning, SARSA and ExpectedSARSA after 4.5×10^5 steps. Since Q-learning yields the lowest percentage number of violations but a similar reward value compared with SARSA and ExpectedSARSA, we computed the energy consumption (EC) rate of other methods with respect to Q-learning (see Figure 5d) to have a more comprehensive comparison:

$$EC_rate = \frac{EC_{method} - EC_{Qlearning}}{EC_{Qlearning}} \times 100. \quad (15)$$

Comparing Figure 5a and Figure 5d we observe that, with respect to SARSA and ExpectedSARSA, Q-learning reduces the violations by 6.5% and 10%, respectively, at the expense of consuming 2.13% and 2.69% more energy.

5.4 Varying latency constraint scenario

To assess how the agent adapts to different latency constraints, we change the threshold value L_{max} at the beginning of each intervals as: $L_{max} = 600, 400, 500ms$. The importance weights are fixed during the whole episode, and set to $\omega_{lat} = 0.93$, $\omega_{eSEW} = 0.03$, $\omega_{ePhone} = 0.02$, $\omega_{5G} = 0$ and $\omega_{rcfg} = 0.02$ (the same weights as the first interval analysed in the previous section).

834
835
836
837
838
839
840
841
842
843
844
845
846
847
848
849
850
851
852
853
854
855
856
857
858
859
860
861
862
863
864
865
866
867
868
869
870
871
872
873
874
875
876
877
878
879
880
881
882

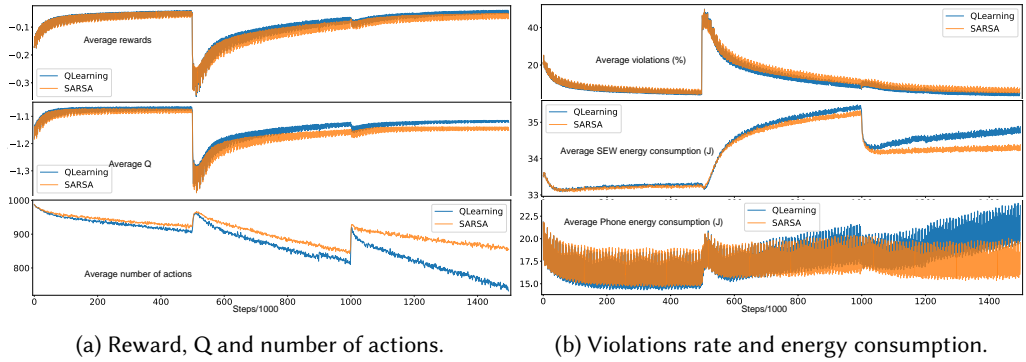


Fig. 6. Experimental results while changing execution time constraint (average across 10 applications and 10 instances each).

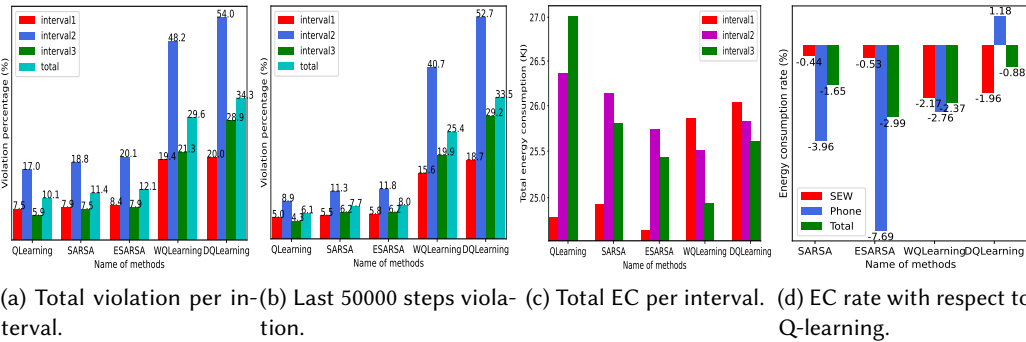


Fig. 7. Violation rates and energy consumption (EC) at different intervals while changing L_{max} (average across 10 applications and 10 instances each).

Figure 6 illustrates the comparison between Q-learning and SARSA in terms of the significant metrics, averaged among all the 10 applications and 10 runs as in Section 5.3. The full comparison among the methods is shown in Appendix B.3. In the first interval, similarly to the previous scenario, the agent explores more during the early steps, but it converges faster and experiences less violations and energy consumption because of a more relaxed latency constraint ($L_{max} = 600$) than in the first interval of Section 5.3.

At the beginning of the second interval, since the latency constraint becomes very strict, the reward experiences a sudden decline and the violation rate increases up to 50%. The agent tries to balance the reward by exploring more actions and increasing the energy consumption: it first explores actions that consume more energy on the phone rather than the SEW (according to the values of the importance weights), but, after some steps (see steps $[5, 5.5] \times 10^5$ in Figure 6b), it realizes that the violations can be compensated only by choosing configurations that process the DNN mostly on SEW, avoiding the communication delay (and its variability) between SEW, phone and cloud. Note that, in the first interval we observe a similar behaviour for SARSA and Q-learning because of the large weight on violations, while in the second interval we observe different behaviours because SARSA uses the current policy to select the next action to update Q-values while Q-learning selects the action that maximizes the Q-value in the next state. This leads Q-learning to select configurations that are mostly run on SEW and phone, increasing the energy consumption




Fig. 8. Distribution of configurations, with changing L_{max} (single instance).

and decreasing violations, while SARSA tends to be less extreme or conservative compared to Q-learning due to its on-policy nature. Finally, when the constraint is relaxed at the beginning of the last interval, the agent can significantly decrease the energy consumption in both SEW and phone, while only slightly decreasing the reward. After some steps, the Q-learning agent learns that, if it runs the DNN mostly on the phone, it can cope with the constraint, consequently consuming more energy on the phone and less on SEW compared with the second interval. SARSA does not necessarily follow the maximum Q-value policy, which leads to a lower energy consumption and to more violations.

Figure 7 compares all the tabular methods in terms of average violations and energy consumption. In particular, Figure 7a highlights that all methods but Weighted Q-learning and Double Q-learning observe the lowest percentage number of violations in the last interval, even though L_{max} is lower than in the first interval. This shows that the agent is trained well particularly in the second interval, because of the strict constraint. Additionally, it illustrates that Q-learning has the lowest violation rate, both per-interval and overall, among all methods. Figure 7b shows the violation percentage during the last 5×10^5 steps of each interval, and the average of the three intervals, referred to as *total* in the legend. This figure exhibits that the agent manages to reduce the violation percentage between 32% and 40% in Q-learning, SARSA and ExpectedSARSA after 4.5×10^5 learning steps per interval, which is an acceptable improvement. In Figure 7c, we analyze the energy consumption of all methods across each interval. For SARSA and Expected SARSA, energy consumption is higher in the second interval due to the stricter constraints imposed during that phase. In contrast, for Q-learning, the maximum energy consumption occurs in the third interval, driven by the policy selection of the maximum Q-value action. Comparing the violation rates in Figure 7a and the energy consumption rate of all methods with respect to Q-learning (Equation 15) in Figure 7d, we observe that Q-learning reduces the violations by 1.3% and 2% while consuming 1.65% and 2.99% more energy compared with SARSA and ExpectedSARSA, respectively. According to all results, Weighted Q-learning and Double Q-learning did not manage to achieve acceptable performance and they need much more steps to converge to a near-to-optimal policy.

Figure 8 shows the five most popular configurations per interval related to the first instance of the first application. Since the constraint in the first interval is relaxed and most of the configurations can satisfy it, the distributions of configurations are very close in both methods (see Figure 8, *interval1*). When the constraint becomes tight, in the second interval, the agent realizes that the best configurations are those between 69 and 84, which correspond to running the DNN mostly on the SEW. For example, configuration 72 is the most popular selected by Q-learning; it runs mostly on the SEW, nothing on the phone and a tiny part on the cloud (see Figure 8a, *interval2*). Finally, in the last interval, when the constraint is relaxed again, the range of the agents decisions expands between 45 and 84, and the violation percentage reaches the lowest possible level (see Figure 8b, first line). The two most popular configurations for both methods are 82 and 78 (see

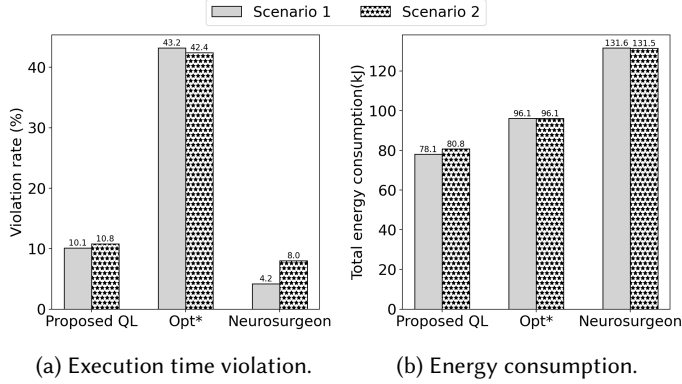


Fig. 9. Comparison of Proposed QL, Neurosurgeon and Opt* considering execution time violation rate and energy consumption for latency-oriented scenario (scenario 1) and energy-aware scenario (scenario 2).

Figure 8, *interval3*), where the former runs the DNN on the SEW and phone (nothing on the cloud), while the latter runs the DNN mostly on the SEW and two tiny parts on the phone and cloud.

5.5 Comparison with Opt* and Neurosurgeon

In this section, we compare our best-performing method, referred to as *Proposed QL*, with the *Opt** and *Neurosurgeon* [23] baselines described in Section 5.2.

Figure 9 presents comparison results of the three approaches considering two scenarios. Scenario 1 is latency-oriented and involves varying execution time thresholds as described in Section 5.4 (i.e., taking values in [600, 400, 500] ms), while setting the importance weights to $\omega_{lat} = 0.93$, $\omega_{eSEW} = 0.03$, $\omega_{e\text{phone}} = 0.02$, $\omega_{5G} = 0$ and $\omega_{rcfg} = 0.02$. Scenario 2 is an energy-aware scenario in which the execution time threshold is fixed at 500 ms, while *Proposed QL* adjusts the importance weights as detailed in Section 5.3. *Neurosurgeon* objective is latency minimization in scenario 1, while it is the minimization of total energy consumption in scenario 2. For *Opt**, as the configuration remains static, we compute the execution time violation rate and energy consumption in both scenarios. Recall that both the execution time violations and energy consumption were computed using the same WiFi and 5G traces, as described in Section 5.1.

For Scenario 1, Figure 9a shows that *Opt** performs the worst in terms of execution time violations, with a violation rate of 43.2%. In contrast, both *Proposed QL* and *Neurosurgeon* demonstrate significantly better performance, with violation rates of 10.1% and 4.2%, respectively. Regarding energy efficiency, Figure 9b reveals that *Proposed QL* is the most efficient, consuming 41% and 19% less energy compared to *Neurosurgeon* and *Opt**, respectively. In Scenario 2, considering both energy consumption and execution time violation rate, *Proposed QL* proves to be the best-performing method. It achieves a violation rate similar to *Neurosurgeon* (10.8% versus 8%), while maintaining a significantly lower violation rate than *Opt** (see Figure 9a). Furthermore, *Proposed QL* ensures substantial energy savings, consuming 39% and 16% less energy than *Neurosurgeon* and *Opt**, respectively (see Figure 9b). In conclusion, *Proposed QL* demonstrates superior performance across both metrics, offering the best trade-off between execution time violations and energy consumption. It effectively captures system dynamics, outperforming the other methods. While *Neurosurgeon* is the most energy-intensive, *Opt** results inadequate for this problem, as it fails to account for system dynamics, inducing higher violation rates.

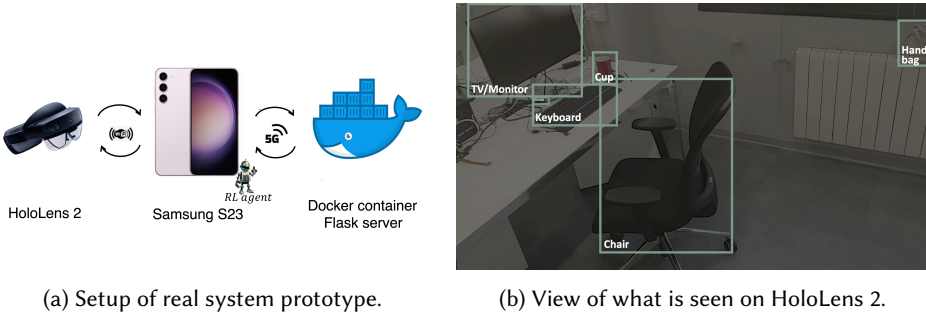


Fig. 10. Setup of real system prototype and example of the HoloLens view.

5.6 Experiments on prototype environment

This section presents the results obtained by executing the trained agent on a real system, to evaluate its performance and ascertain whether it provides the anticipated benefits as described in the previous section. It is important to note that only limited studies have addressed this aspect [12], which is crucial to provide a clear understanding of how the simulation results can be applied to actual systems and how the agent can be utilized in production. We focused on the Q-learning method, identified in Section 5.3 as the optimal approach.

The actual system configuration, illustrated in Figure 10a, precisely aligns with the description provided in Section 5.1. The YOLOv5 object-detection task follows an execution sequence that begins with image capture by the HoloLens 2 camera and ends with the visualization of bounding boxes on its display as seen in Figure 10b.

We used the Samsung S23 WiFi hotspot to establish a wireless connection between different entities. This phone has WiFi6E capability to connect different devices using WiFi6E over the 5GHz band, with a maximum throughput reaching 1000Mbps. To ensure compatibility between Samsung S23 and the third device regarding 5G network throughput, we clipped the connection throughput according to the maximum 5G throughput reported in the real trace presented in the Appendix. Specifically, the 5G throughput value is computed by multiplying the peak value observed in the 5G base trace by the ratio of the actual WiFi6 throughput to its maximum value. Consequently, during data transmission from Samsung S23 to the PC, if the real WiFi6 throughput exceeds the computed value, the server on the PC introduces a delay with a predetermined duration to accommodate the disparity between the real WiFi6 throughput and the computed value. This approach ensures that the experiment accurately reflects the compliance with 5G network specifications in the connection between the second and third computational layers.

We considered an execution time threshold $L_{max} = 500ms$ and adjusted the importance weights to accommodate various setups, following the procedures implemented during the training phase (Section 5.3). As already mentioned, these adjustments account for when there are changes in the battery capacity of the SEW or the phone. To effectively adapt the agent to the new scenario, we profiled the DNN on the real system to evaluate the determinism of execution time across different devices, thereby improving training efficiency. The profiling results revealed that the coefficient of variation in execution time across all partitions consistently remained below 20%. Consequently, we trained the agent using Q-learning, considering all 104 possible configurations of YOLOv5n (YOLOv5 nano), while utilizing the WiFi trace from the profiling and incorporating the 5G trace from [34]. The training process consisted of 1.5 million steps, considering $\epsilon = 0.05$. At the end of the offline training, the total violation achieved is 13.2%.

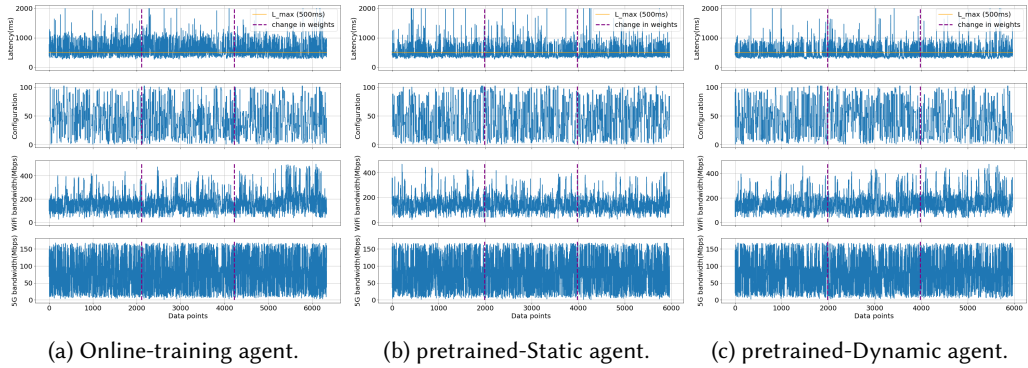


Fig. 11. Latency, reconfigurations, 5G, and WiFi throughputs for *Online-training*, *pretrained-Static*, and *pretrained-Dynamic* agents on real system with 104 configurations in prototype environment. Purple dashed lines indicate the end of an interval, where importance weights are changed.

We evaluated the trained agent in terms of latency constraint violation and energy consumption. The evaluation encompasses four main scenarios: the first one involves a randomly initialized agent, referred to as the *Online-training agent*, which begins the learning process from scratch. The second scenario leverages an agent whose Q-table is initialized with the values from the training process, but which is not updated thereafter (*pretrained-Static agent*). The third scenario follows the same initialization process as the second one, but the agent continues to learn during the experiment, leading to updates in its Q-table (*pretrained-Dynamic agent*). The fourth scenario involves no DNN partitioning, while the entire DNN is executed locally on the HoloLens 2 (*All local*).

For each scenario, the experiment comprises three intervals, each lasting 20 minutes. At the beginning of each interval, the importance weights change as described in Section 5.3. The agent control time (τ) is set to 10 seconds, ensuring that a new action is requested either upon the lapse of the control time or when there are 5 consecutive violations within the control period.

The following sections are organized as follows: Section 5.6.1 and Section 5.6.2 focus on latency constraint violations and energy consumption analyses, respectively. They underscore the *pretrained-Dynamic* and *pretrained-Static* agents capabilities in minimizing latency violations and saving energy, surpassing both the *Online-training* agent and the *All local* configuration. Section 5.6.3 explores how reducing the number of alternative configurations, which also decreases the number of actions and the agent memory requirements, affects the agent performance. This emphasizes the critical importance of configuration optimization in real system deployment.

5.6.1 Assessment of latency constraint violations. Figure 11 provides insights onto the latency constraint violations and the configurations chosen during the experiment for each scenario. Additionally, it includes the WiFi and 5G throughputs. From Figure 11a, we observe that the *Online-training* agent exhibits more violations compared to the others due to the fact that it starts from a random initialization and explores actions that may not be beneficial. The *pretrained-Static* agent, on the other hand, follows the policy learned from the profiled data. As shown in Figure 11b, it experiences fewer violations than the *Online-training* agent since it already possesses some knowledge and selects actions according to the previously learned policy. Finally, Figure 11c demonstrates that the *pretrained-Dynamic* agent exhibits fewer violations compared to the *Online-training* agent and *pretrained-Static* agent.

These observations are supported by the results presented in Table 1, where the *pretrained-Dynamic agent* consistently outperforms the *Online-training* and *pretrained-Static agent* in all three

Table 1. Total and per-interval execution time violation and energy consumption for *Random*, *Static*, *Dynamic*, *All local* agents with 104 configurations in prototype environment.

Agent	Execution time violation (%)				SEW energy consumption (kJ)				Phone energy consumption (kJ)			
	Violation per interval			Total	Energy per interval			Total	Energy per interval			Total
	First	Second	Third		First	Second	Third		First	Second	Third	
Random	45.1	45.8	41.5	44.2	7.2	6.9	6.6	20.7	2.1	2.2	2.1	6.4
Static	22.5	23.7	22.9	23.0	5.0	4.7	6.7	16.4	1.5	1.6	2.5	5.6
Dynamic	17.8	23.7	20.5	20.6	4.8	4.4	4.6	13.8	1.4	1.6	1.5	4.5
All local	0.0	0.0	0.0	0.0	9.2	9.2	9.2	27.6	-	-	-	-

intervals. Overall, the *pretrained-Dynamic agent* incurred a violation rate of 20.6%, compared to 22.9% for the *pretrained-Static agent* and 44.2% for the *Online-training agent*. This indicates that the *pretrained-Dynamic agent* surpasses the *pretrained-Static agent* as it continues to learn from interactions with the real system and can adjust its policy online. In contrast, the violations remain relatively constant across the three intervals for the *pretrained-Static agent* since the policy is fixed and no real system information is learned. *All local* achieves 0% violation rate, as the execution time averaged 450ms and remains consistently below 500 ms. However, this comes at the cost of consuming twice the energy compared to other scenarios, as it will be detailed in [Section 5.6.2](#).

5.6.2 Assessment of energy consumption. We assessed the performance of the trained agent on the real system analyzing the energy consumption of both the Hololens 2 (representing the SEW) and the mobile phone. As in the previous section, we analyzed the performance of the *Online-training*, *pretrained-Static*, *pretrained-Dynamic* agents, and the *All local* configuration.

As discussed in [Section 3.4](#), the energy consumption on a device consists of the energy consumed for computations on the device itself and the energy consumed when transmitting data through the network interface. In this study, we rely on energy profiling on the real system to estimate the device parameters, namely the energy consumed per FLOP (z_{SEW} and z_{phone}) and the power consumed by the wireless network interface (θ_{SEW} and θ_{phone}). The profiling involved determining the energy consumed when executing the partitions of YOLOv5n on the SEW (the phone) without forwarding any data to the subsequent device. From the profiling results, the values of z_{SEW} and z_{phone} are 1085.1 pJ/FLOP and 112.5 pJ/FLOP respectively. Furthermore, the profiling aimed to estimate energy consumption when no operations were performed on the SEW (the phone), but intermediate tensors were transmitted to the subsequent device, facilitating the evaluation of data transfer power consumption. Based on the profiling results, θ_{SEW} and θ_{phone} were estimated to be 7.9 W and 4.5 W, respectively. More details about the profiling are provided in [Appendix B.2](#).

Table 1 presents the energy consumption data (in kilojoules (kJ)) for SEW and phone across different time intervals and scenarios. Notably, we focus on comparing the energy consumption while changing the importance weights and across different scenarios.

We observe that the SEW energy consumption decreases in the second interval, due to a significant increase in the corresponding weight (from $\omega_{e_{SEW}} = 0.03$ in the first interval to $\omega_{e_{SEW}} = 0.5$). In the third interval, the energy consumption varies among the agents: it increases for the *pretrained-Static* and *pretrained-Dynamic* agents but decreases for the *Online-training agent*.

Comparing the other agents with the *All local* scenario, **Table 1** demonstrates that partitioning is significantly more energy-efficient than executing the entire DNN locally. The *pretrained-Dynamic agent* exhibits the highest energy savings on the SEW, reducing energy consumption by 50% compared to the *All local* configuration, effectively doubling the battery life of the Hololens 2. The

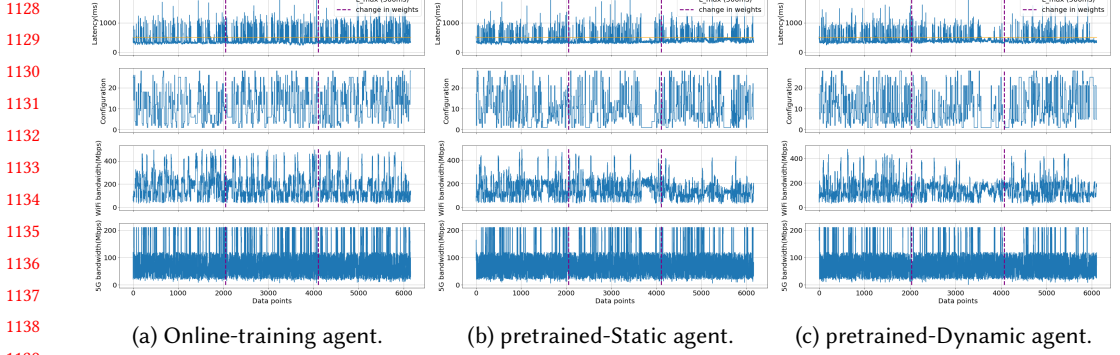


Fig. 12. Latency, configurations, 5G and WiFi throughput of Online-training, pretrained-Static and pretrained-Dynamic agents on real system with 30 configurations in prototype environment.

Online-training and *pretrained-Static* agents achieve savings of 25% and 40.6%, respectively, relative to the *All local* scenario. While the *Online-training* agent outperforms the *All local* configuration in energy savings, its exploratory nature results in more execution time violations, making it less advantageous compared to the other agents.

For the mobile phone, the *pretrained-Dynamic* agent proves to be the most energy-efficient, consuming 4.5kJ of energy, compared to 5.6kJ and 6.4 kJ for the *pretrained-Static* and *Online-training* agents, respectively. Overall, considering the combined energy consumption of the SEW and the mobile phone, the *pretrained-Dynamic* agent emerges as the most energy-efficient option, with total consumption of 18.3 kJ, compared to 22 kJ, 27.1 kJ, and 27.6 kJ for the *pretrained-Static*, *Online-training*, and *All local* scenarios, respectively.

5.6.3 Experiment with reduced number of configurations. While conducting experiments on the real system, we observed that the agents face significant challenges in recovering from violations due to the large number of candidate configurations. To mitigate this issue, we implemented a strategy to reduce the action space by including only those actions that were selected at least 80% of the time during the initial training. This resulted in a reduced action space of 30 candidate configurations (out of the original 104), which not only minimizes violations but also decreases the memory requirements on the mobile phone where the agent runs, as the memory size of the Q-table decreased from 115.2 MB to 9.6 MB.

Subsequently, the agent was retrained using this reduced action space, utilizing the same traces and conditions as in the initial training. As a result, the average latency constraint violations decreased to 8.7% after the offline training. The experiment was then repeated on the prototype system, and the corresponding plots illustrating the violations, chosen configurations, the 5G and WiFi throughputs are reported in [Figure 12](#).

Comparing the performance of the *Online-training agent*, the *pretrained-Static*, and *pretrained-Dynamic* in [Figure 12](#) to the one on the 104 configurations scenario, we observe that the average latency constraint violations decrease for all three agents. Specifically, the *Online-training*, *pretrained-Static*, and *pretrained-Dynamic* agents achieve total violations of 19.0%, 11.8%, and 11.5%, respectively, confirming the superior performance of the *pretrained-Dynamic agent* in this regard.

Moreover, the *pretrained-Dynamic* and *pretrained-Static* agents exhibit similar energy consumption on the SEW, at 14.8 kJ and 14.7 kJ, respectively, while the *Online-training* agent consumes slightly less energy at 13.8 kJ, as shown in [Table 2](#). In comparison to the *All local* scenario, all three alternative scenarios consistently demonstrate superior energy-saving capabilities, aligning with

Table 2. Total and per-interval execution time violation and energy consumption for *Random*, *Static*, *Dynamic*, *All local* agents with 30 configurations in prototype environment.

Agent	Execution time violation (%)				SEW energy consumption (kJ)				Phone energy consumption (kJ)			
	Violation per interval			Total	Energy per interval			Total	Energy per interval			Total
	First	Second	Third		First	Second	Third		First	Second	Third	
Random	18.4	18.1	20.66	19.0	4.6	4.4	4.8	13.8	1.3	1.5	1.6	4.4
Static	11.9	10.5	12.9	11.8	4.7	4.7	5.3	14.7	1.4	1.4	1.5	4.3
Dynamic	13.3	9.5	11.8	11.5	5.1	4.9	4.8	14.8	1.5	1.4	1.4	4.3
All local	0.0	0.0	0.0	0.0	9.2	9.2	9.2	27.6	-	-	-	-

the results observed across 104 configurations. Specifically, the *pretrained-Dynamic* agent achieves a 46.4% reduction in energy consumption compared to *All local*, while the *pretrained-Static* and *Online-training* agents save 46.7% and 50%, respectively. On the mobile device, energy consumption is nearly identical across the three agents.

Overall, when considering both devices, the *pretrained-Dynamic* and *pretrained-Static* agents consume 8.6 kJ less energy than the *All local* scenario. Notably, the 30-configuration scenario offers the best balance between energy consumption and execution time violations, providing an optimal tradeoff compared to the 104-configuration scenario. This underscores the importance of configuration optimization in reducing the violations.

5.6.4 Discussion. To summarize, the findings in this section demonstrate the practical feasibility of employing the trained agent in practical systems. It is also advantageous to start from a pre-learned policy as it diminishes both latency constraint violations and energy consumption. Additionally, utilizing a pretrained-Dynamic agent offers several benefits as it continues to learn and update its knowledge based on the real system state. Furthermore, we noted that reducing the number of configurations by selecting those that are more likely to be chosen effectively reduces violations, since the agent can more promptly recover from actions that result in higher violations.

6 RELATED WORK

Smart devices have enabled the widespread integration of AI applications in various facets of daily life. However, challenges arise when implementing these applications on systems with constrained computational and energy resources. In response to this challenge, researchers have explored techniques such as DNN partitioning and offloading resource-intensive computations to either edge devices or cloud resources. This section briefly reviews the state of the art concerning DNN partitioning (Section 6.1), and resource management and task offloading (Section 6.2, with a particular focus on Reinforcement Learning in Section 6.3).

6.1 DNN partitioning

Authors in [23] present Neurosurgeon, a layer partitioning model based on edge-cloud computing collaboration. This approach involves the deployment of neural networks on both mobile phones and cloud servers, while dynamically adjusting the model partitioning based on the network bandwidth or energy consumption, depending on the specific objective. A similar work in [27] proposes an algorithm to identify the most efficient partitioning point within a multi-device and multi-edge server configuration. This algorithm takes into account the minimum cost associated with each edge server for every mobile device. Authors used a partitioning point retain algorithm to reduce the search space and run a procedure to find the optimal partitioning point considering

the different capabilities of the mobile device and edge server. Differently, authors in [5] considered sparse neural networks and suggested a hypergraph model for partitioning matrices to enhance the scalability and efficiency of parallel sparse matrix-vector multiplications computations. Their approach reduces communication costs and promotes a balanced distribution of computational load across processors which they showed as an advantage for stable training and inference. All the previous studies have predominantly adopted a two-tier partitioning model, where computational tasks are solely distributed between the edge device and an edge/cloud server, and for most of them, the validation is often based on simulation. However, these approaches neglect the potential benefits of a three-tier architecture, where tasks can be dynamically orchestrated across edge devices, edge/fog servers, and a remote cloud server depending on network conditions.

6.2 Resource management and task offloading

Numerous studies tackle the resource allocation and task offloading problems in edge-cloud computing, aiming to achieve equitable distribution among various entities while minimizing expenses, to guarantee acceptable performance levels [8, 13, 21, 48]. The work in [48] introduces a reference system involving multiple users and servers, and proposes an adaptive multi-objective algorithm designed to minimize power consumption, response times, and costs. This is accomplished by implementing a two-tier matching game, which effectively addresses resource optimization.

Authors in [8] propose an alternative solution to solve a similar two-tier matching game. The first tier solves the association problem between the users and edge servers to maximize social welfare, and the second tier focuses on collaboration among the edge servers to minimize the resource cost and transfer delay. Although this study considers the collaboration between different edge servers, a task can only be served by one VM on a single edge server, limiting the problem to a two-tier binary offloading scheme. [21] optimizes computation configuration and bandwidth allocation in Multi-access Edge Computing (MEC) for enhanced user experience, particularly in energy-constrained environments. It presents a joint offloading and resource allocation framework, dynamically allocating resources to mobile devices based on task demands and energy constraints.

Although the mentioned studies yield promising results, the authors consider only resource management for generic tasks with no specific solutions for AI applications. On the other hand, [13] delves into the exploration of DNN model partitioning and offloading strategies for multiple users and proposes an evolutionary pricing strategy and a slot-based model to optimize the efficient distribution of DNN subtasks. Although this work combines DNN partitioning and task offloading, and minimizes the cost on mobile devices, it considers only a two-tier architecture and overlooks the variability of edge servers workload.

6.3 Task offloading RL-based methods

Recently, RL has gained extensive popularity in the realm of task offloading and runtime resource management. For instance, authors in [19] introduce the DROO framework to optimize wireless devices (WD) task offloading decisions and time allocation among WDs in wireless powered MEC networks. The authors use deep reinforcement learning (DRL) to maximize the weighted sum of computation rates for all WDs, improving network performance by leveraging past offloading experiences to refine action generation. In [47], a model-free distributed algorithm leveraging DRL is presented to handle dynamic load variations at edge nodes, enabling mobile devices to independently decide on offloading tasks. Focusing on non-divisible, time-sensitive tasks, the aim is to minimize long-term costs, factoring in task delays and penalties for task drops. However, both [19] and [47] consider a binary offloading strategy and overlook the energy consumption concerns.

Conversely, authors in [16] developed a task-offloading algorithm that relies on DRL and takes into account both logical and data dependencies among subtasks within user applications. Although

Table 3. Summary of the comparison of our approach with related work.

Work	DNN Partitioning	3-tier partitioning	Based on RL	Network Variability	Cloud variability	Energy-aware	Latency-aware	Real system validation
Demirci et al. [5]	✓						✓	
Liao et al. [27]	✓			✓		✓	✓	
Du et al. [8]				✓		✓	✓	
Wang et al. [48]				✓		✓	✓	
Jiang et al. [21]				✓		✓	✓	
Qiu et al. [37]				✓		✓	✓	✓
Gao et al. [13]	✓			✓	✓	✓		
Huang et al. [19]			✓	✓			✓	
Tang et al. [47]			✓		✓		✓	
Gong et al. [16]			✓			✓	✓	
Jain et al. [20]			✓			✓	✓	
Kang et al. [23]	✓			✓	✓	✓	✓	✓
Our approach	✓	✓	✓	✓	✓	✓	✓	✓

designed for scalability and capable of handling complex tasks in multi-user and edge computing environments, this approach overlooks the dynamic nature of server workloads, assuming that a single edge server supports multiple devices without accommodating server workload variations. Furthermore, the binary offloading mechanism restricts task execution to either exclusively local or remote, limiting flexibility and potentially underutilizing available resources. Authors in [20] tackle the fog-cloud task offloading problem using various optimization metrics. They formulated this problem as an MDP to enhance the average resource utility. The considered metrics encompass factors such as task execution time, total service latency, transmission electric power, and task priority. The authors used a DRL-based algorithm and performed an offline training of a multi-agent model on the cloud data center, while the real-time offloading decisions are performed online on the fog. Nonetheless, similar to most of the studies presented in this subsection, the validation results are based on simulation. In [37], authors proposed a DRL-based algorithm to tackle the challenges of computation offloading and channel allocation in MEC environments. Unlike previous methods, this approach utilized a collaborative distributed training scheme and N-steps learning, resulting in faster convergence, stable training, and reduced exploration costs. Their method optimizes long-term performance across diverse applications and is validated through experiments on a real testbed. However, it does not address the partitioning of AI applications and is limited to binary offloading. Table 3 summarizes the key features of the various works presented in this section, emphasizing the objective metrics considered (latency, energy consumption), system factors (network and cloud variability), the adopted approaches (reinforcement learning, DNN partitioning), and the evaluation setup (real system or simulation).

In this work, we proposed a novel approach for runtime offloading relying on a DNN partitioning model that operates at the layer level. Our strategy harnesses the capabilities of several tabular RL methods. The proposed approach considers the fluctuating conditions of 5G and WiFi bandwidth, as well as cloud latency. The key innovation of our work lies on addressing the real-time management of AI tasks, focusing on enhancing the battery life of SEW devices and minimizing 5G-communication cost, all while maintaining an acceptable end-to-end latency to ensure a satisfactory user experience. In particular, this work extends our previous research [22], which focused solely on the energy consumption of the SEW under varying importance weights using Q-learning. In this study, we additionally consider the energy consumption of the phone and conduct extensive comparisons of

multiple tabular methods. The best-performing method is further evaluated against two baselines. We also analyze the agent behavior in scenarios with varying maximum latency constraints to assess how it can adapt to user application constraints. Moreover, the evaluation of the trained RL agent on a real-system prototype utilizing Microsoft HoloLens 2 offers valuable insights into its potential for practical deployment. To the best of our knowledge, and as evidenced by Table 3, this is the first comprehensive study to address these challenges.

7 CONCLUSION

In this study, we tackled the runtime management of AI applications running on next-generation SEWs using various tabular RL methods. Q-learning emerged as the standout performer, exhibiting fewer execution time violations (approximately 10.0% on average across all simulations) and greater stability compared to other algorithms. Additionally, our approach offers a more favorable tradeoff between energy efficiency and execution time violations compared to the two baseline methods considered. We assessed the performance of the Q-learning agent on a real-system prototype, comparing outcomes from running the entire DNN locally with those of random, static, and dynamic agents. Results indicate that the static and dynamic agents are more energy-efficient than the all-local scenario. Experiments on the real-system prototype demonstrated our approach ability to double battery life while maintaining a satisfactory quality of experience, achieving an execution time violation rate of 11%, even under transient conditions (changes in the thresholds or weights) and the significant variability in network traces.

For future work, we aim to explore the potential of DRL algorithms, leveraging their generalization capabilities to potentially yield superior results.

Acknowledgments

This work has been sponsored by the EssilorLuxottica Smart Eyewear Lab.

References

- [1] Danilo Ardagna, Giuliano Casale, Michele Ciavotta, Juan F. Pérez, and Weikun Wang. 2014. Quality-of-service in cloud computing: modeling techniques and their applications. *J. Internet Serv. Appl.* 5, 1 (2014), 11:1–11:17. <https://doi.org/10.1186/s13174-014-0011-3>
- [2] Allysa Kate Brillantes, Edwin Sybingco, Argel Bandala, Robert Kerwin Billones, Alexis Fillone, and Elmer Dadios. 2022. Vehicle Tracking in Low Frame Rate Scenes using Instance Segmentation. In *2022 IEEE 14th International Conference on Humanoid, Nanotechnology, Information Technology, Communication and Control, Environment, and Management (HNICEM)*. 1–5. <https://doi.org/10.1109/HNICEM57413.2022.10109390>
- [3] Ramraj Dangi, Praveen Lalwani, Gaurav Choudhary, Ilsun You, and Giovanni Pau. 2022. Study and Investigation on 5G Technology: A Systematic Review. *Sensors (Basel, Switzerland)* 22 (2022), Issue 1. <https://doi.org/10.3390/S22010026>
- [4] Olivier Debauche, Saïd Mahmoudi, and Adriano Guttadauria. 2022. A New Edge Computing Architecture for IoT and Multimedia Data Management. *Information* 13 (2022), 89. Issue 2.
- [5] Gunduz Vehbi Demirci and Hakan Ferhatosmanoglu. 2021. Partitioning sparse deep neural networks for scalable training and inference. In *Proceedings of the 35th ACM International Conference on Supercomputing (Virtual Event, USA) (ICS '21)*. Association for Computing Machinery, New York, NY, USA, 254–265. <https://doi.org/10.1145/3447818.3460372>
- [6] Cailian Deng, Xuming Fang, Xiao Han, Xianbin Wang, Li Yan, Rong He, Yan Long, and Yuchen Guo. 2020. IEEE 802.11be-Wi-Fi 7: New Challenges and Opportunities. *IEEE Communications Surveys and Tutorials* 22 (2020), 2136–2166. Issue 4. <https://doi.org/10.48550/arxiv.2007.13401>
- [7] X. Dong, B. De Salvo, M. Li, C. Liu, Z. Qu, H. Kung, and Z. Li. 2022. SplitNets: Designing Neural Architectures for Efficient Distributed Computing on Head-Mounted Systems. In *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. IEEE Computer Society, Los Alamitos, CA, USA, 12549–12559. <https://doi.org/10.1109/CVPR52688.2022.01223>
- [8] Yu Du, Jun Li, Long Shi, Tingting Liu, Feng Shu, and Zhu Han. 2022. Two-Tier Matching Game in Small Cell Networks for Mobile Edge Computing. *IEEE Transactions on Services Computing* 15, 1 (2022), 254–265. <https://doi.org/10.1109/TSC.2019.2937777>

- 1373 [9] C. D'Eramo, M. Restelli, and A. Nuara. 2016. Estimating Maximum Expected Value through Gaussian Approximation.
1374 In *ICML Proc.*, Vol. 48. PMLR, NY, USA, 1032–1040. <https://proceedings.mlr.press/v48/deramo16.html>
- 1375 [10] Ersin Elbasi, Nour Mostafa, Zakwan AlArnaout, Aymen I. Zreikat, Elda Cina, Greeshma Varghese, Ahmed Shdefat,
1376 Ahmet E. Topcu, Wiem Abdelbaki, Shinu Mathew, and Chamseddine Zaki. 2023. Artificial Intelligence Technology
1377 in the Agricultural Sector: A Systematic Literature Review. *IEEE Access* 11 (2023), 171–202. [https://doi.org/10.1109/](https://doi.org/10.1109/ACCESS.2022.3232485)
[ACCESS.2022.3232485](https://doi.org/10.1109/ACCESS.2022.3232485)
- 1378 [11] Amir Erfan Eshratifar, Mohammad Saeed Abrishami, and Massoud Pedram. 2021. JointDNN: An Efficient Training
1379 and Inference Engine for Intelligent Mobile Cloud Computing Services. *IEEE Transactions on Mobile Computing* 20
1380 (2021), 565–576. Issue 2. <https://doi.org/10.1109/TMC.2019.2947893>
- 1381 [12] Enrico Galimberti, Bruno Guindani, Federica Filippini, Hamta Sedghani, Danilo Ardagna, Germán Moltó, and Miguel
1382 Caballer. 2023. OSCAR-P and aMLLibrary: Performance Profiling and Prediction of Computing Continua Applications.
1383 In *Companion of the 2023 ACM/SPEC International Conference on Performance Engineering (Coimbra, Portugal) (ICPE*
1384 *'23 Companion)*. Association for Computing Machinery, New York, NY, USA, 139–146. [https://doi.org/10.1145/3578245.](https://doi.org/10.1145/3578245.3584941)
[3584941](https://doi.org/10.1145/3578245.3584941)
- 1385 [13] Mingjin Gao, Rujing Shen, Long Shi, Wen Qi, Jun Li, and Yonghui Li. 2023. Task Partitioning and Offloading in
1386 DNN-Task Enabled Mobile Edge Computing Networks. *IEEE Transactions on Mobile Computing* 22, 4 (2023), 2435–2445.
<https://doi.org/10.1109/TMC.2021.3114193>
- 1387 [14] Maor Gaon and Ronen I. Brafman. 2020. Reinforcement Learning with Non-Markovian Rewards. *AAAI* (2020),
1388 3980–3987. <https://doi.org/10.1609/aaai.v34i04.5814>
- 1389 [15] Ruben Gomez-Ojeda, Francisco Angel Moreno, David Zuñiga-Noël, Davide Scaramuzza, and Javier Gonzalez-Jimenez.
1390 2019. PL-SLAM: A Stereo SLAM System Through the Combination of Points and Line Segments. *IEEE Transactions on*
1391 *Robotics* 35 (2019), 734–746. Issue 3. <https://doi.org/10.1109/TRO.2019.2899783>
- 1392 [16] Bencan Gong and Xiaowei Jiang. 2023. Dependent Task-Offloading Strategy Based on Deep Reinforcement Learning
1393 in Mobile Edge Computing. *Wireless Communications and Mobile Computing* 2023 (2023), 1–12. [https://doi.org/10.](https://doi.org/10.1155/2023/4665067)
[1155/2023/4665067](https://doi.org/10.1155/2023/4665067)
- 1394 [17] H. V. Hasselt. 2010. Double Q-Learning. In *23rd NIPS (Vancouver, British Columbia, Canada)*. Curran Associates Inc.,
1395 Red Hook, NY, USA, 2613–2621.
- 1396 [18] Teresa Hirzle, Florian Müller, Fiona Draxler, Martin Schmitz, Pascal Knierim, and Kasper Hornbæk. 2023. When XR and
1397 AI Meet - A Scoping Review on Extended Reality and Artificial Intelligence. In *Proceedings of the 2023 CHI Conference*
1398 *on Human Factors in Computing Systems (Hamburg, Germany) (CHI '23)*. Association for Computing Machinery, New
1399 York, NY, USA, Article 730, 45 pages. <https://doi.org/10.1145/3544548.3581072>
- 1400 [19] Liang Huang, Suzhi Bi, and Ying-Jun Angela Zhang. 2020. Deep Reinforcement Learning for Online Computation
1401 Offloading in Wireless Powered Mobile-Edge Computing Networks. *IEEE Transactions on Mobile Computing* 19, 11
1402 (2020), 2581–2593. <https://doi.org/10.1109/TMC.2019.2928811>
- 1403 [20] Vibha Jain and Bijendra Kumar. 2023. QoS-Aware Task Offloading in Fog Environment Using Multi-agent Deep
1404 Reinforcement Learning. *Journal of Network and Systems Management* 31, 1 (2023), 7.
- 1405 [21] Hongbo Jiang, Xingxia Dai, Zhu Xiao, and Arun Iyengar. 2023. Joint Task Offloading and Resource Allocation
1406 for Energy-Constrained Mobile Edge Computing. *IEEE Transactions on Mobile Computing* 22, 7 (2023), 4000–4015.
1407 <https://doi.org/10.1109/TMC.2022.3150432>
- 1408 [22] A. W. Kambale, H. Sedghani, F. Filippini, G. Verticale, and D. Ardagna. 2023. Runtime Management of Artificial
1409 Intelligence Applications for Smart Eyewears. In *IEEE/ACM UCC*, Vol. to appear. [https://doi.org/10.1145/3603166.](https://doi.org/10.1145/3603166.3632562)
[3632562](https://doi.org/10.1145/3603166.3632562)
- 1410 [23] Yiping Kang, Johann Hauswald, Cao Gao, Austin Rovinski, Trevor Mudge, Jason Mars, and Lingjia Tang. 2017.
1411 Neurosurgeon: Collaborative intelligence between the cloud and mobile edge. *ACM SIGPLAN Notices* 52 (2017),
1412 615–629. Issue 4. <https://doi.org/10.1145/3037697.3037698>
- 1413 [24] Jyotirmoy Karjee, Praveen Naik S, Kartik Anand, and Vanamala N. Bhargav. 2022. Split computing: DNN inference
1414 partition with load balancing in IoT-edge platform for beyond 5G. *Measurement: Sensors* 23 (2022), 100409. <https://doi.org/10.1016/J.MEASEN.2022.100409>
- 1415 [25] Christian Kerl, Jurgen Sturm, and Daniel Cremers. 2013. Robust odometry estimation for RGB-D cameras. *Proceedings*
1416 *- IEEE ICRA* (2013), 3748–3754. <https://doi.org/10.1109/ICRA.2013.6631104>
- 1417 [26] Hengyi Li, Zhichen Wang, Xuebin Yue, Wenwen Wang, Hiroyuki Tomiyama, and Lin Meng. 2022. An architecture-
1418 level analysis on deep learning models for low-impact computations. *Artificial Intelligence Review* (2022), 1–40.
<https://doi.org/10.1007/S10462-022-10221-5/FIGURES/25>
- 1419 [27] Zhuofan Liao, Weibo Hu, Jiawei Huang, and Jianxin Wang. 2023. Joint multi-user DNN partitioning and task offloading
1420 in mobile edge computing. *Ad Hoc Networks* 144 (2023), 103156. <https://doi.org/10.1016/j.adhoc.2023.103156>
- 1421 [28] Wenhan Luo, Junliang Xing, Anton Milan, Xiaoqin Zhang, Wei Liu, and Tae-Kyun Kim. 2021. Multiple object tracking:
A literature review. *Artificial Intelligence* 293 (2021), 103448. <https://doi.org/10.1016/j.artint.2020.103448>

- 1422 [29] Sultan Javed Majeed and Marcus Hutter. 2018. On Q-learning convergence for non-Markov decision processes. *IJCAI*
 1423 2018-July (2018), 2546–2552. <https://doi.org/10.24963/IJCAI.2018/353>
- 1424 [30] Alexandre Maros, Fabricio Murai, Ana Paula Couto da Silva, Jussara M. Almeida, Marco Lattuada, Eugenio Gianniti,
 1425 Marjan Hosseini, and Danilo Ardagna. 2019. Machine Learning for Performance Prediction of Spark Cloud Applications.
 1426 In *IEEE 12th CLOUD*. 99–106. <https://doi.org/10.1109/CLOUD.2019.00028>
- 1427 [31] Simone Mentasti, Abednego Wamuhindo Kambale, and Matteo Matteucci. 2022. Event-Based Object Detection and
 1428 Tracking - A Traffic Monitoring Use Case -. In *ICIAP Workshops (Lecce, Italy)*. Springer-Verlag, Berlin, Heidelberg,
 1429 95–106. https://doi.org/10.1007/978-3-031-13324-4_9
- 1430 [32] Microsoft. 2024. HoloLens 2. Retrieved June 01, 2024 from [https://www.microsoft.com/en-au/hololens/hardware#](https://www.microsoft.com/en-au/hololens/hardware#document-experiences)
 1431 [document-experiences](https://www.microsoft.com/en-au/hololens/hardware#document-experiences)
- 1432 [33] Zhou Nan. 2021. The application of artificial intelligence technology in the communication engineering industry. In
 1433 *NetCIT*. 252–254. <https://doi.org/10.1109/NetCIT54147.2021.00057>
- 1434 [34] Arvind Narayanan, Xumiao Zhang, Ruiyang Zhu, Ahmad Hassan, Shuwei Jin, Xiao Zhu, Xiaoxuan Zhang, Denis
 1435 Rybkin, Zhengxuan Yang, Zhuoqing Morley Mao, Feng Qian, and Zhi Li Zhang. 2021. A variegated look at 5G in the wild:
 1436 Performance, power, and QoE implications. *ACM SIGCOMM* (2021), 610–625. <https://doi.org/10.1145/3452296.3472923>
- 1437 [35] Arjun Parthasarathy and Bhaskar Krishnamachari. 2022. Partitioning and Placement of Deep Neural Networks on
 1438 Distributed Edge Devices to Maximize Inference Throughput. In *32nd ITNAC*. IEEE, 239–246.
- 1439 [36] Teemu Pulkkinen, Jukka K. Nurminen, and Petteri Nurmi. 2020. Understanding WiFi Cross-Technology Interference
 1440 Detection in the Real World. In *ICDCS*. 954–964. <https://doi.org/10.1109/ICDCS47774.2020.00061>
- 1441 [37] Xiaoyu Qiu, Weikun Zhang, Wuhui Chen, and Zibin Zheng. 2021. Distributed and Collective Deep Reinforcement
 1442 Learning for Computation Offloading: A Practical Perspective. *IEEE Transactions on Parallel and Distributed Systems*
 1443 32, 5 (2021), 1085–1101. <https://doi.org/10.1109/TPDS.2020.3042599>
- 1444 [38] Qualcomm. 2022. *Snapdragon Summit 2022 | Snapdragon Tech Event*. Retrieved March 28, 2024 from [https://www.](https://www.qualcomm.com/company/events/snapdragon-summit)
 1445 [qualcomm.com/company/events/snapdragon-summit](https://www.qualcomm.com/company/events/snapdragon-summit)
- 1446 [39] Qualcomm. 2024. *Snapdragon AR2 Gen 1 Platform*. Retrieved March 28, 2024 from [https://www.qualcomm.com/](https://www.qualcomm.com/products/application/xr-vr-ar/snapdragon-ar2-gen-1-platform)
 1447 [products/application/xr-vr-ar/snapdragon-ar2-gen-1-platform](https://www.qualcomm.com/products/application/xr-vr-ar/snapdragon-ar2-gen-1-platform)
- 1448 [40] Starlin Daniel Raj and Karthiban. 2022. Applications of Artificial Intelligence in Healthcare. In *ICCCI*. 1–2. <https://doi.org/10.1109/ICCCI54379.2022.9741057>
- 1449 [41] Peter Rost and Gerhard Fettweis. 2010. On the transmission-computation-energy tradeoff in wireless and fixed
 1450 networks. *IEEE GC Wkshps* (2010), 1394–1399. <https://doi.org/10.1109/GLOCOMW.2010.5700167>
- 1451 [42] G. Rummery and M. Niranjan. 1994. On-Line Q-Learning Using Connectionist Systems. *Technical Report CUED/F-*
 1452 *INFENG/TR 166* (1994).
- 1453 [43] H. V. Seijen, H. V. Hasselt, S. Whiteson, and M. Wiering. 2009. A theoretical and empirical analysis of Expected Sarsa.
 1454 In *ADPRL*. 177–184. <https://doi.org/10.1109/ADPRL.2009.4927542>
- 1455 [44] Richard S Sutton and Andrew G Barto. 2018. *Reinforcement learning: An introduction*. MIT press.
- 1456 [45] Uma Tadakamalla and Daniel A. Menascé. 2022. Autonomic Resource Management for Fog Computing. *IEEE*
 1457 *Transactions on Cloud Computing* 10, 4 (2022), 2334–2350. <https://doi.org/10.1109/TCC.2021.3064629>
- 1458 [46] U. Tadakamalla and D. A. Menascé. 2022. Autonomic Resource Management for Fog Computing. *IEEE Transactions on*
 1459 *Cloud Computing* 10, 4 (2022), 2334–2350. <https://doi.org/10.1109/TCC.2021.3064629>
- 1460 [47] Ming Tang and Vincent W.S. Wong. 2022. Deep Reinforcement Learning for Task Offloading in Mobile Edge Computing
 1461 Systems. *IEEE Transactions on Mobile Computing* 21, 6 (2022), 1985–1997. <https://doi.org/10.1109/TMC.2020.3036871>
- 1462 [48] Peng Wang, Kenli Li, Senior Member, Bin Xiao, and Keqin Li. 2022. Multiobjective Optimization for Joint Task
 1463 Offloading, Power Assignment, and Resource Allocation in Mobile Edge Computing. *IEEE INTERNET OF THINGS*
 1464 *JOURNAL* 9 (2022), Issue 14. <https://doi.org/10.1109/JIOT.2021.3132080>
- 1465 [49] C. Watkins and P. Dayan. 1992. Autonomic Resource Management for Fog Computing. *Machine Learning* 8 (1992),
 1466 pages279–292. <https://doi.org/10.1007/BF00992698>
- 1467 [50] Wenbo Yao. 2019. The application of artificial intelligence in the internet of things. In *ITCA*. 141–144. <https://doi.org/10.1109/ITCA49981.2019.00038>
- 1468 [51] K Paul Yoon and Ching-Lai Hwang. 1995. *Multiple attribute decision making: an introduction*. Sage publications.
- 1469 [52] Ahmad Zendebudi and Salimur Choudhury. 2022. Designing a Deep Q-Learning Model with Edge-Level Training for
 1470 Multi-Level Task Offloading in Edge Computing Networks. *Applied Sciences* 2022, Vol. 12, Page 10664 12 (2022), 10664.
 Issue 20. <https://doi.org/10.3390/APP122010664>

A PARAMETERS AND DECISION VARIABLES

Table 4 summarises the problem parameters and decision variables adopted in Section 3.

Table 4. Parameters and variables.

Parameters	
Λ	Input workload [frames/second]
\mathcal{D}	Set of computing component indices; $\mathcal{D} = \{1, 2, 3\}$
\mathcal{K}	Set of all the candidate DNN configurations κ^i
p_\emptyset^i	Fictitious partition meaning that nothing is executed
δ_{dk}^i	Per-request data sent by partition p_d^i to p_k^i , $\forall d, k = (d+1) \in \mathcal{D}$
δ_0	Size of the input tensor
μ_d^i	Per-request computational workload for running p_d^i [flops]
t_d^i	Latency for running partition p_d^i on device d
F	Set of the network domains
T_f	Access time of network domain f
ϕ_f	Channel bandwidth of network domain f
r_f	Throughput of network domain f
g	Cost paid to transfer one byte of data over the 5G network
c_{5G}	Time unit cost paid to transfer the data over the 5G network
L_{max}	Maximum time latency allowed for the AI task
θ_{SEW}	Transmission power of the SEW
θ_{phone}	Transmission power of the phone
α	per unit cost of the energy consumption
z_{SEW}	Energy consumption in joules for performing one floating point operation on the SEW
z_{phone}	Energy consumption in joules for performing one floating point operation on the phone
Decision Variables	
x^i	Binary variable specifying that the configuration c^i is selected
P_{exeSEW}	Power consumption of the SEW when running a partition locally
$P_{exePhone}$	Power consumption of the phone when running a partition locally
P_{wnu}	Power consumption of the SEW WNI when uploading data to the smartphone
P_{5gu}	Power consumption of the phone 5G when uploading data to the cloud
P_{SEW}, P_{phone}	Total power consumption of the SEW and phone
R	Reward (this is the negative of the cost)
RL State Variables	
r_{WiFiT}	Data transmission rate between the SEW and the smartphone
r_{5G}	Data transmission rate between the smartphone and the cloud server
l_{SEW}	Latency associated with the execution on the SEW
l_{phone}	Latency associated with the execution on the smartphone
l_{cloud}	Latency associated with the execution on the cloud server
RL Parameters	
τ	Control time period
γ	Discount factor
lr	Learning rate
ϵ	Probability of selecting the next action at random

B TRACES, CONFIGURATIONS AND ADDITIONAL TRAINING DETAILS

This Appendix section provides additional details regarding the experiments outlined in Section 5. In Section B.1 we explain how we generated the 5G, WiFi and cloud latency traces, while we provide information on the applications candidate configurations in Section B.2. In Section B.3 we present the comparison among all the methods, both when varying the importance weights and the latency constraint threshold.

B.1 5G, WiFi and cloud latency traces

In real-world scenarios, the 5G network is an important source of variability. Hence, we used a 5G throughput trace taken from real-time measurements [34] to make the offline training close to the real environment conditions. The dataset contains 11024 rows, and the uplink throughput values range between 0 and 230.75 Mbps. As an episode takes 1.5 Million steps to complete, we replayed the trace after 11024 steps, and, to avoid the periodicity in the trace, we performed some operations as random shifts, insertion of random noise ($\pm 10\%$ of the value at time step t) and trace inversion in its middle. Figure 13a presents a snapshot of the 5G trace for the first 11000 time steps.

1520 On the other hand, when two or more people are using the SEWs in the same environment, the
 1521 problem of interference may arise, and the WiFi5 network throughput may change. Therefore, we
 1522 considered the WiFi5 network throughput as another source of variability. The WiFi5 trace was
 1523 obtained by collecting the throughput during the DNN profiling (the HoloLens was connected to
 1524 the mobile phone through a WiFi6E hotspot). Figure 13b provides a snapshot of the WiFi trace
 1525 for 3000 steps. As for the 5G trace, after 3000 steps the WiFi trace is replayed considering some
 1526 random shifts and the insertion of random noise.

1527 Finally, cloud latency fluctuation is another important source of variability, particularly when
 1528 the cloud is solicited by different entities, which can impact the execution time of the inference
 1529 task coming from the SEW application. To emulate the behavior of real-world applications, we
 1530 treated the cloud server as an M/M/1 queue, where requests follow a Poisson distribution, as it is
 1531 widely assumed in cloud and edge literature [46]. The cloud latency trace was obtained by drawing
 1532 the cloud execution time from exponential distributions whose mean values depend on the chosen
 1533 DNN configuration according to the values reported in Table 5.

1534 A snapshot of the cloud latency samples corresponding to the configuration which runs DNN
 1535 totally on cloud for the first 3000 time steps is given in Figure 13c.

1536

1537

1538

1539


1540

1541

1542

1543

1544



1545

1546

1547

1548

1549

1550

1551

1552

1553

1554

1555

1556

1557

1558

1559

1560

1561

1562

1563

1564

1565

1566

1567

1568

(a) 5G trace in steps [0, 11000]. (b) WiFi5 trace in steps [0, 3000]. (c) Cloud latency in steps [0, 3000].

Fig. 13. The snapshot of traces used for the experiments as the sources of variability in the system.

B.2 YOLOv5 partitioning and configuration parameters

We considered YOLOv5n model in ONNX format to avoid framework dependency, and relied on ONNX Runtime (version 1.13.1) to run different partitions on both the HoloLens 2 and the mobile phone. The goal of the profiling was to determine the execution time of each DNN partition on the three devices we consider. Figure 14 illustrates the different execution times for the DNN partitions considering the partitioning points between the HoloLens (i.e., SEW), the mobile phone, and between the phone and the edge (cloud) server. The end-to-end execution time ranges between 220ms and 620ms across the majority of the configurations.

Note that, after executing the partitioner, the YOLOv5 model exhibited 12 feasible partitioning points. To account for scenarios where nothing is run on a device or the entire DNN model is run on the device, we added two additional partitioning points, denoted as 0 and 13, respectively. The candidate configurations are then generated as follows:

- Two configurations account for the fact that the whole DNN can be executed on the phone or in the cloud;
- 12 candidate configurations can be constructed by choosing one of the possible partitioning points and deploying the first subset of layers on the SEW and the second one on the phone;

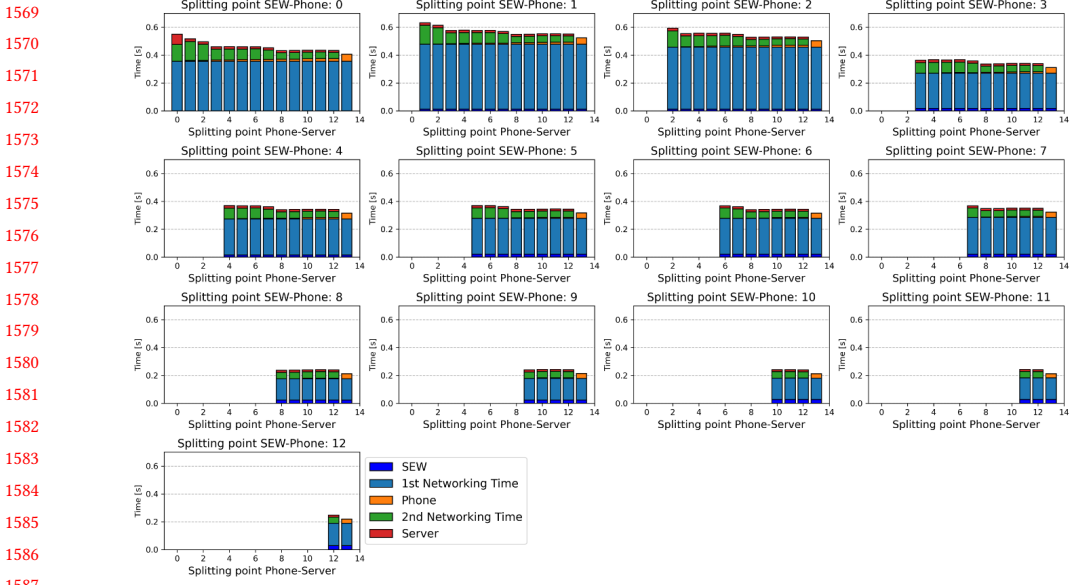


Fig. 14. Profiling of the YOLOv5n on three devices: the SEW (glasses), the phone, and the cloud server.

- Similarly, 12 candidate configurations consider one of the partitioning points and divide the computation between SEW and cloud;
- A third group of 12 candidate configurations selects a possible partitioning point and places the initial DNN layers on the phone and the others in the cloud.
- Finally, 66 candidate configurations can be generated by selecting two partitioning points among the 12 available, so that the computation is split among SEW, phone and cloud.

The total number of configurations can be determined by $\frac{n!}{s!(n-s)!} + n$, where n represents the total number of possible partitioning points (in our case, 14), and s is equal to the number of computation layers (devices) minus one. Overall we have a total of 105 candidate configurations.

Figure 14 reports all the candidate configurations and the corresponding profiling data, listed as detailed in the following: each sub-plot identifies a specific partitioning point between SEW and phone, while its x-axis lists the possible partitioning points between phone and cloud. Therefore, the first sub-plot (corresponding to the synthetic partitioning point 0, which identifies *no splitting*) lists the 14 candidate configurations for which no computation runs on the SEW; the first bar denotes the configuration running the whole DNN on the cloud (also the second partitioning point is 0, see as an example Figure 2a), and thereafter we progressively shift the partitioning point up to 13 (the last bar), which corresponds to the whole DNN running on the phone (see Figure 2b). Similarly, the second sub-plot characterizes the candidate configurations generated by selecting the first candidate partitioning point between SEW and phone, and then considering all the possible partitioning points (from 1 to 13) to further divide the computation between phone and cloud. We proceed as follows in all the remaining sub-plots, up to the last one, which selects the twelfth candidate partitioning point between SEW and phone.

Table 5 presents the range of configuration parameters obtained from the profiling. It gives the range of tensor sizes δ and π (in MB of data) that are transferred from the SEW to the phone, and from the phone to the cloud, respectively. It also provides the range for the latency (in ms) and the

Table 5. Ranges of DNN configuration parameters

Configuration parameters							
δ (MB)	π (MB)	Partition 1		Partition 2		Partition 3	
		latency(ms)	μ (MFLOPs)	latency(ms)	μ (MFLOPs)	latency(ms)	μ (MFLOPs)
(0, 6.25)	(0, 6.25)	(0, 108)	(0, 1680)	(0, 222)	(0, 5427)	(0, 78)	(0,5427)

workload (in MFLOPs) of all partitions when running on specific devices. Recall that Partition 1 runs on the SEW, Partition 2 on the phone, and Partition 3 on the cloud.

Crucial characteristics of the SEW (phone) in the experiments are the electrical energy z_{SEW} (z_{phone}) it needs to perform one FLOP and the electrical power θ_{SEW} (θ_{phone}) required by the network interface while sending data. We assessed these parameters using energy profiling on HoloLens and Samsung S23, executing YOLOv5 model partitions either locally or by transmitting intermediate tensors to a server.

To determine z_{SEW} , we charged the HoloLens to 100%, then ran a YOLOv5n partition until the battery reached 95%. This operation was performed for all the 14 possible partitions running on the SEW. The following step involved computing the energy injected in the HoloLens while recharging it from 95% to 100%. We measured the energy during this process using a TC66² device. The energy consumption of a partition is obtained as $\frac{e_m}{n_{iter}}$, where e_m represents the energy needed to charge the HoloLens from 95% to 100%, and n_{iter} is the number of partition executions (iterations) during the battery discharge from 100% to 95%. This yielded a value of $z_{SEW} \approx 1085.1$ pJ/FLOP through linear regression. Figure 15a and Figure 15b provide views of the real value versus the predicted value of the energy consumption of YOLOv5n on Samsung S23 and HoloLens 2, respectively.


Additionally, to determine θ_{SEW} , no partitions were executed on the HoloLens. Instead, intermediate tensors matching the size of those in YOLOv5 partitions were transmitted to the subsequent device, enabling the evaluation of data-transfer power consumption. Employing a similar methodology as for determining z_{SEW} , the power consumed during data transfer can be computed as $\frac{e_m}{n_{iter} \cdot t}$, where e_m represents the energy needed to charge the phone from 95% to 100%, n_{iter} denotes the iterations performed while transmitting the tensor during the HoloLens discharge from 100% to 95%, and t is the time taken for this discharge. Figure 15c and Figure 15d report the energy consumption while considering the different intermediate tensor sizes on Samsung S23 and HoloLens 2, respectively. Considering the time taken to send the results and the energy consumption, the profiling results revealed a constant power consumption of $\theta_{SEW} \approx 7.9$ W during data transmission.

On Samsung S23, we performed exactly the same energy profiling as described earlier on HoloLens 2 to determine z_{phone} and θ_{phone} . We obtained the values of 112.5 pJ/FLOP and 4.5 W for z_{phone} and θ_{phone} respectively.

B.3 Comparison of all methods

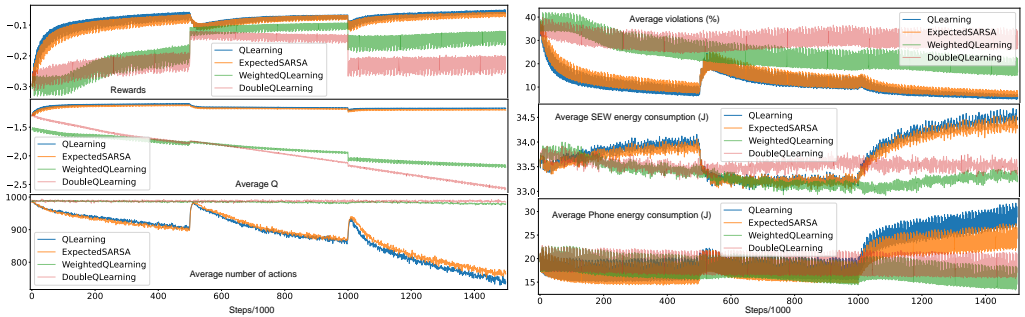
Figure 16 and Figure 17 compare all the methods across all the metrics we consider in the *Varying importance weights* (see Section 5.3) and *Varying latency constraint* (see Section 5.4) scenarios, respectively. Since the corresponding comparison between Q-learning and SARSA is already shown in Figure 3 and Figure 6, we omit SARSA in Figure 16 and Figure 17 to enhance the figures readability. As mentioned in Section 5, Q-learning proved to be the most effective method for solving our problem. Both SARSA and Expected SARSA demonstrated behavior comparable to Q-learning,

²<https://gzhls.at/blob/ldb/2/5/2/8/0e6f880a0a23ad0ef7a16e5c990f2c5b0216.pdf>



(a) Real versus predicted (b) Real versus predicted (c) Energy consumption (d) Energy consumption energy consumption of energy consumption of for YOLOv5n’s interme- for YOLOv5n’s interme- YOLOv5n partitions on YOLOv5n partitions on diate tensors transfer diate tensors transfer from Samsung S23 to next server. from Samsung S23 to next server. HoloLens

Fig. 15. Energy consumption on Samsung S23 and HoloLens 2 for executing YOLOv5n partitions and transferring intermediate tensors.



(a) Reward, Q, and number of actions. (b) Percentage number of violations and energy consumption.

Fig. 16. Methods comparison in varying importance weights scenario (average across 10 applications and 10 instances).

whereas the Weighted Q-learning and Double Q-learning agents were unable to adapt to our specific problem domain, even after hyperparameter exploration, and resulted in a significant number of execution time violations. Moreover, due to the large state space in our problem, the use of a large Q-table is impractical leading to increased memory and computational requirements, so that the training time of one episode for Double Q-learning exceeded that of other methods by more than 20 times.

Received 20 February 2007; revised 12 March 2009; accepted 5 June 2009

1716
1717
1718
1719
1720
1721
1722
1723
1724
1725
1726
1727
1728
1729
1730
1731
1732
1733
1734
1735
1736
1737
1738
1739
1740
1741
1742
1743
1744
1745
1746
1747
1748
1749
1750
1751
1752
1753
1754
1755
1756
1757
1758
1759
1760
1761
1762
1763
1764

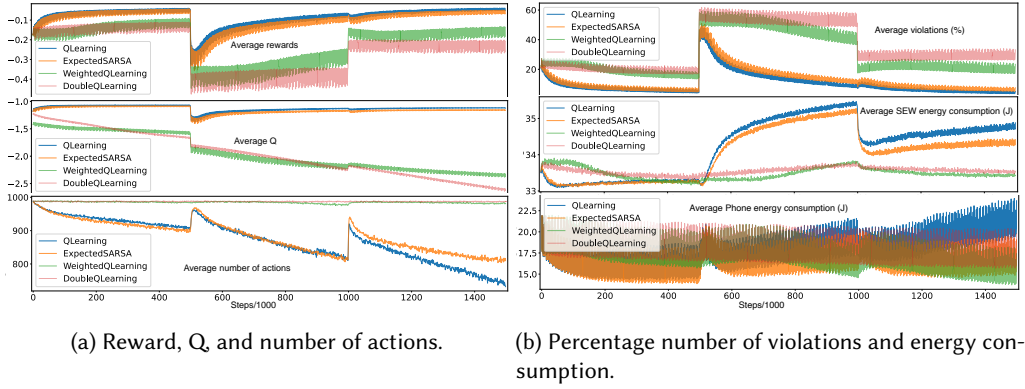


Fig. 17. Methods comparison in varying latency constraint scenario (average across 10 applications and 10 instances).