

# 000 001 002 003 004 005 006 007 008 009 010 011 012 013 014 015 016 017 018 019 020 021 022 023 024 025 026 027 028 029 030 031 032 033 034 035 036 037 038 039 040 041 042 043 044 045 046 047 048 049 050 051 052 053 FROM INTENTS TO ACTIONS: AGENTIC AI IN AUTONOMOUS NETWORKS

Anonymous authors

Paper under double-blind review

## ABSTRACT

Telecommunication networks are increasingly expected to operate autonomously while supporting heterogeneous services with diverse and often conflicting *intents*—that is, performance objectives, constraints, and requirements specific to each service. However, transforming high-level intents—such as ultra-low latency, high throughput, or energy efficiency—into concrete control actions (i.e., low-level actuator commands) remains beyond the capability of existing heuristic approaches. This work introduces an Agentic AI system for intent-driven autonomous networks, structured around three specialized agents. A supervisory *interpreter agent*, powered by language models, performs both lexical parsing of intents into executable optimization templates and cognitive refinement based on feedback, constraint feasibility, and evolving network conditions. An *optimizer agent* converts these templates into tractable optimization problems, analyzes trade-offs, and derives preferences across objectives. Lastly, a preference-driven *controller agent*, based on multi-objective reinforcement learning, leverages these preferences to operate near the Pareto frontier of network performance that best satisfies the original intent. Collectively, these agents enable networks to autonomously interpret, reason over, adapt to, and act upon diverse intents and network conditions in a scalable manner.

## 1 INTRODUCTION

Radio access networks (RANs) are large-scale, real-time distributed systems that must operate reliably in highly dynamic and uncertain radio environments, while serving a broad range of connectivity services and applications. Currently, these systems rely heavily on manual intervention for configuration optimization and functional fine-tuning. This dependence on human expertise limits scalability, slows adaptation to environmental changes, and increases operational costs.

The next generation of communication networks is expected to address these limitations by becoming increasingly autonomous. This evolution—already underway in 5G-Advanced through standardized intent management frameworks, e.g., 3GPP (2025c) and TMForum (2021)—envisions self-configuring, self-optimizing, and self-healing systems guided by high-level *network intents*. Intents specify performance objectives, requirements, and constraints for a connectivity service or management workflow 3GPP (2025c), allowing operators to express *what* the network should achieve rather than *how*. For example, an operator may specify a goal as “*maximize user coverage while minimizing energy consumption*,” leaving the network to autonomously determine the appropriate actions, such as antenna tilt adjustments to improve coverage or carrier deactivation to save energy. In this context, intents act as directives, while the network abstracts away the implementation details, much like a compiler translates high-level code into machine-executable instructions.

Converting intents into network actions is fundamentally a problem of planning and reasoning across multiple abstraction layers—from natural-language specifications to optimization formulations, and ultimately to control policies executed at the RAN. These requirements exceed the capabilities of current heuristic and rule-based approaches. Bridging this gap calls for a new class of artificial intelligence (AI) systems that move beyond perception and prediction, linking abstract objectives with dynamic decision-making through iterative reasoning and planning.

Agentic AI has recently emerged as a promising paradigm for building autonomous, goal-driven systems capable of interpreting objectives, planning multi-step actions, and adapting to dynamic envi-

054 environments with minimal human oversight. Unlike traditional AI approaches based on fixed heuristics  
 055 or monolithic models, Agentic AI structures intelligence into specialized agents that interact and co-  
 056 operate through well-defined workflows Sapkota et al. (2025). Central to this paradigm are large-scale  
 057 generative models—particularly large language models (LLMs)—which enable agents to understand  
 058 and generate natural language, decompose goals, generalize across tasks, invoke specialized tools, and  
 059 reason in open-ended contexts Liu et al. (2024). As such, Agentic AI offers a compelling architectural  
 060 foundation for autonomous and intent-driven network management and optimization.

061 This paper takes a step toward realizing this vision by introducing an Agentic AI system comprising  
 062 an **interpreter**, an **optimizer**, and a **controller**. Our contributions are:  
 063

- 064 **1. Cognitive intent processing.** The interpreter is a supervisory cognitive agent with two core  
 065 functions: converting high-level intents into structured templates and recursively refining them  
 066 on a slow timescale by reasoning over network observations and feedback on intent fulfillment.  
 067 To meet RAN compute and memory constraints, we adopt a dual-SLM architecture that separates  
 068 intent translation and in-context reasoning among two small language models (SLMs).
- 069 **2. Preference optimization.** The optimizer agent transforms optimization template models (OTMs)  
 070 into constrained optimization problems over a preference space, performs preference planning  
 071 via Bayesian optimization to dynamically adapt preferences to network conditions, and steers  
 072 the controller policy to satisfy the service intents expressed by the OTM.
- 073 **3. Multi-objective control.** The controller leverages multi-objective reinforcement learning  
 074 (MORL) to realize adaptive policies that operate near the Pareto front of network performance. A  
 075 central technical contribution is distributed envelope Q-learning (D-EQL), a scalable distributed  
 076 variant of envelope Q-learning (EQL) Yang et al. (2019) that: (i) decouples learner–actors with  
 077 sharded prioritized replay for high-throughput training; (ii) distributes the exploration of the  
 078 preference simplex across actors while learning a single preference-conditioned network; (iii)  
 079 uses envelope updates with vector TD targets plus a cosine-stability loss; and (iv) refreshes  
 080 priorities with hindsight preference relabeling. Together, these extensions improve scalability,  
 081 accuracy and exploration over established MORL art Yang et al. (2019); Basaklar et al. (2023).
- 082 **4. Proof of concept.** We showcase the agentic system through an intent-aware radio resource  
 083 management (RRM) use case combining interpreter and optimizer agents with a novel MORL-  
 084 based link adaptation (LA), and adapt its policy on the fly to diverse connectivity service goals.  
 085 Our approach outperforms traditional reinforcement learning (RL)—which cannot adapt a single  
 086 policy across goals—and exceeds the state-of-the-art LA baseline of 5G/5G-A systems.

087 Results from high-fidelity system-level simulations of a 5G-compliant network suggest that Agentic  
 088 AI can transform high-level human intents into self-optimizing control mechanisms for next-  
 089 generation networks, thereby paving the way toward scalable network autonomy.

## 090 2 RELATED WORK

093 **Agentic AI:** Agentic AI is an emerging paradigm that structures intelligence as a modular network  
 094 of specialized agents collaborating to achieve complex, high-level goals (Hughes et al., 2025). Recent  
 095 surveys highlight recurring design patterns and challenges related to reliability and evaluation (Guo  
 096 et al., 2024; Li et al., 2024). A central mechanism is *goal decomposition*, whereby broad objectives  
 097 are divided into subtasks handled by agents with distinct functions. Prior work has demonstrated that  
 098 agents can integrate reasoning and action in recursive loops (Yao et al., 2023), improve performance  
 099 through reflective memory (Shinn et al., 2023), and operate collectively via structured communication  
 100 (Wu et al., 2024). To coordinate distributed intelligence, orchestration layers or meta-agents  
 101 assign roles, manage life cycles and task dependencies, and resolve conflicts using centralized or  
 102 decentralized mechanisms (Qian et al., 2024). Furthermore, persistent goals and memory enable  
 103 adaptation over long time horizons (Wang et al., 2024; Agashe et al., 2025). Domain-specific systems,  
 104 such as MAGIS (Tao et al., 2024), illustrate how these principles scale to collaborative workflows.

105 **Bayesian optimization:** Zhan & Xing (2020) reviews the evolution of expected improvement (EI)  
 106 as an acquisition function for surrogate-based optimization, detailing its extensions to parallel, multi-  
 107 objective, constrained, noisy, multi-fidelity, and high-dimensional settings, analyzing their theoretical  
 108 properties, and highlighting future research directions. Zhao et al. (2024) shows that the performance

108 of high-dimensional Bayesian optimization is strongly limited by poor random initialization of  
 109 acquisition function maximizers and proposes AIBO, a simple framework that uses past evaluations  
 110 and heuristic search to generate better starting points, significantly boosting optimization efficiency.  
 111

112 **Multi-objective reinforcement learning:** MORL addresses control problems in which optimality  
 113 is defined by a Pareto front of policies, each capturing different trade-offs among multiple objectives.  
 114

115 Early approaches to multi-objective optimization (Kim & de Weck, 2005; Konak et al., 2006; Yoon  
 116 et al., 2009) reduced the problem to scalar optimization—typically via utility functions with fixed  
 117 weights across objectives—followed by standard RL. These methods are tied to a single preference  
 118 setting and cannot adapt when goals or constraints change (Liu et al., 2015), thereby necessitating  
 119 retraining. To improve generality, subsequent work sought to approximate the entire Pareto front by  
 120 learning multiple optimal policies over the preference space (Natarajan & Tadepalli, 2005; Barrett &  
 121 Narayanan, 2008; Mossalam et al., 2016). However, training a separate policy for each preference  
 122 combination quickly becomes computationally infeasible in large domains.  
 123

124 A more scalable approach is to learn a single universal policy conditioned on preferences (Yang et al.,  
 125 2019; Xu et al., 2020; Abdolmaleki et al., 2020), enabling adaptation across tasks without retraining.  
 126 For instance, Yang et al. (2019) proposed envelope Q-learning, which generalizes the Bellman  
 127 equation to optimize the convex envelope of multi-objective Q-values under linear preferences  
 128 using deep networks. Extensions such as those in Basaklar et al. (2023) introduced parallelization  
 129 to improve sample efficiency and Pareto approximation. Nonetheless, efficiently exploring the  
 130 preference space and learning universal MORL policies remain open challenges (Hayes et al., 2022).  
 131

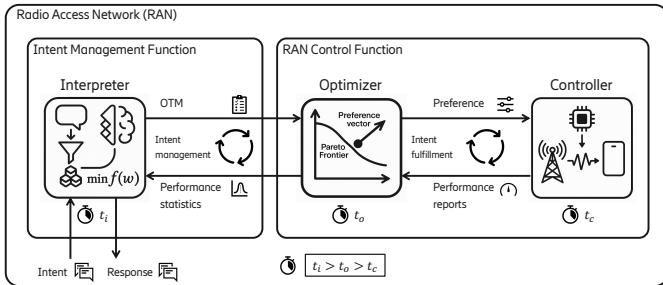
132 **Agentic AI in Communication Systems:** Intent-based management is already part of modern  
 133 5G-Advanced systems 3GPP (2025c), and its extension toward 6G is strongly supported in current  
 134 standardization efforts 3GPP (2025g). Concurrently, academic and industrial interest in Agentic  
 135 AI is rapidly growing, positioning it as a key enabler of next-generation autonomous networks,  
 136 particularly for intent-driven operations Bimo et al. (2025); ZTE (2025); Intel & NEC (2025). Recent  
 137 work on agent-based and LLM-guided control frameworks for network optimization and service  
 138 management Qayyum et al. (2025); Jolicoeur-Martineau (2025); Bimo et al. (2025) highlights a  
 139 shift toward systems capable of reasoning, adaptation, and collaboration. This trajectory is reflected  
 140 across 3GPP, Open RAN, and TM Forum. For example, 3GPP TR 22.870 3GPP (2025a) identifies  
 141 AI-agent-enabled service coordination, LLM-assisted interactions, and agent-supported UE–network  
 142 cooperation as 6G use cases, while IETF (2025) defines protocols for AI-agent communication.  
 143 Furthermore, the 3GPP SA5 workgroup has identified *intent-driven agentic autonomous management*  
 144 as a priority areas for 6G 3GPP (2025i;h) while SA2 is examining agentic mechanisms for the 6G  
 145 core network 3GPP (2025f). Together, these developments indicate that agentic and intent-based  
 146 paradigms are increasingly viewed as foundational elements of future 6G architectures.  
 147

148 **Differentiation from Prior Agentic AI Work:** Existing Agentic AI systems have largely been  
 149 applied to reasoning, planning, and tool use, where control loops operate over long timescales in  
 150 relatively stable environments. By contrast, we integrate agentic AI into the fast control loops of  
 151 RRM, where sub-millisecond decisions must adapt to fading channels, mobility, and heterogeneous  
 152 service requirements. To our knowledge, this is among the first applications of Agentic AI in highly  
 153 dynamic, stochastic environments, extending its reach to performance-critical autonomous networks.  
 154

155 We demonstrate the workflow with an end-to-end, cognitively guided intent-aware RRM design  
 156 for supporting different connectivity services, where control policies adapted by reasoning over  
 157 individual service goals and network observations are then executed in time-varying, frequency-  
 158 selective environments to meet the goals. Our results show superior performance compared to  
 159 traditional RL and the state-of-the-art LA algorithm adopted in 5G/5G-A systems.  
 160

### 161 3 AGENTIC AI SYSTEM FOR RAN CONTROL

162 At its core, the proposed *Agentic AI system* comprises three specialized agents—interpreter, optimizer,  
 163 and controller—whose interactions form an *agentic workflow* consisting of two loops: an *intent*  
 164 *management loop*, executed by the interpreter–optimizer pair, and an *intent fulfillment loop*, executed  
 165 by the optimizer–controller pair. Each loop operates on a distinct timescale, forming a two-timescale  
 166 control architecture analogous to Kahneman’s dual-process theory (Kahneman, 2011), with a slower,  
 167 deliberative outer System 2 and a faster, reactive inner System 1.  
 168

162  
163  
164  
165  
166  
167  
168  
169  
170  
171172 Figure 1: Agentic AI system for intent and resource management in autonomous networks.  
173174 The *interpreter* is a supervisory cognitive agent that converts high-level intents into structured  
175 templates and adaptively refines them on a slow timescale using network states and fulfillment  
176 feedback. The *optimizer* recursively plans and adjusts the downstream controller configurations to  
177 satisfy the intent, aggregating controller feedback into slower-timescale statistical summaries returned  
178 to the interpreter. The *controller* executes real-time decision-making, collects observations, and  
179 provides periodic performance reports to the optimizer.180 This triadic workflow provides a blueprint for a broader Agentic AI system for autonomous man-  
181 agement and optimization of communication networks. Its realization, however, requires a twofold  
182 extension. *Horizontally*, the interpreter may coordinate with multiple optimizer–controller pairs  
183 supporting different RAN functions within a single architectural layer. *Vertically*, the workflow can  
184 be embedded across different layers of the RAN protocol stack, whose operational timescales range  
185 from slow (for network management at higher layers) to very fast (for RRM at lower layers).186  
187

## 3.1 TIMESCALES SEPARATION

188  
189  
190  
191  
192  
193  
194  
195  
196  
197  
198The workflow separates responsibilities across three timescales. The controller handles real-time  
decisions and thereby establishes the system’s reference timescale  $t_c$ . Because this agent replaces an  
existing RAN control function, it inherits that function’s native latency budget, which may range from  
sub-millisecond operation for RRM functions (e.g., link adaptation) to minutes or hours for network  
optimization tasks (e.g., cell shaping). The optimizer adjusts the controller’s policy at a deliberately  
slower timescale  $t_o$ , spanning hundreds of milliseconds to seconds for fast RRM functions and up to  
hours for RAN management functions, ensuring that its decisions do not interfere with the primary  
control loop. The interpreter operates on the slowest supervisory cadence  $t_i$ , which spans seconds  
to minutes for RRM supervision and up to hours for RAN-wide management. At this timescale,  
the interpreter evaluates intent feasibility, reasons over observed key performance indicator (KPI)  
deviations, and generates refined intents without imposing timing constraints on downstream agents.199  
200  
201  
202Decoupling long-term reasoning and intermediate adaptation from real-time control ensures that (a)  
the interpreter supervisory role is non-latency-critical; (b) latency-critical operations are confined to  
the controller—for any RAN control function involved; and (c) the fast control loop remains stable.203  
204

## 4 LANGUAGE-GUIDED INTENT MANAGEMENT

205  
2064.1 INTERPRETER AGENT  
The interpreter is a language-guided supervisory agent aligned with the scope of an intent management  
function (IMF) (TMForum, 2024). It performs two complementary functions: (a) *transforming intents*  
into structured OTMs, and (b) *cognitive reasoning* for recursive intent adaptation.210  
211  
212  
213  
214  
215The interpreter agent must integrate domain awareness, intent stabilization, and adherence to the  
computational and memory constraints of the RAN system. Domain awareness includes understanding  
which control agents operate within each sub-domain, their capabilities, parameters, and timescales,  
as well as the KPIs they influence. This knowledge enables the interpreter to produce feasible OTM  
formulations for a given intent, route each intent to the appropriate RAN control agent, and ensure  
intent stabilization by reasoning over system observations, optimizer feedback, and network dynamics  
to perform safe, explainable OTM refinements when required.

Model	Schema accuracy		OTM accuracy	
	Objectives	Constraints	Objectives	Constraints
Qwen-2.5-7B-Instruct (Before fine-tuning)	100.0%	45.00%	21.50%	11.30%
Qwen-2.5-7B-Instruct (After fine-tuning)	100.0%	<b>100.0%</b>	<b>98.00%</b>	<b>98.00%</b>

Table 1: Schema and OTM accuracy for interpreters using the Qwen-2.5-7B-Instruct model.

Meeting these requirements within current 5G/5G-A RAN hardware necessitates a design that is both computationally efficient and functionally modular. Deploying a single large general-purpose LLM is impractical due to compute and memory constraints in current RAN platforms, and integrating dedicated accelerators is neither scalable nor cost-effective. To address this, we adopt a dual-SLM architecture that separates the interpreter’s two core functions—intent translation and cognitive reasoning—across two lightweight, complementary SLMs, as detailed in Appendix A.

**Intent translation.** This module is the workflow entry point. It interprets the intent, decomposes it into sub-intents, selects the appropriate downstream control agent, and initiates the intent-fulfillment loop. A fine-tuned SLM renders the intent as a structured, schema-compliant OTM by disambiguating objectives, constraints, requirements, and metadata. This step extends beyond lexical parsing: the model must map high-level intents into optimization structures grounded in domain knowledge. Using a fine-tuned SLM ensures low-complexity generation of machine-readable OTMs that reflect RAN semantics and remain robust to linguistic variability. Appendix B discusses the generality of the OTM schema, while Appendix C outlines the fine-tuning of a Qwen-2.5-7B-Instruct model Qwen et al. (2025), which achieves the high schema validity and OTM accuracy shown in Table 1.

**Cognitive reasoning and adaptation.** Complementing the translator, a lightweight general-purpose SLM performs supervisory reasoning via in-context learning. It evaluates feasibility, diagnoses constraint violations, and refines OTMs when strict requirements cannot be met, proposing alternative trade-offs or adapting objectives to evolving network conditions. Intent stabilization is achieved through structured monitoring, advisory evaluation, and guarded execution (see Appendix A). This supervisory closed-loop reasoning extends beyond static templates or rule-based logic and is essential for autonomous, intent-driven, network management under real-world network dynamics.

This division of labor preserves contextual knowledge and ensures adaptability for intent handling, while remaining compatible with practical constraints of contemporary RAN deployments. The dual-SLM interpreter—built from small-scale models—and the infrequent, non-latency-critical nature of SLM inference within the agents’ timescale separation allow the system to maintain low compute and energy overhead. As a result, the overall design is feasible on current 5G/5G-Advanced hardware.

## 4.2 OPTIMIZER AGENT

The optimizer agent performs three key tasks: (i) decoding the OTM received from the interpreter, (ii) recursively solving the associated optimization problem to align the controller’s policy with the intent, and (iii) coordinating the two feedback loops within the workflow. Upon receiving an OTM, the optimizer formulates a constrained optimization problem aligned with the specified intent, such as

$$\begin{aligned} & \underset{\omega \in \Omega}{\text{minimize}} \quad f(\omega) \\ & \text{subject to} \quad g_i(\omega) \leq b_i, \quad i = 1, \dots, p, \end{aligned} \tag{1}$$

where  $f(\omega)$  quantifies the system performance (e.g., energy, latency, throughput), and the decision variable  $\omega$  belongs to a feasible set  $\Omega \subseteq \mathbb{R}^m$ . The inequality constraints  $g_i(\omega) \leq b_i$  capture operational limitations—e.g., bandwidth, latency, or power—or service requirements. Since both objective and constraints are often non-convex, the solution landscape may contain multiple local optima, making the identification of feasible or optimal solutions challenging.

The decision variables  $\omega$  link the optimizer to the controller by representing hyperparameters that tune the controller’s policy. In our framework, the controller follows a MORL approach (Section 5), so  $\omega$  corresponds directly to the preference weights in its multi-dimensional reward function.

Since the explicit forms of  $f$  and  $g_i$  are unknown and their evaluations are computationally expensive, the optimizer employs Bayesian optimization (BO), leveraging surrogate models trained on RAN performance data (e.g., throughput, spectral efficiency, block error rate (BLER)) relevant to the intent.

270 These models guide the exploration of preference weights  $\omega$  (i.e., decision variables), which steer the  
 271 controller’s actions. Additional details of the BO design are provided in Appendix D.  
 272

273 **4.2.1 PAX-BO: PREFERENCE-ALIGNED EXPLORATION BAYESIAN OPTIMIZATION**  
 274

275 We next address the preference-based constrained BO problem (1) in the multi-service case, where  $S$   
 276 connectivity services must be jointly optimized under  $p$  constraints that capture requirements such as  
 277 data rate, latency, and reliability. The optimization problem (1) becomes

$$\begin{aligned} 278 \quad & \underset{\mathbf{W} \in \Omega^S}{\text{minimize}} \quad f(\mathbf{W}) \\ 279 \quad & \text{subject to} \quad g_i(\mathbf{W}) \leq b_i, \quad i = 1, \dots, p, \end{aligned} \quad (2)$$

281 where  $\mathbf{W} = [\omega^{(1)}, \dots, \omega^{(S)}]$  collects the service-specific preference vectors  $\omega^{(s)} \in \Omega$  ( $\Omega = \Delta^{m-1}$ )  
 282 on the probability simplex. The objective  $f(\mathbf{W})$  quantifies system-wide performance, while the  
 283 constraints  $g_i(\mathbf{W}) \leq b_i$  enforce joint service requirements. Problem (2) reduces to (1) when  $S = 1$ .  
 284

285 **PAX-BO**, shown in Algorithm 2, solves Problem (2) by optimizing preference vectors on the  
 286 simplex through BO in an unconstrained internal space. Let  $U = [u^{(1)}, \dots, u^{(S)}] \in \mathbb{R}^{m \times S}$  and  
 287  $\bar{u} = \text{vec}(U)$ . Each service  $s$  has a projected simplex weight  $\omega^{(s)} = \Pi_{\Delta}(u^{(s)}) \in \Delta^{m-1}$ , and  
 288  $\mathbf{W}(U) = [\omega^{(1)}, \dots, \omega^{(S)}] \in (\Delta^{m-1})^S$ . At each iteration, we fit surrogate models that approximate  
 289 the system objective and constraints as  $\mathcal{F}(\bar{u}) \approx f(\mathbf{W}(U))$  and  $\mathcal{G}_i(\bar{u}) \approx g_i(\mathbf{W}(U))$ , and build a  
 290 constraint-aware acquisition  $\alpha(\bar{u})$  (e.g., Log-EI times a feasibility probability).

291 A *trust region* (TR)—an  $\ell_\infty$  box with center  $s_c$  and radius  $L \in [L_{\min}, L_{\max}]$ —constrains local  
 292 exploration. At each iteration, the acquisition function is maximized within the TR, and the solution  
 293 is projected back onto the simplex:

$$294 \quad \bar{u}_t = \arg \max_{\|\bar{v} - s_c\|_\infty \leq L} \alpha(\bar{v}), \quad U_t = \text{mat}(\bar{u}_t), \quad \mathbf{W}_t = \Pi_{\Delta}(U_t).$$

297 After evaluating  $o_t = f(\mathbf{W}_{t-1})$  and  $c_t^{(i)} = g_i(\mathbf{W}_{t-1})$ , we declare success if  $c_t^{(i)} \leq 0$  for all  $i$  and  
 298  $o_t \geq f_{t-1}^* + \epsilon$ , with  $\epsilon \ll 1$ . On success, we set  $f_t^* \leftarrow o_t$ ,  $s_c \leftarrow \bar{u}_{t-1}$ , and expand  $L$  after  $s_{\text{th}}$   
 299 consecutive successes; otherwise,  $L$  is shrunk after  $f_{\text{th}}$  failures, clamped to  $[L_{\min}, L_{\max}]$ . If the TR  
 300 stalls at  $L_{\min}$  for  $w$  rounds, a *reset* is triggered:  $n$  candidates are sampled from  $(\Delta^{m-1})^S$ , scored by  
 301 (acquisition)  $\times$  (feasibility)  $\times$  (novelty), and the best candidate reinitializes  $s_c$  with  $L \leftarrow L_0$ .  
 302

302 Overall, PAX-BO is a lift-and-project BO method with TR safeguards and reset mechanisms, tailored  
 303 to simplex-valued preferences that jointly influence a constrained system objective.  
 304

305 **5 PREFERENCE-GUIDED INTENT FULFILLMENT**  
 306

307 The optimizer and controller agents operate in a closed loop to achieve intent fulfillment. The  
 308 optimizer recursively adapts the preference vector  $\omega$  based on performance feedback from the  
 309 controller. The optimal (or near-optimal) vector  $\omega^*$ , obtained by solving (1), is then passed to the  
 310 controller, which aligns network actions with the original intent.  
 311

312 **5.1 CONTROLLER AGENT**  
 313

314 The controller implements a policy trained via D-EQL, a distributed extension of EQL (Yang et al.,  
 315 2019). D-EQL learns a single policy/value network conditioned on a linear preference vector  $\omega \in \Omega$   
 316 (the probability simplex) and scales exploration through a learner–actor architecture with prioritized  
 317 replay (cf. Horgan et al. (2018)).  
 318

318 During training, actors are assigned to distinct strata of the simplex defined by a simplex-lattice  
 319 partition. Each actor samples preferences uniformly within its stratum using barycentric sampling,  
 320 executes an  $\epsilon$ -greedy policy with the scalarization

$$321 \quad Q_\omega(s, a; \theta) = \omega^\top Q(s, a, \omega; \theta),$$

322 and initializes replay priorities by drawing an independent preference  $\tilde{\omega}$  to compute a scalar temporal-  
 323 difference error. Transitions and priorities are batched locally and sent to sharded replay buffers.  
 324

Algorithm	Partition	Hindsight	Replay memory			Actor	CFR1	Hypervol.
			Sampling	Update	Sharded	Distrib.	Comm.	Improv.
Yang et al. (2019)	No	Yes	Prioritized	No	No	No	–	–
Basaklar et al. (2023)	Yes	Yes	Uniform	No	No	Yes	Synch.	12.33%
D-EQL (ours)	Yes	Yes	Prioritized	Yes	Yes	Yes	Asynch.	<b>22.10%</b>
								<b>89.37%</b>

Table 2: Comparison of D-EQL with Yang et al. (2019) and Basaklar et al. (2023) in terms of design features and achieved CFR1 performance in Fruit Tree Navigation with depth 7.

The learner assigns strata of the simplex to actors for distributed exploration, retrieves prioritized minibatches from all shards, samples preferences from a Dirichlet distribution, and forms a Cartesian product so that each transition is evaluated under every sampled preference. The learner performs envelope backups by maximizing over actions and supporting preferences, updates parameters using a regression loss with an optional cosine-alignment term, refreshes priorities, and periodically synchronizes the target network. Updated weights are then broadcast to all actors.

The envelope backup is expressed as

$$\mathbf{y} = \mathbf{r} + \gamma(1 - d) Q(s', a^*, \tilde{\omega}^*; \theta^-), \quad (a^*, \tilde{\omega}^*) = \arg \max_{a', \omega' \in \Omega} \omega^\top Q(s', a', \omega'; \theta).$$

Compared with state-of-the-art MORL algorithms such as Yang et al. (2019) and Basaklar et al. (2023), D-EQL introduces (i) a hindsight replay memory with prioritized sampling and priority updates, (ii) partitioned exploration of the preference space across distributed asynchronous actors, and (iii) a sharded replay memory. This architecture improves scalability in environments with large state-action-preference spaces by enabling systematic simplex exploration, dense preference supervision, and high-throughput stable learning. As shown in Table 2, D-EQL achieves a 22.1% performance CFR1 improvement over Yang et al. (2019) and an additional 8% gain over Basaklar et al. (2023) in the Fruit Tree Navigation environment with depth 7, as well as 89.37% hypervolume improvement over Yang et al. (2019) and an extra 6.05% gain over Basaklar et al. (2023). Additional design details and extended comparisons are provided in Appendix F.

## 6 CASE STUDY: AGENTIC RADIO RESOURCE MANAGEMENT

RRM encompasses some of the most demanding and dynamic control functions in RANs, including user scheduling, resource allocation, link adaptation, power control, and beamforming. These mechanisms operate on sub-millisecond timescales and must continuously adapt to the stochastic nature of the wireless channel to maintain reliable and efficient over-the-air communications.

As proof of concept, we apply our Agentic AI system to support differentiated connectivity services using a MORL-based controller agent for LA—a key function that tunes modulation and coding scheme (MCS) parameters to the radio link capacity. The detailed description of the MORL LA controller agent is provided in Appendix G. Here, we note that the reward is a vector  $\mathbf{r} = [r_1, r_2]^\top \in \mathbb{R}^2$  with two competing components:  $r_1$  measures the number of information bits successfully delivered per packet, and  $r_2$  captures the time-frequency resources consumed per packet transmission.

In our agentic system, the MORL LA controller agent defines the fastest operational timescale, running on a sub-millisecond cadence. This cadence sets the reference timescale for dimensioning the optimizer and interpreter. The optimizer updates the preference weights of the MORL controller once per second, based on performance reports and observed network conditions. This update rate is fast enough to steer the controller toward MCS selections aligned with the intent goals, yet slow enough not to interfere with the stability of the LA decision loop.

At the same cadence, the optimizer agent provides feedback to the interpreter agent for supervisory monitoring of intent fulfillment. However, the interpreter’s cognitive loop is triggered only on an event-driven basis. Upon receiving an alert message from the optimizer, the interpreter leverages its general-purpose SLM to perform cognitive reasoning over KPIs deviations, intent-fulfillment, and evolving network conditions to determine whether the intent must be refined. In our case study, such intervention occurs when changing network conditions render the service requirements temporarily infeasible.

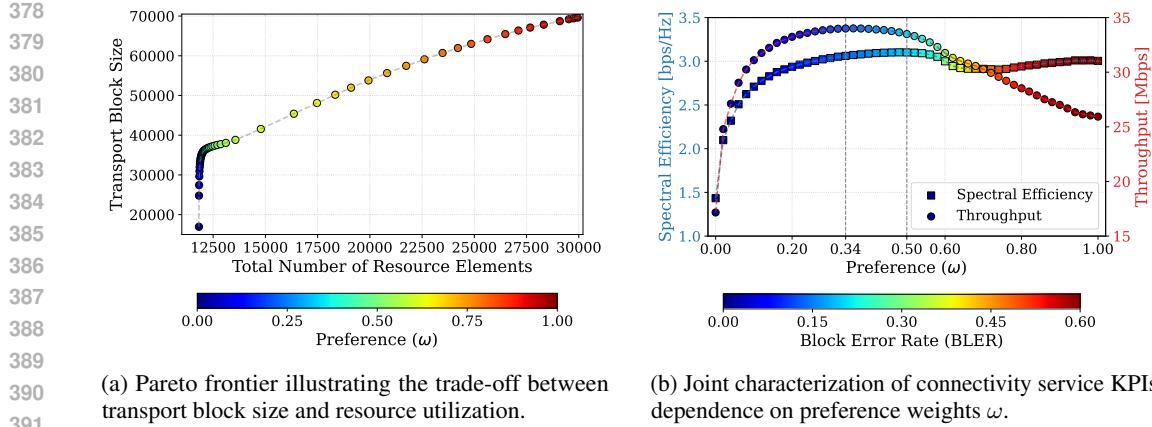


Figure 2: Characterization of preference-guided LA using MORL to satisfy service intents.

## 7 EXPERIMENT

This section evaluates the empirical performance of our Agentic AI system for intent-aware RRM using a 5G-compliant event-driven network simulator. We validate our approach in three steps using a multi-cell setup described in Appendix H: First we validate the MORL controller agent design; secondly, we evaluate the optimizer-controller loop; and lastly we benchmark the overall workflow.

### 7.1 MORL CONTROLLER AGENT FOR LINK ADAPTATION

Figure 2 illustrates how the preference-guided MORL controller for LA steers trade-offs among service KPIs, like spectral efficiency, throughput, and BLER, assuming long communication sessions (e.g., streaming services). Figure 2a shows the Pareto frontier for the two reward components, while Figure 2b maps each point on the frontier to link-level KPIs. When  $\omega \approx 0$ , the controller selects conservative MCS values, resulting in resource efficient and high-reliable transmissions (with near-zero BLER), but at the cost of low throughput (i.e., due to small transport block sizes (TBSs)) and spectral efficiency. At the other extreme,  $\omega \approx 1$  drives aggressive MCS choices that exploit retransmissions to target a spectral efficiency beyond the channel capacity, inducing resource-hungry and unreliable transmissions (with BLER  $\approx 60\%$ ). The best operating points emerge for intermediate preferences, with  $\omega \approx 0.34$  maximizing throughput and  $\omega \approx 0.5$  maximizing spectral efficiency. Appendix H extends the analysis to examples with multiple connectivity services.

### 7.2 INTENT-FULFILLMENT LOOP VALIDATION

Next, we evaluate *only* the optimizer–controller loop, assuming a single forward interaction with the interpreter to obtain an OTM. That is, when stochastic changes in the RAN environment render the OTM specifications infeasible, the interpreter’s cognitive refinement loop is not triggered. While the optimizer–controller pair cannot resolve temporary infeasibility caused by evolving RAN conditions.

We illustrate this by considering an intent that combines two contrasting connectivity services:

*Maximize cell throughput while serving mobile broadband users on a best-effort basis, and guaranteeing 99.99% reliability for a ultra-reliable traffic.*

This intent reflects quality of service (QoS) requirements for streaming and reliable services. In the agentic workflow, the interpreter constructs an OTM that (a) identifies the two services, (b) defines an overall objective based on their achieved throughput, and (c) formulates a reliability constraint for the reliable service. The optimizer then instantiates an optimization problem to adapt the two vectors,  $\omega_{\text{mbb}} = [\omega_{\text{mbb}}, 1 - \omega_{\text{mbb}}]^\top$  and  $\omega_{\text{rel}} = [\omega_{\text{rel}}, 1 - \omega_{\text{rel}}]^\top$ , each aligned to a service, by maximizing the aggregate throughput  $f(\omega_{\text{mbb}}, \omega_{\text{rel}}) = f_{\text{mbb}}(\omega_{\text{mbb}}, \omega_{\text{rel}}) + f_{\text{rel}}(\omega_{\text{mbb}}, \omega_{\text{rel}})$  subject to the reliability constraint  $g_{\text{rel}}(\omega_{\text{mbb}}, \omega_{\text{rel}}) \geq 0.9999$ .

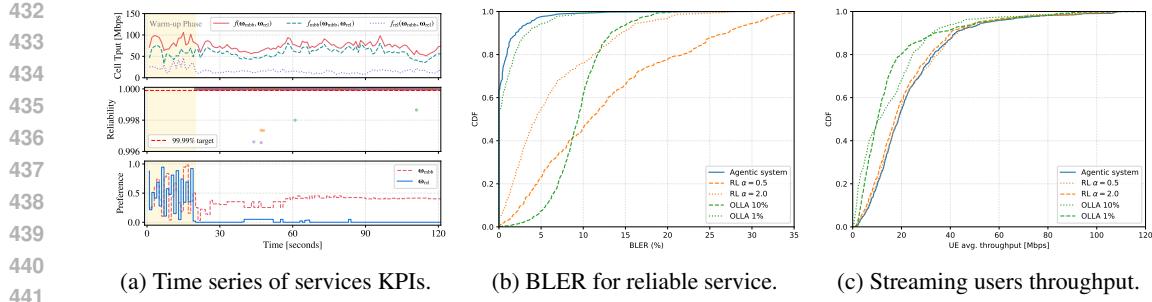


Figure 3: Validation of the intent fulfillment loop between optimizer-controller for two examples.

Figure 3a shows the optimizer–controller dynamics over a two-minute simulation. After an initial warm-up phase, the PAX-BO optimizer steers  $\omega_{\text{mbb}}$  and  $\omega_{\text{rel}}$  so that the D-EQL controller applies Pareto-optimal policies matched to each service’s requirements under varying network conditions. For reliable services, the optimizer converges to  $\omega_{\text{rel}} \approx 0$  (consistent with Figure 2b), driving the controller toward conservative MCS selections that deliver ultra-reliable performance throughout the simulation—exceeding 99.99% reliability in 94% of the run. Only a few packets are lost during isolated deep-fading episodes; under persistent fading, the interpreter could be invoked to relax the reliability target. For enhanced-streaming traffic, the optimizer converges to  $\omega_{\text{mbb}} \approx 0.45$ , prioritizing higher mean user throughput. Appendix H provides additional analysis and results.

Figure 3b and Figure 3c show that our agentic system outperforms both the state-of-the-art outer-loop link adaptation (OLLA) used in 5G systems and the traditional RL-based LA of Demirel et al. (2025). Unlike our approach—which adapts a single D-EQL model on-the-fly to different connectivity requirements and radio conditions—both OLLA and traditional RL require separate configurations optimized for each service type. For OLLA, we consider a standard target BLER of 10% for maximizing throughput in streaming services and 1% for highly reliable transmissions. Traditional RL similarly requires distinct models with reward functions tailored to each service; following Demirel et al. (2025), we use robustness parameters  $\alpha = 0.5$  for throughput and  $\alpha = 2$  for reliability. Figure 3b shows that our agentic system achieves substantially lower BLER for reliable services than both OLLA and the RL baseline with  $\alpha = 2$ , yielding more reliable transmissions. Figure 3c further shows that the same D-EQL model also attains throughput comparable to an RL model explicitly trained for throughput optimization. While D-EQL handles both services with a single model, using multiple RL models is impractical: inference must complete within a few hundred microseconds for all users, making rapid model switching across services infeasible.

### 7.3 TRIADIC AGENT WORKFLOW VALIDATION

We next evaluate the complete agentic AI system, with both intent management and intent fulfillment loops working in unison to provide a continuous solution to an intent formulation that combines a primary system objective (i.e., cell throughput) with QoS requirements of a connectivity service:

*Maximize cell throughput and serve streaming users with a minimum average data rate of 7 Mbps whenever possible.*

The peculiarity of this problem stems from the highly likelihood of the QoS requirements to become infeasible for users with poor channel conditions (such as cell-edge and high mobility users). When such an event occurs, persisting with a rigid QoS requirement would induce the system to over-provision users with poor channel regardless of their inability to meet the QoS goal, at the expense of users with a better channel quality. In turns, this may induce users with better channel to achieve lower throughput (due to less resources) and therefore compromise the primary intent objective.

Figure 4 compares the agentic AI system with two settings: (a) a formulation with rigid QoS requirements; and (b) a formulation with flexible QoS requirements. In the latter case, when the optimizer agent alerts the interpreter agent of a consistent violation of the service constraint, the interpreter reasons over the cause of the problem and plans a solution to relax the QoS requirements.

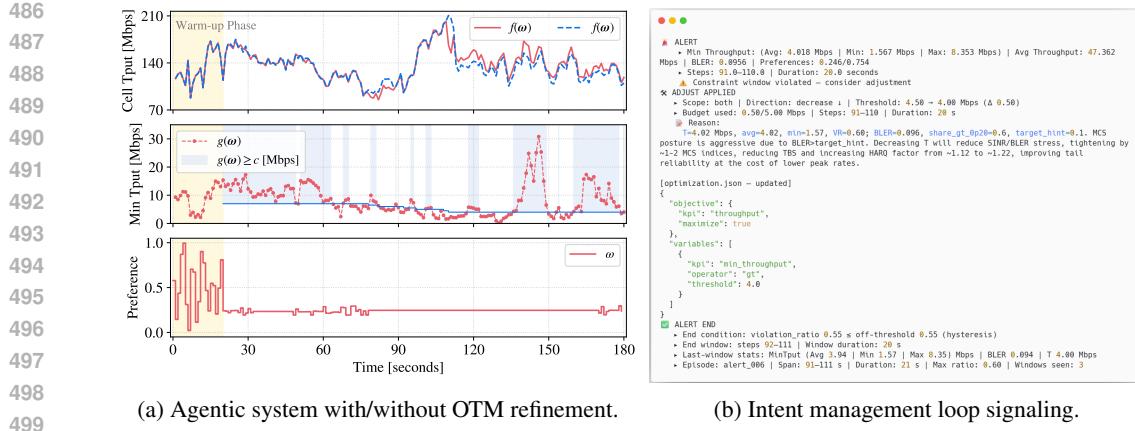


Figure 4: Validation of the full agentic workflow, with intent management loop and intent fulfillment loop working in unison. We compare two formulations with rigid and flexible service requirements.

Figure 4b shows an instance of this intent management loop between the interpreter-optimizer agents, where the latter reacts to the constraint violation by relaxing the service threshold, in an attempt to improve the primary objective, and providing a revised OTM. This choice allows the optimizer agent to choose an  $\omega$  setting that guides the MORL controller towards a less aggressive MCS selection policy for LA, making packets transmissions more reliable for users with poor channel conditions.

Despite the interpreter’s recursive adaptation of QoS requirements, infeasibility may still persist. This occurs because (a) the adaptor module includes guardrails that prevent abrupt QoS changes during OTM refinement (cf. Appendix A); and (b) prolonged poor channel conditions—such as deep fading, high pathloss, or shadowing—may yield spectral efficiencies too low to satisfy the QoS constraints, regardless of how the interpreter adjusts them. Nonetheless, adapting the OTM still yields tangible system-level benefits. By relaxing QoS targets for users in persistently poor channel conditions, the system frees radio resources that can be reallocated to users with better channel quality, thus with higher spectral efficiency. This redistribution increases the primary intent objective (cell throughput), even if some individual QoS constraints remain infeasible. As illustrated in Figure 4a, once OTM adaptation begins in the second half of the simulation, the cell throughput improves by a 4.79%.

## 8 CONCLUSIONS

We presented an Agentic AI system for intent-driven control in autonomous networks, structured around three cooperating agents: interpreter, optimizer, and controller. Their coordinated interaction links high-level service intents to concrete network actions, enabling continuous reasoning, trade-off resolution, and real-time adaptation across multiple timescales of autonomous network control.

Our contributions span the full intent-to-control pipeline. The interpreter uses a lightweight dual-SLM architecture to convert natural-language intents into structured optimization templates, assess feasibility, diagnose constraint violations, and refine templates using optimizer feedback. The optimizer performs preference planning via BO, dynamically adjusting the downstream controller’s policy to meet the service requirements encoded in the template. The controller builds on MORL to execute fast-timescale actions and adapt policies to evolving network conditions. To support this role, we introduce a distributed MORL algorithm that integrates envelope Q-learning with actor–learner decoupling, preference-space exploration, and prioritized hindsight replay, improving scalability, exploration coverage, and performance over state-of-the-art MORL approaches.

Proof-of-concept experiments in a high-fidelity, 5G-compliant RAN simulator demonstrate that the proposed system reconciles heterogeneous service requirements—including throughput and reliability—while operating near the Pareto front of network performance and adapting effectively to dynamic conditions, exceeding traditional RL and state-of-the-art functions of in 5G/5G-A systems.

Looking ahead, a key challenge is scaling this workflow across hierarchical layers of the RAN—from cell-level control to cluster-level coordination and end-to-end service orchestration—while ensuring intent consistency, agent interoperability, and robustness to uncertainty at each level.

540 9 LLM USAGE STATEMENT  
541542 In this paper, the authors used LLMs to check grammar, spelling, punctuation, and style compliance.  
543544 REFERENCES  
545546 3GPP. Technical report (TR) 22.870. technical specification group tsg sa; study on 6Guse cases  
547 and service requirements. stage 1 (release 20) v0.4.1 (2025-10). Technical report, 3rd Generation  
548 Partnership Project (3GPP), October 2025a.549 3GPP. Technical specification (TS) 23.501. System architecture for the 5G system (5GS), v19.4.0,  
550 June 2025b.  
551552 3GPP. Technical specification (TS) 28.312. Management and orchestration; intent driven management  
553 services for mobile networks, v19.2.1, June 2025c.554 3GPP. Technical specification (TS) 38.211. NR; Physical channels and modulation, v18.7.0, June  
555 2025d.557 3GPP. Technical specification (TS) 38.214. NR; Physical layer procedures for data, v19.0.0, June  
558 2025e.

559 3GPP. S2-2507223 [wt3, network ai agent] agentic core network for 6g, October 2025f.

561 3GPP. S5-250861 - Study on intent driven management services for mobile network phase 4. In  
562 *3GPP TSG SA Meeting 108*, June 2025g. URL [https://www.3gpp.org/ftp/tsg\\_sa/TSG\\_SA/TSGS\\_108\\_Prague\\_2025-06/Docs/SP-250861.zip](https://www.3gpp.org/ftp/tsg_sa/TSG_SA/TSGS_108_Prague_2025-06/Docs/SP-250861.zip).

564 3GPP. S5-255123 - Study on 6G management and orchestration, October 2025h.

566 3GPP. SA5 NWM discussion for rel-20 6G OAM work areas, August 2025i. URL  
567 [https://www.3gpp.org/ftp/Email\\_Discussions/SA5/OAM%20rapporteur%20calls/Rapporteur%20call%20%23161/SA5\\_NWM\\_Discussion\\_for\\_Rel-20\\_6G\\_OAM\\_Work\\_Areas-v0.0.7.pdf](https://www.3gpp.org/ftp/Email_Discussions/SA5/OAM%20rapporteur%20calls/Rapporteur%20call%20%23161/SA5_NWM_Discussion_for_Rel-20_6G_OAM_Work_Areas-v0.0.7.pdf).570 Abbas Abdolmaleki, Sandy H. Huang, Leonard Hasenclever, Michael Neunert, H. Francis Song,  
571 Martina Zambelli, Murilo F. Martins, Nicolas Heess, Raia Hadsell, and Martin Riedmiller. A  
572 distributional view on multi-objective policy optimization. In *Proceedings of the 37th International  
573 Conference on Machine Learning*, ICML'20. JMLR.org, 2020.575 Saaket Agashe, Jiuzhou Han, Shuyu Gan, Jiachen Yang, Ang Li, and Xin Eric Wang. Agent s:  
576 An open agentic framework that uses computers like a human. In *The Thirteenth International  
577 Conference on Learning Representations*, 2025. URL <https://openreview.net/forum?id=1IVRgt4nLv>.579 Leon Barrett and Srinivas Narayanan. Learning all optimal policies with multiple criteria. In *International  
580 Conference on Machine Learning*, 2008. URL <https://api.semanticscholar.org/CorpusID:14572936>.582 Toygun Basaklar, Suat Gumussoy, and Ümit Y. Ogras. PD-MORL: preference-driven multi-objective  
583 reinforcement learning algorithm. *International Conference on Learning Representations (ICLR)*,  
584 2023. doi: 10.48550/ARXIV.2208.07914. URL <https://doi.org/10.48550/arXiv.2208.07914>.587 Fransiscus Asisi Bimo, Maria Amparo Canaveras Galdon, Chun-Kai Lai, Ray-Guang Cheng, and  
588 Edwin K. P. Chong. Intent-based network for ran management with large language models, 2025.  
589 URL <https://arxiv.org/abs/2507.14230>.590 Burak Demirel, Yu Wang, Cristian Tatino, and Pablo Soldati. Generalization in reinforcement learning  
591 for radio access networks, 2025. URL <https://arxiv.org/abs/2507.06602>.593 Stefan Elfwing, Eiji Uchibe, and Kenji Doya. Sigmoid-weighted linear units for neural network  
function approximation in reinforcement learning. *Neural Networks*, 107:3–11, November 2018.

594 Taicheng Guo, Xiuying Chen, Yaqi Wang, Ruidi Chang, Shichao Pei, Nitesh V. Chawla, Olaf Wiest,  
 595 and Xiangliang Zhang. Large language model based multi-agents: A survey of progress and  
 596 challenges. In Kate Larson (ed.), *Proceedings of the Thirty-Third International Joint Conference*  
 597 *on Artificial Intelligence, IJCAI-24*, pp. 8048–8057. International Joint Conferences on Artificial  
 598 Intelligence Organization, 8 2024. doi: 10.24963/ijcai.2024/890. URL <https://doi.org/10.24963/ijcai.2024/890>.

600 601 Conor F. Hayes, Roxana Rădulescu, Eugenio Bargiacchi, Johan Källström, Matthew Macfarlane,  
 602 Mathieu Reymond, Timothy Verstraeten, Luisa M. Zintgraf, Richard Dazeley, Fredrik Heintz, Enda  
 603 Howley, Athirai A. Irissappane, Patrick Mannion, Ann Nowé, Gabriel Ramos, Marcello Restelli,  
 604 Peter Vamplew, and Diederik M. Roijers. A practical guide to multi-objective reinforcement  
 605 learning and planning. *Autonomous Agents and Multi-Agent Systems*, 36(1), April 2022. ISSN  
 606 1573-7454. doi: 10.1007/s10458-022-09552-y. URL <http://dx.doi.org/10.1007/s10458-022-09552-y>.

607 608 Dan Horgan, John Quan, David Budden, Gabriel Barth-Maron, Matteo Hessel, Hado van Hasselt, and  
 609 David Silver. Distributed prioritized experience replay. In *International Conference on Learning*  
 610 *Representations*, 2018. URL <https://openreview.net/forum?id=H1Dy---0Z>.

611 612 Laurie Hughes, Yogesh K. Dwivedi, Tegwen Malik, Mazen Shawosh, Mousa Ahmed Albashrawi,  
 613 Il Jeon, Vincent Dutot, Mandanna Appanderanda, Tom Crick, Rahul De', Mark Fenwick,  
 614 Senali Madugoda Gunaratnege, Paulius Jurcys, Arpan Kumar Kar, Nir Kshetri, Keyao Li, Sashah  
 615 Mutasa, Spyridon Samothrakis, Michael Wade, and Paul Walton. AI agents and agentic systems:  
 616 A multi-expert analysis. *Journal of Computer Information Systems*, 65(4):489–517, 2025. doi:  
 617 10.1080/08874417.2025.2483832. URL <https://doi.org/10.1080/08874417.2025.2483832>.

618 619 IETF. Ai agent protocols for 6g systems, October 2025. URL <https://datatracker.ietf.org/doc/draft-stephan-ai-agent-6g/>.

620 621 Maximilian Igl, Kamil Ciosek, Yingzhen Li, Sebastian Tschiatschek, Cheng Zhang, Sam Devlin,  
 622 and Katja Hofmann. Generalization in reinforcement learning with selective noise injection and  
 623 information bottleneck. In *Advances in Neural Information Processing Systems*, volume 32. Curran  
 624 Associates, Inc., 2019.

625 626 Intel and NEC. Our views on ai/ml management enhancements for 6g system. In *3GPP TSG SA5*  
 627 *6G workshop*, June 2025. URL [https://www.3gpp.org/ftp/Email\\_Discussions/SAC/SA5-SA5-level%20discussions/SA5\\_Workshop\\_on\\_6G\\_Rel20/Rel-20%206G%20workshop%20Our%20views%20on%20AIML%20Management%20enhancement%20for%206G%20system%20\(updated%20version\).pdf](https://www.3gpp.org/ftp/Email_Discussions/SAC/SA5-SA5-level%20discussions/SA5_Workshop_on_6G_Rel20/Rel-20%206G%20workshop%20Our%20views%20on%20AIML%20Management%20enhancement%20for%206G%20system%20(updated%20version).pdf).

628 629 Alexia Jolicoeur-Martineau. Less is more: Recursive reasoning with tiny networks, 2025. URL  
 630 <https://arxiv.org/abs/2510.04871>.

631 632 Daniel Kahneman. *Thinking, Fast and Slow*. Farrar, Straus and Giroux, 2011.

633 634 I. Y. Kim and O. L. de Weck. Adaptive weighted sum method for multiobjective optimization: a new  
 635 method for pareto front generation. *Structural and multidisciplinary optimization*, 31(2), 2005.

636 637 Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization, 2017. URL  
 638 <https://arxiv.org/abs/1412.6980>.

639 640 Abdullah Konak, David W. Coit, and Alice E. Smith. Multi-objective optimization using genetic  
 641 algorithms: A tutorial. *Reliability Engineering & System Safety*, 91(9):992–1007, 2006. ISSN 0951-  
 642 8320. doi: <https://doi.org/10.1016/j.ress.2005.11.018>. URL <https://www.sciencedirect.com/science/article/pii/S0951832005002012>. Special Issue - Genetic Algorithms  
 643 and Reliability.

644 645 Xinyi Li, Sai Wang, Siqi Zeng, Yu Wu, and Yi Yang. A survey on llm-based multi-agent  
 646 systems: Workflow, infrastructure, and challenges. *Vicinagearth*, 1(9), 2024. doi: 10.1007/s44336-024-00009-2. URL <https://link.springer.com/article/10.1007/s44336-024-00009-2>.

648 Chunming Liu, Xin Xu, and Dewen Hu. Multiobjective reinforcement learning: A comprehensive  
 649 overview. *IEEE Transactions on Systems, Man, and Cybernetics: Systems*, 45(3):385–398, 2015.  
 650 doi: 10.1109/TSMC.2014.2358639.

651

652 Zuxin Liu, Thai Quoc Hoang, Jianguo Zhang, Ming Zhu, Tian Lan, Shirley Kokane, Juntao Tan,  
 653 Weiran Yao, Zhiwei Liu, Yihao Feng, Rithesh R N, Liangwei Yang, Silvio Savarese, Juan Carlos  
 654 Niebles, Huan Wang, Shelby Heinecke, and Caiming Xiong. APIGen: Automated Pipeline for  
 655 generating verifiable and diverse function-calling datasets. In *The Thirty-eighth Conference on*  
 656 *Neural Information Processing Systems Datasets and Benchmarks Track*, 2024. URL <https://openreview.net/forum?id=Jfg3vw2bjx>.

657

658 Kurt Van Moffaert and Ann Nowé. Multi-objective reinforcement learning using sets of pareto  
 659 dominating policies. *Journal of Machine Learning Research*, 15:3483–3512, 2014. URL <https://jmlr.org/papers/volume15/vanmoffaert14a/vanmoffaert14a.pdf>.

660

661

662 Hossam Mossalam, Yannis M. Assael, Diederik M. Roijers, and Shimon Whiteson. Multi-objective  
 663 deep reinforcement learning. *CoRR*, abs/1610.02707, 2016. URL <http://arxiv.org/abs/1610.02707>.

664

665 Sriraam Natarajan and Prasad Tadepalli. Dynamic preferences in multi-criteria reinforcement learning.  
 666 In *Proceedings of the 22nd International Conference on Machine Learning*, ICML '05, pp. 601–608,  
 667 New York, NY, USA, 2005. Association for Computing Machinery. ISBN 1595931805. doi:  
 668 10.1145/1102351.1102427. URL <https://doi.org/10.1145/1102351.1102427>.

669

670 Thanh Thi Nguyen, Ngoc Duy Nguyen, and Saeid Nahavandi. A multi-objective deep reinforcement  
 671 learning framework. In *Proceedings of the 2018 International Joint Conference on Neural Networks*  
 672 (IJCNN), pp. 1–8, 2018. URL <https://arxiv.org/pdf/1803.02965.pdf>.

673

674 Klaus I. Pedersen, Guillaume Monghal, Istvan Z. Kovacs, Troels E. Kolding, Akhilesh Pokhariyal,  
 675 Frank Frederiksen, and Preben Mogensen. Frequency domain scheduling for OFDMA with limited  
 676 and noisy channel feedback. In *IEEE 66th Vehicular Technology Conference*, pp. 1792–1796,  
 677 2007.

678

679 Martin L. Puterman. *Markov Decision Processes: Discrete Stochastic Dynamic Programming*. Wiley,  
 680 New York, 1994.

681

682 Adnan Qayyum, Abdullatif Albaseer, Junaid Qadir, Ala Al-Fuqaha, and Mohamed Abdallah. LLM-  
 683 driven multi-agent architectures for intelligent self-organizing networks. *IEEE Networks*, 2025.

684

685 Chen Qian, Wei Liu, Hongzhang Liu, Nuo Chen, Yufan Dang, Jiahao Li, Cheng Yang, Weize  
 686 Chen, Yusheng Su, Xin Cong, Juyuan Xu, Dahai Li, Zhiyuan Liu, and Maosong Sun. Chatdev:  
 687 Communicative agents for software development, 2024. URL <https://arxiv.org/abs/2307.07924>.

688

689 Qwen, :, An Yang, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chengyuan  
 690 Li, Dayiheng Liu, Fei Huang, Haoran Wei, Huan Lin, Jian Yang, Jianhong Tu, Jianwei Zhang,  
 691 Jianxin Yang, Jiaxi Yang, Jingren Zhou, Junyang Lin, Kai Dang, Keming Lu, Keqin Bao, Kexin  
 692 Yang, Le Yu, Mei Li, Mingfeng Xue, Pei Zhang, Qin Zhu, Rui Men, Runji Lin, Tianhao Li, Tianyi  
 693 Tang, Tingyu Xia, Xingzhang Ren, Xuancheng Ren, Yang Fan, Yang Su, Yichang Zhang, Yu Wan,  
 694 Yuqiong Liu, Zeyu Cui, Zhenru Zhang, and Zihan Qiu. Qwen2.5 technical report, 2025. URL  
<https://arxiv.org/abs/2412.15115>.

695

696 Diederik M. Roijers, Peter Vamplew, Shimon Whiteson, and Richard Dazeley. A survey of multi-  
 697 objective sequential decision-making. *Journal of Artificial Intelligence Research*, 48(1):67–113,  
 698 October 2013. URL <https://arxiv.org/abs/1402.0590>.

699

700 Diederik M. Roijers, Peter Vamplew, Shimon Whiteson, and Richard Dazeley. Computing convex  
 701 coverage sets for faster multi-objective coordination. *Journal of Artificial Intelligence Research*, 52:  
 702 399–443, 2015. URL <https://www.jair.org/index.php/jair/article/view/10933>.

702 Ranjan Sapkota, Konstantinos I. Roumeliotis, and Manoj Karkee. AI agents vs. agentic AI: A  
 703 conceptual taxonomy, applications and challenges. *Information Fusion*, 126:103599, feb 2025.  
 704 ISSN 1566-2535. doi: 10.1016/j.inffus.2025.103599. URL <http://dx.doi.org/10.1016/j.inffus.2025.103599>.

705

706 Noah Shinn, Federico Cassano, Edward Berman, Ashwin Gopinath, Karthik Narasimhan, and  
 707 Shunyu Yao. Reflexion: Language agents with verbal reinforcement learning, 2023. URL  
 708 <https://arxiv.org/abs/2303.11366>.

709

710 Richard S. Sutton and Andrew G. Barto. *Reinforcement Learning: An Introduction*. MIT Press, 2nd  
 711 edition, 2018. URL <https://incompleteideas.net/book/the-book-2nd.html>.

712

713 Wei Tao, Yucheng Zhou, Yanlin Wang, Wenqiang Zhang, Hongyu Zhang, and Yu Cheng. MAGIS:  
 714 LLM-based multi-agent framework for github issue resolution. In *The Thirty-eighth Annual  
 715 Conference on Neural Information Processing Systems*, 2024. URL <https://openreview.net/forum?id=qevq3FZ63J>.

716

717 TMForum. Ig1251x autonomous networks – reference architecture v1.0.0, 2021.

718

719 TMForum. Tr292a intent management elements v3.6.0, 2024.

720

721 Guanzhi Wang, Yuqi Xie, Yunfan Jiang, Ajay Mandlekar, Chaowei Xiao, Yuke Zhu, Linxi  
 722 Fan, and Anima Anandkumar. Voyager: An open-ended embodied agent with large lan-  
 723 guage models. *Transactions on Machine Learning Research*, 2024. ISSN 2835-8856. URL  
 724 <https://openreview.net/forum?id=ehfRif0R3a>.

725

726 Qingyun Wu, Gagan Bansal, Jieyu Zhang, Yiran Wu, Beibin Li, Erkang Zhu, Li Jiang, Xiaoyun  
 727 Zhang, Shaokun Zhang, Jiale Liu, Ahmed Hassan Awadallah, Ryen W White, Doug Burger, and  
 728 Chi Wang. Autogen: Enabling next-gen LLM applications via multi-agent conversations. In *First  
 729 Conference on Language Modeling*, 2024. URL <https://openreview.net/forum?id=BAakY1hNKS>.

730

731 Jie Xu, Yunsheng Tian, Pingchuan Ma, Daniela Rus, Shinjiro Sueda, and Wojciech Matusik.  
 732 Prediction-guided multi-objective reinforcement learning for continuous robot control. In *Proceed-  
 733 ings of the 37th International Conference on Machine Learning*, ICML’20. JMLR.org, 2020.

734

735 Runzhe Yang, Xingyuan Sun, and Karthik Narasimhan. A generalized algorithm for multi-objective  
 736 reinforcement learning and policy adaptation. In *Advances in Neural Information Processing  
 737 Systems*, volume 32. Curran Associates, Inc., 2019.

738

739 Shunyu Yao, Jeffrey Zhao, Dian Yu, Nan Du, Izhak Shafran, Karthik Narasimhan, and Yuan Cao.  
 740 React: Synergizing reasoning and acting in language models. In *International Conference on  
 741 Learning Representations (ICLR)*, 2023. URL <https://arxiv.org/abs/2210.03629>.

742

743 Min Yoon, Yeboon Yun, and Hirotaka Nakayama. Sequential approximate multiobjective op-  
 744 timization using computational intelligence. In *Vector Optimization*, 2009. URL <https://api.semanticscholar.org/CorpusID:2036814>.

745

746 Dawei Zhan and Huanlai Xing. Expected improvement for expensive optimization: a review. *J.  
 747 of Global Optimization*, 78(3):507–544, November 2020. ISSN 0925-5001. doi: 10.1007/s10898-020-00923-x. URL <https://doi.org/10.1007/s10898-020-00923-x>.

748

749 Jiayu Zhao, Renyu Yang, Shenghao Qiu, and Zheng Wang. Unleashing the potential of acquisition  
 750 functions in high-dimensional bayesian optimization, 2024. URL <https://arxiv.org/abs/2302.08298>.

751

752 ZTE. R3-256538 - initial discussion on airan for 6g network. In *3GPP TSG-RAN WG3 Meeting  
 753 129bis*, October 2025. URL [https://www.3gpp.org/ftp/TSG\\_RAN/WG3\\_Iu/TSGR3\\_129-bis/Docs/R3-256538.zip](https://www.3gpp.org/ftp/TSG_RAN/WG3_Iu/TSGR3_129-bis/Docs/R3-256538.zip).

754

755

756	CONTENTS OF APPENDIX	
757		
758		
759	<b>1 Introduction</b>	<b>1</b>
760		
761	<b>2 Related Work</b>	<b>2</b>
762		
763	<b>3 Agentic AI System for RAN Control</b>	<b>3</b>
764	3.1 Timescales Separation . . . . .	4
765		
766		
767	<b>4 Language-Guided Intent Management</b>	<b>4</b>
768	4.1 Interpreter Agent . . . . .	4
769	4.2 Optimizer Agent . . . . .	5
770	4.2.1 PAX-BO: Preference-Aligned eXploration Bayesian Optimization . . . . .	6
771		
772		
773	<b>5 Preference-Guided Intent Fulfillment</b>	<b>6</b>
774	5.1 Controller Agent . . . . .	6
775		
776		
777	<b>6 Case Study: Agentic Radio Resource Management</b>	<b>7</b>
778		
779	<b>7 Experiment</b>	<b>8</b>
780	7.1 MORL Controller Agent for Link Adaptation . . . . .	8
781	7.2 Intent-fulfillment loop validation . . . . .	8
782		
783	7.3 Triadic Agent Workflow Validation . . . . .	9
784		
785		
786	<b>8 Conclusions</b>	<b>10</b>
787		
788	<b>9 LLM Usage Statement</b>	<b>11</b>
789		
790	<b>A Interpreter Agent: Responsibilities, Design, Implementation</b>	<b>18</b>
791		
792	A.1 Scope and Responsibilities . . . . .	18
793	A.2 Architectural Overview . . . . .	18
794	A.2.1 Translator . . . . .	18
795	A.2.2 Sliding-Window Monitor . . . . .	19
796	A.2.3 Advisor (Advisory Layer) . . . . .	20
797	A.2.4 Adaptor (Magnitude, Safety, Persistence) . . . . .	21
798		
799	A.3 Algorithmic Summary and Interfaces . . . . .	21
800	A.3.1 Interfaces . . . . .	21
801		
802	A.4 Models . . . . .	22
803		
804	A.5 Stability, Safety, and Complexity . . . . .	22
805		
806	A.6 Failure Modes and Mitigations . . . . .	22
807		
808	<b>B Optimization Template Model</b>	<b>24</b>
809	B.1 OTM Schema and Domain Semantics . . . . .	24

810	B.2 Example of OTM adaptation. . . . .	26
811		
812	<b>C Translator SLM Fine-Tuning</b>	<b>28</b>
813		
814	C.1 Dataset Curation . . . . .	28
815		
816	C.2 Training Methodology . . . . .	29
817		
818	C.3 Comparative Evaluation . . . . .	30
819		
820	<b>D Optimizer Agent Design</b>	<b>32</b>
821		
822	D.1 Bayesian Optimization . . . . .	32
823		
824	D.1.1 Gaussian Process Priors . . . . .	32
825		
826	D.1.2 Acquisition Functions . . . . .	32
827	D.2 PAX-BO: Preference-Aligned eXploration Bayesian Optimization . . . . .	33
828		
829	<b>E Multi-Objective Reinforcement Learning</b>	<b>36</b>
830		
831	E.1 Multi-Objective Markov Decision Process . . . . .	36
832		
833	E.2 Convex Coverage Set . . . . .	37
834		
835	E.3 Envelope Q-Learning . . . . .	37
836		
837	<b>F Distributed Envelope Q-Learning</b>	<b>40</b>
838		
839	F.1 D-EQL Architecture . . . . .	40
840		
841	F.2 Distributed Actors . . . . .	41
842		
843	F.3 Centralized Learner . . . . .	42
844		
845	F.4 Stratified Sampling on the Probability Simplex . . . . .	44
846		
847	F.4.1 Deterministic Equal-Volume Strata via a Simplex Lattice . . . . .	44
848		
849	F.4.2 Uniform Sampling Within a Stratum . . . . .	45
850		
851	F.4.3 Assigning Strata to Actors . . . . .	45
852		
853	F.4.4 Discussion and alternatives . . . . .	45
854	F.5 Experiments . . . . .	45
855		
856	F.5.1 Environments and Setup . . . . .	45
857		
858	F.5.2 Metrics and Results . . . . .	46
859		
860	F.5.3 Hyperparameters . . . . .	47
861		
862	<b>G Case Study</b>	<b>48</b>
863		
864	G.1 Link adaptation . . . . .	48
865		
866	G.2 MOMDP Design for Link Adaptation . . . . .	48
867		
868	G.3 Action space . . . . .	49
869		
870	G.4 Reward vector and preference space . . . . .	49
871		
872	G.5 State Design . . . . .	50
873		
874	<b>H Extended Experimental Evaluation</b>	<b>51</b>
875		
876	H.1 Network Simulator Environment . . . . .	51
877		

864	H.2 Training Setup . . . . .	51
865		
866	H.3 Testing the MORL LA Controller Agent . . . . .	52
867	H.3.1 Single Connectivity Service . . . . .	52
868	H.3.2 Multi Connectivity Services with QoS Differentiation . . . . .	54
869		
870	H.4 Online Preference Optimization . . . . .	60
871		
872	<b>I Compute resources and Hyperparameters</b>	<b>61</b>
873		
874		
875		
876		
877		
878		
879		
880		
881		
882		
883		
884		
885		
886		
887		
888		
889		
890		
891		
892		
893		
894		
895		
896		
897		
898		
899		
900		
901		
902		
903		
904		
905		
906		
907		
908		
909		
910		
911		
912		
913		
914		
915		
916		
917		

918 A INTERPRETER AGENT: RESPONSIBILITIES, DESIGN, IMPLEMENTATION  
919920 A.1 SCOPE AND RESPONSIBILITIES  
921922 The *interpreter agent* is the gateway from high-level intent to optimization-ready control. It fulfills  
923 two primary responsibilities: (1) translating intents expressed in natural language into an initial  
924 structured OTM; and (2) recursively reasoning over system observations and optimizer feedback  
925 to stabilize intent fulfillment by revising the OTM when required (e.g., when constraints become  
926 infeasible).927 **Division of Labour (Dual-SLM).** To address these responsibilities under tight computational  
928 budgets, we employ two complementary SLMs: (1) a **fine-tuned SLM** for intent-to-OTM translation;  
929 and (2) an **in-context learning (ICL) based SLM** for adaptive intent management, which reasons  
930 over structured prompts and windowed KPI statistics, refines intent requirements when needed, and  
931 provides an explicit textual rationale.932 While alternative realizations of an interpreter agent are possible, our design enables the use of  
933 lightweight SLMs that adhere to the compute and memory constraints of 4G/5G RAN systems  
934 (see Appendix A.5).936 A.2 ARCHITECTURAL OVERVIEW  
937

938 The interpreter agent architecture, showed in Figure 5, consists of four tightly coupled modules:

939  
940 • **Translator (Appendix A.2.1)** uses a fine-tuned SLM to convert an incoming intent into a  
941 structured, machine-readable OTM that specifies objectives, constraints, aggregation units, and  
942 provenance for different connectivity services and operational goals.  
943 • **Monitor (Appendix A.2.2)** subscribes to optimizer telemetry, aligns the telemetry stream to the  
944 OTM-defined intent-management timescale, extracts per-window summaries, and bridges short  
945 gaps.  
946 • **ICL-based Advisor (Appendix A.2.3)** uses an ICL-based SLM to reason over window sum-  
947 maries and active policy thresholds, selects an advisory direction  
948

949 
$$a \in \{\text{increase, decrease, no\_change}\}$$
  
950

951 and generates a compact rationale  $\mathcal{R}$  grounded in RRM. It proposes only a direction, not a  
952 magnitude.953 • **Adaptor (Appendix A.2.4)** converts the advisory action  $a$  into a bounded threshold update  
954  $\Delta b$  under guardrails (e.g., caps, lifetime budget, floor/ceiling, cooldown), persists the updated  
955 threshold atomically into the OTM, and emits an audit record.956 During the intent-management loop, the OTM is treated as a *living document* jointly maintained  
957 by the interpreter and optimizer agents. The optimizer continuously solves against the current  
958 OTM snapshot and reports telemetry (e.g., windowed KPI statistics) to the monitor. Guided by this  
959 feedback, the ICL-based advisor recommends adjustments when intent requirements become overly  
960 tight or infeasible under the current network state. The adaptor then applies bounded updates to the  
961 corresponding OTM constraints, yielding a refreshed OTM for the optimizer.962  
963 A.2.1 TRANSLATOR  
964965 The translator employs a fine-tuned SLM to convert intents into deterministic, schema-compliant  
966 OTM instances. Its role extends well beyond lexical parsing: it must interpret natural-language  
967 intents into meaningful optimization structures grounded in domain knowledge, and identify the  
968 appropriate downstream control agent to execute them. For example, a service intent requesting high  
969 reliability—such as the case in Section 7.2—may translate into a non-obvious constraint formulated  
970 in terms of BLER.971 Further details on the translator design are provided in Appendix B and Appendix C, which discuss  
972 the OTM schema and the supervised fine-tuning and evaluation of the translator SLM, respectively.

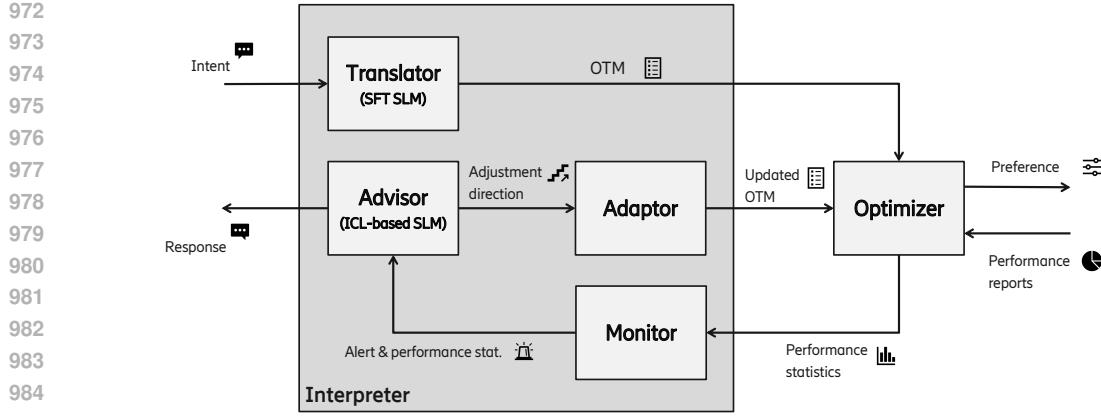


Figure 5: **Dual-SLM interpreter agent.** A (supervised) fine-tuned SLM generates the OTM; an intent monitor aligns telemetry; an ICL-based advisory module outputs discrete adjustment directions with rationale; and an adaptor applies bounded updates and persists them atomically. The optimizer then solves against the latest OTM snapshot, with telemetry closing the loop.

#### A.2.2 SLIDING-WINDOW MONITOR

Consider a single constraint  $k^A = \langle \text{kpi}, \odot, b, A, \text{unit} \rangle$  of an OTM, where  $\odot \in \{\leq, \geq\}$  and  $A$  is the per-step aggregation operator declared in the OTM (e.g., mean, min, max, p95). To simplify notation, we refer to the constraint function  $k^A(\cdot)$  as  $y(\cdot)$ , and let  $y_t$  denote the KPI value at step  $t$  after applying  $A$  over the telemetry bin of length  $\Delta$  (e.g., 10 s). With window length  $W$ , the monitor maintains a ring buffer over  $\{y_i\}_{i=t-W+1}^t$  and computes a *signed margin*:

$$m_i = s(y)(y_i - b) \quad \text{where} \quad s(y) = \begin{cases} +1, & \odot \in \{\geq\} \text{ (lower bound)} \\ -1, & \odot \in \{\leq\} \text{ (upper bound).} \end{cases}$$

A step  $i$  is a violation if and only if  $m_i < 0$  (negative margin). The window statistics are then

$$\text{violation\_ratio}(t) = \frac{1}{W} \sum_{i=t-W+1}^t \mathbf{1}[m_i < 0],$$

$$\bar{y} = \frac{1}{W} \sum_{i=t-W+1}^t y_i, \quad y_{\min} = \min_i y_i, \quad y_{\max} = \max_i y_i,$$

and the average *shortfall/slack* (useful for controllers and prompts):

$$\text{shortfall\_avg} = \frac{1}{W} \sum_{i=t-W+1}^t \max\{0, -m_i\},$$

$$\text{slack\_avg} = \frac{1}{W} \sum_{i=t-W+1}^t \max\{0, m_i\}.$$

**Hysteresis and alerting.** Hysteresis prevents chattering: An ALERT\_START event is declared when  $VR > \rho_{\text{on}}$  and an ALERT\_END event when  $VR < \rho_{\text{off}}$  with  $\rho_{\text{on}} > \rho_{\text{off}}$ . At each window end (the decision point), if an alert is active, the monitor produces a compact, constraint-centric context:

```

1015
1016
1017
1018
1019
1020
1021
1022
1023
1024
1025

```

$$\left\{ \begin{array}{l} \text{window} : \{W, t-W+1 \dots t, VR\}, \\ \text{constraint\_metric} : \{\bar{y}, y_{\min}, y_{\max}, b, \text{shortfall\_avg}, \text{slack\_avg}, \text{unit}\}, \\ \text{constraint\_id} : \text{id}, \end{array} \right.$$

Optionally, the context may be augmented with domain-specific auxiliaries (e.g., aux\_kpis) if available.

Table 3: Key hyperparameters of the interpreter agent.

Symbol	Name	Description
$W$	Window size	Number of samples used to compute moving averages and the violation ratio.
$b$	Threshold	Current target value for the monitored KPI.
$\rho_{\text{on}}$	Alert-on ratio	Violation ratio above which an alert episode is initiated.
$\rho_{\text{off}}$	Alert-off ratio	Violation ratio below which an alert episode is terminated.
$d$	Step size	Base increment or decrement applied to threshold updates.
$g_{\uparrow}, g_{\downarrow}$	Guardrail gains	Maximum upward or downward adjustment permitted per update.
$s_{\text{max}}$	Smoothing cap	Maximum smoothing applied across consecutive updates.
$B$	Budget	Maximum number of updates allowed within a single alert episode.
$b_{\min}, b_{\max}$	Bounds	Minimum and maximum permissible threshold values.
$C$	Cooldown	Minimum number of steps that must elapse before another update can be applied.

**Complexity.** The monitor executes in  $O(1)$  time per step through the use of a fixed-size ring buffer and incremental summary updates, with no rescans required. Memory usage grows linearly with the window size, i.e.,  $O(W)$ .

### A.2.3 ADVISOR (ADVISORY LAYER)

The advisory layer determines the *direction of adaptation* and supplies a textual justification  $\mathcal{R}$ . It does *not* specify the magnitude of change. Two modes are supported:

- (a) **Rule-based.** Thresholds on summary statistics (e.g., violation ratio, mean deviation from the target, minimum deviation from the target, auxiliary posture indicators) determine an advisory action  $a$ .
- (b) **ICL-based SLM** A structured prompt encodes (i) the set of allowed actions, (ii) the decision policy, (iii) domain-specific guardrails, and (iv) a strict JSON output schema. The SLM produces an advisory adjustment

`{"action": "...", "justification": "..."}`

conditioned on the parsed telemetry payload from the intent monitor.

**Guardrails in Prompting.** Schema fidelity and reproducibility are enforced through:

- (a) JSON-only outputs;
- (b) end-of-sentence token fences;
- (c) banned tokens (e.g., URLs, markdown code fences); and
- (d) near-deterministic decoding with low-variance sampling to avoid verbatim repetition while maintaining stability.

The justification must cite explicit numerical values extracted from the payload (e.g., target  $b$ , mean  $\bar{y}$ , minimum  $y_{\min}$ , violation ratio VR) and must classify posture relative to a domain-specific auxiliary metric (e.g., “aggressive” vs. “conservative”).

**Prompt Contract (Abridged).** Allowed actions are `{increase, decrease, no_change}`. The required output format is strictly JSON:

`{"action": "...", "justification": "..."}`

The justification must reference the relevant statistics and the auxiliary posture label. Domain-specific instantiations (e.g., using BLER as the auxiliary metric) appear in examples in Section 7.

1080 A.2.4 ADAPTOR (MAGNITUDE, SAFETY, PERSISTENCE)  
10811082 Given an advisory action  $a \in \{\text{increase, decrease, no\_change}\}$ , the adaptor computes a  
1083 candidate step size  $\Delta b$  using a deadband  $d$  and asymmetric gains ( $g_{\uparrow}, g_{\downarrow}$ ):  
1084

1085 
$$\Delta b = \begin{cases} g_{\downarrow} \max(0, (b - \bar{x}) - d), & a = \text{decrease}, \\ g_{\uparrow} \max(0, (\bar{x} - b) - d), & a = \text{increase}, \\ 0, & a = \text{no\_change}. \end{cases}$$
  
1086

1087 Safety guardrails limit the actuation:  
1088

1089 
$$\Delta b \leftarrow \min\{\Delta b, s_{\max}, B_{\text{left}}, b - b_{\min}, b_{\max} - b\}, \quad b \leftarrow \text{clip}(b \pm \Delta b, b_{\min}, b_{\max}).$$
  
1090

1091 Budgets and cooldown counters are updated after each actuation. Final thresholds are written  
1092 atomically to the OTM, ensuring that the optimizer and monitor operate on consistent snapshots.  
1093 Section 7 illustrates with concrete examples (e.g., throughput maximization with minimum guarantees  
1094 per user, or bounds on BLER) how this generic mechanism applies across KPIs.  
10951096 A.3 ALGORITHMIC SUMMARY AND INTERFACES  
10971098 The closed-loop operation of the interpreter agent—integrating monitoring, advisory, and adaptation—  
1099 is summarized in Algorithm 1. The procedure shows how the agent detects constraint violations,  
1100 issues advisory actions, and applies bounded adaptations under guardrails.  
11011102 **Algorithm 1** Interpreter Agent (Monitor  $\rightarrow$  Advisor  $\rightarrow$  Adaptor)  
1103

```

1: Input: window size  $W$ ; thresholds  $(b, \rho_{\text{on}}, \rho_{\text{off}})$ ; guardrails  $(d, g_{\uparrow}, g_{\downarrow}, s_{\max}, B, b_{\min}, b_{\max}, C)$ 
2: for each step  $t$  do
3:   Push observation  $y_t$  into ring buffer; update  $(\bar{y}, y_{\min}, \text{VR})$ 
4:   if  $\text{VR} > \rho_{\text{on}}$  and not in alert then
5:     Start episode; reset budget and cooldown
6:   end if
7:   if in alert then
8:     Build parsed telemetry payload; select action  $a$  via rules or ICL SLM; log rationale  $\mathcal{R}$ 
9:     if  $a \neq \text{no\_change}$  and cooldown expired and  $B_{\text{left}} > 0$  then
10:      Compute  $\Delta b$ ; apply guardrails; update  $b$ ; persist OTM; decrement budget; reset
11:        cooldown  $C$ 
12:      end if
13:      if  $\text{VR} < \rho_{\text{off}}$  then
14:        End episode; log summary
15:      end if
16:   end if
17: end for

```

1119 A.3.1 INTERFACES  
11201121 **(i) From Monitor to Advisor.** Upon receiving telemetry from the optimizer, the intent monitor  
1122 produces a compact summary aligned to the OTM timescale. This parsed payload becomes the sole  
1123 input to the ICL-based advisory module. An example summary from our experiments is:  
1124

```

1: {
2:   "window": {
3:     "start": 1020, "end": 1139, "W": 12, "violation_ratio": 0.60
4:   },
5:   "constraint_metric": {
6:     "name": "throughput",
7:     "avg": 6.92,
8:     "min": 3.08,
9:     "monitor_threshold": 7.00,
10:    "unit": "Mbps"
11:  },
12:   "radio_kpis": {"bler": {"avg": 0.14, "target_hint": 0.10}}
13: }

```

1134 **(ii) From Advisor to Adaptor.** The advisory module returns only an *adjustment direction* along  
 1135 with a textual justification, both constrained by the *current OTM* used by the optimizer. It never  
 1136 proposes numeric magnitudes. Example output:

```
1138 {
1139   "action": "decrease",
1140   "justification": "relax to reduce MCS pressure and HARQ overhead."
1141 }
```

1142 **(iii) From Adaptor to OTM (atomic).** The adaptor converts the advisory direction into a bounded  
 1143 step  $\Delta b$ , applies guardrails (e.g., clipping to  $[b_{\min}, b_{\max}]$ ), persists the updated threshold atomically,  
 1144 and records the rationale:

```
1146 {
1147   "kpi": "throughput",
1148   "aggregation": "min"
1149   "old_threshold": 7.00,
1150   "new_threshold": 6.92,
1151   "delta": -0.08,
1152   "episode": "alert_002",
1153   "rationale": "VR=0.60; BLER aggressive"
1154 }
```

#### A.4 MODELS

1155 **Fine-tuned SLM (Intent-to-OTM).** A domain-specialized causal SLM is fine-tuned to generate  
 1156 OTM JSON directly from natural-language intents. Training uses instruction-style pairs of the form  
 1157 (intent, OTM) that adhere to domain schemas (objective, KPI, operator, threshold). The model is  
 1158 evaluated using exact-match accuracy and schema validity. This component is implemented using the  
 1159 Qwen-2.5-7B-Instruct model (Qwen et al., 2025) with supervised fine-tuning; additional details are  
 1160 provided in Appendix C.

1161 **ICL-based SLM (Constraint Adaptation).** A general-purpose SLM—also based on Qwen-2.5-  
 1162 7B-Instruct (Qwen et al., 2025) but without task-specific weight updates—is prompted with: (i)  
 1163 the allowed actions and guardrails, (ii) policy rules governing the violation ratio (VR) and KPI  
 1164 slack/shortfall, (iii) BLER posture rules with target hints, and (iv) a strict JSON schema. Outputs are  
 1165 assessed for schema validity, internal consistency (e.g., adherence to policy rules), and justification  
 1166 quality.

#### A.5 STABILITY, SAFETY, AND COMPLEXITY

1168 Guardrails constrain actuation by ensuring that the target parameter  $b$  remains within the safe interval  
 1169  $[b_{\min}, b_{\max}]$ . A hysteresis mechanism further prevents rapid oscillations caused by frequent threshold  
 1170 updates. The computational overhead of the method is minimal: each control step requires constant  
 1171 time  $O(1)$ , and memory usage grows linearly with the window size  $O(W)$ . This design minimizes  
 1172 the impact on RAN compute and memory resources.

1173 To evaluate the practical performance of the agentic AI system for intent management, we report the  
 1174 following metrics: (i) reduction in violation ratio relative to baseline operation; (ii) percentage of  
 1175 observation windows that request a change; (iii) percentage of updates clipped by guardrails; (iv)  
 1176 validity rate of JSON payloads against the schema; (v) observed episode lengths; and (vi) adaptation  
 1177 latency per update.

#### A.6 FAILURE MODES AND MITIGATIONS

1181 Despite these safeguards, the system remains susceptible to several failure modes. The corresponding  
 1182 mitigation strategies are:

- 1183 **Prompt sensitivity:** Malformed or ambiguous payloads may arise from language model outputs.  
 1184 This risk is mitigated through strict schema enforcement, exclusion of unsafe tokens, and  
 1185 regression testing on canonical telemetry payloads.

1188

- **Distribution shift:** Variations in traffic or channel conditions can create discrepancies between

1189 training and deployment distributions. The system addresses this through window normalization

1190 and by providing BLER posture hints to the model. In extreme cases, the controller can revert to

1191 a rules-only mode to preserve stability.

1192

- **Over-actuation:** Excessive threshold adjustments may cause oscillations or instability. To

1193 prevent this, the system enforces lifetime update budgets, per-step update caps, cooldown

1194 intervals, and explicit floor/ceiling bounds on  $b$ .

1195

- **Explainability drift:** Generated rationales may deviate from the underlying numerical evidence.

1196 The advisory module  $\mathcal{R}$  must cite explicit numerical values, and all rationale cards are logged

1197 and checked against policy expectations to ensure traceability and consistency.

1198

1199 This appendix outlines how the interpreter agent determines *when to act, how to act and why, and to*

1200 *what extent to act*. These behaviors are realized through dual SLMs, classical control guardrails, and

1201 auditable OTM persistence.

1202  
1203  
1204  
1205  
1206  
1207  
1208  
1209  
1210  
1211  
1212  
1213  
1214  
1215  
1216  
1217  
1218  
1219  
1220  
1221  
1222  
1223  
1224  
1225  
1226  
1227  
1228  
1229  
1230  
1231  
1232  
1233  
1234  
1235  
1236  
1237  
1238  
1239  
1240  
1241

## 1242 B OPTIMIZATION TEMPLATE MODEL

1244 **Purpose.** The OTM defines the contract between the interpreter agent and the downstream optimizer.  
 1245 It (i) specifies the optimization *objective* and the associated *constraints*, including explicit units and  
 1246 aggregation semantics; (ii) records provenance for auditability (*origin*, *modified\_by*); and (iii)  
 1247 serves as a *living document* that can be safely updated by the adaptor during execution.

1249 **Formal view.** Let  $\mathcal{X}$  denote the optimizer’s decision space, and let  $k(\cdot)$  be a network KPI evaluated  
 1250 under an aggregation operator  $A$  (e.g., mean, min,  $p95$ ). We define an OTM instance as

$$1252 \max_{x \in \mathcal{X}} k_{\text{obj}}^{A_{\text{obj}}}(x) \quad \text{s.t. } \forall i \in \{1, \dots, m\} : \begin{cases} k_i^{A_i}(x) \leq b_i & \text{if operator} \in \{\text{lt, le}\} \\ k_i^{A_i}(x) \geq b_i & \text{if operator} \in \{\text{gt, ge}\} \end{cases} \quad (3)$$

1254 where each constraint  $i$  specifies *service*, *kpi*, *operator*, threshold  $b_i$ , aggregation  $A_i$ , units,  
 1255 and scope. In essence, this formulation revisits the optimization (1) by rewriting the objective  $f(\cdot)$   
 1256 and the constraints  $g_i(\cdot)$  in terms of a more generic KPI construct  $k(\cdot)$  used in the OTM schema.

### 1258 B.1 OTM SCHEMA AND DOMAIN SEMANTICS

1260 The OTM schema is a minimal versioned JSON contract comprising four blocks, *objective*,  
 1261 *constraints*, and *metadata*, *version*, characterizing the OTM formalism in equation 3.

1263 Listing 1: Generic OTM schema applicable to different RAN control problems.

```

1264 {
1265   "objective": {
1266     "service": <service_name>,           // {"mbb", "urllc", "gaming", "streaming", "slice", ...},
1267     "kpi": <kpi_name>,                 // {"throughput", "reliability", "latency", ...}
1268     "scope": <scope_name>,             // {"per_user", "per_cell", "per_slice", ...}
1269     "aggregation": <aggr_name>,       // {"mean", "min", "max", "p95", "sum", ...}
1270     "unit": <unit_name>,               // {"Mbps", "Gbps", "ms", "s", "%", ...}
1271     "maximize": <value>,               // boolean value {true, false}
1272   },
1273   "constraints": [
1274     {
1275       "id": <value>,                  // string value (e.g., "C1", "C2")
1276       "service": <service_name>,       // {"mbb", "urllc", "gaming", "streaming", ...},
1277       "kpi": <kpi_name>,             // {"throughput", "reliability", "latency", ...}
1278       "scope": <scope_name>,           // {"per_user", "per_cell", "per_slice", ...}
1279       "aggregation": <aggr_name>,       // {"mean", "min", "max", "p95", "sum", ...}
1280       "unit": <unit_name>,             // {"Mbps", "Gbps", "ms", "s", "%", ...}
1281       "operator": <operator_type>,       // {"lt", "le", "ge", "gt"}
1282       "threshold": <value>,             // float expressed in "unit"
1283       "modified": <value>,               // boolean value {true, false}
1284     },
1285     {
1286       ...
1287     }
1288   ],
1289   "metadata": {
1290     "otm": {
1291       "id": <value>,                  // string value (e.g., "O1", "O2")
1292       "created_by": <model_id>,       // string value {"SFT_LLM", ...}
1293       "timestamp": <value>,             // formatted as iso-8601
1294       "timescale": <value>,             // string with window value (e.g., "10s_window")
1295     }
1296     "episode": {
1297       "id": <value>,                  // string value (e.g., "E1", "E2")
1298       "episode_type": <type_name>,       // {"alert", "alert_resolved", ...}
1299       "modified_by": <model_id>,       // string value {"ICL_LLM", ...}
1300       "timestamp": <value>,             // formatted as iso-8601
1301     }
1302     "adaptation_log": []
1303   },
1304   "version": "1.0",
1305 }

```

1293 *Note:* Listings include // comments for readability; they are illustrative and not strict JSON.

1294 This structure is simple, yet generic enough to accommodate a wide range of problems, from simple  
 1295 single-service policies and more complex multi-service optimization directives spanning typical

mobile traffic types—e.g., URLLC, mMTC, streaming, web, gaming, and voice—each associated with domain-appropriate KPIs (e.g., reliability and latency for URLLC, throughput and jitter for gaming and streaming). Specifically, the OTM blocks define:

- **objective**: This block specifies the primary goal to optimize for a plurality of services, including KPIs, their aggregation level (mean/min/max/p95/sum), scope, unit, and optimization sense (i.e., maximize or minimize).
- **constraints**: This block encodes optional service-specific KPI bounds, each expressed as an inequality  $k^A \odot b$ . To this end, it shares the same fields of the objective block, and additionally includes an operator (lt/le/ge/gt) that specifies the relation to a threshold  $b$  expressed in the stated unit. Optionally, it includes fields indicating modification to a service constraint (modified, modified\_by, id).
- **metadata**: This block records OTM static information, such as provenance, time of origination, etc. and dynamic information related to the last episode event that triggered a modification of the OTM to the and the aggregation timescale, a timestamp, an episode identifier, and an append-only adaptation\_log.
- **version**: specifies the OTM version.

**OTM fields semantics.** The OTM schema currently requires only 15 fields, some of which are common across the schema blocks:

- **service**: Specifies a service class (e.g., mbb, urllc, gaming, streaming, slice).
- **kpi**: Indicates canonical KPIs key resolvable by both the telemetry layer and the optimizer (e.g., throughput, reliability, latency, bler).
- **scope**: Indicates a spatial or logical domain related to a KPI scope (e.g., per\_user, per\_user, per\_slice, per\_user\_group, per\_cell\_group, etc.)
- **aggregation**: Indicates an operator defining how raw samples are aggregated to optimize, evaluate or compare a KPI (e.g., mean, min, max, sum, p95, etc.).
- **unit**: Indicates the type of unit used for KPI or a threshold value (e.g., Mbps, Gbps, ms, s, %, etc.).
- **operator**: Defines relational semantics in a constraint like  $\geq$ ,  $\leq$ ,  $=$  etc. (e.g., le, ge, ge, gt, eq etc.).
- **threshold**: Defines the threshold value  $b_i$  associated to a constraint stated unit; for modified constraints, the value is updated atomically by the adaptor.
- **maximize**: Defines the direction of an optimization (can be true or false)
- **id**: Indicate an identifier associated with OTM, a constraint, an event, etc.
- **episode\_type**: Indicate the type of event that caused a revision of the OTM.
- **created\_by**: Identifies the model or module that originated the OTM
- **modified/modified\_by**: Indicate whether an OTM has been modified and by which model or module
- **timestamp**: Records events times, such as OTM creation and modification...
- **timescale**: Indicates monitoring window

The field **metadata.adaptation\_log** is append-only and used to trace updates with  $\langle$ old, new,  $\Delta$ , rationale, episode, time $\rangle$ . Table 4 exemplifies how typical lexical descriptions of service goals or requirements are mapped into the OTM schema fields:

**Lifecycle and updates.** The fine-tuned LLM creates the initial OTM combining connectivity service intents and network operational intents. This includes verifying the OTM schema validity prior to hand-off to the optimizer through a set of rules: (i) All KPIs must declare units and aggregation; (ii) the operator must be consistent with KPI directionality; (iii) thresholds must lie within domain bounds. During the execution, the ICL-based advisor may propose an update direction with rationale based on telemetry data. The adaptor then computes thresholds adjustments  $\Delta b$  under guardrails, updates the target threshold, and persists the new OTM snapshot atomically. Each episode produces a versioned OTM with a growing adaptation\_log.

1350	Lexical description	KPI	Unit	Aggregation	Scope	Maximize	Operator	Threshold
1351							-	-
1352	Maximize mean cell throughput	throughput	Mbps	mean	per_cell	true	-	-
1353	Minimum user rate above 7Mbps	throughput	Mbps	min	per_user	-	ge	7
1354	Mean users BLER smaller than 10%	bler	%	mean	per_user	-	le	10
1355	95%-tile users latency less than 10ms	latency	ms	p95	per_user	-	le	10

Table 4: Examples of OTM schema values for typical service definitions.

## B.2 EXAMPLE OF OTM ADAPTATION.

Listing 2 illustrates an OTM produced by the fine-tuned LLM. The objective is to maximise mean throughput (Mbps). Three constraints are active: **C1** enforces mean BLER  $\leq 0.10$  over a per-cell window (unitless ratio); **C2** caps user-level latency at 20 ms using the p95 aggregator; and **C3** requires a per-cell minimum user throughput of at least  $b = 7.00$  Mbps. Provenance marks **C3** as modified\_by: ICL\_LLM, indicating that its threshold may be adjusted online. The metadata block specifies the aggregation timescale (10s\_window) and records a snapshot timestamp/episode.

Listing 2: Illustrative OTM with multiple constraints, before adaptation.

```

1  {
2    "version": "1.0",
3    "objective": {
4      "service": "mbb",
5      "kpi": "throughput",
6      "aggregation": "mean",
7      "unit": "Mbps",
8      "maximize": true
9    },
10   "constraints": [
11     {
12       "id": "C1",
13       "service": "mbb",
14       "kpi": "bler",
15       "operator": "le",
16       "threshold": 0.10,
17       "aggregation": "mean",
18       "unit": "",
19       "scope": "per_cell_window",
20       "origin": "fine_tuned_LLM"
21     },
22     {
23       "id": "C2",
24       "service": "mbb",
25       "kpi": "latency_ms",
26       "operator": "le",
27       "threshold": 20,
28       "aggregation": "p95",
29       "unit": "ms",
30       "scope": "per_user_window",
31       "origin": "fine_tuned_LLM"
32     },
33     {
34       "id": "C3",
35       "service": "mbb",
36       "kpi": "tpt_min_mbps",
37       "operator": "ge",
38       "threshold": 7.00,
39       "aggregation": "min",
40       "unit": "Mbps",
41       "scope": "per_cell_window",
42       "origin": "fine_tuned_LLM",
43       "adapted_by": "ICL_LLM"
44     }
45   ],
46   "metadata": {
47     "timescale": "10s_window",
48     "timestamp": "2025-09-22T10:20:00Z",
49     "episode": "alert_001",
50     "adaptation_log": []
51   }
52 }
```

At runtime, the *sliding*-window monitor observes a violation ratio  $VR = 0.60$  for `tpt_min_mbps`, together with  $BLER = 0.14$  (classified as *aggressive* relative to the 0.10 target). The ICL-based advisory selects action `decrease`; the adaptor computes a clipped update  $\Delta b = -0.08$  (subject to caps, budgets, and bounds) and persists the new threshold. Listing 3 shows the resulting *living* OTM: only **C3** changes ( $b : 7.00 \rightarrow 6.92 \text{ Mbps}$ ), while **C1** and **C2** remain unchanged. An `adaptation_log` entry documents the update with `(old_threshold, new_threshold,  $\Delta b$ , rationale, episode, timestamp)`.

Listing 3: Same OTM after one adaptation of constraint C3.

```
1413 1 {  
1414 2   "version": "1.0",  
1415 3   "objective": {  
1416 4     "service": "mbb",  
1417 5     "kpi": "throughput",  
1418 6     "aggregation": "mean",  
1419 7     "unit": "Mbps",  
1420 8     "maximize": true  
1421 9   },  
1422 10  "constraints": [  
1423 11    {  
1424 12      "id": "C1",  
1425 13      "service": "mbb",  
1426 14      "kpi": "bler",  
1427 15      "operator": "le",  
1428 16      "threshold": 0.10,  
1429 17      "aggregation": "mean",  
1430 18      "unit": "",  
1431 19      "scope": "per_cell_window",  
1432 20      "origin": "fine_tuned_LLM"  
1433 21    },  
1434 22    {  
1435 23      "id": "C2",  
1436 24      "service": "mbb",  
1437 25      "kpi": "latency_ms",  
1438 26      "operator": "le",  
1439 27      "threshold": 20,  
1440 28      "aggregation": "p95",  
1441 29      "unit": "ms",  
1442 30      "scope": "per_user_window",  
1443 31      "origin": "fine_tuned_LLM"  
1444 32    },  
1445 33    {  
1446 34      "id": "C3",  
1447 35      "service": "mbb",  
1448 36      "kpi": "tpt_min_mbps",  
1449 37      "operator": "ge",  
1450 38      "threshold": 6.92,  
1451 39      "aggregation": "min",  
1452 40      "unit": "Mbps",  
1453 41      "scope": "per_cell_window",  
1454 42      "origin": "fine_tuned_LLM",  
1455 43      "adapted_by": "ICL_LLM"  
1456 44    }  
1457 45  ],  
1458 46  "metadata": {  
1459 47    "timescale": "10s_window",  
1460 48    "timestamp": "2025-09-22T10:28:00Z",  
1461 49    "episode": "alert_002",  
1462 50    "adaptation_log": [  
1463 51      {  
1464 52        "id": "C3",  
1465 53        "old_threshold": 7.00,  
1466 54        "new_threshold": 6.92,  
1467 55        "delta": -0.08,  
1468 56        "rationale": "VR=0.60; avg=6.92<b=7.00; BLER posture aggressive; relax b to stabilize  
1469 57          HARQ."  
1470 58      }  
1471 59    ]  
1472 60  }  
1473 61 }
```

**1455 Design rationale.** The OTM is deliberately minimal (objective, constraints, metadata) yet extensible  
1456 (aggregation, scope, provenance). This ensures interface stability across RAN domains while enabling  
1457 adaptive operation and full auditability of constraint updates.

1458 C TRANSLATOR SLM FINE-TUNING  
14591460 C.1 DATASET CURATION  
14611462 This section describes the methodology used to construct the supervised corpus for training the Intent-  
1463 to-OTM translator, together with a statistical characterization of the resulting dataset. The objective  
1464 of the curation process is to create a corpus that captures the semantic breadth of natural-language  
1465 QoS intents encountered in operational networks while ensuring strict adherence to the OTM schema  
1466 required for structured policy generation. The design integrates domain knowledge from 5G/6G  
1467 communication systems, QoS-engineering practice, service semantics, and the linguistic variability  
1468 typical of operator-to-system interactions. The final dataset comprises 90,000 samples derived from  
1469 30,000 distinct OTM structures, each paired with three three paraphrased intent utterances.  
14701471 The construction process is guided by four principles: *schema consistency, domain realism, linguistic*  
1472 *diversity, and multi-service generality*. Every instance conforms to the prescribed OTM JSON  
1473 structure to eliminate structural ambiguity. Services, KPIs, thresholds, aggregation functions, and  
1474 operator semantics are selected to reflect realistic RAN-engineering practice rather than arbitrary  
1475 sampling. Multiple paraphrases express the same underlying intent using different linguistic styles,  
1476 while both single-service and multi-service formulations are included to reflect realistic optimization  
1477 scenarios such as cross-slice coordination or heterogeneous multi-tenant workloads. Together, these  
1478 principles ensure that the model learns not only syntactically correct outputs but also semantically  
1479 grounded mappings aligned with operational decision-making.  
14801481 Seven KPIs central to QoS and quality of experience (QoE) optimization are represented: latency,  
1482 packet delay budget, jitter, packet error rate, block error rate, throughput, and spectral efficiency.  
1483 Each KPI is characterized by its physical unit, optimization orientation (minimize or maximize),  
1484 and a plausible operational range. Service-specific threshold distributions are used to maintain  
1485 realism. URLLC thresholds, for example, are drawn from tight low-delay intervals consistent with  
1486 ultra-reliable low-latency requirements; gaming jitter values are sampled from moderate-sensitivity  
1487 ranges; and streaming throughput thresholds reflect bandwidth levels typical of video services. These  
1488 calibrated ranges ensure that the model encounters thresholds reflective of actual RAN-optimization  
1489 tasks rather than arbitrary numeric values.  
14901491 Real-world QoS requirements frequently rely on percentile-based performance metrics, and the  
1492 dataset reflects this by including mean, minimum, maximum, and percentile aggregations from p25  
1493 to p99. Sampling is intentionally biased toward domain-appropriate usage: reliability-sensitive KPIs  
1494 such as latency, jitter, and error rates predominantly use high percentiles (p95 or p99), whereas  
1495 throughput-oriented KPIs typically rely on mean values. This probabilistic, domain-aware selection  
1496 encourages the model to internalize the relationship between service reliability expectations and  
1497 suitable aggregation choices. Constraint operators are chosen in accordance with KPI orientation,  
1498 with minimization KPIs paired with “ $\leq$ ” constraints and maximization KPIs paired with “ $\geq$ ”.  
1499 Semantically invalid combinations, such as lower bounds on error rates, are excluded to prevent the  
1500 model from learning physically implausible relations.  
15011502 Each OTM instance is paired with three natural-language paraphrases produced from four stylistic  
1503 registers: operator-style technical phrasing, 3GPP-inspired formal language, casual expressions,  
1504 and terse imperative commands. These stylistic variants emulate the diverse ways in which human  
1505 operators, analysts, and automated systems articulate QoS intents. The paraphrases incorporate  
1506 synonyms for KPIs and services, linguistic variations in percentile expressions, and syntactic diversity  
1507 ranging from multi-sentence descriptions to compact directives. This controlled diversity promotes  
1508 robustness to real-world phrasing while preserving semantic consistency across paraphrases.  
15091510 To reflect realistic optimization scenarios in multi-slice and multi-tenant RAN deployments, a  
1511 controlled fraction of OTMs include multi-service dependencies in which the optimization objective  
1512 applies to one service while constraints reference another. These cases emulate common operational  
1513 patterns such as managing cross-service interference or guaranteeing simultaneous user-experience  
1514 requirements across heterogeneous traffic types. Their presence strengthens the model’s ability to  
1515 process complex interdependencies and to generate coherent, jointly feasible policies.  
15161517 All generated samples include metadata fields such as an ISO-8601 timestamp and an episode  
1518 identifier. The episode field is fixed to “unspecified” to avoid introducing unintended temporal  
1519

1512 semantics while maintaining compatibility with future policy-orchestration workflows requiring  
 1513 contextual metadata.

1514 The statistical structure of the corpus reflects these design choices. Each OTM specifies one optimi-  
 1515 zation objective and between zero and three constraints consistent with QoS-engineering practices  
 1516 in 5G and 6G networks. Because each template is associated with three paraphrases, the full cor-  
 1517 pus contains 90,000 samples. Constraint cardinality follows a non-uniform distribution chosen to  
 1518 represent operator practice: approximately 10% of OTMs contain no constraints, 45% contain one,  
 1519 30% contain two, and 15% contain three. Consequently, roughly 90% of the corpus includes at least  
 1520 one constraint, and nearly half include multiple constraints. This distribution exposes the model to a  
 1521 broad range of multi-constraint optimization scenarios rather than biasing it toward oversimplified  
 1522 workflows.

1523 Service representation spans eight canonical categories—gaming, streaming, web, messaging,  
 1524 URLLC, mMTC, VoLTE, and VoIP. Sampling is intentionally skewed toward services with stringent  
 1525 QoS requirements. URLLC and gaming each account for approximately 20–25% of OTMs, streaming  
 1526 contributes about 15%, voice (VoLTE and VoIP combined) contributes roughly another 15%, and the  
 1527 remainder corresponds to web, messaging, and mMTC use cases. This distribution ensures adequate  
 1528 coverage across both throughput-oriented and latency-critical traffic classes.

1529 KPI coverage is similarly broad: all seven KPIs appear throughout the dataset following the service-  
 1530 specific domain profiles described above. Over 95% of reliability-related constraints use high-  
 1531 percentile aggregations (p90–p99), preserving realism in statistical QoS modeling. Approximately  
 1532 12–18% of OTM instances include multi-service dependencies, providing inductive signals for joint  
 1533 optimization patterns common in next-generation RAN automation. Finally, linguistic variation  
 1534 reflects the stylistic sampling weights: operator style (40%), 3GPP-inspired formal expressions  
 1535 (30%), casual phrasing (20%), and terse directives (10%). This variation enhances generalization to  
 1536 heterogeneous real-world intent expressions while maintaining semantic consistency across samples.

1537 Collectively, the dataset provides extensive coverage of service semantics, KPI behavior, constraint  
 1538 types, and linguistic variation. Its strict structural consistency, calibrated numerical modeling, and  
 1539 broad paraphrastic diversity make it well suited for supervised fine-tuning of models tasked with  
 1540 translating diverse natural-language intents into precise, schema-compliant OTM structures aligned  
 1541 with RAN-optimization practice.

1542

1543

1544

## C.2 TRAINING METHODOLOGY

1545

1546

1547 The Intent-to-OTM translator is trained using supervised fine-tuning on the curated corpus described  
 1548 in Section C. The task is formulated as a conditional sequence-generation problem: given a natural-  
 1549 language intent and a fixed system prompt, the model must produce a complete and structurally  
 1550 valid OTM in JSON format. Since the mapping between intents and OTMs is deterministic and  
 1551 schema-constrained, the training objective emphasizes exact reproduction of field names, values,  
 1552 ordering, and hierarchical structure.

1553

1554

1555

1556

1557

A transformer-based, instruction-tuned model (Qwen-2.5-7B-Instruct Qwen et al. (2025)) serves as  
 the underlying architecture. Fine-tuning proceeds in a left-to-right autoregressive manner in which  
 each token is generated conditioned on both the input intent and the previously generated output. This  
 preserves the strengths of the pre-trained model while enabling specialization toward domain-specific  
 reasoning over services, KPIs, and QoS constraints.

1558

1559

1560

1561

1562

The optimization pipeline follows established practices for adapting large language models. Stability  
 and parameter-efficiency mechanisms are incorporated, and the specific optimizer settings, learning-  
 rate schedule, and adaptation configuration are reported in Table 13 of Appendix I. These components  
 ensure reliable convergence when generating long, nested JSON structures that are sensitive to  
 single-token variations.

1563

1564

1565

Fine-tuning is performed for a small number of epochs, as the deterministic target format and internal  
 consistency of the dataset enable rapid convergence without significant risk of overfitting. A held-out  
 validation set is used to monitor generalization performance and to detect potential memorization of  
 stylistic artifacts present in the synthetic paraphrases.

1566 Model evaluation combines syntactic and semantic criteria. Token-level accuracy measures fidelity  
 1567 to the target JSON sequence, while a schema-validity check verifies exact compliance with the  
 1568 required OTM specification. In addition, semantic alignment metrics assess whether the model  
 1569 correctly identifies the optimization objective, reproduces the appropriate constraints, and matches  
 1570 the complete ground-truth OTM. Together, these metrics provide a comprehensive assessment of  
 1571 structured intent translation accuracy.

1572

1573 Table 5: Comparison of Evaluation Metrics Between the Fine-Tuned and Baseline Models  
 1574

Metric	Fine-Tuned Model	Baseline Model
Total evaluation examples	1000	1000
JSON valid rate	1.000	1.000
Objective match rate	<b>1.000</b>	0.450
Constraints match rate	<b>0.98</b>	0.215
Full OTM match rate	<b>0.98</b>	0.113
Number of constraint mismatches	20	785
Number of objective mismatches	0	550
Full-match examples	980	113

1586  
 1587  
 1588 C.3 COMPARATIVE EVALUATION  
 1589

1590 This section presents a comparative evaluation of two models for translating natural-language QoS  
 1591 intents into structured OTM representations. The *fine-tuned model* is a supervised LoRA-adapted  
 1592 Qwen-2.5-7B-Instruct trained on a curated corpus of 90,000 intent-OTM pairs (see Section C.2),  
 1593 whereas the *baseline model* is the unmodified Qwen-2.5-7B-Instruct. Both models were evaluated on  
 1594 a held-out test set of 1,000 examples under a strict schema-constrained matching protocol.

1595 As summarized in Table 5, the fine-tuned model demonstrates near-perfect structural and semantic  
 1596 adherence to the OTM schema, achieving 100% JSON validity, 100% objective correctness, and a  
 1597 98% full OTM match rate. In contrast, the baseline model—although also producing syntactically  
 1598 valid JSON in every case—achieves only 45% objective match and 11.3% full OTM match. These  
 1599 results indicate that prompt-only use of a generic instruction-tuned model is insufficient for reliable  
 1600 schema-grounded semantic parsing.

1601 The most pronounced disparity appears in the reconstruction of constraint sets. The fine-tuned  
 1602 model correctly predicts the complete constraint set—including KPI type, operator, aggregation  
 1603 level, threshold, and service scope—in 98% of cases. The baseline model succeeds in only 21.5%  
 1604 of examples, frequently selecting incorrect KPIs, operators, or units despite emitting valid JSON.  
 1605 As shown in Table 6, the baseline model exhibits severe error inflation, particularly in the “missing  
 1606 constraint” (724 occurrences) and “extra constraint” (737 occurrences) categories. It is worth  
 1607 noting that a single erroneous sample may contribute to multiple error categories simultaneously.  
 1608 In contrast, the fine-tuned model’s residual errors remain modest and are concentrated primarily in  
 1609 scope mismatches (50% of erroneous samples), with all other categories occurring infrequently.

1610 Two complementary mechanisms explain the fine-tuned model’s performance advantage. First,  
 1611 supervised adaptation aligns the model’s internal representation with the deterministic structure of  
 1612 the OTM schema. Although the baseline model possesses broad linguistic and domain knowledge, it  
 1613 lacks incentives to prioritize schema-specific conventions such as canonical KPI naming, service–KPI  
 1614 associations, valid threshold ranges, and operator semantics (e.g.,  $\leq$  for reliability-oriented KPIs).  
 1615 Fine-tuning effectively anchors the model’s output distribution to the OTM schema.

1616 Second, the curated dataset encodes domain priors that are internalized during training. For example,  
 1617 URLLC latency values cluster in the 1–10 ms range; gaming intents commonly include jitter con-  
 1618 straints with p95 aggregation; streaming intents typically optimize for throughput; and error-rate KPIs  
 1619 almost always appear with  $\leq$  operators. Lacking these priors, the baseline model frequently produces  
 semantically plausible but non-canonical KPI selections, percentile aggregations, or threshold values.

1620  
1621  
1622  
1623  
1624  
1625  
1626  
1627  
1628  
1629  
1630  
1631  
1632  
1633  
1634  
1635 Table 6: Comparison of Error Categories Between the Fine-Tuned and Baseline Models  
1636  
1637  
1638  
1639  
1640  
1641  
1642  
1643  
1644  
1645  
1646  
1647  
1648  
1649  
1650  
1651  
1652  
1653  
1654  
1655  
1656  
1657  
1658  
1659  
1660  
1661  
1662  
1663  
1664  
1665  
1666  
1667  
1668  
1669  
1670  
1671  
1672  
1673

Error Category	Fine-Tuned Model		Baseline Model	
	Count	Percentage	Count	Percentage
Scope mismatch	10	50%	76	9.7%
Missing constraint	4	25%	724	92.2%
Extra constraint	5	20%	737	93.9%
Threshold mismatch	2	10%	60	7.6%
Operator mismatch	2	10%	68	8.7%
Aggregation mismatch	2	10%	20	2.5%
Count mismatch	1	5%	72	9.2%

Importantly, the fine-tuned model produces no objective-field errors (0 occurrences in Table 6), indicating complete mastery of service–objective alignment. Its remaining errors occur almost exclusively within the constraint block and are typically minor: substituting p90 for p95, small numerical deviations around thresholds, or occasional variations in service scope. These contrast sharply with the baseline model’s structurally inconsistent outputs, which reflect a lack of grounding in the semantics encoded by the OTM schema.

Taken together, Tables 5–6 show that supervised, domain-specific fine-tuning is essential for this task. Although both models generate syntactically valid JSON, only the fine-tuned model functions as a high-precision compiler from natural-language intents to machine-interpretable OTM structures. For practical deployment in autonomous RAN-optimization pipelines, relying solely on prompting a pretrained instruction-tuned model is insufficient; structured domain adaptation is required to ensure correctness, robustness, and operational safety.

1674 **D OPTIMIZER AGENT DESIGN**  
16751676 **D.1 BAYESIAN OPTIMIZATION**  
16771678 We consider the problem of minimizing an unknown objective function  $f : \mathcal{X} \rightarrow \mathbb{R}$ , over a domain  
1679  $\mathcal{X} \subseteq \mathbb{R}^d$ . The goal is to identify

1680 
$$x^* = \arg \min_{x \in \mathcal{X}} f(x).$$
  
1681

1682 In classical numerical optimization, one distinguishes between *global optimization*, where the absolute  
1683 minimum of  $f$  is sought, and *local optimization*, where the search is restricted to neighborhoods of  
1684 an initial point. If  $f$  is convex and  $\mathcal{X}$  is a convex set, the problem reduces to convex optimization, for  
1685 which efficient algorithms exist.1686 However, in many modern applications,  $f$  is a *black-box function*, meaning that its closed-form  
1687 expression is unavailable and evaluations are costly (e.g., expensive simulations or training machine  
1688 learning models). In such cases, *Bayesian optimization (BO)* provides an efficient framework for  
1689 global optimization by maintaining a probabilistic model of  $f$  and selecting query points via a  
1690 surrogate criterion known as an *acquisition function*.1691 **D.1.1 GAUSSIAN PROCESS PRIORS**  
16921693 Bayesian optimization typically employs a *Gaussian process (GP)* prior to model the unknown  
1694 function  $f$ . A GP is defined by a mean function  $\mu(x)$  and covariance function  $K(x, x')$ :

1695 
$$p(f) = \mathcal{GP}(f; \mu, K).$$
  
1696

1697 Given a set of noiseless observations  
1698

1699 
$$\mathcal{D} = \{(x_i, f(x_i))\}_{i=1}^n,$$

1700 the posterior distribution over  $f$  is again a GP with updated mean and covariance functions  $\mu_{f|\mathcal{D}}(x)$   
1701 and  $K_{f|\mathcal{D}}(x, x')$ . This posterior provides both a predictive mean (exploitation) and predictive  
1702 uncertainty (exploration), which form the basis of acquisition functions.1704 **D.1.2 ACQUISITION FUNCTIONS**  
17051706 An *acquisition function*  $a(x)$  encodes the utility of evaluating  $f$  at a candidate point  $x$ . Since  
1707 acquisition functions are cheap to evaluate, the optimization problem is reduced to

1708 
$$x_{t+1} = \arg \max_{x \in \mathcal{X}} a(x),$$
  
1709

1710 or its minimization equivalent. Common acquisition functions include:  
17111712 **Probability of improvement (PI):** Selects points with the highest probability of improving upon  
1713 the best observed value  $f'$ :

1714 
$$a_{\text{PI}}(x) = \Phi\left(\frac{f' - \mu(x)}{\sigma(x)}\right),$$
  
1715

1716 where  $\Phi$  is the Gaussian cumulative distribution function (CDF).  
17171718 **Expected improvement (EI):** Accounts for the *magnitude* of improvement:  
1719

1720 
$$a_{\text{EI}}(x) = (f' - \mu(x)) \Phi\left(\frac{f' - \mu(x)}{\sigma(x)}\right) + \sigma(x) \phi\left(\frac{f' - \mu(x)}{\sigma(x)}\right),$$
  
1721

1722 where  $\phi$  is the Gaussian probability density function (PDF). This criterion naturally balances  
1723 exploration ( $\sigma(x)$ ) and exploitation ( $\mu(x)$ ).  
17241725 **Entropy search (ES)** Reduces uncertainty about the optimizer’s location by minimizing the entropy  
1726 of the distribution  $p(x^*|\mathcal{D})$ . While analytically intractable, approximations make this approach  
1727 feasible.

1728 **Upper confidence bound (UCB):** Promotes exploration via an optimism-in-the-face-of-uncertainty  
 1729 principle:

$$a_{\text{UCB}}(x; \beta) = \mu(x) - \beta\sigma(x),$$

1730 where  $\beta > 0$  is a tunable trade-off parameter. Despite lacking an expected-utility interpretation, UCB  
 1731 has strong theoretical guarantees for asymptotic convergence to the global optimum.  
 1732

1733 This Bayesian decision-theoretic framework provides a principled way to trade off exploration and  
 1734 exploitation, making BO a powerful tool for solving expensive black-box optimization problems.  
 1735

## 1736 D.2 PAX-BO: PREFERENCE-ALIGNED EXPLORATION BAYESIAN OPTIMIZATION

1737 We have

$$\Delta^{d-1} = \{\omega \in \mathbb{R}_{\geq 0}^d : \mathbf{1}^\top \omega = 1\}$$

1738 **Decision space and projection.** We have  $S$  services  $\mathcal{S} = \{1, \dots, S\}$  and preference dimension  $d$ .  
 1739 The internal (unconstrained) variable is

$$U = [u^{(1)}, \dots, u^{(S)}] \in \mathbb{R}^{d \times S}, \quad \bar{u} = \text{vec}(U) \in \mathbb{R}^{dS}.$$

1740 Each service  $s$  uses a probability-simplex preference

$$\mathbf{w}^{(s)} \in \Delta^{d-1} \quad \text{with} \quad \Delta^{d-1} = \{\mathbf{w} \in \mathbb{R}_{\geq 0}^d : \mathbf{1}^\top \mathbf{w} = 1\},$$

1741 obtained by the columnwise Euclidean simplex projection

$$\mathbf{w}^{(s)} = \Pi_{\Delta}(u^{(s)}), \quad \mathbf{W}(U) = [\mathbf{w}^{(1)}, \dots, \mathbf{w}^{(S)}] \in (\Delta^{d-1})^S.$$

1742 Closed form for  $\Pi_{\Delta}$ : sort  $u$  in descending order, find the threshold  $\theta$ , and set  $\Pi_{\Delta}(u) = \max(u - \theta \mathbf{1}, 0)$   
 1743 with  $\mathbf{1}^\top \Pi_{\Delta}(u) = 1$ .

1744 **Objective, constraints, and data.** We optimize a single objective  $f : (\Delta^{d-1})^S \rightarrow \mathbb{R}$  subject to  
 1745 scalar constraints  $\{g_i : (\Delta^{d-1})^S \rightarrow \mathbb{R}\}_{i=1}^p$ . A configuration  $U$  is feasible iff

$$g_i(\mathbf{W}(U)) \leq 0, \quad i = 1, \dots, p.$$

1746 At iteration  $t$ , we evaluate at  $U_{t-1}$  and observe

$$o_t = f(\mathbf{W}(U_{t-1})), \quad c_t^{(i)} = g_i(\mathbf{W}(U_{t-1})) \quad (i = 1, \dots, p),$$

1747 forming the dataset

$$\mathcal{D}_t = \left\{ (\bar{u}_k, o_k, c_k^{(1)}, \dots, c_k^{(p)}) \right\}_{k=1}^t.$$

1748 **Surrogates and acquisition in  $U$ -space.** Fit surrogates for the compositions

$$\mathcal{F}(\bar{u}) \approx f(\mathbf{W}(U)), \quad \mathcal{G}_i(\bar{u}) \approx g_i(\mathbf{W}(U)) \quad (i = 1, \dots, p).$$

1749 Let the incumbent best feasible value be

$$f_{t-1}^* = \max\{o_k : c_k^{(i)} \leq 0 \forall i, k \leq t-1\} \quad (\text{use } -\infty \text{ if none}).$$

1750 Define the acquisition on  $\bar{u}$  (e.g., constrained Log-EI) by

$$\alpha(\bar{u}) = \text{ACQ}(\mathcal{F}, \{\mathcal{G}_i\}; \bar{u}, f_{t-1}^*) \approx \text{LogEI}(\mu_{\mathcal{F}}(\bar{u}), \sigma_{\mathcal{F}}(\bar{u}); f_{t-1}^*) \times \prod_{i=1}^p \Phi\left(-\mu_{\mathcal{G}_i}(\bar{u})/\sigma_{\mathcal{G}_i}(\bar{u})\right).$$

1751 **Single trust region (TR) in  $U$ -space.** Maintain center  $s_c \in \mathbb{R}^{dS}$  and radius  $L > 0$  (half side-length  
 1752 in  $\ell_\infty$ ), giving the box

$$\mathcal{B}_t = \{\bar{v} \in \mathbb{R}^{dS} : \|\bar{v} - s_c\|_\infty \leq L\}.$$

1753 Let  $\kappa_s, \kappa_f, \kappa_\ell$  be the success, failure, and “stuck-at-floor” counters. Parameters:  $L_0$  (initial),  $L_{\min}$   
 1754 (floor),  $L_{\max}$  (cap); thresholds  $s_{\text{th}}$  (expand),  $f_{\text{th}}$  (shrink); tolerance  $\epsilon > 0$ .

1782 **Local proposal (inner step).** Choose the next internal point by maximizing  $\alpha$  within the TR:  
 1783

$$1784 \quad \bar{u}_t \in \arg \max_{\bar{v} \in \mathcal{B}_t} \alpha(\bar{v}), \quad U_t = \text{mat}(\bar{u}_t), \quad \mathbf{W}_t = \mathbf{W}(U_t) = \Pi_{\Delta}(U_t).$$

1786 **Success test and TR adaptation.** After executing  $\bar{u}_{t-1}$ , define the success flag  
 1787

$$1788 \quad \text{SF}_t = \left( \forall i : c_t^{(i)} \leq 0 \right) \wedge \left( o_t \geq f_{t-1}^* + \epsilon \right).$$

1790 If  $\text{SF}_t$ :

$$1791 \quad f_t^* \leftarrow o_t, \quad s_c \leftarrow \bar{u}_{t-1}, \quad \kappa_s \leftarrow \kappa_s + 1, \quad \kappa_f \leftarrow 0, \quad \kappa_\ell \leftarrow 0,$$

1792 and if  $\kappa_s \geq s_{\text{th}}$  then

$$1793 \quad L \leftarrow \min(2L, L_{\text{max}}), \quad \kappa_s \leftarrow 0, \quad \kappa_f \leftarrow 0.$$

1794 Else (failure):

$$1795 \quad f_t^* \leftarrow f_{t-1}^*, \quad \kappa_f \leftarrow \kappa_f + 1, \quad \kappa_s \leftarrow 0,$$

1796 and if  $\kappa_f \geq f_{\text{th}}$  then

$$1797 \quad L \leftarrow \max(L/2, L_{\text{min}}), \quad \kappa_s \leftarrow 0, \quad \kappa_f \leftarrow 0, \quad \kappa_\ell \leftarrow \kappa_\ell + \mathbf{1}\{L = L_{\text{min}}\}.$$

1800 **Smart reset (escape when stuck).** If  $L = L_{\text{min}}$  and  $\kappa_\ell \geq w$ , draw  $n$  candidates on the product  
 1801 simplex:  $\{\mathbf{W}^{(j)} \in (\Delta^{d-1})^S\}_{j=1}^n$  (e.g., Dirichlet/QMC per column), and set  $\bar{u}^{(j)} = \text{vec}(\mathbf{W}^{(j)})$ . For  
 1802 each  $j$  compute

$$1803 \quad \alpha^{(j)} = \alpha(\bar{u}^{(j)}), \quad P_{\text{feas}}^{(j)} = \prod_{i=1}^p \Phi\left(-\mu_{\mathcal{G}_i}(\bar{u}^{(j)})/\sigma_{\mathcal{G}_i}(\bar{u}^{(j)})\right), \quad d^{(j)} = \min_{k \leq t} \|\mathbf{W}^{(j)} - \mathbf{W}(U_k)\|_F.$$

1806 Normalize  $z \in \{\alpha, P_{\text{feas}}, d\}$  to  $\tilde{z}^{(j)} \in [0, 1]$  and score  
 1807

$$1808 \quad \text{score}^{(j)} = \tilde{\alpha}^{(j)} \tilde{P}_{\text{feas}}^{(j)} (\tilde{d}^{(j)})^\beta.$$

1810 Choose  $j^* = \arg \max_j \text{score}^{(j)}$  and reset  
 1811

$$1812 \quad s_c \leftarrow \bar{u}^{(j^*)}, \quad L \leftarrow L_0, \quad \kappa_s, \kappa_f, \kappa_\ell \leftarrow 0.$$

1813 **Action selection (vector-valued  $Q$ ).** Given state  $s_t$  and requested service  $\sigma_t$ , act by preference-  
 1814 aligned scalarization:  
 1815

$$1816 \quad a_t \in \arg \max_{a \in \mathcal{A}} \langle Q(s_t, a), \mathbf{w}_t^{(\sigma_t)} \rangle, \quad \mathbf{w}_t^{(\sigma_t)} = \Pi_{\Delta}(u_t^{(\sigma_t)}).$$

1818 **Remarks.** (i) All modeling and optimization happens in the unconstrained  $\bar{u}$ -space; feasibility  
 1819 is enforced by the projection  $\mathbf{W}(U) = \Pi_{\Delta}(U)$ . (ii) The single TR stabilizes steps; expand/shrink  
 1820 is governed by  $(s_{\text{th}}, f_{\text{th}})$  and improvement tolerance  $\epsilon$ . (iii) The smart reset proposes diverse,  
 1821 high-acquisition, high-feasibility candidates directly on  $(\Delta^{d-1})^S$ .  
 1822

1823

1824

1825

1826

1827

1828

1829

1830

1831

1832

1833

1834

1835

1836

1837

1838 **Algorithm 2** PAX-BO — Preference-Aligned eXploration Bayesian Optimization1839 **Objects.**

1840

- Services  $\mathcal{S} \triangleq \{1, \dots, S\}$ , preference dimension  $d$ .
- Internal (unconstrained) variables  $U = [u^{(1)}, \dots, u^{(S)}] \in \mathbb{R}^{d \times S}$ ; vectorization  $\bar{u} = \text{vec}(U) \in \mathbb{R}^{dS}$ .
- Columnwise simplex projection:  $p^{(s)} = \Pi_{\Delta}(u^{(s)}) \in \Delta^{d-1}$ ; stack  $P(U) = [p^{(1)}, \dots, p^{(S)}] \in (\Delta^{d-1})^S$ .
- Single objective  $f : (\Delta^{d-1})^S \rightarrow \mathbb{R}$ ; constraints are *scalar* functions  $\{g_i : (\Delta^{d-1})^S \rightarrow \mathbb{R}\}_{i=1}^p$ ; feasible iff  $g_i(P) \leq 0 \forall i$ .
- Surrogates (e.g., GPs) model compositions  $\mathcal{F}(\bar{u}) \approx f(P(U))$  and  $\mathcal{G}_i(\bar{u}) \approx g_i(P(U))$  for  $i = 1, \dots, p$ .
- Best feasible value  $f_t^* = \max\{o_j : g_i(P_j) \leq 0 \forall i\}$  (use  $-\infty$  if none).
- Acquisition on  $U$ :  $\alpha(\bar{u}) = \text{ACQ}(\mathcal{F}, \{\mathcal{G}_i\}_{i=1}^p; \bar{u}, f_t^*)$  (e.g., constrained Log-EI).

**Projection.** For  $u \in \mathbb{R}^d$ :  $\Pi_{\Delta}(u) = \arg \min_{p \in \Delta^{d-1}} \|p - u\|_2$ , closed form: sort/threshold;  $p = \max(u - \theta \mathbf{1}, 0)$  with  $\mathbf{1}^\top p = 1$ .

**Single trust region in  $u$ -space.** Center  $s_c \in \mathbb{R}^{dS}$ , radius  $L > 0$ , box  $\mathcal{B} \triangleq \{\bar{v} : \|\bar{v} - s_c\|_\infty \leq L\}$ . Counters: successes  $\kappa_s$ , failures  $\kappa_f$ , stuck-at- $L_{\min} \kappa_\ell$ .

- 1: **Given:** window  $W$ ; trust-region radii  $L_0$  (initial),  $L_{\min}$  (shrink floor),  $L_{\max}$  (expansion cap); thresholds  $s_{\text{th}}$  (successes to expand),  $f_{\text{th}}$  (failures to shrink); tolerance  $\epsilon$ ; reset window  $w$ ; candidate count  $n$ ; diversity exponent  $\beta$
- 2: **Init:** choose  $U_0$ ;  $P_0 = \Pi_{\Delta}(U_0)$ ;  $s_c \leftarrow \text{vec}(U_0)$ ;  $L \leftarrow L_0$ ;  $\kappa_s \leftarrow 0$ ;  $\kappa_f \leftarrow 0$ ;  $\kappa_\ell \leftarrow 0$ ;  $\mathcal{D}_0 \leftarrow \emptyset$
- 3: **for**  $t = 1, 2, \dots$  **do**
- 4:   **Evaluate** at  $\bar{u}_{t-1}$ :  $P_{t-1} = \Pi_{\Delta}(U_{t-1})$ ; observe  $o_t = f(P_{t-1})$  and  $c_t^{(i)} = g_i(P_{t-1})$  for  $i = 1, \dots, p$
- 5:   **Update data**  $\mathcal{D}_t = \mathcal{D}_{t-1} \cup \{(\bar{u}_{t-1}, o_t, (c_t^{(i)})_{i=1}^p)\}$ ; keep last  $W$ ; refit  $\mathcal{F}$ ,  $\{\mathcal{G}_i\}$  on  $(\bar{u}, o, (c^{(i)})_{i=1}^p)$ ; compute  $f_{t-1}^*$
- 6:   **Suggest next  $U$  (local TR maximization)**

$$\bar{u}_t \in \arg \max_{\bar{v} \in \mathcal{B}} \alpha(\bar{v}), \quad U_t = \text{mat}(\bar{u}_t), \quad P_t = \Pi_{\Delta}(U_t).$$

- 7:   **TR update** with success flag  $\text{SF} \leftarrow \left(\forall i : c_t^{(i)} \leq 0\right) \wedge \left(o_t \geq f_{t-1}^* + \epsilon\right)$
- 8:   **if**  $\text{SF}$  **then**
- 9:      $f_t^* \leftarrow o_t$ ;  $s_c \leftarrow \bar{u}_{t-1}$ ;  $\kappa_s \leftarrow \kappa_s + 1$ ;  $\kappa_f \leftarrow 0$ ;  $\kappa_\ell \leftarrow 0$
- 10:     **if**  $\kappa_s \geq s_{\text{th}}$  **then**  $L \leftarrow \min(2L, L_{\max})$ ;  $\kappa_s \leftarrow 0$ ;  $\kappa_f \leftarrow 0$
- 11:     **end if**
- 12:   **else**
- 13:      $\kappa_f \leftarrow \kappa_f + 1$ ;  $\kappa_s \leftarrow 0$
- 14:     **if**  $\kappa_f \geq f_{\text{th}}$  **then**
- 15:        $L \leftarrow \max(L/2, L_{\min})$ ;  $\kappa_s \leftarrow 0$ ;  $\kappa_f \leftarrow 0$
- 16:       **if**  $L = L_{\min}$  **then**  $\kappa_\ell \leftarrow \kappa_\ell + 1$
- 17:       **end if**
- 18:     **end if**
- 19:   **end if**
- 20:   **Smart reset (if stuck).** If  $L = L_{\min}$  and  $\kappa_\ell \geq w$ :
- 21:     Sample  $n$  candidates  $P^{(j)} \in (\Delta^{d-1})^S$  (Dirichlet/QMC per service); set  $\bar{u}^{(j)} = \text{vec}(P^{(j)})$
- 22:     For each  $j$ : compute  $\alpha(\bar{u}^{(j)})$ ;  $P_{\text{feas}}^{(j)} = \prod_{i=1}^p \Phi\left(-\mu_{\mathcal{G}_i}(\bar{u}^{(j)})/\sigma_{\mathcal{G}_i}(\bar{u}^{(j)})\right)$ ;  $d^{(j)} = \min_{(\bar{u}_k, \cdot, \cdot) \in \mathcal{D}_t} \|P^{(j)} - P_k\|_F$
- 23:     Normalize  $\tilde{z} = \frac{z - \min z}{\max z - \min z + \epsilon}$  for  $z \in \{\alpha, P_{\text{feas}}, d\}$ ;  $\text{score}^{(j)} = \tilde{\alpha}^{(j)} \tilde{P}_{\text{feas}}^{(j)} (\tilde{d}^{(j)})^\beta$
- 24:     Set  $s_c \leftarrow \bar{u}^{(j^*)}$  where  $j^* = \arg \max_j \text{score}^{(j)}$ ;  $L \leftarrow L_0$ ;  $\kappa_s \leftarrow 0$ ;  $\kappa_f \leftarrow 0$ ;  $\kappa_\ell \leftarrow 0$
- 25:   **Action (vector-valued  $Q$ ).** For state  $s_t$  and requested service  $\sigma_t$ :
- 26:      $a_t \in \arg \max_{a \in \mathcal{A}} \langle Q(s_t, a, p_t^{(\sigma_t)}), p_t^{(\sigma_t)} \rangle$  with  $p_t^{(\sigma_t)} = \Pi_{\Delta}(u_t^{(\sigma_t)})$ .

1887 **26: end for**

1888

1889

1890 

## E MULTI-OBJECTIVE REINFORCEMENT LEARNING

1891 

### E.1 MULTI-OBJECTIVE MARKOV DECISION PROCESS

1892 A multi-objective Markov decision process (MOMDP) extends the traditional Markov decision  
1893 process (MDP) framework (Puterman, 1994; Sutton & Barto, 2018) by considering not just one,  
1894 but multiple objectives, which may conflict with each other (Rojers et al., 2013; Yang et al., 2019).  
1895 Within this framework, an agent seeks to optimize several reward functions simultaneously, each  
1896 corresponding to a different objective. These objectives can either conflict or complement each other;  
1897 thus, improvements in one may adversely affect another or contribute positively to shared goals. The  
1898 primary goal of an MOMDP is to derive a policy that achieves an optimal balance among multiple  
1899 objectives. This trade-off is typically represented by a Pareto front comprising a set of optimal  
1900 policies such that no policy can outperform another across all objectives simultaneously, thereby  
1901 making them non-dominated (Rojers et al., 2013; Moffaert & Nowé, 2014).  
1902

1903 Formally, an MOMDP is defined by the tuple  $\langle \mathcal{S}, \mathcal{A}, \mathcal{P}, \mathbf{r}, \Omega, f_\omega, \gamma \rangle$ , where  $\mathcal{S}$  denotes the state  
1904 space,  $\mathcal{A}$  the action space,  $\mathcal{P}(s' \mid s, a)$  the transition probabilities, and  $\gamma \in [0, 1)$  a discount factor.  
1905 The vector  $\mathbf{r} = [r_1, r_2, \dots, r_m]^\top$  represents the  $m$ -dimensional reward vector, and we assume the  
1906 preference space  $\Omega$  to be the standard  $(m - 1)$ -dimensional simplex (probability simplex), defined as  
1907

1908 
$$\Delta^{m-1} = \left\{ \omega \in \mathbb{R}^m : \sum_{i=1}^m \omega_i = 1 \text{ and } \omega_i \geq 0 \text{ for } i = 1, \dots, m \right\}. \quad (4)$$

1909 Here, each preference vector  $\omega \in \Omega$  assigns a normalized non-negative weight to each objective,  
1910 reflecting its relative importance. We focus on a class of MOMDPs with a *linear preference function*,  
1911 in which a scalarization function  $f_\omega(\mathbf{r}) = \omega^\top \mathbf{r}$  converts the reward vector into a scalar return using  
1912 the preference vector  $\omega$  (Rojers et al., 2013; Hayes et al., 2022). The cumulative expected return  
1913 under a policy  $\pi$  is then given by  
1914

1915 
$$\hat{r}^\pi = \mathbb{E} \left[ \sum_{t=0}^{\infty} \gamma^t \mathbf{r}(s_t, a_t) \mid \pi \right].$$

1916 **Observation 1** When the preference vector  $\omega \in \Omega$  is fixed, an MOMDP reduces to a standard MDP  
1917 with scalar rewards.  
1918

1919 **Remark 1 (Interpretation of Linear Scalarization)** The scalarization function  $f_\omega(\mathbf{r}) = \omega^\top \mathbf{r}$   
1920 can be interpreted as an expectation. If  $\mathbf{r} = [r_1, r_2, \dots, r_m]^\top$  is the reward vector and  $\omega =$   
1921  $[w_1, w_2, \dots, w_m]^\top$  is a weight vector with  $\sum_{i=1}^m w_i = 1$  and  $w_i \geq 0$ , then  
1922

1923 
$$r_s = \omega^\top \mathbf{r} = \sum_{i=1}^m w_i r_i = \mathbb{E}_{i \sim \omega}[r_i].$$

1924 Thus,  $\omega$  can be interpreted as a probability distribution over objectives, and the function  $f_\omega$  corre-  
1925 sponds to the expected reward under this distribution (Rojers et al., 2013).  
1926

1927 **Example 1 (Multi-Objective Q-Learning)** In traditional single-objective Q-learning, the agent  
1928 estimates a scalar Q-value for each state-action pair;  
1929

1930 
$$Q(s, a) = \mathbb{E} \left[ \sum_{t=0}^{\infty} \gamma^t r(s_t, a_t) \mid s_0 = s, a_0 = a \right].$$

1931 In contrast, in multi-objective reinforcement learning, each action may yield a vector of rewards  
1932 corresponding to different objectives (Rojers et al., 2013; Nguyen et al., 2018).  
1933

1934 For example, in a self-driving car scenario, the agent may consider:  
1935

1936 


1937 - $r_{\text{speed}}$ : how fast the car goes,
1938 - $r_{\text{fuel}}$ : fuel efficiency,
1939 - $r_{\text{safety}}$ : safety score.

1944 Thus, the  $Q$ -value function becomes vector-valued:  
 1945

$$1946 \quad 1947 \quad 1948 \quad Q(s, a) = \begin{bmatrix} Q_{\text{speed}}(s, a) \\ Q_{\text{fuel}}(s, a) \\ Q_{\text{safety}}(s, a) \end{bmatrix} \in \mathbb{R}^3.$$

1949 Each decision-maker may have different trade-offs between these objectives, expressed as a preference  
 1950 vector  $\omega \in \Omega = \Delta^2$ . For instance,  $\omega = [0.7, 0.2, 0.1]^\top$  indicates 70% priority on speed, 20% on  
 1951 fuel efficiency, and 10% on safety. The corresponding scalarized utility is given by  $\omega^\top Q(s, a)$ .  
 1952

## 1953 E.2 CONVEX COVERAGE SET

1954 In multi-objective optimization, the *Pareto front*  $\mathcal{F}^*$  contains all Pareto optimal solutions (Rojers  
 1955 et al., 2013; Moffaert & Nowé, 2014), defined as

$$1956 \quad \mathcal{F}^* \triangleq \{ \mathbf{r} \in \mathbb{R}^m : \nexists \mathbf{r}' \in \mathbb{R}^m \text{ such that } \mathbf{r}'_i \geq \mathbf{r}_i \text{ for all } i \text{ and } \mathbf{r}'_j > \mathbf{r}_j \text{ for at least one } j \} .$$

1957 However, not all Pareto-optimal points are relevant when preferences are restricted to linear scalariza-  
 1958 tions. In this case, only a subset of the Pareto front—the *convex coverage set* (CCS)—is sufficient  
 1959 (Rojers et al., 2015).

1960 The CCS is a subset of  $\mathcal{F}^*$  consisting of non-dominated solutions that are optimal under some linear  
 1961 preference vector. Mathematically, it is defined as

$$1962 \quad \mathcal{C} \triangleq \{ \hat{\mathbf{r}} \in \mathcal{F}^* : \exists \omega \in \Delta^{m-1} \text{ such that } \omega^\top \hat{\mathbf{r}} \geq \omega^\top \hat{\mathbf{r}}' \text{ for all } \hat{\mathbf{r}}' \in \mathcal{F}^* \} .$$

1963 Thus, the CCS comprises those points on the outer convex boundary of  $\mathcal{F}^*$  that maximize utility for  
 1964 at least one preference vector  $\omega$  (with  $\Omega = \Delta^{m-1}$  as defined above) (Rojers et al., 2015).

1965 **Example 2** Consider a bi-objective optimization problem with the objectives of maximizing accuracy  
 1966 and maximizing interpretability. The CCS would include those points on the Pareto front that  
 1967 maximize a weighted sum of the two objectives for some given trade-off between them; these points  
 1968 lie on the convex outer boundary of the feasible set in the objective space.

1969 Therefore, the CCS represents the minimal subset of the Pareto front that guarantees optimality under  
 1970 some linear preference. It is particularly valuable in decision-making scenarios where preferences  
 1971 may vary, since it identifies exactly those solutions that are relevant for all possible linear trade-offs  
 1972 (Rojers et al., 2013).  
 1973

## 1974 E.3 ENVELOPE Q-LEARNING

1975 **Scope.** Learn a *single* preference-conditioned action-value function  $Q : \mathcal{S} \times \mathcal{A} \times \Delta^{m-1} \rightarrow \mathbb{R}^m$  such  
 1976 that, for any linear preference  $\omega \in \Delta^{m-1}$ , the scalar projection  $\omega^\top Q(s, a, \omega)$  equals the optimal  
 1977 scalarized value for acting under  $\omega$  (Yang et al., 2019). The induced policy is  
 1978

$$1979 \quad \pi_\omega(s) = \arg \max_{a \in \mathcal{A}} \omega^\top Q(s, a, \omega; \theta). \quad (5)$$

1980 **Envelope maximizer selection (double DQN (DDQN) style).** Given a transition  $(s, a, r, s')$  with  
 1981 reward vector  $\mathbf{r}(s, a) \in \mathbb{R}^m$ , EQL (Yang et al., 2019) bootstraps from the *envelope* at the next state  
 1982 by selecting the action–preference pair that maximizes the  $\omega$ -projection using the *online* network  
 1983 (parameters  $\theta$ ):  
 1984

$$1985 \quad (a^*, \omega^*) = \arg \max_{a' \in \mathcal{A}, \omega' \in \Delta^{m-1}} \omega'^\top Q(s', a', \omega'; \theta). \quad (6)$$

1986 This couples preferences because  $\omega^*$  is not necessarily equal to the current  $\omega$ .  
 1987

1988 **Vector temporal difference (TD) target (envelope bootstrap).** Evaluate the selected pair with the  
 1989 *target* network  $\theta^-$  to obtain a *vector* target  
 1990

$$1991 \quad y = r(s, a) + \gamma Q(s', a^*, \omega^*; \theta^-) \in \mathbb{R}^m, \quad (7)$$

1992 where the expectation over  $s' \sim \mathcal{P}(\cdot | s, a)$  is approximated by sampling  $s'$  from the replay buffer.  
 1993

---

**Algorithm 3** Envelope Q-learning (EQL) with Preference-Guided Replay

---

**Require:** Discount factor  $\gamma$ , prioritized buffer  $\mathcal{B}$ , target period  $C$ , Dirichlet prior  $\alpha$ , preferences-per-sample  $K$ , minibatch size  $B$

1: Initialize  $\theta$ ; set  $\theta^- \leftarrow \theta$  ▷ Interaction (Figure 6, steps 1–5)

2: **for** each environment step **do** ▷ (1)

3:   observe  $s$  ▷ (2)

4:   sample  $\omega \sim \text{Dir}(\alpha)$  ▷ (3)

5:   choose  $a$  by  $\varepsilon$ -greedy on  $\omega^\top Q(s, a, \omega; \theta)$  ▷ (4)

6:   execute  $a$ ; observe  $r, s'$  ▷ (4)

7:   push  $(s, a, r, s')$  into  $\mathcal{B}$  with initial priority  $p_{\max}$  ▷ (5)

8: **end for** ▷ Learning (Figure 6, steps 6–11)

9: **for** each gradient step **do** ▷ (7)

10:   sample  $\{(s_i, a_i, r_i, s'_i)\}_{i=1}^B \sim \mathcal{B}$  by priority ▷ (7)

11:   **for**  $i = 1$  to  $B$  **do** ▷ (6), (8)

12:     sample  $\mathcal{W}_i = \{\omega'_{ij}\}_{j=1}^K$  ▷ (6), (8)

13:      $(a_i^*, \omega_i^*) \leftarrow \arg \max_{a', \omega' \in \mathcal{W}_i} \omega_i^\top Q(s'_i, a', \omega'; \theta)$  ▷ (6), (8)

14:     **Vector target:**  $y_i \leftarrow r_i + \gamma Q(s'_i, a_i^*, \omega_i^*; \theta^-)$  ▷ Equation (7)

15:      $\delta_i^A \leftarrow y_i - Q(s_i, a_i, \omega_i; \theta)$  ▷ vector TD

16:      $\delta_i^B \leftarrow \omega_i^\top y_i - \omega_i^\top Q(s_i, a_i, \omega_i; \theta)$  ▷ scalarized TD

17:   **end for** ▷ (10)

18:   minimize  $\mathcal{L}(\theta) = \frac{1}{B} \sum_i [(1 - \lambda) |\delta_i^B| + \lambda \|\delta_i^A\|_2^2]$  ▷ Equations (8) to (10)

19:   update priorities in  $\mathcal{B}$  using  $\|\delta_i^A\|_1$  or  $|\delta_i^B|$  ▷ Equation (12)

20:   **if** step mod  $C = 0$  **then** ▷ (12)

21:      $\theta^- \leftarrow \theta$  ▷ (12)

22:   **end if** ▷ (12)

23: **end for** ▷ (12)

---

**Losses.** The primary objective regresses the full vector target (cf. Eq. (6) in (Yang et al., 2019)):

$$\mathcal{L}^A(\theta) = \mathbb{E}_{(s, a, r, s'), \omega} \left[ \|y - Q(s, a, \omega; \theta)\|_2^2 \right]. \quad (8)$$

Because the optimal frontier contains many discrete extreme points (a nonsmooth landscape), an auxiliary *scalarized* loss improves optimization stability (cf. Eq. (7) in (Yang et al., 2019)):

$$\mathcal{L}^B(\theta) = \mathbb{E}_{(s, a, r, s'), \omega} \left[ \left\| \omega^\top y - \omega^\top Q(s, a, \omega; \theta) \right\|_2 \right]. \quad (9)$$

The neural network is trained by employing a simple homotopy:

$$\mathcal{L}(\theta) = (1 - \lambda) \mathcal{L}^B(\theta) + \lambda \mathcal{L}^A(\theta), \quad \lambda \in [0, 1], \lambda \uparrow 1. \quad (10)$$

**Approximating the inner maximization.** The maximization over  $\omega' \in \Delta^{m-1}$  in Equation (6) is approximated by sampling a small candidate set  $\mathcal{W} = \{\omega'_j\}_{j=1}^K \subset \Delta^{m-1}$  (e.g., from a Dirichlet distribution) and computing

$$(a^*, \omega^*) \approx \arg \max_{a' \in A, \omega' \in \mathcal{W}} \omega^\top Q(s', a', \omega'; \theta). \quad (11)$$

Each transition is *relabelled* with multiple sampled preferences (hindsight preference relabeling), which couples learning across the preference space and greatly improves sample efficiency. The complete training loop—with preference sampling, hindsight relabeling, prioritized replay, and the envelope bootstrap—is summarized in Figure 6 and Algorithm 3.

**Replay priority.** Priorities can be derived from vector or scalarized TD errors, e.g.,

$$p \propto \|y - Q(s, a, \omega; \theta)\|, \quad \text{or} \quad p \propto |\omega^\top y - \omega^\top Q(s, a, \omega; \theta)|. \quad (12)$$

**Theory and intuition.** The *envelope Bellman operator* (induced by Equation (7)) has a unique fixed point and is a  $\gamma$ -contraction under a suitable metric; hence EQL converges in tabular settings (Yang et al., 2019).

2052  
2053 **Geometric interpretation.** The selection of the envelope maximizer, seen in Equation (6), *backs*  
2054 *up from* the upper convex hull of the next state returns. Define  
2055

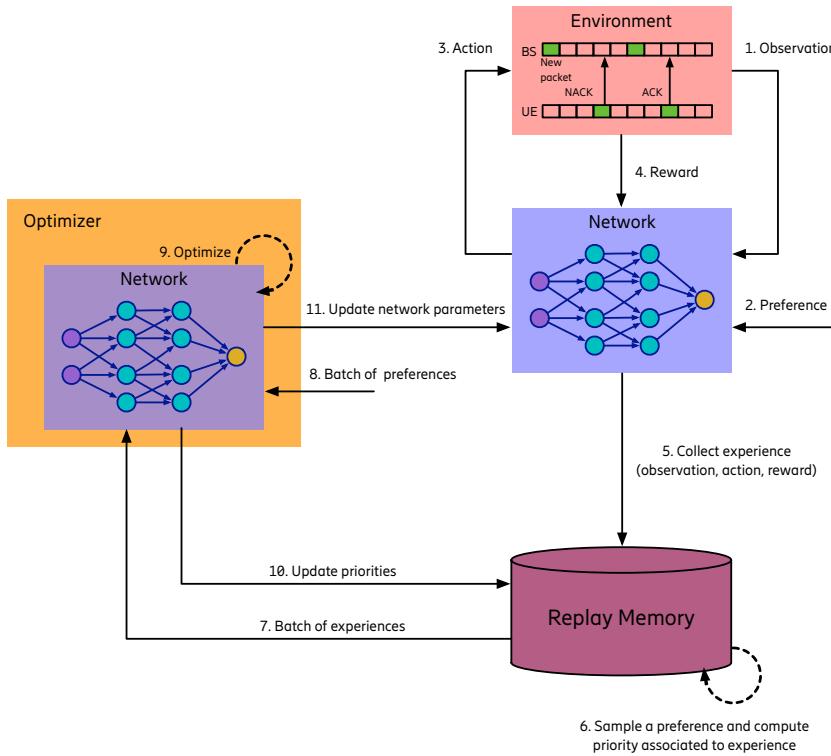
$$\mathcal{V}(s') \triangleq \{ Q(s', a', \omega') : a' \in \mathcal{A}, \omega' \in \Omega \} \subset \mathbb{R}^m.$$

2056 Since  $\omega^\top(\cdot)$  is linear, maximizing it over a set equals maximizing it over that set’s convex hull  
2057 (support-function invariance):  
2058

$$\max_{a', \omega'} \omega^\top Q(s', a', \omega') = \max_{v \in \text{conv } \mathcal{V}(s')} \omega^\top v. \quad (13)$$

2061 Consequently, Equation (6) selects a supporting extreme point of  $\text{conv } \mathcal{V}(s')$ , and the target vector  
2062 in Equation (7) *bootstraps* from this convex envelope of the solution frontier; hence dominated  
2063 trade-offs are not reinforced and EQL effectively targets the convex coverage set (CCS).<sup>1</sup>  
2064

2065 **Adaptation.** At runtime, the trained policy  $\pi_\omega$  can be executed with any desired preference vector  
2066  $\omega$  without retraining.  
2067



2092 **Figure 6: Multi-objective RL with preference-guided optimization (EQL training loop).** (1)  
2093 Observe state; (2) sample preference  $\omega$ ; (3) act; (4) receive vector reward  $r$ ; (5) store  $(s, a, r, s')$ ; (6)  
2094 memory samples auxiliary preferences and computes envelope-based priorities; (7) return prioritized  
2095 batch; (8) return batch of preferences; (9) optimize using the vector target in Equation (7) and  
2096 losses Equations (8) to (10); (10) update priorities; (11) update network parameters.  
2097

2098 **EQL shortcomings:** Two design aspects limit the scalability of EQL Yang et al. (2019) in large  
2099 state–action spaces. First, EQL relies on a singleton architecture Algorithm 3, where a single actor  
2100 must explore the full joint state–action–preference space, leading to poor coverage and inefficient  
2101 learning. Second, sample priorities are assigned only once at generation and never updated during  
2102 training (unlike, e.g., (Horgan et al., 2018)), which slows convergence. Since RAN control problems  
2103 involve vast state–action spaces, we propose a distributed EQL variant where multiple actors share  
2104 the exploration load that leads to improved coverage and performance appendix F.5.  
2105

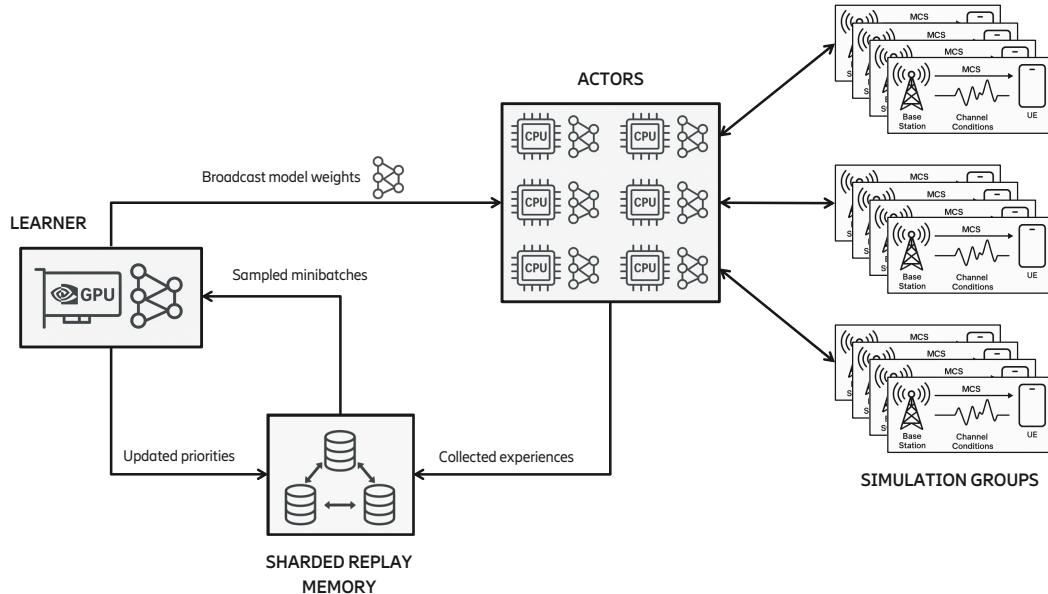
<sup>1</sup>With finite  $\mathcal{A}$  (and discretized  $\Omega$  in practice),  $\mathcal{V}(s')$  is finite, so  $\text{conv } \mathcal{V}(s')$  is a polytope and the maximum in equation 13 is attained at a vertex.

## 2106 F DISTRIBUTED ENVELOPE Q-LEARNING

2108 We propose D-EQL, a distributed MORL algorithm that extends EQL (Yang et al., 2019) with an  
 2109 APE-X-style distributed architecture (Horgan et al., 2018) for faster and more efficient exploration of  
 2110 the preference space. As in vanilla EQL, our method optimizes a single policy/value network over  
 2111 preferences for multiple competing objectives. Unlike the original setting, we employ learner–actor  
 2112 decoupling (Horgan et al., 2018) and distribute exploration over the preference space across multiple  
 2113 parallel actors. Specifically, we partition the preference simplex into subspaces and allocate different  
 2114 actors to explore different subspaces in parallel. While the partitioning of the preference space is  
 2115 inspired by (Xu et al., 2020), distributing exploration across actors improves coverage and exploration  
 2116 efficiency. Furthermore, we employ distributed prioritized experience replay with hindsight to  
 2117 improve sample efficiency: prioritized replay selects the most informative experiences at each training  
 2118 step, while hindsight relabeling increases reuse by updating priorities under multiple preferences  $\omega$ .  
 2119

### 2120 F.1 D-EQL ARCHITECTURE

2121 Our framework follows a scalable distributed (multi-objective) reinforcement learning architecture  
 2122 that decouples data collection, storage, and learning; see Figure 7. A set of CPU-based actors, each  
 2123 running a replica of the policy network, interact in parallel with multiple simulation environments  
 2124 to generate trajectories of state, action, reward, and next state tuples. Actors generate experiences  
 2125 by exploring *only* a partition the preference simplex. These experiences are first stored locally and  
 2126 then pushed to a sharded replay memory, where data is distributed either via load-balancing or fixed  
 2127 actor-to-shard mappings. Each shard operates as an independent replay buffer, enabling parallel  
 2128 writes and prioritized sampling. A GPU-based learner initially allocates a subspaces of the preference  
 2129 simplex to each actor for distributed exploration. The learner then periodically samples mini-batches  
 2130 from all shards, performs gradient updates on the policy network, and returns updated priorities  
 2131 to maintain efficient replay. Updated network weights are then broadcast to all actors, ensuring  
 2132 consistent synchronization across distributed processes. This design allows the system to scale  
 2133 efficiently with the number of actors and environments, achieving high-throughput of experience  
 2134 collection while preserving stability in training.



2154 **Figure 7: Overview of the distributed (multi-objective) reinforcement learning architecture.**  
 2155 CPU-based actors interact in parallel with multiple simulation groups to generate trajectories, which  
 2156 are stored in a sharded replay memory, while exploring different subspaces of the preference simplex.  
 2157 A GPU-based learner samples mini-batches from the shards, updates the policy network, and broad-  
 2158 casts the updated weights back to the actors while sending updated priorities to the replay memory.  
 2159 This design enables scalable experience collection and stable policy learning.

2160 F.2 DISTRIBUTED ACTORS  
2161

2162 We consider a MOMDP with state space  $\mathcal{S}$ , action set  $\mathcal{A}$ , transition kernel  $\mathcal{P}(s' | s, a)$ , discount  
2163 factor  $\gamma \in [0, 1]$ , and vectorial reward  $\mathbf{r}(s, a) \in \mathbb{R}^m$ . Preferences lie on the probability simplex  
2164  $\Omega = \Delta^{m-1}$ . The learner maintains vectorial action-value functions  $Q(s, a, \omega; \theta) \in \mathbb{R}^m$  (online) and  
2165  $Q(s, a, \omega; \theta^-) \in \mathbb{R}^m$  (target) and periodically shares them with all actors. For any  $\omega \in \Omega$ , we define  
2166 the scalarization as

$$2167 Q_\omega(s, a; \theta) \triangleq \omega^\top Q(s, a, \omega; \theta) \in \mathbb{R}. \quad (14)$$

2168 Each actor (i) generates experiences under an  $\varepsilon$ -greedy policy conditioned on a preference  $\omega$ , (ii)  
2169 computes a *scalarized* DDQN TD error to initialize prioritized experience replay priorities, and (iii)  
2170 sends batched transitions and priorities to its assigned replay memory shard. Each actor is assigned  
2171 to a stratum  $\Omega_L^{(u)}$  from a simplex lattice (see Algorithm 6) and samples  $\omega$  uniformly in that simplex  
2172 via barycentric weights (see Algorithm 7).  
2173

2174 BEHAVIOR POLICY AND DATA GENERATION  
2175

2176 At the beginning of episode  $e$ , the actor draws a preference with support on the actor's stratum  $\Omega_L^{(u)}$   
2177 (see Appendix F.4) and keeps it fixed:

$$2178 \omega^{(e)} \sim p_u(\cdot), \quad \omega_t \equiv \omega^{(e)} \text{ for all } t \text{ in episode } e.$$

2179 At the environment step  $t \in \mathbb{N}$ , the exploration rate is linearly annealed from  $\varepsilon_{\max}$  to  $\varepsilon_{\min}$  over  $T_{\text{decay}}$   
2180 steps:

$$2182 \varepsilon_t = \max \left\{ \varepsilon_{\min}, \varepsilon_{\max} - \frac{\varepsilon_{\max} - \varepsilon_{\min}}{T_{\text{decay}}} t \right\}. \quad (15)$$

2184 Here,  $\varepsilon_{\max} \in (0, 1]$  is the initial exploration rate,  $\varepsilon_{\min} \in [0, \varepsilon_{\max})$  is the floor, and  $T_{\text{decay}} \in \mathbb{N}$  is the  
2185 annealing horizon (after which  $\varepsilon_t$  is clamped).

2186 We define the greedy action under the fixed preference (using Equation 14) as

$$2187 2188 a_\omega^*(s) = \arg \max_{a \in \mathcal{A}} Q_\omega(s, a; \theta). \quad (16)$$

2189 Let  $\xi_t \sim \text{Unif}(0, 1)$  and let  $\text{UnifAct}(\mathcal{A})$  denote a single uniform draw from  $\mathcal{A}$  used only when  
2190 exploring. As a result, the executed action becomes:

$$2192 2193 a_t = \begin{cases} \text{UnifAct}(\mathcal{A}), & \text{if } \xi_t < \varepsilon_t, \\ a_\omega^*(s_t), & \text{otherwise.} \end{cases} \quad (17)$$

2194 Applying  $a_t$  yields a transition  $(s_t, a_t, \mathbf{r}_t, s_{t+1}, d_t)$  with a terminal flag  $d_t \in \{0, 1\}$ .  
2195

2196 INITIAL PRIORITY COMPUTATION  
2197

2198 To initialize prioritized replay, the actor computes a *scalar* DDQN TD error using a *fresh* preference  
2199  $\tilde{\omega} \sim p_u(\cdot)$  (independent of the behavior preference) to diversify the priorities:

$$2200 2201 a_{\tilde{\omega}}^*(s_{t+1}) = \arg \max_{a' \in \mathcal{A}} \tilde{\omega}^\top Q(s_{t+1}, a', \tilde{\omega}; \theta), \quad (18)$$

$$2202 2203 \delta_{\text{act}} = \underbrace{\tilde{\omega}^\top \mathbf{r}_t}_{r_{\tilde{\omega}}} + \gamma(1 - d_t) \tilde{\omega}^\top Q(s_{t+1}, a_{\tilde{\omega}}^*(s_{t+1}), \tilde{\omega}; \theta^-) - \tilde{\omega}^\top Q(s_t, a_t, \tilde{\omega}; \theta), \quad (19)$$

$$2204 2205 p_{\text{init}} = |\delta_{\text{act}}| + \epsilon_0, \quad \epsilon_0 \ll 1. \quad (20)$$

2206 The pair  $((s_t, a_t, \mathbf{r}_t, s_{t+1}, d_t), p_{\text{init}})$  is buffered locally and flushed to the assigned replay shard.  
2207

2208 LOCAL BATCHING AND BATCHED COMMUNICATION  
2209

2210 Let  $u \in \{0, \dots, U-1\}$  denote the actor id and  $K$  be the number of replay shards. The actor accumulates  
2211 transitions in a local circular buffer of capacity  $C$  and, when full, sends the batch to a designated  
2212 shard via a single remote procedure call (RPC). Shard selection uses the deterministic mapping  
2213

$$2214 k(u) = u \bmod K. \quad (21)$$

2215 The actor  $u$  then transmits the batch  $\{(s_n, a_n, \mathbf{r}_n, s'_n, d_n), p_{\text{init}, n}\}_{n=1}^C$  to shard  $k(u)$ .  
2216

---

2214 **Algorithm 4** D-EQL Actor

---

2215 **Require:** actor id  $u$ , shards  $\{\text{SHARD}_k\}_{k=0}^{K-1}$ , local buffer capacity  $C$ , discount factor  $\gamma$ , schedule

2216    params  $(\varepsilon_{\max}, \varepsilon_{\min}, T_{\text{decay}})$ , stratum  $\Omega_L^{(u)}$  with sampler  $p_u(\cdot)$ , online  $Q(\cdot; \theta)$ , target  $Q(\cdot; \theta^-)$

2217    1:  $k \leftarrow u \bmod K$  ▷ assigned shard, cf. Equation 21

2218    2:  $\mathcal{B} \leftarrow \emptyset, \mathcal{P} \leftarrow \emptyset, t \leftarrow 0$

2219    3: **for** episode  $e = 1, 2, \dots$  **do**

2220    4:    **Sample episodic preference:**  $\omega^{(e)} \sim p_u(\cdot)$ ; set  $\omega_t \equiv \omega^{(e)}$  for this episode

2221    5:    **while** episode not terminated **do**

2222    6:     observe  $s_t$

2223    7:     compute  $\varepsilon_t$  via the linear schedule Equation 15

2224    8:     draw  $\xi_t \sim \text{Unif}(0, 1)$

2225    9:     **if**  $\xi_t < \varepsilon_t$  **then**

2226    10:        $a_t \leftarrow \text{UnifAct}(\mathcal{A})$

2227    11:     **else**

2228    12:        $a_t \leftarrow \arg \max_{a \in \mathcal{A}} Q_{\omega_t}(s_t, a; \theta)$  ▷ greedy map Equation 16

2229    13:     **end if**

2230    14:     execute  $a_t$ ; observe  $(\mathbf{r}_t, s_{t+1}, d_t)$

2231    15:     **Sample priority preference:**  $\tilde{\omega} \sim p_u(\cdot)$

2232    16:      $a^* \leftarrow \arg \max_{a' \in \mathcal{A}} \tilde{\omega}^\top Q(s_{t+1}, a'; \tilde{\omega}; \theta)$  ▷ selection Equation 18

2233    17:      $\delta_{\text{act}} \leftarrow \tilde{\omega}^\top \mathbf{r}_t + \gamma(1 - d_t) \tilde{\omega}^\top Q(s_{t+1}, a^*, \tilde{\omega}; \theta^-) - \tilde{\omega}^\top Q(s_t, a_t, \tilde{\omega}; \theta)$  ▷ TD error

2234    18:      $p_{\text{init}} \leftarrow |\delta_{\text{act}}| + \epsilon_0$  ▷ priority Equation 20

2235    19:     append  $((s_t, a_t, \mathbf{r}_t, s_{t+1}, d_t), p_{\text{init}})$  to  $(\mathcal{B}, \mathcal{P})$

2236    20:     **if**  $|\mathcal{B}| = C$  **then**

2237    21:       ADDEXPERIENCES( $\text{SHARD}_k, \mathcal{B}, \mathcal{P}$ ); reset  $\mathcal{B}, \mathcal{P} \leftarrow \emptyset$

2238    22:     **end if**

2239    23:     PERIODICALLY( $\theta \leftarrow \text{LEARNER.PARAMETERS}()$ )

2240    24:      $t \leftarrow t + 1$

2241    25: **end while**

2242    26: **end for**

---

### F.3 CENTRALIZED LEARNER

The learner (i) assembles prioritized mini-batches from  $K$  replay shards, (ii) samples a mini-batch of preferences from a Dirichlet and forms a *Cartesian product* with the transitions, (iii) for each (transition, preference) pair computes an *envelope* DDQN target by maximizing over both actions and a finite set of supporting preferences, (iv) updates the online parameters  $\theta$  using a vector regression objective plus a cosine similarity term, (v) refreshes per-transition priorities on the shards, and (vi) periodically synchronizes the target network  $\theta^-$  and publishes the latest online parameters to actors.

## MINI-BATCH ASSEMBLY FROM REPLAY SHARDS

Let assume shard  $k \in \{0, \dots, K-1\}$  stores  $N_k$  items with priorities  $\{p_{k,i}\}$ , with a total running value  $Z_k \triangleq \sum_{i=1}^{N_k} p_{k,i}^\alpha$ , where  $\alpha \in [0, 1]$ . A pair  $(k, i)$  is sampled with probability

$$\Pr((k, i)) = \frac{Z_k}{\sum_{\ell=0}^{K-1} Z_\ell} \cdot \frac{p_{k,i}^\alpha}{Z_k} = \frac{p_{k,i}^\alpha}{\sum_\ell \sum_j p_{\ell,j}^\alpha}. \quad (22)$$

Given  $N \triangleq \sum_k N_k$ , experience replay importance weights  $w_{k,i}$  (with exponent  $\beta \in [0, 1]$ ) are

$$w_{k,i} = \left( \frac{1}{N} \cdot \frac{1}{\Pr((k,i))} \right)^\beta \Bigg/ \max_{k',i'} \left( \frac{1}{N} \cdot \frac{1}{\Pr((k',i'))} \right)^\beta. \quad (23)$$

The learner queries the shards values  $\{Z_k\}_{k=1}^K$ , allocates per-shard batch sizes proportionally, fetches tuples (transition, index,  $w_k$ , shard id), and aggregates them into a transition batch of size  $B$ .

2268 PREFERENCE SAMPLING  
2269

2270 For each training step, we draw an i.i.d. mini-batch of preferences from a Dirichlet distribution:

2271 
$$\{\omega_j\}_{j=1}^P \subset \Omega = \Delta^{m-1}, \quad \omega_j \stackrel{\text{i.i.d.}}{\sim} \text{Dir}(\alpha), \quad \alpha \in (0, \infty)^m. \quad (24)$$

2272 The choice  $\alpha = \mathbf{1}_m$  yields the uniform distribution on the simplex;  $\alpha < \mathbf{1}$  emphasizes corners, while  
2273  $\alpha > \mathbf{1}$  emphasizes the interior. To couple transitions and preferences, we form the *Cartesian product*  
2274 index set

2275 
$$\mathcal{I} = \{1, \dots, B\} \times \{1, \dots, P\},$$

2276 so every transition is paired with every preference. For  $(i, j) \in \mathcal{I}$ , we write  $(s_i, a_i, \mathbf{r}_i, s'_i, d_i, \omega_j)$  and  
2277 define the scalarization  $Q_{\omega_j}(s, a; \theta) = \omega_j^\top Q(s, a, \omega_j; \theta)$ .

2278 ENVELOPE DDQN SELECTION AND VECTOR TARGET  
2279

2280 The EQL's envelope backup requires maximizing over both actions and supporting preferences.  
2281 Directly optimizing over  $\Omega$  is intractable, so we *approximate* the inner maximization by searching  
2282 over a sampled set  $\mathcal{W} = \{\omega_j\}_{j=1}^P$ . For each pair  $(i, j) \in \mathcal{I}$ , we get:

2283 
$$(a_{i,j}^*, \tilde{\omega}_{i,j}^*) = \arg \max_{a' \in \mathcal{A}, \omega' \in \mathcal{W}} \omega_j^\top Q(s'_i, a', \omega'; \theta), \quad (25)$$

2284 Thus each *query preference*  $\omega_j$  selects a *supporting preference*  $\tilde{\omega}_{i,j}^* \in \mathcal{W}$  and action  $a_{i,j}^*$  that together  
2285 realize the envelope along direction  $\omega_j$ . The online prediction is  $Q_{\text{pred},i,j} = Q(s_i, a_i, \omega_j; \theta) \in \mathbb{R}^m$ .

2286 TRAINING LOSS  
2287

2288 For each  $(i, j) \in \mathcal{I}$ , we define:

2289 
$$\mathcal{L}_{\text{mmse}}(i, j; \theta) = \|\mathbf{y}_{i,j} - Q(s_i, a_i, \omega_j; \theta)\|_2^2, \quad (27)$$

2290 
$$\mathcal{L}_{\text{cos}}(i, j; \theta) = 1 - \frac{\omega_j^\top Q(s_i, a_i, \omega_j; \theta)}{\|\omega_j\|_2 \|Q(s_i, a_i, \omega_j; \theta)\|_2}. \quad (28)$$

2291 With tradeoff  $\lambda \geq 0$  and PER weights  $w_{k,i}$  tied to the *transition* (replicated over its  $P$  preference  
2292 copies), the learner minimizes

2293 
$$\mathcal{L}(\theta) = \frac{1}{BP} \sum_{i=1}^B \sum_{j=1}^P w_{k(i), \text{idx}(i)} (\mathcal{L}_{\text{mmse}}(i, j; \theta) + \lambda \mathcal{L}_{\text{cos}}(i, j; \theta)), \quad (29)$$

2294 where  $k(i)$  and  $\text{idx}(i)$  are the shard id and local index of transition  $i$ .

2295 PRIORITY REFRESH (LEARNER SIDE)  
2296

2297 To refresh priority weights, we first compute scalarized residuals per pair  $(i, j)$ ,

2298 
$$\delta_{i,j} = \omega_j^\top (\mathbf{y}_{i,j} - Q(s_i, a_i, \omega_j; \theta)), \quad (30)$$

2299 then aggregate to a single priority per *original transition*  $i$ ,

2300 
$$p_{\text{new}}(i) = \max_{1 \leq j \leq P} |\delta_{i,j}| + \epsilon_0, \quad \epsilon_0 > 0, \quad (31)$$

2301 and return  $(\text{indices}(i), p_{\text{new}}(i))$  to the corresponding shards to update PER totals.

2302 TARGET UPDATES AND PARAMETER BROADCAST  
2303

2304 Every  $C_{\text{tgt}}$  steps, the learner updates the target network either by a hard copy

2305 
$$\theta^- \leftarrow \theta \quad (32)$$

2306 or a soft update with factor  $\tau \in (0, 1]$ :

2307 
$$\theta^- \leftarrow (1 - \tau) \theta^- + \tau \theta. \quad (33)$$

2308 Every  $C_{\text{push}}$  steps, the latest online parameters are published to the shared model used by all actors.

43

---

2322 **Algorithm 5** D-EQL Learner

---

2323 **Require:** shards  $k \in \{0, \dots, K-1\}$ , discount  $\gamma$ , PER exponents  $(\alpha, \beta)$ , target period  $C_{\text{tgt}}$ , push  
 2324 period  $C_{\text{push}}$ , preference batch size  $P$ , tradeoff  $\lambda$ , Dirichlet parameter  $\alpha$

2325 1: Initialize online  $\theta$ ; set target  $\theta^- \leftarrow \theta$

2326 2: **while** training **do**

2327 3:   Query  $\{Z_k\}$ ; allocate per-shard batch sizes; fetch prioritized transitions with  
 2328   ( $\text{indices}(i), w_{k(i), \text{idx}(i)}$ ) ▷ Equation 22–Equation 23

2329 4:   Sample preferences  $\{\omega_j\}_{j=1}^P \stackrel{\text{i.i.d.}}{\sim} \text{Dir}(\alpha)$  ▷ Equation 24

2330 5:   Form Cartesian product  $\mathcal{I} = \{1..B\} \times \{1..P\}$  (replicate transitions across all  $\omega_j$ )

2331 6:   **for all**  $(i, j) \in \mathcal{I}$  **do**

2332 7:      $(a_{i,j}^*, \tilde{\omega}_{i,j}^*) \leftarrow \arg \max_{a' \in \mathcal{A}, \tilde{\omega} \in \{\omega_1, \dots, \omega_P\}} \omega_j^\top Q(s'_i, a', \tilde{\omega}; \theta)$  ▷ Equation 25

2333 8:      $y_{i,j} \leftarrow r_i + \gamma(1 - d_i) Q(s'_i, a_{i,j}^*, \tilde{\omega}_{i,j}^*; \theta^-)$  ▷ Equation 26

2334 9:   **end for**

2335 10:   Compute  $\mathcal{L}(\theta)$  via Equation 29; take a gradient step on  $\theta$

2336 11:   For each  $i \in \{1..B\}$ , compute  $p_{\text{new}}(i)$  via Equation 30–Equation 31; send updates to shards

2337 12:   **if** step mod  $C_{\text{tgt}} = 0$  **then**

2338     Update  $\theta^-$  via Equation 32 or Equation 33

2339   **end if**

2340 15:   **if** step mod  $C_{\text{push}} = 0$  **then**

2341     Publish  $\theta$  to actors

2342   **end if**

2343 18: **end while**

---

2344 **Notes.** (i) The envelope selection (Equation 25) is the finite-set approximation of the bi-level inner  
 2345 maximization over preferences; letting  $P \uparrow \infty$  densifies the approximation. (ii) The Cartesian product  
 2346 ensures that *every* transition is trained under *every* sampled preference each step (dense supervision).  
 2347 (iii) Priorities are defined per transition by aggregating scalarized TD residuals over  $P$  preference  
 2348 replicas.

2349  
 2350 **Communication summary.** The learner *pulls* batches proportional to  $Z_k$ , *pushes* refreshed priorities  
 2351  $p_{\text{new}}$  (updating shard totals), and periodically *broadcasts* the latest online parameters. Preference  
 2352 expansion couples updates across  $\Omega$ , while DDQN selection/evaluation preserves stability.

2353  
 2354 **F.4 STRATIFIED SAMPLING ON THE PROBABILITY SIMPLEX**

2355  
 2356 We want to sample preferences  $\omega \in \Omega := \Delta^{m-1}$  so that *all regions* of the simplex are adequately  
 2357 covered during data generation, thereby reducing variance and avoiding mode collapse toward a few  
 2358 scalarizations.

2359  
 2360 **Baseline (no stratification).** When no partition is imposed, we draw independent and identically  
 2361 distributed (i.i.d.) preferences from a Dirichlet distribution,

2362 
$$\omega \sim \text{Dir}(\alpha), \quad \alpha \in (0, \infty)^m. \quad (34)$$

2363 Choosing  $\alpha = \mathbf{1}_m$  yields the uniform distribution on  $\Omega$ .

2364  
 2365 **F.4.1 DETERMINISTIC EQUAL-VOLUME STRATA VIA A SIMPLEX LATTICE**

2366 We partition  $\Omega$  into congruent  $(m-1)$ -simplices by using a barycentric lattice with resolution  $L \in \mathbb{N}$ .  
 2367 We define the lattice vertices as:

2368  
 2369 
$$\mathcal{V}_L := \left\{ \frac{\mathbf{k}}{L} \in \Omega : \mathbf{k} = (k_1, \dots, k_m) \in \mathbb{N}^m, \sum_{i=1}^m k_i = L \right\}, \quad (35)$$

2370 and we consider a fixed permutation  $\pi$  of  $\{1, \dots, m\}$ . For each *base* point  $\mathbf{k} \in \mathbb{N}^m$  with  $\sum_i k_i =$   
 2371  $L-1$ , we form the  $m$  lattice points as:

2372  
 2373 
$$\mathbf{v}_0 = \frac{\mathbf{k}}{L}, \quad \mathbf{v}_r = \frac{\mathbf{k} + \mathbf{e}_{\pi(1)} + \dots + \mathbf{e}_{\pi(r)}}{L}, \quad r = 1, \dots, m-1, \quad (36)$$

---

2376 **Algorithm 6** Simplex–Lattice Stratification (build strata at resolution  $L$ )  
2377  
2378 **Require:** dimension  $m$ , resolution  $L$ , a fixed permutation  $\pi$  of  $\{1, \dots, m\}$   
2379 1:  $\mathcal{S} \leftarrow \emptyset$  ▷ list of strata (each as  $m$  vertices in  $\mathbb{R}^m$ )  
2380 2: **for all**  $\mathbf{k} \in \mathbb{N}^m$  with  $\sum_{i=1}^m k_i = L - 1$  **do**  
2381 3:    $\mathbf{v}_0 \leftarrow (\mathbf{k})/L$   
2382 4:   **for**  $r = 1$  to  $m - 1$  **do**  
2383 5:      $\mathbf{v}_r \leftarrow (\mathbf{k} + \mathbf{e}_{\pi(1)} + \dots + \mathbf{e}_{\pi(r)})/L$   
2384 6:   **end for**  
2385 7:   append  $\text{conv}\{\mathbf{v}_0, \dots, \mathbf{v}_{m-1}\}$  to  $\mathcal{S}$   
2386 8: **end for**  
2387 9: **return**  $\mathcal{S}$  ▷  $|\mathcal{S}| = L^{m-1}$  equal-volume strata  
2388

---

2388 **Algorithm 7** Sample uniformly from a stratum  $\Omega_L(\mathbf{k})$   
2389

---

2390 **Require:** vertices  $\{\mathbf{v}_0, \dots, \mathbf{v}_{m-1}\}$  of  $\Omega_L(\mathbf{k})$   
2391 1: draw  $\mathbf{z} \sim \text{Dir}(\mathbf{1}_m)$   
2392 2: **return**  $\boldsymbol{\omega} = \sum_{r=0}^{m-1} z_r \mathbf{v}_r$  ▷ uniform in the stratum  
2393

---

2394 and define the micro-simplex (stratum):

$$\Omega_L(\mathbf{k}) := \text{conv}\{\mathbf{v}_0, \mathbf{v}_1, \dots, \mathbf{v}_{m-1}\} \subset \Omega. \quad (37)$$

2395 The collection  $\{\Omega_L(\mathbf{k}) : \mathbf{k} \in \mathbb{N}^m, \sum_i k_i = L - 1\}$  tiles  $\Omega$  into  $L^{m-1}$  equal-volume strata.  
2396

#### 2398 F.4.2 UNIFORM SAMPLING WITHIN A STRATUM

2400 Let  $\Omega_L(\mathbf{k}) = \text{conv}\{\mathbf{v}_0, \dots, \mathbf{v}_{m-1}\}$  be any stratum. We draw barycentric weights  $\mathbf{z} \sim \text{Dir}(\mathbf{1}_m)$  and  
2401 map affinely:

$$\boldsymbol{\omega} = \sum_{r=0}^{m-1} z_r \mathbf{v}_r \in \Omega_L(\mathbf{k}). \quad (38)$$

2402 This yields a sample *uniform* in the stratum.  
2403

#### 2404 F.4.3 ASSIGNING STRATA TO ACTORS

2405 Index the  $L^{m-1}$  strata in a fixed order as  $\{\Omega_L^{(u)}\}_{u=0}^{U-1}$  with  $U = L^{m-1}$  (or group them when  $U$   
2406 exceeds the number of actors). Actor  $u$  repeatedly samples  $\boldsymbol{\omega} \in \Omega_L^{(u)}$  via Algorithm 7, ensuring  
2407 non-overlapping coverage across actors.  
2408

#### 2409 F.4.4 DISCUSSION AND ALTERNATIVES

2410 **Coverage and variance.** Compared to i.i.d. Dirichlet sampling shown in Equation 34, the lattice  
2411 partition yields systematic coverage of the entire simplex and reduces estimator variance by ensuring  
2412 that each subregion is represented.  
2413

2414 **Resolution.** Larger  $L$  gives finer strata ( $L^{m-1}$  pieces) and smoother coverage at the cost of more  
2415 partitions to manage.  
2416

2417 **Clustering alternative.** When equal-volume strata are unnecessary, a simple alternative is to draw  
2418 a large pilot set  $\{\boldsymbol{\omega}^{(n)}\}_{n=1}^{N_0} \sim \text{Dir}(\boldsymbol{\alpha})$  and run  $k$ -means on the  $(m-1)$ -dimensional simplex (with  
2419 cosine or Euclidean distance); the Voronoi cells of the cluster centers define strata. Sampling within a  
2420 cell can be done by re-running Dirichlet draws and accepting points whose nearest center matches  
2421 the cell (approximately uniform within each cell).  
2422

### 2423 F.5 EXPERIMENTS

#### 2424 F.5.1 ENVIRONMENTS AND SETUP

2425 We evaluate our approach using two well-established MORL benchmarks: deep sea treasure (DST)  
2426 and fruit tree navigation (FTN). Both environments are widely used in the literature (Yang et al.,  
2427

2430 2019; Basaklar et al., 2023), providing standardized testbeds for assessing Pareto front coverage and  
 2431 preference generalization.

2432 The DST environment is a grid-world wherein an agent controls a submarine that must navigate from  
 2433 the surface to collect one of several treasures placed at different depths. Each treasure yields a two-  
 2434 dimensional reward: a positive value and a negative time penalty. The task is inherently multi-  
 2435 objective, requiring the agent to balance collecting high-value treasures against minimizing travel time.  
 2436 The Pareto front is well understood and serves as a reliable benchmark for coverage and accuracy.

2437 The FTN environment generalizes this idea to a tree-structured setting. Starting from the root, the  
 2438 agent makes sequential decisions until it reaches a leaf node, where it receives a multi-dimensional  
 2439 reward corresponding to the chosen fruit. The tree depth controls task complexity. At depth five,  
 2440 the agent makes five sequential decisions; at depth seven, the number of possible outcomes grows  
 2441 exponentially, producing a much larger and more diverse Pareto front. This makes FTN with higher  
 2442 depth a significantly more challenging benchmark, particularly for algorithms that must adapt to  
 2443 unseen preferences or maintain wide coverage.

2444 In all experiments, algorithms are trained across a range of randomly sampled linear preference  
 2445 vectors, following the setup in Yang et al. (2019). At test time, additional preference vectors are  
 2446 sampled to assess generalization. All results are averaged over multiple random seeds to account for  
 2447 variance.

### 2449 F.5.2 METRICS AND RESULTS

2450 We evaluate performance using three widely adopted metrics in MORL:

- 2453 1. **Coverage Ratio F1 (CRF1):** A harmonic mean of precision and recall that captures both  
 2454 accuracy and coverage of the Pareto front.
- 2455 2. **Hypervolume:** The volume dominated by the obtained solutions with respect to a reference  
 2456 point, reflecting both the quality and diversity of the Pareto front.
- 2457 3. **Sparsity:** The average distance between neighboring solutions, indicating how uniformly the  
 2458 Pareto front is covered.

2460  
 2461 Table 7: We compare the performance of the proposed distributed Q-learning algorithm with the  
 2462 original Q-learning approach Yang et al. (2019) and Basaklar et al. (2023) in terms of CRF1,  
 2463 hypervolume and sparsity, showing superior performance in more complex scenarios.

	Deep Sea Treasure			Fruit Tree Nav. (d=5)		Fruit Tree Nav. (d=7)	
	CRF1	Hyperv.	Sparsity	CRF1	Hyperv.	CRF1	Hyperv.
Yang et al. (2019)	0.994	227.39	2.62	1.0	6920.58	0.819	6395.27
Basaklar et al. (2023)	1.0	241.73	1.14	1.0	6920.58	0.920	11419.58
D-EQL (ours)	<b>1.0</b>	<b>241.73</b>	<b>1.14</b>	<b>1.0</b>	<b>6920.58</b>	<b>1.0</b>	<b>12110.74</b>

2471 Table 7 compares D-EQL with prior approaches. On the simpler DST domain, all methods achieve  
 2472 near-perfect coverage. Our method, on par with Basaklar et al. (2023), attains CRF1 = 1.0 and  
 2473 simultaneously achieves the highest hypervolume and lowest sparsity, indicating comprehensive and  
 2474 well-distributed solutions along the Pareto front.

2475 On FTN with depth five, performance saturates across all methods with perfect coverage and identical  
 2476 hypervolume, reflecting the relative simplicity of this setting.

2477 The advantage of D-EQL becomes evident in the more complex FTN with depth seven. Prior methods  
 2478 show a notable drop in coverage (CRF1 = 0.819 for Yang et al. (2019) and 0.920 for Basaklar et al.  
 2479 (2023)), whereas D-EQL maintains perfect coverage (CRF1 = 1.0). Moreover, D-EQL achieves the  
 2480 highest hypervolume (12110.74), showing 22.1% and 8.69% improvements over Yang et al. (2019)  
 2481 and Basaklar et al. (2023), respectively, and demonstrating broader Pareto front coverage and superior  
 2482 solution diversity. These results show that our distributed training scheme preserves accuracy while  
 2483 scaling effectively to environments with exponentially growing outcome spaces.

2484 Overall, D-EQL achieves state-of-the-art performance: it matches existing methods on simpler tasks  
 2485 and clearly outperforms them in complex scenarios, highlighting the benefits of distributed training  
 2486 for multi-objective reinforcement learning. This makes D-EQL a more suitable for handling the vast  
 2487 dimensions of state-action spaces in RAN control functions.  
 2488

2489 **F.5.3 HYPERPARAMETERS**

2490 The hyperparameters used in D-EQL training for DST and FTN environment are listed in Table 8.  
 2491

2492 Table 8: Hyperparameters used for D-EQL on deep sea treasure and fruit tree navigation.  
 2493

2494 <b>Hyperparameter</b>	2495 <b>DST</b>	2495 <b>FTN</b>
<b>Model</b>		
2497 Hidden feature (units per layer)	256	512
2498 Activation	SiLU	SiLU Elfwing et al. (2018)
2499 Number of layers	3	3
<b>Actor</b>		
2502 Number of actors	10	10
2503 $\epsilon$ -greedy (linear, start $\rightarrow$ final)	$0.8 \rightarrow 0.1$	$0.8 \rightarrow 0.1$
2504 Anneal timesteps	$1 \times 10^6$	$1 \times 10^6$
2505 Local buffer capacity	125	125
2506 Max environment step	$1 \times 10^6$	$1 \times 10^6$
<b>Learner</b>		
2508 Learning rate	$3.5 \times 10^{-4}$	$3.5 \times 10^{-4}$
2509 Target update period (gradient step)	1	1
2510 Model sync period (gradient step)	250	250
2511 Discount factor $\gamma$	0.99	0.99
2512 Prefetched batches	16	16
2513 Transition batch size	128	128
2514 Preference batch size	128	128
<b>Replay Memory (PER)</b>		
2516 Number of shards	1	1
2517 Capacity	$5 \times 10^5$	$5 \times 10^5$
2518 Priority exponent $\alpha$	0.7	0.7
2519 Importance sampling $\beta$ (linear, start $\rightarrow$ final)	$0.4 \rightarrow 1.0$	$0.4 \rightarrow 1.0$
2520 $\beta$ anneal timesteps	$2 \times 10^4$	$2 \times 10^4$

2522  
 2523  
 2524  
 2525  
 2526  
 2527  
 2528  
 2529  
 2530  
 2531  
 2532  
 2533  
 2534  
 2535  
 2536  
 2537

2538 **G CASE STUDY**  
25392540  
2541 We design a controller agent for LA, a crucial functionality of modern wireless communication  
2542 systems that employs adaptive coding and modulation to optimize the spectral efficiency of the radio  
2543 link between transmitter and receiver. By adopting a MORL approach, the LA controller agent can  
2544 adjust transmission parameters to meet connectivity service intents expressed in terms of data rate,  
2545 reliability, and latency requirements for individual users.  
25462547 **G.1 LINK ADAPTATION**  
25482549  
2550 LA adapts the modulation order and code rate of individual packet transmissions to match the capacity  
2551 of the radio link capacity, given the radio link state. The LA parameters are encoded into a unique  
2552 value, referred to as MCS index in 3GPP (2025e), that is provided to the receiver for packet decoding.  
25532554 The 3GPP 5<sup>th</sup> Generation (5G) New Radio (NR) system, rely on an OLLA approach inspired  
2555 to Pedersen et al. (2007) to maximize the link spectral efficiency while adhering to a predefined  
2556 BLER target using receiver-side channel state information (CSI), such as channel quality indicator  
2557 (CQI) 3GPP (2025e), and hybrid automatic repeat request (HARQ) feedback—a 1-bit information  
2558 indicating whether a prior packet transmission was successful or not. While this approach suits  
2559 best-effort traffic, its reliance on long communication sessions to converge makes is suboptimal to  
2560 address connectivity service intents under more general conditions, such as short bursty traffic, fast  
2561 channel aging, medium-high user mobility, etc.  
25622563 A MORL approach instead enables to dynamic LA toward selection transmission parameters that best  
2564 align with different service intents. For example, selecting MCS conservatively—e.g., lower modula-  
2565 tion orders such as Quadri-Phase Shift Keying (QPSK) or reduced code rates—favor robustness by  
2566 lowering the probability of decoding errors. This enables to achieve highly reliable transmissions at  
2567 the cost of throughput, since more time-frequency resource element (RE) are required per information  
2568 bit. Conversely, an aggressive MCS selection can push spectral efficiency closer to or even beyond  
2569 the instantaneous link capacity, exploiting retransmissions to increase data rate and reduce latency  
2570 for best-effort traffic. However, overly aggressive choices may lead to excessive retransmissions and  
2571 throughput degradation. By explicitly balancing these conflicting objectives, MORL allows LA to  
2572 adapt beyond fixed BLER-driven policies, supporting a wider range of connectivity intents.  
25732574 **G.2 MOMDP DESIGN FOR LINK ADAPTATION**  
25752576 Our goal is to train a single pareto efficient uniform model (PEUMO) for LA to learn the Pareto  
2577 frontier outlining the optimal trade-off between the utilization of radio resources and the amount of  
2578 information bits delivered by a packet transmission.  
25792580 As LA and HARQ operate on a per-user equipment (UE) and per-packet transmission basis, we  
2581 formulate this problem as an episodic MOMDP  $\mathcal{M} = \langle \mathcal{S}, \mathcal{A}, p, \mathbf{r}, \Omega, \gamma, \rho_0 \rangle$ , where  $\mathcal{S}$  denotes the  
2582 state space,  $\mathcal{A}$  the action space,  $p(s' | s, a)$  the transition dynamics,  $\mathbf{r} : \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}^K$  a multi-  
2583 dimensional reward vector,  $\Omega \subseteq \mathbb{R}^K$  the reward preference space,  $\gamma \in [0, 1]$  a discount factor, and  $\rho_0$   
2584 the initial state distribution.  
25852586 An episode models the lifespan of a UE packet in the HARQ process—from its first transmission to  
2587 either a successful reception or the packet being dropped upon  $N$  transmission attempts, as illustrated  
2588 in Figure 8. This enables us to train a single RL policy from the collective experience generated  
2589 by any UEs across the network. A transition in the episode represents the duration of a packet  
2590 transmission in the HARQ process, from the selection of LA parameters (i.e., the action) to the  
2591 reception of the associated HARQ feedback, i.e., an positive acknowledgment (ACK) or negative  
2592 acknowledgment (NACK) for successful or failed transmission, respectively. For instance, the 3GPP  
2593 5G NR system, used in our evaluations, supports at most four packet retransmissions. Hence, the  
2594 episode length  $N$  may range from one to five steps. Each step is characterized by a state, an action,  
2595 and an associated reward and preference vectors, as presented next.  
2596

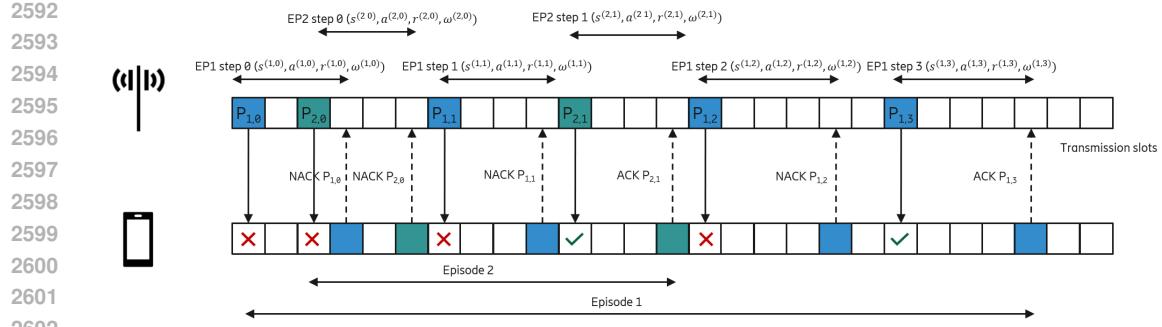


Figure 8: Example of MOMDP episodes modelling the downlink LA and HARQ process. Here, the notation  $P_{i,n}$  denotes the  $n$ -th transmission attempt of the  $i$ -th data packet.

### G.3 ACTION SPACE

The action space consists of the set of MCS index values supported by a communication standard, i.e.,  $\mathcal{A} = \{a_m \mid a_m = m, m = 0, \dots, M - 1\}$ . Therefore, an action  $a_m \in \mathcal{A}$  implicitly provides a combination of modulation order, code rate, and spectral efficiency to be used to transmit a packet. The 5G NR system used in our evaluations in Section 7 supports  $M = 28$  MCS index values as specified in Table 5.1.3.1-1 or Table 5.1.3.1-2 in 3GPP (2025e), corresponding to modulation orders up to 64QAM and 256QAM, respectively.

For a new packet transmission, the selection of an MCS index, combined with the time-frequency resources allocated by a scheduler, determines the amount of information bits, i.e., the TBS, to be transmitted. Packet re-transmissions, however, reuse the TBS value of the original transmission, as no new information is transmitted. A packet re-transmission, however, may occur with a different MCS index therefore resulting in possibly a different amount of radio resources.

### G.4 REWARD VECTOR AND PREFERENCE SPACE

We design a two-dimensional reward function  $\mathbf{r} = [r_1, r_2]^\top \in \mathcal{S} \times \mathcal{A} \in \mathbb{R}^2$  with two competing components:  $r_1$  representing the amount of information bits successfully carried by a packet; and  $r_2$  denoting the cost, in terms of time-frequency resource, incurred in each individual transmission of the packet. Specifically, for each transmission attempt  $n$  of a packet, we define the reward function as

$$\mathbf{r}^{(n)}(s, a) = \begin{cases} \begin{bmatrix} 0 \\ -\frac{N_{RE}^{(n)}}{N_{RE}^{\max}} \end{bmatrix} & \text{if transmission fails at } n\text{-th attempt,} \\ \begin{bmatrix} \frac{TBS}{N_{RE}^{\max}} \\ -\frac{N_{RE}^{(n)}}{N_{RE}^{\max}} \end{bmatrix} & \text{if transmission succeeds at } n\text{-th attempt,} \end{cases} \quad (39)$$

where the TBS and  $N_{RE}^{(n)}$  denote the number of information bits carried by the packet and the number of RE used for the  $n$ -th transmission attempt, respectively, and  $N_{RE}^{\max}$  is the maximum number of REs available, given the system bandwidth. Scaling the reward components by  $N_{RE}^{\max}$  has a twofold purpose: Firstly, it keeps each component within similar range of values, while preserving the functional relation between the MCS index selected to transmit TBS information bits and the required number of time-frequency RE. This relation is specified by communication standards, as in the 3<sup>rd</sup> Generation Partnership Project (3GPP) technical specification (TS) 38.211 3GPP (2025d). Secondly, it makes the reward design agnostic to the system bandwidth, with  $\frac{TBS}{N_{RE}^{\max}}$  representing the spectral efficiency for transmitting TBS bits using the entire system bandwidth. This allows us to employ domain randomization in training (cf. appendix H.1) to improve model generalization over the RAN environment.

Therefore, for each packet transmission attempt  $n$ , the first reward component takes value  $r_1^{(n)} = \frac{TBS}{N_{RE}^{\max}}$  if the transmission is successful or  $r_1^{(n)} = 0$  otherwise. The second reward component, on the

2646 other hand, always indicates the resource cost incurred at the current transmission attempt  $n$ , i.e.,  
 2647  $r_2^{(n)} = -\frac{N_{RE}^{(n)}}{N_{RE}^{\max}}$ , regardless of whether the transmission attempt succeeds or fails.  
 2648

2649 **G.5 STATE DESIGN**

2650 A key goal of our design is to achieve model generalization across diverse RAN environments,  
 2651 enabling a single MORL model to operate reliably under different deployments and radio conditions.  
 2652 To this end, we construct a rich state space  $\mathcal{S} \subseteq \mathbb{R}^K$  and apply domain randomization in training Igl  
 2653 et al. (2019). To model the state space  $\mathcal{S}$  for link adaptation, we follow Demirel et al. (2025)  
 2654 which considers a deep Q-network (DQN) approach for LA with a single, fixed reward design based  
 2655 on the link spectral efficiency. In particular, we model  $\mathcal{S}$  using two types of features: (a) semi-  
 2656 static information characterizing the network deployment surrounding the UE; (b) and information  
 2657 describing observable link dynamics relevant to infer LA parameters.

2658 Semi-static information characterizing the network deployment may include, for instance, deployment  
 2659 type (e.g., rural, urban, dense urban, etc.), location, orientation, relationships among network sites  
 2660 or radio cells, as well as technology configurations, such as whether the system operated in time-  
 2661 duplex or full duplex mode, carrier frequency, system bandwidth, transmit power, antenna array  
 2662 type, etc. On the other hand, information characterizing the dynamics of LA consist of real-time  
 2663 observation (measured in a milliseconds timescale), such as channel state information, HARQ  
 2664 feedback, measurement of path loss, data buffer state, historical actions, and more. We refer  
 2665 to Demirel et al. (2025) for a complete description of the state features.

2666  
 2667  
 2668  
 2669  
 2670  
 2671  
 2672  
 2673  
 2674  
 2675  
 2676  
 2677  
 2678  
 2679  
 2680  
 2681  
 2682  
 2683  
 2684  
 2685  
 2686  
 2687  
 2688  
 2689  
 2690  
 2691  
 2692  
 2693  
 2694  
 2695  
 2696  
 2697  
 2698  
 2699

## 2700 H EXTENDED EXPERIMENTAL EVALUATION

2701  
 2702 This Appendix extends the discussion and empirical evaluation presented in Section 7 with additional  
 2703 results. We organized the material as follows: Appendix H.1 and H.2 describe the network simulator  
 2704 environment and training setup for the MORL controller agent; Appendix H.3 extend the analysis of  
 2705 the controller agent presented in Section 7.1, including an additional scenario with two communication  
 2706 services. Finally, Appendix H.4 extends our analysis of the intent fulfillment loop.

### 2708 H.1 NETWORK SIMULATOR ENVIRONMENT

2709  
 2710 We train and evaluate the MORL controller agent using a high-fidelity, event-driven system-level  
 2711 simulator compliant with the 3GPP 5G NR specifications. Each rollout simulation models a heteroge-  
 2712 neous multi-cell RAN operating in time division duplexing (TDD) mode with single-user multiple  
 2713 input multiple output (SU-MIMO) transmission. The carrier frequency is set to 3.5 GHz, and the  
 2714 physical layer follows the orthogonal frequency division multiplexing (OFDM) numerology  $\mu = 0$   
 2715 specified in 3GPP TS 38.211 (cf. Table 4.2-1 3GPP (2025d)).

2716 To improve model generalization across diverse RAN deployments and radio environments, we apply  
 2717 domain randomization across multiple network characteristics, summarized in Table 9. Each sim-  
 2718 ulation consists of three tri-sector radio sites, randomly configured as either conventional multiple  
 2719 input multiple output (MIMO) or massive multiple input multiple output (mMIMO), with antenna  
 2720 attributes defined in Table 9. Site-level parameters such as location, cell radius, system bandwidth,  
 2721 and downlink transmit power are also randomized by sampling values from the same parameter set.

2722 The training scenario is further diversified by randomizing cell load, traffic type, UEs, and receiver  
 2723 configuration. UEs are generated with a mixture of full buffer (FB) and enhanced mobile broadband  
 2724 (eMBB) traffic, randomly placed in the simulated area according to one of the indoor/outdoor  
 2725 probability distributions in Table 9. Each eMBB UE generates traffic with variable packet size and  
 2726 inter-arrival times, modeled using empirical distributions derived from field measurement campaigns.

2727 Finally, individual UEs are randomized in terms of antenna configuration, mobility (speed), and  
 2728 receiver implementation. The latter accounts for manufacturer-specific differences in hardware (e.g.,  
 2729 antenna arrays and chipsets) and internal algorithms (e.g., CSI estimation), which influence perceived  
 2730 radio conditions. This randomized environment ensures that the MORL controller agent is trained  
 2731 under various realistic network conditions, thus improving its ability to generalize to unseen scenarios.

### 2732 H.2 TRAINING SETUP

2733  
 2734 We train the MORL LA controller agent using our D-EQL algorithm, described in Appendix F,  
 2735 with a single GPU and 560 CPU cores. The learner uses Adam optimizer (Kingma & Ba, 2017)  
 2736 with a learning rate of  $5 \times 10^{-5}$ , weight decay of  $0.02/512$ , and default momentum terms  $(\beta_1, \beta_2) =$   
 2737  $(0.9, 0.999)$ , and mean squared error (MSE) loss. He initialization is used for all network parameters.  
 2738 A soft target update policy is applied with an update factor of 0.001 and a period of one timestep. The  
 2739 model synchronization period is 200 gradient iterations, and training begins after 50,000 timesteps. To  
 2740 reduce communication overhead between the learner and replay memory, 16 batches are prefetched  
 2741 per cycle. Experience and preference batches contain 512 and 128 samples, respectively.

2742 The actor subsystem consists of 40 CPU-based rollout workers, each interacting with 14 parallel  
 2743 simulations (one CPU core per simulation), resulting in efficient experience generation. Each actor  
 2744 collects about 112 samples per second, for a total of roughly 279 million over the training horizon.  
 2745 The learner processes about 27,500 samples per second for gradient updates. Each actor maintains a  
 2746 local buffer of 2,500 samples and follows a linear epsilon-greedy strategy, decaying  $\epsilon$  from 0.8 to  
 2747 0.05 over 5.5 million timesteps. The agent operates with a discount factor of 1.0. Training throughput  
 2748 is about 53.8 batches per second, with each batch containing 65,536 samples.

2749 Replay memory is organized as a single module with four independent shards, each capable of storing  
 2750 four million samples. Each shard has a fixed communication path to a designated learner shard,  
 2751 minimizing cross-shard delays. Prioritized experience replay is employed with parameters  $\alpha = 0.6$   
 2752 and  $\beta = 0.4$  to improve sample efficiency. In total, the system runs 11,200 simulations under different  
 2753 random seeds to ensure reproducibility across diverse network conditions. Communication details of  
 the distributed system are further discussed in Appendix I.

2754 Table 9: RAN environment simulation parameters for domain randomization during training.  
2755

2756 <b>Parameter</b>	2757 <b>Value range</b>	2758 <b>Description</b>
Duplexing type	TDD	Fixed
Carrier frequency	3.5 GHz	Fixed
Deployment type	3-site 9-sector	
Site type	{MIMO, mMIMO}	Randomized
Antenna array	1x2x2 MIMO (4) 8x4x2 mMIMO (64)	Fixed Fixed
Cell radius	{166, 300, 600, 900, 1200} m	Randomized
Bandwidth	{20, 40, 50, 80, 100} MHz	Randomized
Number of sub-bands	{20, 106, 133, 217, 273}	Randomized
DL TX power	{20, 40, 50, 80, 100} W	Randomized
UE antennas	{2, 4}	Randomized
Maximum TX rank	{2, 4}	As per UE ant.
Maximum DL TX	5	Fixed
UE traffic type	{FB, eMBB}	Randomized
Number FB UEs	{1, 5, 10}	Randomized
Number eMBB UEs	{0, 10, 25, 50, 100, 200, 300}	Randomized
Speed UE FB	{0.67, 10, 15, 30} m/s	Randomized
Speed UE eMBB	{0.67, 1.5, 3} m/s	Randomized
UE receiver types	{type0, type1, type2, type3}	Randomized
Indoor probability	{0.2, 0.4, 0.8}	Randomized

2777 Furthermore, we explore a preference space  $\Omega = \Delta^1 \triangleq \{\omega \mid \omega = [\omega, 1 - \omega]^\top, \omega \in [0, 1]\}$  defined  
2778 for the two-dimensional reward in (39). The preference space is partitioned into strata, and each  
2779 actor is assigned to explore a different stratum. Preferences are then sampled from the corresponding  
2780 strata following the procedure in Algorithm 6 (stratum construction) and Algorithm 7 (stratum-based  
2781 sampling). Further details on all hyperparameters used in training are summarized in Appendix I.  
2782

### 2783 H.3 TESTING THE MORL LA CONTROLLER AGENT

#### 2784 H.3.1 SINGLE CONNECTIVITY SERVICE

2785 We extend the analysis presented in Section 7.1 by further evaluating the MORL-based LA controller  
2786 for a single connectivity service: video streaming users. Focusing on a single user class simplifies the  
2787 analysis of the Pareto front achievable by the MORL controller agent.  
2788

2789 Figure 9 shows the Pareto front defined by the two-dimensional reward function in Equation (39)  
2790 for a 3-cell deployment with 10 streaming users. Each point on the frontier is obtained from 480  
2791 independent simulations, where each radio cell employs the MORL-based LA controller with a fixed  
2792 preference value  $\omega \in [0, 1]$ . The parameter  $\omega$  determines the trade-off between minimizing radio  
2793 resources required for packet transmission and maximizing the transmitted payload size. Moving  
2794 along the frontier results in different performance trade-offs in network KPIs, as detailed in Figure 10.  
2795

2796 At the top-right corner of Figure 9 and 10, large values of  $\omega$  prioritize payload maximization (i.e.,  
2797 high TBS), but at the cost of excessive radio resource consumption. In this case, the controller agent  
2798 selects overly aggressive MCS values relative to the channel state (see Figure 11a), targeting spectral  
2799 efficiencies beyond channel capacity. This leads to frequent transmission failures (BLER  $\approx 60\%$ ,  
2800 see Figure 11b) and numerous retransmissions, yielding suboptimal throughput and spectral efficiency.  
2801

2802 In contrast, when  $\omega \approx 0$  (bottom-left corner of Figure 9 and 10), the controller favors conservative  
2803 MCS choices (see Figure 11a), targeting spectral efficiencies well below channel capacity. Although  
2804 this results in low resource utilization and highly reliable transmissions (BLER  $\approx 0\%$ , see Figure 11b),  
2805 it under-utilizes favorable channel conditions by employing low modulation orders and code rates.  
2806 Thus, the system fails to deliver higher payloads, limiting throughput and spectral efficiency.  
2807

2808 Overall, throughput and spectral efficiency peak at  $\omega \approx 0.34$  and  $\omega \approx 0.5$ , respectively. Beyond  
2809 these values, throughput declines more rapidly than spectral efficiency due to the rising BLER, which

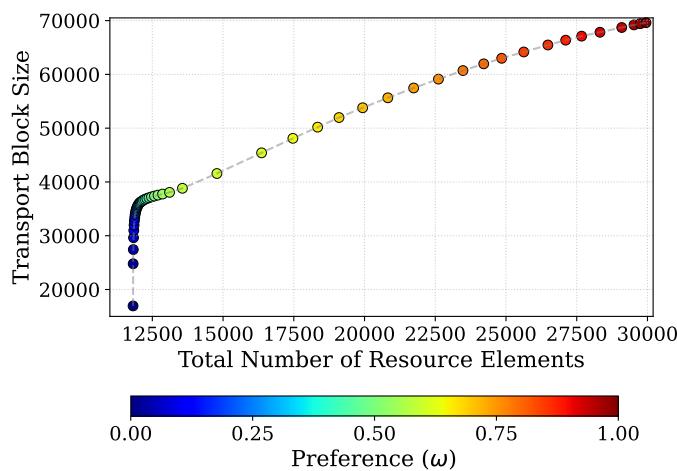


Figure 9: **Pareto front illustrating the trade-off between transport block size and resource utilization.** The Pareto front captures the relationship between the transport block size (vertical axis) and the total number of resource elements (horizontal axis) across a range of system configurations. Each point represents an outcome from 480 independent simulations, computed using distinct preference vectors  $\omega \in [0, 1]$ , and is color-coded by the corresponding preference weight.

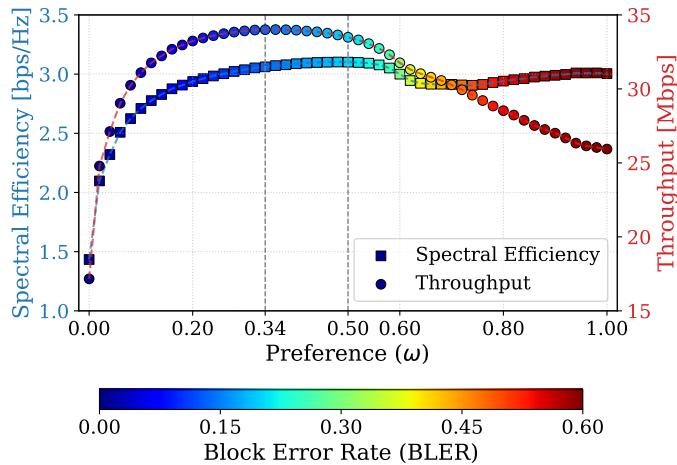
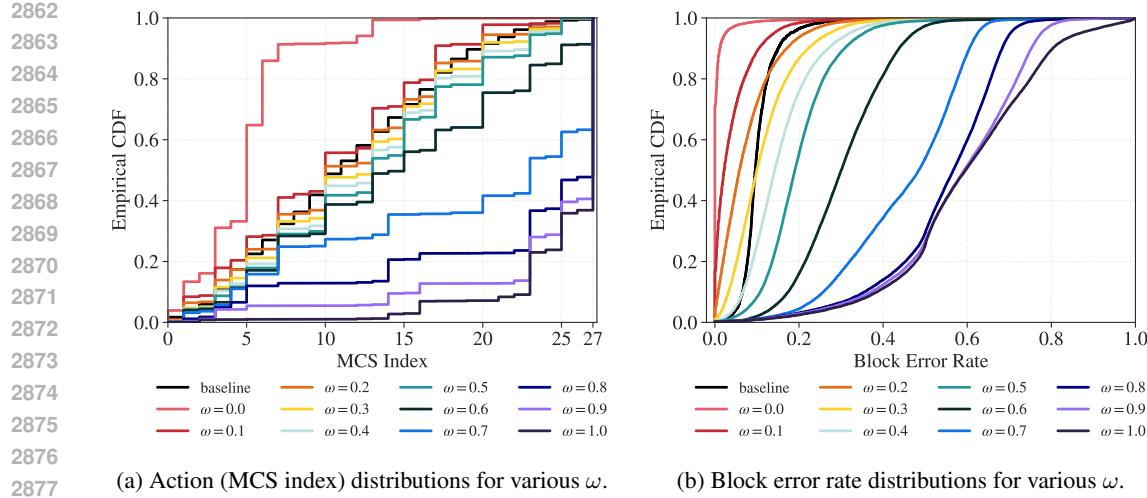


Figure 10: **Joint characterization of spectral efficiency and throughput under varying preference weights  $\omega$ , with BLER-encoded performance.** The figure presents the trade-off between spectral efficiency (squares, left axis) and throughput (circles, right axis) as a function of the preference weight  $\omega$ , which governs the optimization objective. Data points are color-coded based on BLER, with cooler hues indicating lower error rates. Dashed vertical lines denote peaks in the performance trends.

reduces transmission reliability. These results highlight the intrinsic tension between maximizing data rate and maintaining reliability in link adaptation, highlighting how a MORL controller agent can be deployed to provide differentiated connectivity services.

Figure 11 further illustrates how the controller policy changes by selecting different preference values. In particular, Figure 11a illustrate the action (i.e., MCS index) distribution induced by different preference values  $\omega$ . For example, it clearly highlights how small values of  $\omega$  induce a link adaptation policy that selects overly conservative MCS index values relative to the channel state, thus aiming for transmissions with low spectral efficiency (i.e., characterized by low modulation order and code rate). Although this makes the transmission very robust, as demonstrated by the corresponding BLER distribution in Figure 11b, such policy leads to low data throughput.

Figure 11: Controller agent behavior for different preference values  $\omega$ .

Conversely, Figure 11a also shows how large preference values  $\omega \approx 1$  induce a link adaptation policy that selects overly aggressive MCS index values relative to the channel state, thus aiming for transmissions with too high spectral efficiency (i.e., characterized by high modulation order and code rates). This makes the transmissions over-the-air unreliable, as demonstrated by the corresponding BLER distribution in Figure 11b.

### H.3.2 MULTI CONNECTIVITY SERVICES WITH QoS DIFFERENTIATION

We consider a practical scenario with two service applications with distinct QoS profiles concurrently sharing the resources of a radio cell: *real-time gaming* and *web browsing* users. Table 10 characterizes their QoS profile in terms of purpose, service type, differentiated services code point (DSCP) value, 5G QoS identifier (5QI) value and QoS features.

Table 10: QoS Profile Comparison: Real-time gaming vs Web Browsing

Aspect	Real-time gaming	Web browsing
<b>Purpose</b>	Real-time, delay-sensitive traffic	Delay-tolerant, no bandwidth guarantees
<b>Service type</b>	Expedited forwarding (EF)	Best effort (BE)
<b>DSCP value</b>	EF (46)	BE (0)
<b>5QI</b>	3	9
<b>QoS features</b>	Guaranteed bit rate (GBR) Ultra-low latency Low jitter Packet delay budget (PDB) $\approx 50$ ms packet error rate (PER) $\approx 1 \times 10^{-3}$	Non-guaranteed bit rate (non-GBR) No strict latency No strict jitter PDB $\approx 300$ ms

**Real-time gaming** traffic consists of continuous, high-frequency bidirectional streams, often transmitted over user datagram protocol (UDP) based protocols to support real-time video rendering and user input feedback. This type of traffic demands substantially higher bitrates (ranging from 5 to 25 Mbps), ultra-low latency, and minimal jitter to maintain responsive and seamless game-play. As such, real-time gaming is classified as GBR traffic and expedited forwarding service, necessitating stringent QoS settings, including 5QI = 3 and DSCP values like EF (46), corresponding to a PDB of  $\approx 50$  ms and PER  $\approx 10^{-3}$  (cf. Table 5.7.4-1, 3GPP (2025b)).

**Web browsing** traffic is instead elastic, delay-tolerant, and bursty, following a request-response model (like the HyperText Transfer Protocol (HTTP)) over reliable transmission control protocol (TCP) connections. It generally demands low to moderate bitrates (typically below 1 Mbps) and is relatively

Table 11: RAN environment simulation parameters.

Load scenario	Number of gaming users		Web users arrival rate		Performance KPIs
	Indoor	Outdoor	Indoor	Outdoor	
<b>Low</b>	12	6	3.15	1.35	Figure 12
<b>Medium</b>	24	12	6.3	2.7	Figure 13
<b>High</b>	48	24	12.6	5.4	Figure 14
<b>Very high</b>	72	36	18.9	8.1	Figure 15

insensitive to latency and jitter, making it tolerant of network delays. Web browsing is typically classified as non-GBR traffic, associated with 5QI = 9 and DSCP values such as BE (0), corresponding to a PDB of  $\approx 300$  ms and PER  $\approx 10^{-6}$ .

**The evaluation scenario** consists of a dense-urban deployment comprising three 3-sector sites operating at 3.5 GHz with a 100 MHz bandwidth with inter-site distance of 167 meters to ensure full uplink coverage across the simulation area. Traffic is predominantly downlink-oriented, with minimal uplink activity. Real-time gaming users remain active throughout the simulation duration, whereas web browsing users follow a Poisson arrival process with a distribution modeled to fit realistic field data patterns and depart the simulation upon completing their downloads (e.g., webpage, email, etc.).

Unlike the single-service application considered in Appendix H.3.1, the controller agent here applies a different preference vector to each service application:  $\omega_g = [\omega_g, 1 - \omega_g]^T$  and  $\omega_w = [\omega_w, 1 - \omega_w]^T$ . Like before, we analyze how shifting the MORL controller policy along Pareto front defined by the two reward components in Equation (39), by tuning  $\omega_g$  or  $\omega_w$ , produces different trade-offs in various performance KPIs. Under these settings, the values achievable for a performance KPI  $g(\cdot)$  of each service application depends on both preference vectors, i.e.,  $g_g = g_g(\omega_g, \omega_w)$   $g_w = g_w(\omega_g, \omega_w)$ .

Figure 12 to Figure 15 present the achievable user experience for real-time gaming and web browsing services in terms of three KPIs that closely relate to their QoS profile: user throughput, latency, and BLER. Each figure refers to one of the four traffic load scenarios, with a mixture of indoor and outdoor users, summarized in Table 11. Each figure also depicts the average MCS value selected by the MORL controller, showing how the controller agent applies a different policy to each service application for different combinations of preference vectors  $\omega_g$  or  $\omega_w$  and network load conditions.

For example, let us analyze the throughput distributions for real-time gaming uses (Figure 12a to Figure 15a) and web browsing user (Figure 12b to Figure 15b) for the various scenarios. For low and medium low load conditions, cf. Figure 12a-12b and Figure 13a-13b, respectively the mean throughput distribution of the two services shows similarities due to the abundance of radio resources compared to traffic load. The difference in mean throughput magnitude between the two type of services (i.e., Mbps vs Kbps) can be explained by the difference in traffic: continuous video streaming vs sporadic downloads of small packets.

At high and very high load conditions, cf. Figure 14a-14b and Figure 15a-15b, respectively, the two distributions of throughput start showing significant differences, clearly revealing how each service achieves the best mean throughput with different combinations of preference values  $(\omega_g, \omega_w)$ . As the traffic load becomes very high, the region of preference values  $(\omega_g, \omega_w)$  that optimizes the throughput of each service shrinks into a smaller and well defined area. Furthermore, since in these scenarios more users share the same amount of radio resources, both services achieve lower throughput.

Similar trends can be observed for latency (expressed as round-trip time (RTT) for real-time gaming users and as webpage load time for web browsing users, respectively), and block error rate. The conditions observed in different service KPI in Figure 12 to Figure 15 can be related to the constraints  $g_i(\omega) \leq b_i$  that can be required to be fulfilled by a service intent in the optimization problem (1) solved by the optimizer agent to dynamically adapt the preference vectors for each service applications. For instance, in Section 7.3 we presented an example with a video streaming service requiring a minimum of 7 Mbps per active user (i.e.,  $g_{i,thr}(\omega) \geq 7$ ). This threshold is rated as good for most real-time gaming applications at 720p and 1080p resolutions, and excellent for video streaming, given that typical requirements range from 5 Mbps for HD to 15 Mbps for 4K content.

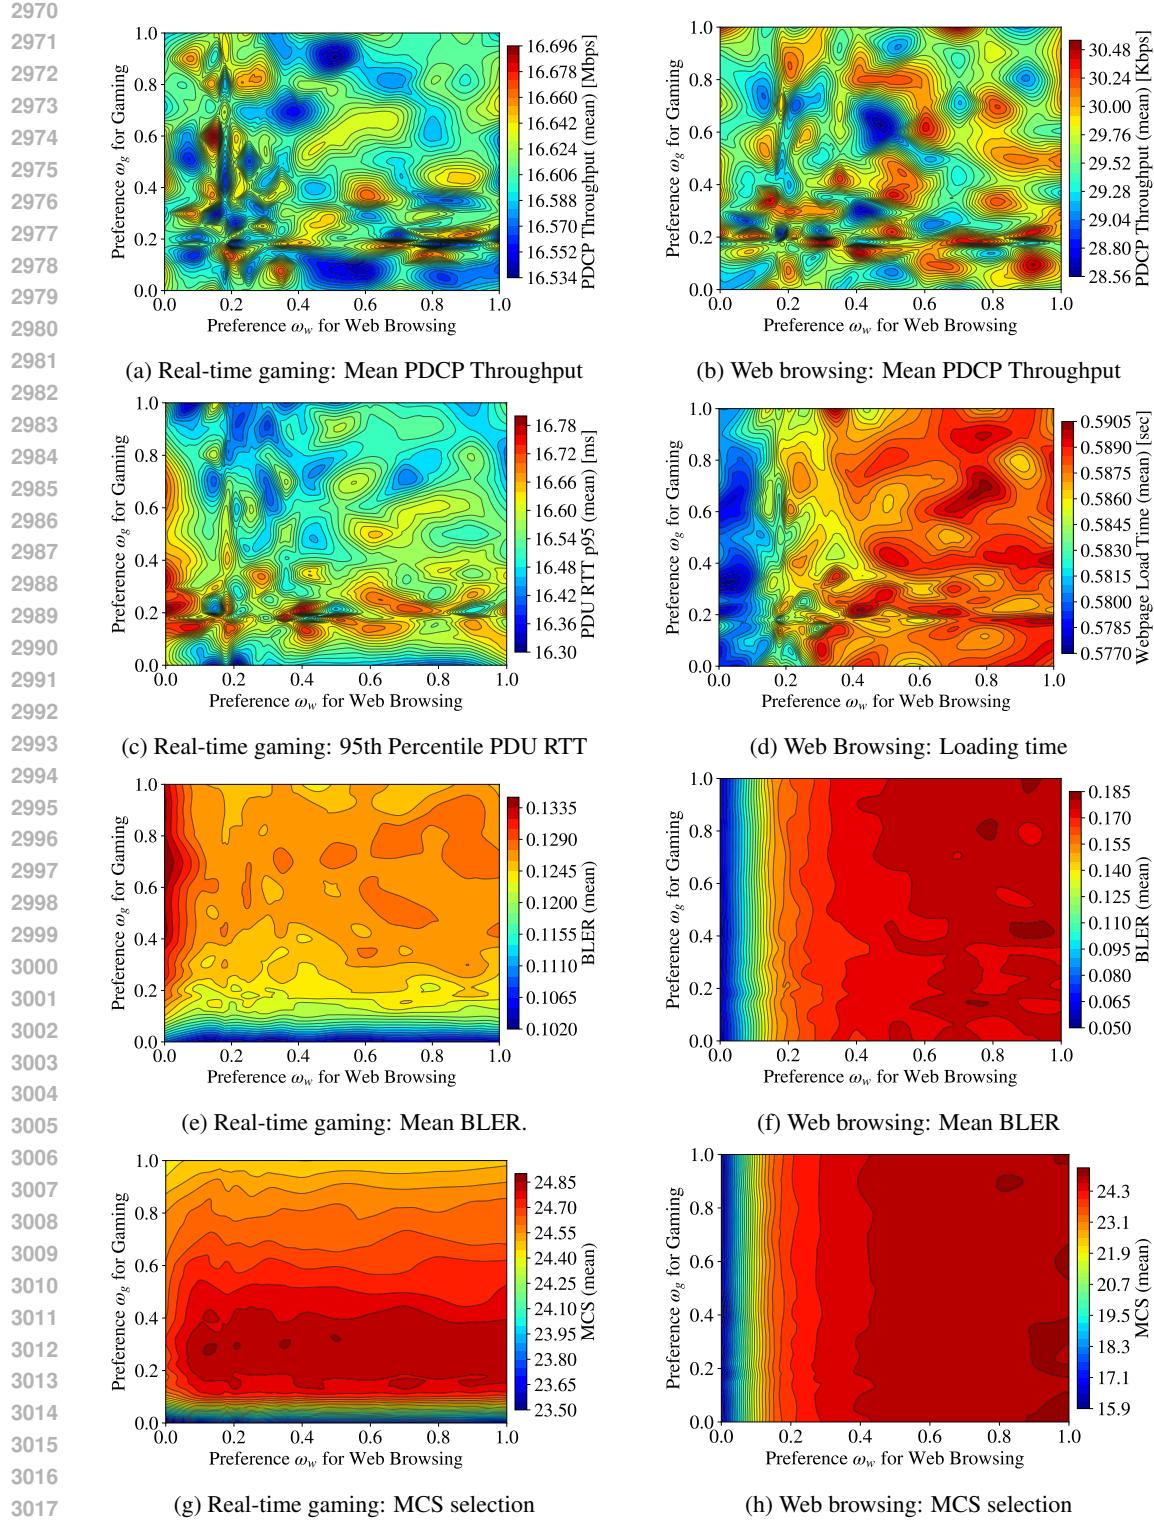
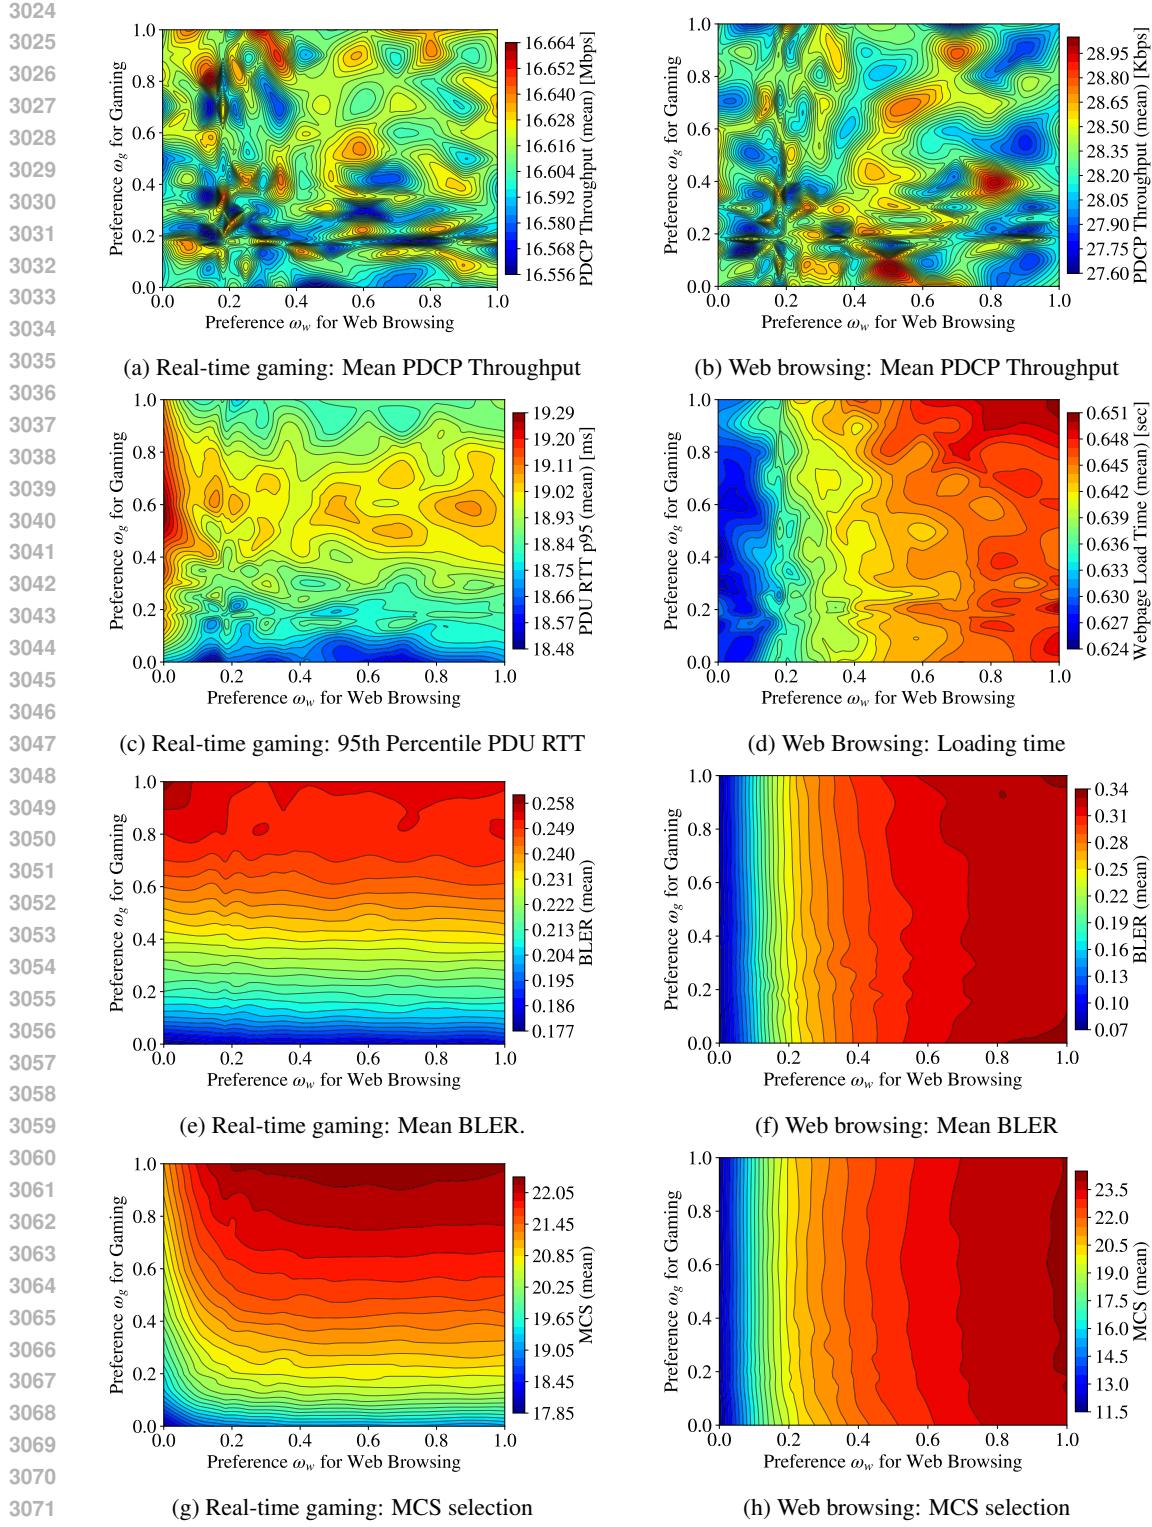


Figure 12: **Impact of user preference weights on the performance of real-time gaming and browsing users in low network load conditions.** Each subplot shows a distinct QoS metric for real-time gaming users (left column) and for web browsing users (right column) under varying preference weights ( $\omega_g, \omega_w$ ) reflecting resource allocation priorities for the two connectivity services. Metrics include: (a)-(b) mean user throughput, (c)-(d) mean latency (defined according to the service), (e)-(f) mean BLER. Furthermore, (g)-(h) show the action (mean MCS) distribution under  $(\omega_g, \omega_w)$ .



**Figure 13: Impact of user preference weights on the performance of real-time gaming and browsing users in medium network load conditions.** Each subplot shows a distinct QoS metric for real-time gaming users (left column) and for web browsing users (right column) under varying preference weights ( $\omega_g, \omega_w$ ) reflecting resource allocation priorities for the two connectivity services. Metrics include: (a)-(b) mean user throughput, (c)-(d) mean latency (defined according to the service), (e)-(f) mean BLER. Furthermore, (g)-(h) show the action (mean MCS) distribution under  $(\omega_g, \omega_w)$ .

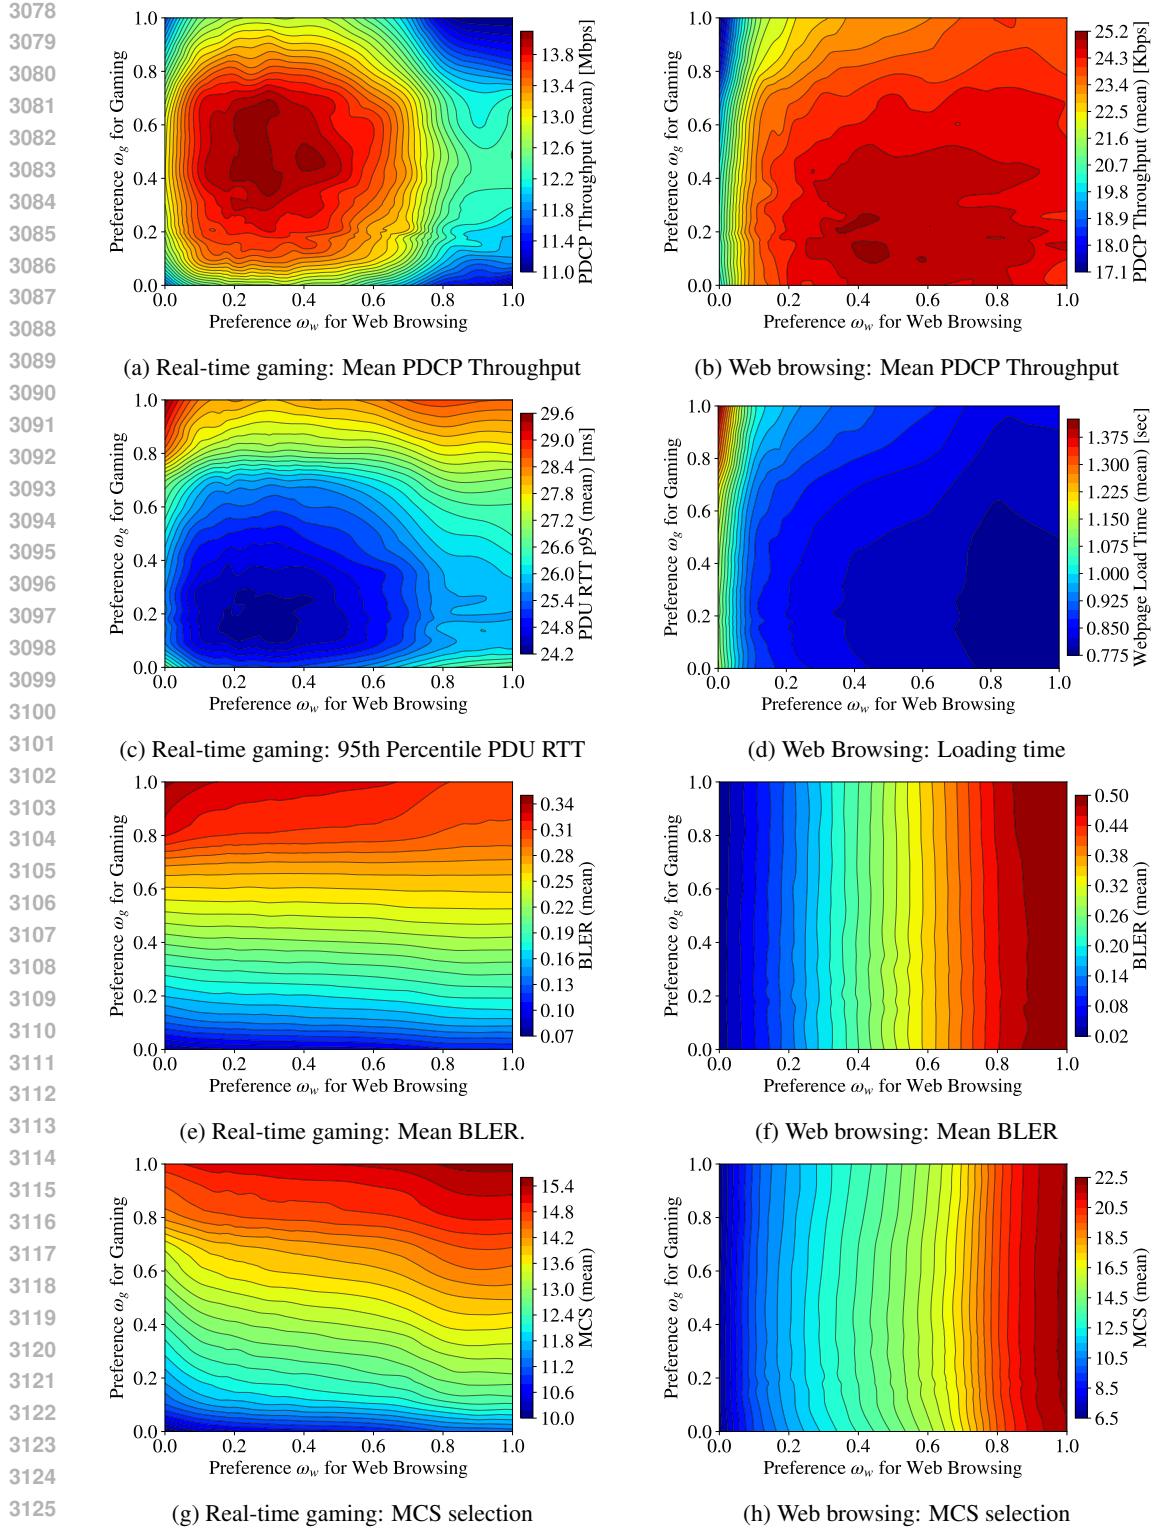


Figure 14: **Impact of user preference weights on the performance of real-time gaming and browsing users in high network load conditions.** Each subplot shows a distinct QoS metric for real-time gaming users (left column) and for web browsing users (right column) under varying preference weights ( $\omega_g, \omega_w$ ) reflecting resource allocation priorities for the two connectivity services. Metrics include: (a)-(b) mean user throughput, (c)-(d) mean latency (defined according to the service), (e)-(f) mean BLER. Furthermore, (g)-(h) show the action (mean MCS) distribution under  $(\omega_g, \omega_w)$ .

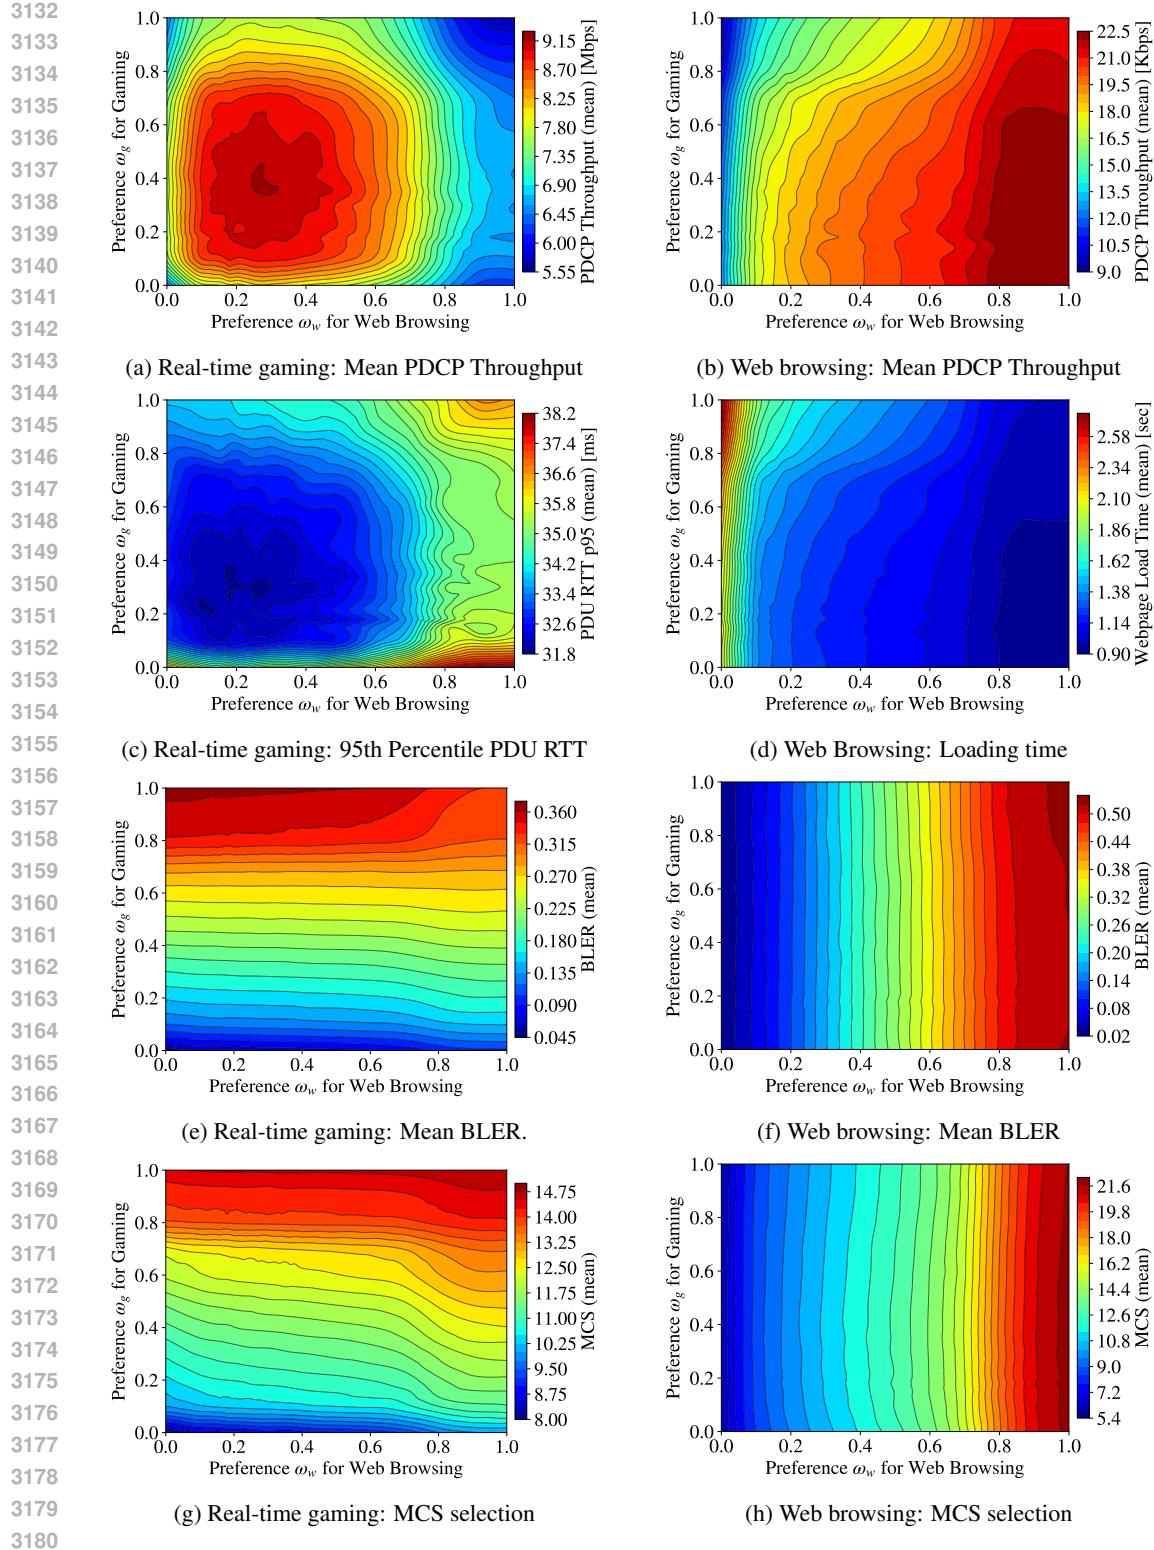


Figure 15: **Impact of user preference weights on the performance of real-time gaming and browsing users in very high network load conditions.** Each subplot shows a distinct QoS metric for real-time gaming users (left column) and for web browsing users (right column) under varying preference weights ( $\omega_g, \omega_w$ ) reflecting resource allocation priorities for the two connectivity services. Metrics include: (a)-(b) mean user throughput, (c)-(d) mean latency (defined according to the service), (e)-(f) mean BLER. Furthermore, (g)-(h) show the action (mean MCS) distribution under  $(\omega_g, \omega_w)$ .

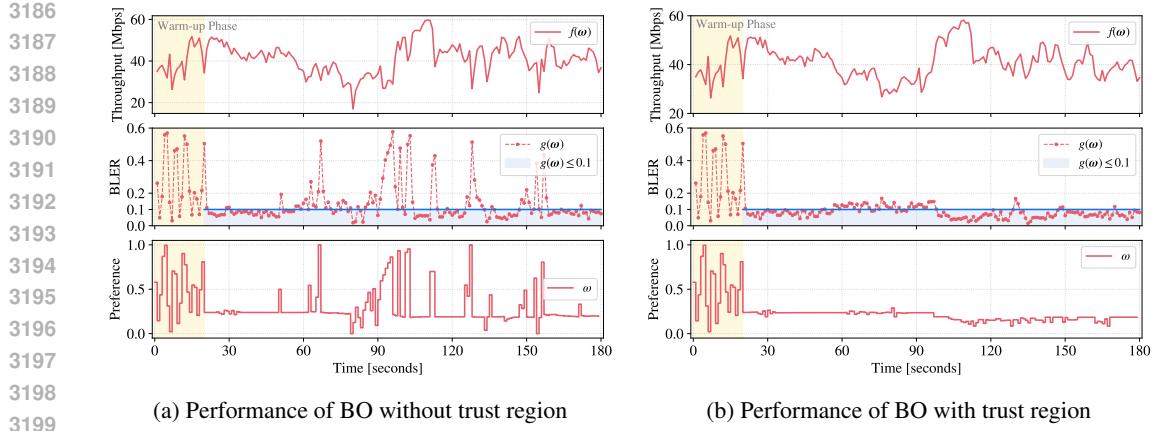


Figure 16: Performance comparison of optimizer agent when using the PAX-BO algorithm (a) without trust region and (b) with trust region enabled.

#### H.4 ONLINE PREFERENCE OPTIMIZATION

Appendix D presented the Preference-Aligned eXploration Bayesian Optimization (PAX-BO) algorithm for the optimizer agent (cf. Algorithm 2). A key design feature of PAX-BO is the integration of trust regions to stabilize the selection of the preference values  $\omega$  for the downstream controller agent.

Figure 16 compares the performance of PAX-BO considering an intent definition that requires to maximize the aggregate system throughput while keeping the BLER of each user below 10% – which corresponds to the typical configuration of link adaptation in 4G/5G RAN systems. Figure 16a shows the performance of PAX-BO without trust region, whereas in Figure 16b we enabled the trust region. The top panel of both figures shows the aggregate system throughput  $f(\omega)$ , the middle panel shows BLER constraint evaluations  $g(\omega)$  with threshold  $g(\omega) \leq 0.1$ , and the bottom panel shows the evolution of the preference parameter  $\omega$ . Compared to Figure 16a, the trust region stabilizes the optimization, reducing constraint violations and leading to smoother preference adaptation and improved system throughput (with smoother degradations).

## 3240 I COMPUTE RESOURCES AND HYPERPARAMETERS

3241  
 3242 All MORL training runs were performed on a high-performance computing (HPC) cluster. The main  
 3243 training node was equipped with an NVIDIA A100-PCIE-40GB GPU and 48 CPU cores, which hosted  
 3244 the learner, actors, and replay memory. The replay buffer was partitioned into four independently  
 3245 prioritized shards, each pinned to a dedicated CPU core to support parallelized access. Co-locating  
 3246 the learner, actors, and replay shards on the same node minimized intra-node communication latency.

3247 We used 40 actors for each experiment, and each actor launched two threads that interacted with  
 3248 14 simulator instances in parallel. The simulators were distributed across multiple compute nodes,  
 3249 totaling 560 CPU cores. Each simulator ran in a separate process and communicated with its assigned  
 3250 actor via ZeroMQ, enabling scalable multi-node environment interaction.

3251 Cluster job scheduling and resource management were handled by the Load Sharing Facility (LSF),  
 3252 which managed job queueing, monitoring, and node allocation according to the experiments’ resource  
 3253 specifications.

3255 Table 12: Hyperparameters Used for Adaptor in Interpreter.  
 3256

3257 <b>Monitor</b>	
3259 Window size (W)	12
3260 Alert-on ratio ( $\rho_{on}$ )	0.55
3261 Alert-off ratio ( $\rho_{off}$ )	0.45
3262 <b>Adjust</b>	
3263 Step (Mbps)	0.08
3264 Lifetime (Mbps)	0.40
3265 Floor	5.00
3266 Ceiling	9.00
3267 Cooldown steps	2
3268 Deadband	0.05
3269 Gain (up)	1.0
	1.0

3271 Table 13: Hyperparameters Used for Supervised Fine-Tuning of the Intent-to-OTM Translator.  
 3272

3273 <b>Component</b>	3274 <b>Setting</b>
3275 Base model	Qwen2.5-7B-Instruct
3276 Parameter-efficient tuning	LoRA (rank 64, $\alpha = 16$ , dropout = 0.05)
3277 LoRA target modules	q_proj, k_proj, v_proj, o_proj
3278 Precision	bfloat16
3279 Epochs	2
3280 Batch size (per device)	2
3281 Gradient accumulation steps	8
3282 Optimizer	AdamW (Torch fused)
3283 Learning rate	$2 \times 10^{-4}$
3284 Scheduler	Cosine decay
3285 Warmup ratio	0.03
3286 Weight decay	0.01
3287 Gradient clipping	1.0
3288 Gradient checkpointing	Enabled
3289 Max sequence format	Qwen chat template (intent $\rightarrow$ OTM pair)
3290 Evaluation frequency	Every 200 steps
3291 Checkpoint frequency	Every 200 steps (max 5 checkpoints)

3294  
 3295  
 3296  
 3297  
 3298  
 3299  
 3300  
 3301  
 3302  
 3303  
 3304  
 3305  
 3306  
 3307  
 3308  
 3309

Table 14: Bayesian Optimization Hyperparameters Used in the Actor.

Hyperparameter	Value / Description
Acquisition function	qLogEI
MC samples for acquisition	256
Raw samples for optimization	512
Number of restarts	10
Batch size ( $q$ )	1
GP refit frequency	Every 1 observation
Training window size	60 most recent samples
Input scaling (normalize)	Yes (Normalize transform)
Output scaling	Standardize outcomes
Trust region initial radius	0.15
Minimum trust region radius	0.05
Trust region shrink factor	0.7
Infeasible patience	2 consecutive infeasible samples
No-improvement patience	5 evaluations
Initial preference samples	20 Sobol samples (fixed list)
Preference domain	$[0, 1]^2$

3333  
 3334  
 3335  
 3336  
 3337  
 3338  
 3339  
 3340  
 3341  
 3342  
 3343  
 3344  
 3345  
 3346  
 3347

3348  
 3349  
 3350  
 3351  
 3352  
 3353  
 3354  
 3355  
 3356  
 3357

Table 15: Hyperparameters used for Multi-Objective Reinforcement Learning (MORL).

3358  
 3359  
 3360  
 3361  
 3362  
 3363  
 3364  
 3365  
 3366  
 3367  
 3368  
 3369  
 3370  
 3371  
 3372  
 3373  
 3374  
 3375  
 3376  
 3377  
 3378  
 3379  
 3380  
 3381  
 3382  
 3383  
 3384  
 3385  
 3386  
 3387  
 3388  
 3389  
 3390  
 3391  
 3392  
 3393  
 3394  
 3395  
 3396  
 3397  
 3398  
 3399  
 3400  
 3401

<b>Learner</b>	
Optimizer	Adam Kingma & Ba (2017)
Learning rate	$5 \times 10^{-5}$
$\beta_1$ (Adam momentum term)	0.9
$\beta_2$ (Adam second moment term)	0.999
$\epsilon$ (Adam numerical stability)	$1.5 \times 10^{-4}$
Weight decay	0.02/512
Gradient norm	20
Target update period	Every 1 gradient updates
Target update policy	Soft
Target update factor	$1.0 \times 10^{-3}$
Model update interval	Every 200 gradient updates
Prefetched batches	16
Batch size (experience)	512
Batch size (preference)	128
Warm-up phase	50,000 samples
Loss function	MSE
<b>Actor</b>	
Number of actors	40
Local buffer capacity	2,500
Discount factor ( $\gamma$ )	1.0
$\epsilon$ -greedy (linear decay)	$0.8 \rightarrow 0.05$
Timesteps	5,500,000
<b>Replay Memory</b>	
Number of shards	4
Capacity of each shard	4,000,000
Prioritization exponent ( $\alpha$ )	0.6
Importance sampling exponent ( $\beta$ )	0.4
<b>Model</b>	
Activation function	ReLU
Number of blocks	6
Number of layers per block	2
units per layer	128
Dropout probability	0.1
Layer normalization	True