

Bayesian Optimization for Online Management in Dynamic Mobile Edge Computing

Jia Yan¹, *Member, IEEE*, Qin Lu², *Member, IEEE*, and Georgios B. Giannakis³, *Life Fellow, IEEE*

Abstract—Recent years have witnessed the emergence of mobile edge computing (MEC), on the premise of a cost-effective enhancement in the computational ability of hardware-constrained wireless devices (WDs) comprising the Internet of Things (IoT). In a general multi-server multi-user MEC system, each WD has a computational task to execute and has to select binary (off)loading decisions, along with the analog-amplitude resource allocation variables in an online manner, with the goal of minimizing the overall energy-delay cost (EDC) with dynamic system states. While past works typically rely on the explicit expression of the EDC function, the present contribution considers a practical setting, where in lieu of system state information, the EDC function is not available in analytical form, and instead only the function values at queried points are revealed. Towards tackling such a challenging online combinatorial problem with only bandit information, novel Bayesian optimization (BO) based approaches are put forth by leveraging the multi-armed bandit (MAB) framework. Per time slot, the discrete offloading decisions are first obtained via the MAB method, and the analog resource allocation variables are subsequently optimized using the BO selection rule. By exploiting both temporal and contextual information, two novel BO approaches, termed time-varying BO and contextual time-varying BO, are developed. Numerical tests validate the merits of the proposed BO approaches compared with contemporary benchmarks under different MEC network sizes.

Index Terms—Mobile edge computing, Bayesian optimization, online learning, task offloading, resource allocation, Internet of Things.

I. INTRODUCTION

THE era of massive connectivity is brought into being by the Internet of Things (IoT), where tens of billions of wireless devices (WDs) are ubiquitously connected to the Internet through cellular networks. Constrained by limited batteries and low-power on-chip computing units, the WDs face challenges to support latency-sensitive applications in the current IoT paradigms such as autonomous driving, online gaming and virtual reality. To meet the intensive computation demands far beyond the WDs' capacities, mobile edge computing (MEC) has emerged as a promising technology by releasing and distributing computing resources to the edge servers within the radio access networks to facilitate real-time services. Capitalizing on the MEC architecture, WDs in the IoT are able to carry out high-performance computation by offloading tasks to the servers located at the network edge [2]. Compared with traditional mobile cloud computing, the MEC no longer suffers from high overhead and long backhaul latency.

Due to the time-varying wireless channel conditions and the heterogeneity in both the WDs and edge servers, judiciously offloading computations can offer significant performance enhancement. In general, MEC has two computation offloading models, referred to as binary and partial offloading [2]. Binary offloading requires each task to be either executed locally or offloaded to the edge server as a whole [3]. On the other hand, a task under partial offloading model is allowed to be partitioned and computed both locally and at the edge server [4], [5]. In this work, we focus on binary computation task offloading, which is commonly used in IoT to process indivisible simple tasks such as face recognition and temperature monitoring in smart home [2]. Prior works on offloading computations typically focus on offline algorithms by adopting either convex [3], [4] or non-convex (e.g., convex relaxation [6] and heuristic local search [7], [8]) optimization methods, which assume that the system states are known a priori, even though such knowledge is challenging to acquire beforehand.

With unknown system dynamics, online computational task offloading approaches have been extensively investigated. Building on the assumption of stationarity, a class of online algorithms rely on stochastic optimization methods such as Lyapunov optimization to determine the task

Manuscript received 3 February 2023; revised 25 June 2023; accepted 12 August 2023. Date of publication 29 August 2023; date of current version 11 April 2024. This work was supported by NSF Grant 2102312, Grant 2103256, Grant 1901134, Grant 2126052, Grant 2128593, and Grant 2220292. The work of Jia Yan was supported in part by the Guangzhou Municipal Science and Technology Project under Grant 2023A03J0011. An earlier version of this paper was presented in part at the IEEE 56th Asilomar Conference on Signals, Systems, and Computers, Pacific Grove, California, USA, October 2022 [DOI: 10.1109/IEEECONF56349.2022.10051868]. The associate editor coordinating the review of this article and approving it for publication was X. Gong. (Corresponding author: Jia Yan.)

Jia Yan was with the Department of Electrical and Computer Engineering, University of Minnesota, Twin Cities, Minneapolis, MN 55455 USA. He is now with the Intelligent Transportation Thrust, The Hong Kong University of Science and Technology (Guangzhou), Guangzhou 511453, China, and also with the Guangdong Provincial Key Laboratory of Integrated Communication, Sensing and Computation for Ubiquitous Internet of Things, Guangzhou 511453, China (e-mail: jasonjiayan@hkust-gz.edu.cn).

Qin Lu was with the Department of Electrical and Computer Engineering, University of Minnesota, Twin Cities, Minneapolis, MN 55455 USA. She is now with the School of Electrical and Computer Engineering, University of Georgia, Athens, GA 30602 USA (e-mail: Qin.Lu@uga.edu).

Georgios B. Giannakis is with the Department of Electrical and Computer Engineering, University of Minnesota, Twin Cities, Minneapolis, MN 55455 USA (e-mail: georgios@umn.edu).

Color versions of one or more figures in this article are available at <https://doi.org/10.1109/TWC.2023.3307875>.

Digital Object Identifier 10.1109/TWC.2023.3307875

offloading decisions within each time slot without future information [9], [10], [11]. Nevertheless, the nonstationarity introduced by the human participation in IoT makes the stochastic optimization impractical. Targeting at the nonstationary system dynamics, existing works focus on the online convex optimization (OCO) algorithms [12], [13], [14], where the sequence of convex task offloading costs changes in an unknown and possibly adversarial manner. Yet, the OCO approaches necessitate the availability of explicit cost function forms or their gradients.

In practice though, the unpredictable WD preferences (e.g., service latency, reliability or privacy) render it prohibitive to model the objective function analytically in dynamic IoT environment. In fact, the IoT controller can only have available objective function values at queried points. In this context, the OCO has been extended to the bandit setting by leveraging only point-wise values of objective functions for the gradient estimations, which is referred to as bandit convex optimization (BCO) [15], [16], [17]. Tailored for partial task offloading strategies among multiple edge servers, BCO with both time-varying costs and constraints was studied in [18]. On the other hand, aiming at binary computational offloading strategies with such a bandit feedback, multi-armed bandit (MAB) based methods have been popular in MEC systems [19], [20], [21], [22]. An online combinatorial bandit upper confidence bound algorithm was proposed in [19] for the task scheduling to asymptotically minimize the computing delay. The security-aware server selection strategies based on MAB were reported in [20]. The MAB-based task offloading approach was further adapted to the vehicular edge computing systems in [21].

Although achieving promising results, the aforementioned BCO or MAB based works deal only with either continuous or discrete decision variables. In many practical settings though, the analog-amplitude communication and computation resource allocation variables (e.g., transmit power and local computing speed) need to be jointly optimized with discrete variables that capture offloading decisions for optimum MEC performance. Finely discretizing the analog action space (or relaxing the discrete task offloading decisions), renders the existing MAB methods (or the BCO approaches) inaccurate and computationally prohibitive. In addition, the convexity of objective functions commonly assumed in BCO algorithms may not hold in practice [2], [3], [6], [7], [8]. Although dealing with arbitrary objective functions, MAB methods require to explore every single arm at least once to accumulate sufficient statistics, which may incur sudden performance drops and slow down the learning processes for large MEC networks [19], [20], [21], [22].

Alleviating these limitations, we advocate a novel approach based on Bayesian optimization (BO) [23]. BO is a promising methodology for black-box derivative-free (i.e., only function value observations at queried points are available without derivative information) global optimization with well-documented merits, including sample efficiency, uncertainty quantification, and safe exploration [23], [24]. The key idea of BO is to build a Bayesian surrogate model (typically, the Gaussian process [25], [26], [27], [28]) for the black-

box objective function, guided by which an acquisition function is designed to decide the next function evaluation point. Apart from the applications such as hyperparameter tuning in machine learning [29], drug discovery [30], and robotics [31], BO has been applied to several problems in the context of wireless networks, including radio resource allocation [32], coverage and capacity optimization in cellular networks [33], as well as beam alignment in mmWave MIMO systems [34]. Very recently, targeting video analysis in MEC, a BO-based approach is put forth for edge server and frame resolution selection in [35], where the case of analog-amplitude communication and computation resource allocation is not accounted for.

The existing BO-based approaches to resource allocation for MEC focus on optimizing black-box objective functions with either continuous (e.g., transmit power [32], [33] and beam alignment [34]) or integer (e.g., frame resolution [35]) variables. The present work, on the other hand, deals with the mixture of categorical and continuous variables, which poses nontrivial challenges in kernel design of the Gaussian process (GP) based surrogate model and acquisition rule relative to previous works. Further, the sought objective function here is modeled as time-varying under the unknown system state dynamics (e.g., wireless channel conditions and edge server computing speeds), what deteriorates the performance of existing BO approaches where stale data is exploited as equally important as fresh data. Towards addressing the aforementioned challenges, the current work puts forth a novel kernel that incorporates the temporal and contextual information, and further devises the associated acquisition rule.

In a multi-server multi-user MEC, time-varying system state information, including wireless channel conditions, edge server computing speeds, task workloads, and input data sizes, has significant impact on task offloading and resource allocation decisions. For example, a WD with larger task workload at current time slot may prefer edge computing, rather than local execution, so as to leverage higher computing speed at edge servers. Also, the WD may choose the edge server with best channel condition to offload its task with small transmit power for energy saving. Therefore, exploiting such temporal and contextual information (i.e., partially observed system state) will intuitively yield enhanced performance than the vanilla BO approach for MEC systems with bandit feedback.

Relative to the aforementioned existing works, the present work is the first attempt to develop novel BO-based approaches for the joint optimization of discrete task offloading decisions and analog-amplitude resource allocation strategies in time-varying multi-server multi-user MEC systems with bandit feedback. Specifically, our main contributions are summarized as follows.

- 1) Building on the BO framework for online bandit optimization of categorical and continuous decision variables, a GP-based surrogate model is adopted for the sought objective function with novel kernel design. The resultant kernel function not only leverages a weighted combination of sum and product compositions of individual kernels over categorical and continuous

variables in order to allow for more expressive coupling, but also capitalizes on a temporal kernel to account for unknown dynamics in the black-box function.

- 2) With the GP-based surrogate model, an innovative acquisition rule is developed in the time-varying BO scheme to select new optimization variables per iteration. Specifically, given the categorical offloading decisions obtained by the MAB-based method, the analog-amplitude resource allocation variables are determined using the conventional BO-based selection rule.
- 3) Under the scenario where each WD reveals its task characterization variables (task computational workload and input data size) at the beginning of each time slot, a generalized contextual time-varying BO scheme is further devised by incorporating the contextual kernel in the GP surrogate model.
- 4) Numerical simulations under various MEC network sizes demonstrate that our proposed BO approaches benefit from both temporal and contextual information, and exhibit superior performance compared with traditional BO and other representative benchmarks.

The rest of the paper is organized as follows. The system model and problem formulation are presented in Sec. II, following which a novel time-varying BO algorithm for online joint optimization of task offloading and resource allocation under bandit setting is proposed in Sec. III. Further leveraging observed state information, Sec. IV develops the contextual time-varying BO approach for dynamic MEC management. In Sec. V, the performance of the proposed BO methods is evaluated on synthetic tests. Finally, concluding remarks are made in Sec. VI.

Notation: $(\cdot)^\top$ and $(\cdot)^{-1}$ denote transpose and matrix inverse, respectively, and $\|\mathbf{x}\|$ stands for the l_2 -norm of a vector $\mathbf{x}$. Besides, $\mathbf{0}_t$, $\mathbf{1}_t$ and $\mathbf{I}_t$ denote the $t \times 1$ all-zero vector, the $t \times 1$ all-one vector and the $t \times t$ identity matrix, respectively. Inequalities for vector $\mathbf{x} > \mathbf{0}$ are entry-wise. $\mathbb{I}(x = x')$ denotes the indicator function taking the value of 1 if $x = x'$, and 0 otherwise. $\mathcal{N}(\mathbf{x}; \boldsymbol{\mu}, \mathbf{K})$ stands for the probability density function (pdf) of a Gaussian random vector $\mathbf{x}$ with mean $\boldsymbol{\mu}$ and covariance $\mathbf{K}$.

II. SYSTEM MODEL AND PROBLEM FORMULATION

Consider a MEC system with M WDs, and N base stations (BSs). Each BS $n \in \mathcal{N} := \{1, \dots, N\}$ is the gateway of edge servers to provide MEC services to the power-limited WDs indexed by $m \in \mathcal{M} := \{1, \dots, M\}$. Per slot $t \in \mathcal{T} := \{1, \dots, T\}$, the m -th WD has a computational task characterized by the pair (I_t^m, L_t^m) , where I_t^m denotes the size of input data in bits, and L_t^m represents the workload in terms of the total number of CPU cycles to execute the aforementioned task. This WD could either execute its task locally or offload it to one of the BSs, a choice that is henceforth captured by the categorical variable $c_t^m \in \{0, 1, \dots, N\}$. Specifically, $c_t^m = 0$ indexes local computing, and $c_t^m = n, n \in \mathcal{N}$, stands for offloading task to BS n , i.e.,

$$c_t^m = \begin{cases} 0, & \text{local computing} \\ n, & \text{offloading task to BS } n \end{cases} \quad \forall m \in \mathcal{M}, n \in \mathcal{N}, t \in \mathcal{T}. \quad (1)$$




Fig. 1. The considered mobile edge computing (MEC) system with M wireless devices (WDs) and N base stations (BSs).

For both scenarios, the computational overhead per task consists of the execution delay and energy consumption, which will be elaborated as follows.

A. Local Computing

If WD m chooses to execute its task locally (i.e., $c_t^m = 0$) per slot t , it has to select the local CPU frequency f_t^m , based on which the task computing time is given by

$$\tau_{l,t}^m = \frac{L_t^m}{f_t^m} \quad (2)$$

and the corresponding energy consumption is

$$\epsilon_{l,t}^m = \xi L_t^m (f_t^m)^2 \quad (3)$$

where ξ denotes the effective switched capacitance parameter [2].

B. Edge Computing

If WD m alternatively goes for edge computing at BS n per slot t , that is, $c_t^m = n$, it must first offload the task using transmit power p_t^m . Suppose that the wireless channel coefficient between WD m and BS n for task offloading is $h_t^{m,n}$, and the receiver is corrupted by additive white Gaussian noise (AWGN) with mean zero and variance σ^2 . Here, the wireless channel is assumed to be invariant within each slot and may change across different slots. Then, the uplink transmission data rate for the sought offloading task is

$$R_t^{m,n} = W \log_2 \left(1 + \frac{p_t^m |h_t^{m,n}|^2}{\sigma^2} \right) \quad (4)$$

where W is the identical bandwidth of the dedicated spectral resource block allocated to each WD. Accordingly, the offloading transmission time is

$$\tau_{u,t}^m = \sum_{n=1}^N \frac{\mathbb{I}(c_t^m = n) I_t^m}{R_t^{m,n}} \quad (5)$$

and the transmission energy consumption of WD m is

$$\epsilon_{u,t}^m = p_t^m \tau_{u,t}^m. \quad (6)$$

For edge computing at BS n , the total computation resource per slot t is signified by the CPU frequency $f_{c,t}^n$. Upon receiving all the offloaded tasks, the edge server

generates multiple virtual machines (VMs) to execute the tasks in parallel, and equally partitions $f_{c,t}^n$ to yield $f_{c,t}^n/(1 + \sum_{m' \in \mathcal{M}/m} \mathbb{I}(c_t^{m'} = n))$ per task. The edge execution time for WD m 's task is thus

$$\tau_{c,t}^m = \sum_{n=1}^N \mathbb{I}(c_t^m = n) \frac{L_t^m (1 + \sum_{m' \in \mathcal{M}/m} \mathbb{I}(c_t^{m'} = n))}{f_{c,t}^n}. \quad (7)$$

It is worth mentioning that the time delay for downloading the task output from the BS to the WD is ignored given the relatively small output data size and strong downlink transmit power of the BS.

C. Problem Formulation

Accounting for both local and edge computing, the total time delay for executing the task at WD m per slot t is given by

$$D_t^m = \mathbb{I}(c_t^m = 0) \tau_{l,t}^m + \mathbb{I}(c_t^m \neq 0) (\tau_{u,t}^m + \tau_{c,t}^m). \quad (8)$$

Here, D_t^m is equal to the local execution time $\tau_{l,t}^m$ if WD m chooses local computing (i.e., $c_t^m = 0$). Otherwise, D_t^m in (8) equals the sum of offloading transmission time $\tau_{u,t}^m$ and the edge computing time $\tau_{c,t}^m$.

Similarly, the energy consumption of WD m per slot t is given by

$$E_t^m = \mathbb{I}(c_t^m = 0) \epsilon_{l,t}^m + \mathbb{I}(c_t^m \neq 0) \epsilon_{u,t}^m \quad (9)$$

which is $\epsilon_{l,t}^m$ for local computing ($c_t^m = 0$) and $\epsilon_{u,t}^m$ otherwise.

Taking a weighted sum of task execution time delay D_t^m and energy consumption E_t^m yields the energy-delay cost (EDC) per WD m as

$$EDC_t^m(c_t^m, f_t^m, p_t^m) = \beta_d D_t^m + \beta_e E_t^m \quad (10)$$

where β_d, β_e are positive scalars that balance these two costs. For notational brevity, collect the optimization variables in $\mathbf{c}_t := [c_t^1, \dots, c_t^M]^\top$, $\mathbf{p}_t := [p_t^1, \dots, p_t^M]^\top$, and $\mathbf{f}_t := [f_t^1, \dots, f_t^M]^\top$. The objective is to choose online (at the beginning of each slot t) the categorical task offloading decisions (i.e., $\mathbf{c}_t$) and analog-amplitude resource allocation strategies (i.e., $\mathbf{p}_t, \mathbf{f}_t$) minimizing the accumulated EDC across all WDs, that is

$$\begin{aligned} \text{(P1)} \quad & \min_{\{\mathbf{c}_t, \mathbf{p}_t, \mathbf{f}_t\}_t} \sum_{t=1}^T \sum_{m=1}^M EDC_t^m(c_t^m, f_t^m, p_t^m), \\ \text{s.t.} \quad & c_t^m \in \{0, 1, 2, \dots, N\}, \\ & 0 < p_t^m \leq P_{\text{peak}}, \\ & 0 < f_t^m \leq f_{\text{peak}}, \forall m \in \mathcal{M}, t \in \mathcal{T} \end{aligned}$$

where f_{peak} and P_{peak} are the peak local CPU frequency and transmit power of the WDs, respectively. By further introducing $\mathbf{x}_t := [\mathbf{p}_t^\top, \mathbf{f}_t^\top]^\top$ and the reward function $\varphi_t(\mathbf{c}_t, \mathbf{x}_t) := -\sum_{m=1}^M EDC_t^m$ at slot t , (P1) can be equivalently expressed as

$$\begin{aligned} \text{(P2)} \quad & \max_{\{\mathbf{c}_t, \mathbf{x}_t\}_t} \sum_{t=1}^T \varphi_t(\mathbf{c}_t, \mathbf{x}_t), \\ \text{s.t.} \quad & \mathbf{c}_t \in \{0, 1, 2, \dots, N\}^M, \\ & \mathbf{0} < \mathbf{x}_t \leq \mathbf{x}_{\text{peak}}, \forall t \in \mathcal{T} \end{aligned}$$

where $\mathbf{x}_{\text{peak}} := [P_{\text{peak}} \mathbf{1}_M^\top, f_{\text{peak}} \mathbf{1}_M^\top]^\top$, and $\mathbf{1}_M$ is the M -dimensional all-one column vector.

A major challenge facing (P2) (equivalently (P1)) is that the wireless channels $\{h_t^{m,n}\}$, the edge computing capacities $\{f_{c,t}^n\}$, the computational task characterization $\{I_t^m, L_t^m\}$ are not available; thus, the explicit form of the time-varying EDC function is unknown when making the task offloading and resource allocation decisions $\{\mathbf{c}_t, \mathbf{x}_t\}$ per slot. After performing $\{\mathbf{c}_t, \mathbf{x}_t\}$, only noisy EDC function value (equivalently the realization of $\varphi_t(\mathbf{c}_t, \mathbf{x}_t)$) at that queried point can be acquired at the end of slot t . The difficulty of such a bandit setup is further exacerbated by its combinatorial nature that calls for the joint optimization of the categorical $\mathbf{c}_t$ and continuous $\mathbf{x}_t$. To tackle this bandit mix-integer program, novel BO-based approaches will be pursued in the following sections.

III. TIME-VARYING BO FOR DYNAMIC MEC MANAGEMENT

BO has well-documented merits in optimizing black-box functions that arise in several settings [23]. To account for the temporal variation arising from unknown system dynamics (e.g., changing channel conditions and computing capacities of the edge servers), the slot index t is augmented as an additional input of the sought black-box function, i.e., $\varphi(\mathbf{c}_t, \mathbf{x}_t, t) := \varphi_t(\mathbf{c}_t, \mathbf{x}_t)$. In short, BO seeks to maximize the black-box $\varphi(\mathbf{z}_t)$ with $\mathbf{z}_t := [\mathbf{c}_t^\top, \mathbf{x}_t^\top, t]^\top$ by sequentially acquiring function observations using a surrogate model. Collect all the acquired data up to slot t in $\mathcal{D}_t := \{(\mathbf{z}_\tau, y_\tau)\}_{\tau=1}^t$ with y_τ denoting the possibly noisy observation of $\varphi(\mathbf{z}_\tau)$. Each BO iteration consists of i) obtaining the function posterior pdf $p(\varphi(\mathbf{z})|\mathcal{D}_t)$ based on the chosen surrogate model using $\mathcal{D}_t$; and, ii) selecting $\mathbf{z}_{t+1}$ to evaluate at the beginning of slot $t+1$, whose observation y_{t+1} will be acquired at the end of slot $t+1$. In the following, we will introduce the GP-based surrogate model and the acquisition rule for $\mathbf{z}_{t+1}$, respectively.

A. GP-Based Surrogate Model for Time-Varying Function φ and Kernel Design

As an established Bayesian nonparametric approach, the GP can learn black-box functions with quantifiable uncertainty and sample efficiency, making it suitable for surrogate modeling in BO. Specifically, given data $\mathcal{D}_t$, the goal is to learn the function $\varphi(\cdot)$ that links the input $\mathbf{z}_\tau$ with the scalar output y_τ as $\mathbf{z}_\tau \rightarrow \varphi(\mathbf{z}_\tau) \rightarrow y_\tau$. Towards this, a GP prior is assumed on the unknown φ as $\varphi \sim \mathcal{GP}(0, \kappa(\mathbf{z}, \mathbf{z}'))$, where $\kappa(\cdot, \cdot)$ is a kernel (covariance) function measuring pairwise similarity of any two inputs. Then, the joint prior pdf of any t function evaluations $\varphi_t := [\varphi(\mathbf{z}_1), \dots, \varphi(\mathbf{z}_t)]^\top$ at inputs $\mathbf{Z}_t := [\mathbf{z}_1, \dots, \mathbf{z}_t]^\top$ is jointly Gaussian distributed as [25]

$$p(\varphi_t|\mathbf{Z}_t) = \mathcal{N}(\varphi_t; \mathbf{0}_t, \mathbf{K}_t), \forall t \quad (11)$$

where $\mathbf{K}_t$ is a $t \times t$ covariance matrix with (τ, τ') -th entry $[\mathbf{K}_t]_{\tau, \tau'} = \text{cov}(\varphi(\mathbf{z}_\tau), \varphi(\mathbf{z}_{\tau'})) := \kappa(\mathbf{z}_\tau, \mathbf{z}_{\tau'})$. The estimation of φ relies on the observed outputs $\mathbf{y}_t := [y_1, \dots, y_t]^\top$ that are linked with φ_t through the Gaussian conditional likelihood $p(\mathbf{y}_t|\varphi_t, \mathbf{Z}_t) = \mathcal{N}(\mathbf{y}_t; \varphi_t, \sigma_o^2 \mathbf{I}_t)$, where σ_o^2 is the

noise variance. Along with the GP prior in (11), one can readily obtain the function posterior pdf $p(\varphi(\mathbf{z})|\mathcal{D}_t)$ via Bayes' rule as

$$p(\varphi(\mathbf{z})|\mathcal{D}_t) = \mathcal{N}(\varphi(\mathbf{z}); \mu_t(\mathbf{z}), \sigma_t^2(\mathbf{z})) \quad (12)$$

where its mean and variance have the following closed-form expressions

$$\mu_t(\mathbf{z}) = \mathbf{k}_t^\top(\mathbf{z})(\mathbf{K}_t + \sigma_o^2 \mathbf{I}_t)^{-1} \mathbf{y}_t \quad (13)$$

$$\sigma_t^2(\mathbf{z}) = \kappa(\mathbf{z}, \mathbf{z}) - \mathbf{k}_t^\top(\mathbf{z})(\mathbf{K}_t + \sigma_o^2 \mathbf{I}_t)^{-1} \mathbf{k}_t(\mathbf{z}) \quad (14)$$

where $\mathbf{k}_t(\mathbf{z}) := [\kappa(\mathbf{z}_1, \mathbf{z}), \dots, \kappa(\mathbf{z}_t, \mathbf{z})]^\top$. Notice that the posterior mean $\mu_t(\mathbf{z})$ is a weighted average of the observed function values $\mathbf{y}_t$, with the weights determined by evaluations of the kernel function at the input values. Besides, the posterior variance $\sigma_t^2(\mathbf{z})$ is equal to the prior covariance $\kappa(\mathbf{z}, \mathbf{z})$ minus the term corresponding to the variance reduction by observing $\mathbf{y}_t$.

Clearly, the performance of this GP predictor (13)-(14) highly hinges on the design of the kernel function $\kappa(\cdot, \cdot)$ over the input space. Accounting for both the continuous $\mathbf{x}_\tau$ for resource allocation and the categorical $\mathbf{c}_\tau$ for task offloading in the function input $\mathbf{z}_\tau$, as well as temporal variations across slots, three separate kernels are considered, which are $\kappa_x(\mathbf{x}_\tau, \mathbf{x}_{\tau'})$ over continuous inputs, $\kappa_c(\mathbf{c}_\tau, \mathbf{c}_{\tau'})$ over categorical inputs, and the temporal kernel $\kappa_{temp}(\tau, \tau')$.

Various kernel functions are available for continuous inputs; see [25]. A popular choice is the class of *Matérn* kernels

$$\kappa_x^{MT}(\mathbf{x}_\tau, \mathbf{x}_{\tau'}) = \frac{2^{1-\nu}}{\Gamma(\nu)} \left(\frac{\sqrt{2\nu} \|\mathbf{x}_\tau - \mathbf{x}_{\tau'}\|}{l} \right)^\nu \times B_\nu \left(\frac{\sqrt{2\nu} \|\mathbf{x}_\tau - \mathbf{x}_{\tau'}\|}{l} \right) \quad (15)$$

with parameter $\nu > 0$ controlling the smoothness of the learning function. The smaller ν is, the less smooth the sought function is assumed to be. In (15), l is the characteristic lengthscale, B_ν is a modified Bessel function, and Γ is the gamma function. Specifically, as $\nu \rightarrow \infty$, the kernel (15) boils down to the well-known radial basis function (RBF) $\kappa_x^{RBF}(\mathbf{x}_\tau, \mathbf{x}_{\tau'}) := \alpha \exp(-\frac{\|\mathbf{x}_\tau - \mathbf{x}_{\tau'}\|^2}{2l^2})$, where the pairwise similarity grows exponentially as a function of the squared distance between any two continuous inputs.

As for categorical variables, we follow [36] to adopt the kernel function $\kappa_c(\mathbf{c}_\tau, \mathbf{c}_{\tau'})$ as

$$\kappa_c(\mathbf{c}_\tau, \mathbf{c}_{\tau'}) = \frac{\omega}{M} \sum_{m=1}^M \mathbb{I}(c_\tau^m = c_{\tau'}^m) \quad (16)$$

where ω is the categorical kernel variance. Note that the categorical kernel defined in (16) is a special case of the RBF kernel with $\alpha = 1$ and $l \rightarrow 0$. To allow for a richer set of couplings between the continuous and categorical domains, a mixture of the sum and product compositions of the two kernels κ_x and κ_c is proposed for the kernel function $\kappa_{x,c}$ over continuous and categorical variables [36], i.e.,

$$\begin{aligned} \kappa_{x,c}([\mathbf{x}_\tau^\top, \mathbf{c}_\tau^\top]^\top, [\mathbf{x}_{\tau'}^\top, \mathbf{c}_{\tau'}^\top]^\top) \\ = (1-\lambda)[\kappa_c(\mathbf{c}_\tau, \mathbf{c}_{\tau'}) + \kappa_x(\mathbf{x}_\tau, \mathbf{x}_{\tau'})] \\ + \lambda \kappa_c(\mathbf{c}_\tau, \mathbf{c}_{\tau'}) \kappa_x(\mathbf{x}_\tau, \mathbf{x}_{\tau'}) \end{aligned} \quad (17)$$

where $\lambda \in [0, 1]$ weighs the contributions from the sum and product compositions of κ_c and κ_x . When $\lambda = 0$, only the sum composition exists in (17), leading to independence of the black-box function φ over the continuous and categorical domains with limited expressiveness. On the other hand, the pure product composition with $\lambda = 1$ will take the value of 0 if there is no pairwise overlap between two categorical variables $\mathbf{c}_\tau$ and $\mathbf{c}_{\tau'}$, that is, $\kappa_c(\mathbf{c}_\tau, \mathbf{c}_{\tau'}) = 0$ according to (16), thus preventing the GP model from learning. Towards overcoming the aforementioned two limitations, one can leverage a weighted combination of the sum and product components with $0 < \lambda < 1$ in (17).

To further capture the temporal variation of the black-box function φ due to the unknown system dynamics, the following temporal kernel function $\kappa_{temp}(\tau, \tau')$ is adopted based on [37]

$$\kappa_{temp}(\tau, \tau') = (1 - \rho)^{\frac{|\tau - \tau'|}{2}} \quad (18)$$

where $\rho \in [0, 1]$ is the hyperparameter that controls the level of temporal dynamics in the learning function φ . The larger the value of ρ , the more frequently φ varies over time. In particular, when $\rho = 0$, $\kappa_{temp}(\tau, \tau') = 1$ for any (τ, τ') , thus inducing no dynamics in φ .

Henceforth, applying the product composition of $\kappa_{x,c}$ (17) and κ_{temp} (18) yields the overall kernel function given by

$$\kappa(\mathbf{z}_\tau, \mathbf{z}_{\tau'}) = \kappa_{temp}(\tau, \tau') \kappa_{x,c}([\mathbf{x}_\tau^\top, \mathbf{c}_\tau^\top]^\top, [\mathbf{x}_{\tau'}^\top, \mathbf{c}_{\tau'}^\top]^\top). \quad (19)$$

It can be observed that the temporal kernel imposes a scaling factor on $\kappa_{x,c}$ based on the time separation of any pair of inputs. This agrees well with intuition that inputs that are well separated in time (i.e., large $|\tau - \tau'|$) yield less correlated function values for $\rho \neq 0$.

Remark 1 (Learning the GP Hyperparameters): The GP hyperparameters, collected in θ that consists of the characteristic length-scale l , categorical kernel variance ω , and the noise variance σ_o^2 , are optimized by maximizing the log marginal likelihood [25]

$$\begin{aligned} \mathcal{L}(\theta) &:= \log p(\mathbf{y}_t | \mathbf{Z}_t) = \log \left(\int p(\mathbf{y}_t | \varphi_t, \mathbf{Z}_t) p(\varphi_t | \mathbf{Z}_t) d\varphi_t \right) \\ &= -\frac{1}{2} \mathbf{y}_t^\top (\mathbf{K}_t + \sigma_o^2 \mathbf{I}_t)^{-1} \mathbf{y}_t \\ &\quad - \frac{1}{2} \log |\mathbf{K}_t + \sigma_o^2 \mathbf{I}_t| - \frac{t}{2} \log 2\pi \end{aligned} \quad (20)$$

where the first term involving the observations represents the data-fit; the second term indicates the complexity penalty; and, the last term is a normalization constant. Accordingly, the gradient of the $\mathcal{L}(\theta)$ with respect to the hyperparameters θ is given by

$$\begin{aligned} \frac{\partial \mathcal{L}(\theta)}{\partial \theta} &= \frac{1}{2} \mathbf{y}_t^\top (\mathbf{K}_t + \sigma_o^2 \mathbf{I}_t)^{-1} \frac{\partial (\mathbf{K}_t + \sigma_o^2 \mathbf{I}_t)}{\partial \theta} (\mathbf{K}_t + \sigma_o^2 \mathbf{I}_t)^{-1} \mathbf{y}_t \\ &\quad - \frac{1}{2} \text{tr} \left((\mathbf{K}_t + \sigma_o^2 \mathbf{I}_t)^{-1} \frac{\partial (\mathbf{K}_t + \sigma_o^2 \mathbf{I}_t)}{\partial \theta} \right) \end{aligned} \quad (21)$$

based on which the gradient-based optimizer is adopted to learn θ every δ time slots.

B. Acquisition for $\mathbf{z}_{t+1}$ Based on GP Surrogate Model

Having available GP-based posterior function model (12) with the form of kernel function specified by (19) at slot t , one is ready to select the next decisions $\mathbf{z}_{t+1}$. Coping with both categorical and continuous variables, this is certainly a nontrivial task, but can fortunately be handled by relying on the MAB framework. Since the cardinality of the categorical variables is exponential with respect to the number M of WDs, a scalable multi-agent MAB approach will be leveraged with each WD m acting as an agent simultaneously and independently determining its local task offloading decision $c_t^m \in \{0, 1, \dots, N\}$. As the overall reward function in the resultant MAB framework does not follow any statistical distribution, it is more sensible to rely on the adversarial MAB framework and adopt as the action selection rule the well-known exponential-weight algorithm for exploration and exploitation (EXP3) [38]. Per slot t , EXP3 maintains an unnormalized weight vector $\mathbf{w}_t^m := [w_t^m(0), w_t^m(1), \dots, w_t^m(N)]^\top$ for each WD m to guide the selection of its action. Next, we will delineate how each acquisition step of the time-varying BO selects categorical $\mathbf{c}_{t+1}$ and continuous $\mathbf{x}_{t+1}$ with the help of EXP3.

1) *Acquisition for Categorical Task Offloading Decisions:* Given $\mathbf{w}_t^m$ from the end of slot t , each agent m in EXP3 draws its action c_{t+1}^m randomly according to the probability vector $\mathbf{q}_t^m := [q_t^m(0), q_t^m(1), \dots, q_t^m(N)]^\top$ with [38]

$$q_t^m(k) = (1 - \gamma) \frac{w_t^m(k)}{\sum_{k'=0}^N w_t^m(k')} + \frac{\gamma}{N+1}, \forall k \in \{0, 1, \dots, N\} \quad (22)$$

where $\gamma \in (0, 1]$ is the coefficient that balances *exploitation* given by the normalized weight in the first factor and *exploration* from the uniform probability in the second term. Specifically, by including the uniform distribution, EXP3 allows all $N+1$ decisions to be explored per agent (WD) so as to get good reward estimates.

2) *Acquisition for Analog-Amplitude Resource Allocation Decisions:* With the categorical task offloading decisions $\mathbf{c}_{t+1}$ at hand, the analog-amplitude resource allocation decisions $\mathbf{x}_{t+1}$ are selected by finding the maximizer of the celebrated upper confidence bound (UCB)-based acquisition function as [39]

$$\begin{aligned} \mathbf{x}_{t+1} &= \arg \max_{0 < \mathbf{x} \leq \mathbf{x}_{peak}} u_{t+1}(\mathbf{x} | \mathcal{D}_t, \mathbf{c}_{t+1}, t+1) \\ &:= \mu_t(\mathbf{x}, \mathbf{c}_{t+1}, t+1) + \sqrt{\zeta_{t+1} \sigma_t^2(\mathbf{x}, \mathbf{c}_{t+1}, t+1)} \end{aligned} \quad (23)$$

where the coefficient $\zeta_{t+1} \geq 0$ nicely balances the exploitation and exploration that are signified by the posterior mean μ_t (13) and variance σ_t^2 (14), respectively. With closed-form expressions of μ_t and σ_t^2 at hand, one can readily solve (23) via off-the-shelf gradient-based solvers.

3) *Weight Update in EXP3:* Upon deploying $(\mathbf{c}_{t+1}, \mathbf{x}_{t+1})$ into the MEC system to yield the observed reward y_{t+1} , EXP3 capitalizes on the importance sampling rule to obtain an unbiased estimate of the reward value as

$$\hat{\varphi}_{t+1}^m(k) = \frac{y_{t+1} \mathbb{I}(c_{t+1}^m = k)}{q_t^m(k)}, \forall k \in \{0, 1, \dots, N\}, m \in \mathcal{M} \quad (24)$$

based on which the corresponding weight is updated using the exponential rule as

$$\begin{aligned} w_{t+1}^m(k) &= w_t^m(k) \exp \left(\frac{\gamma \hat{\varphi}_{t+1}^m(k)}{N+1} \right) \\ &= w_0^m(k) \exp \left(\frac{\gamma \sum_{\tau=1}^{t+1} \hat{\varphi}_\tau^m(k)}{N+1} \right), \end{aligned} \quad \forall k \in \{0, 1, \dots, N\}, m \in \mathcal{M}. \quad (25)$$

It is evident that $w_{t+1}^m(k)$ summarizes the cumulative rewards up to slot $t+1$ for action k under WD m , and thus represents the effect of exploitation in (22). Relying on the multi-agent MAB framework with the EXP3 to select the categorical variables, a preliminary regret bound of the proposed BO approach is given by the following lemma.

Lemma 1: Selecting $\gamma = \min \left\{ 1, \sqrt{\frac{(N+1) \ln(N+1)}{(e-1)T}} \right\}$ and having the kernel function satisfy some regularity conditions, the cumulative regret of the BO-based approach is upper bounded by

$$\mathcal{R}(T) \leq \mathcal{O}(M \sqrt{T(N+1) \ln(N+1)}). \quad (26)$$

Proof: For a single agent m , the cumulative regret $\mathcal{R}_m(T)$ when the categorical decisions of all the other agents are fixed to $\mathbf{c}_{\setminus m}$ is given by

$$\begin{aligned} \mathcal{R}_m(T) &= \sup_{\mathbf{c}_{\setminus m}} \left\{ \max_{c_m^* \in \{0, 1, \dots, N\}} \sum_{t=1}^T \varphi_t(c_m^*, \mathbf{c}_{\setminus m}) - \mathbb{E} \left[\sum_{t=1}^T \varphi_t(c_t^m, \mathbf{c}_{\setminus m}) \right] \right\}, \end{aligned} \quad (27)$$

where c_m^* is the best single action over all time slots. According to [38], Theorem 3.1], for the bounded reward $\varphi_t \leq 1, \forall t$ and parameter $\gamma = \min \left\{ 1, \sqrt{\frac{(N+1) \ln(N+1)}{(e-1)T}} \right\}$ in EXP3, we have

$$\begin{aligned} \mathcal{R}_m(T) &\leq \sup_{\mathbf{c}_{\setminus m}} \left\{ 2.63 \sqrt{T(N+1) \ln(N+1)} \right\} \\ &\leq 2.63 \sqrt{T(N+1) \ln(N+1)}. \end{aligned} \quad (28)$$

Then, the cumulative regret of the multi-agent EXP3 approach has the bound

$$\mathcal{R}(T) = \sum_{m=1}^M \mathcal{R}_m(T) \leq 2.63M \sqrt{T(N+1) \ln(N+1)}. \quad (29)$$

Such cumulative regret is sub-linear as

$$\lim_{T \rightarrow \infty} \frac{\mathcal{R}(T)}{T} = 0. \quad (30)$$

□

The sublinear cumulative regret bound in Lemma 1 dictates the convergence of the proposed BO approach.

The pseudo-code of the overall time-varying BO approach is summarized in Algorithm 1.

The main computational complexity of the proposed BO approach comes from updating the GP-based surrogate model using the acquired data $\mathcal{D}_t$ (cf. (13)-(14)), and is of order $\mathcal{O}(|\mathcal{D}_t|^3)$ with $|\mathcal{D}_t|$ denoting the cardinality of $\mathcal{D}_t$. In the BO context, $|\mathcal{D}_t|$ is usually pretty small (a few hundred),

Algorithm 1 Time-Varying BO for Dynamic MEC Management

```

1: Initialization:  $\mathcal{D}_0, w_0^m(k) = 1, \forall k \in \{0, 1, \dots, N\}, m \in \mathcal{M}$ ;
2: for  $t = 0 : T - 1$  do
3:   if  $t \bmod \delta = 1$  then
4:     Learn GP hyperparameters  $\theta$  via multi-started
       gradient descent using (21);
5:   end if
6:   Calculate the posterior mean  $\mu_t$  and variance  $\sigma_t^2$ 
       according to (13)–(14) given  $\mathcal{D}_t$ ;
7:   Compute the action distribution  $\mathbf{q}_t^m, \forall m \in \mathcal{M}$ 
       according to (22);
8:   Draw the discrete task offloading decision  $c_{t+1}^m$ 
       randomly according to  $\mathbf{q}_t^m, \forall m \in \mathcal{M}$ ;
9:   Acquire the analog-amplitude resource allocation deci-
       sions  $\mathbf{x}_{t+1}$  by solving (23);
10:  Deploy decisions  $\mathbf{z}_{t+1} := [\mathbf{c}_{t+1}^\top, \mathbf{x}_{t+1}^\top, t+1]^\top$  to MEC
       system to observe  $y_{t+1}$ ;
11:   $\mathcal{D}_{t+1} = \mathcal{D}_t \cup \{(\mathbf{z}_{t+1}, y_{t+1})\}$  and update  $w_{t+1}^m(k)$ 
       via (25),  $\forall k \in \{0, 1, \dots, N\}, m \in \mathcal{M}$ ;
12: end for

```

thus rendering the computational complexity of no concern. Even in cases when $|\mathcal{D}_t|$ is rather large, one can rely on a replay memory [40] to store the C most recent data samples, reducing the complexity to $\mathcal{O}(C^3)$ per slot. Alternatively, one can effect scalability via the scalable GP approaches, of which the two most prevalent ones are the inducing points-based framework [41] and random feature-based paradigm [42].

IV. CONTEXTUAL TIME-VARYING BO FOR DYNAMIC MEC MANAGEMENT

So far, we have introduced a time-varying BO approach for dynamic MEC management under the bandit setting, where the temporal dynamics of the black-box reward function φ is captured by incorporating the temporal kernel. In some scenarios, in addition to the observed reward value, one could have access to a subset of the system state. Here, each WD m could report its task characterization variables (I_t^m, L_t^m) to the central controller per slot t . The goal of this section is then to generalize the time-varying BO approach for more informed decision-making by leveraging such state information, which will also be termed as “context” hereafter.

In the resultant contextual time-varying BO approach, the black-box reward function $\varphi(\bar{\mathbf{z}}_t)$ has the augmented input $\bar{\mathbf{z}}_t := [\mathbf{z}_t^\top, \mathbf{s}_t^\top]^\top$, where $\mathbf{s}_t$ is the context vector that collects the observed state information as $\mathbf{s}_t := [I_t^1, \dots, I_t^M, L_t^1, \dots, L_t^M]^\top$. As with the time-varying BO approach in the previous section, the generalized counterpart here still consists of two steps per iteration, namely, GP-based surrogate model learning and the acquisition of new decisions.

For the former, a GP prior is postulated for φ as $\varphi \sim \mathcal{GP}(0, \bar{\kappa}(\bar{\mathbf{z}}, \bar{\mathbf{z}}'))$, where the kernel function $\bar{\kappa}$ has to be adapted to capture correlation from the contextual input. Inspired by [43], $\bar{\kappa}(\bar{\mathbf{z}}_\tau, \bar{\mathbf{z}}_{\tau'})$ is proposed as the product combination

of three separate kernels given by

$$\bar{\kappa}(\bar{\mathbf{z}}_\tau, \bar{\mathbf{z}}_{\tau'}) = \kappa_s(\mathbf{s}_\tau, \mathbf{s}_{\tau'}) \kappa_{temp}(\tau, \tau') \kappa_{x,c}([\mathbf{x}_\tau^\top, \mathbf{c}_\tau^\top]^\top, [\mathbf{x}_{\tau'}^\top, \mathbf{c}_{\tau'}^\top]^\top) \quad (31)$$

where $\kappa_s(\mathbf{s}_\tau, \mathbf{s}_{\tau'})$ is the contextual kernel over the observed context variables, and $\kappa_{x,c}$ and κ_{temp} are given by (17) and (18). Given the GP prior and a set of input-output data pairs $\bar{\mathcal{D}}_t := \{(\bar{\mathbf{z}}_\tau, y_\tau)\}_{\tau=1}^t$, the posterior pdf for the reward function is given by (cf. (12))

$$p(\varphi(\bar{\mathbf{z}}) | \bar{\mathcal{D}}_t) = \mathcal{N}(\varphi(\bar{\mathbf{z}}); \bar{\mu}_t(\bar{\mathbf{z}}), \bar{\sigma}_t^2(\bar{\mathbf{z}})) \quad (32)$$

where the closed-form expressions of the mean $\bar{\mu}_t$ and variance $\bar{\sigma}_t^2$ can be obtained similarly as in (13)–(14) by including the context vectors in the input, i.e.,

$$\bar{\mu}_t(\bar{\mathbf{z}}) = \bar{\mathbf{k}}_t^\top(\bar{\mathbf{z}})(\bar{\mathbf{K}}_t + \sigma_o^2 \mathbf{I}_t)^{-1} \mathbf{y}_t \quad (33)$$

$$\bar{\sigma}_t^2(\bar{\mathbf{z}}) = \bar{\kappa}(\bar{\mathbf{z}}, \bar{\mathbf{z}}) - \bar{\mathbf{k}}_t^\top(\bar{\mathbf{z}})(\bar{\mathbf{K}}_t + \sigma_o^2 \mathbf{I}_t)^{-1} \bar{\mathbf{k}}_t(\bar{\mathbf{z}}). \quad (34)$$

Here $\bar{\mathbf{k}}_t(\bar{\mathbf{z}}) := [\bar{\kappa}(\bar{\mathbf{z}}_1, \bar{\mathbf{z}}), \dots, \bar{\kappa}(\bar{\mathbf{z}}_t, \bar{\mathbf{z}})]^\top$ and $\bar{\mathbf{K}}_t$ is the $t \times t$ covariance matrix with (τ, τ') -th entry $[\bar{\mathbf{K}}_t]_{\tau, \tau'} := \bar{\kappa}(\bar{\mathbf{z}}_\tau, \bar{\mathbf{z}}_{\tau'})$. Similar as Remark 1 in Sec. III-A, the GP hyperparameters $\bar{\theta}$ are optimized every δ slots via log marginal likelihood maximization using (21).

As for the acquisition of task offloading and resource allocation decisions for slot $t+1$, contextual time-varying BO proceeds as in Sec. III-B by first selecting the categorical $\mathbf{c}_{t+1}$ via the EXP3 approach based on the multi-agent MAB framework, and then choosing the continuous $\mathbf{x}_{t+1}$ using the UCB rule. Here, the latter has to take into account the observed context vector $\mathbf{s}_{t+1}$, thus yielding $\mathbf{x}_{t+1}$ given by

$$\begin{aligned} \mathbf{x}_{t+1} = \arg \max_{0 < \mathbf{x} \leq \mathbf{x}_{peak}} \bar{u}_{t+1}(\mathbf{x} | \bar{\mathcal{D}}_t, \mathbf{c}_{t+1}, \mathbf{s}_{t+1}, t+1) \\ := \bar{\mu}_t(\mathbf{x}, \mathbf{c}_{t+1}, \mathbf{s}_{t+1}, t+1) + \sqrt{\bar{\zeta}_{t+1}} \\ \times \bar{\sigma}_t^2(\mathbf{x}, \mathbf{c}_{t+1}, \mathbf{s}_{t+1}, t+1) \end{aligned} \quad (35)$$

where $\bar{\zeta}_{t+1} \geq 0$ is the coefficient that balances exploration and exploitation. Please refer to Algorithm 2 for the detailed implementation of the contextual time-varying BO approach.

V. SIMULATION RESULTS

In this section, numerical tests were conducted to evaluate the performance of the proposed BO approaches for dynamic MEC management. In the multi-user multi-server MEC system with M WDs and N BSs, the time-varying wireless channel $h_{t,n}^{m,n}$ from WD m to BS n is modelled as Rician fading channel

$$h_t^{m,n} = \sqrt{\frac{K}{K+1}} h_{t,LoS}^{m,n} + \sqrt{\frac{1}{K+1}} h_{t,NLoS}^{m,n}, \forall m, n, t \quad (36)$$

where $h_{t,LoS}^{m,n}$ denotes the deterministic line of sight (LoS) component determined by the locations of BS n and WD m ; $h_{t,NLoS}^{m,n}$ stands for the non-LoS component following the independent and identically distributed (i.i.d.) standard Gaussian distribution; and $K \geq 0$ is the Rician factor representing the ratio of the power in the LoS component to the power in the non-LoS component. Note that a larger K implies

Algorithm 2 Contextual Time-Varying BO for Dynamic MEC Management

- 1: **Initialization:** observation dataset $\bar{\mathcal{D}}_0$ and $w_0^m(k) = 1, \forall k \in \{0, 1, \dots, N\}, m \in \mathcal{M}$;
- 2: **for** $t = 0 : T - 1$ **do**
- 3: **if** $t \bmod \delta = 1$ **then**
- 4: Learn GP Hyperparameters $\bar{\theta}$ via multi-started gradient descent using (21);
- 5: **end if**
- 6: Calculate the mean $\bar{\mu}_t$ and variance $\bar{\sigma}_t^2$ in the posterior pdf (32) according to (33)–(34);
- 7: Observe the contextual information $\mathbf{s}_{t+1}$;
- 8: Compute the action distribution $\mathbf{q}_t^m, \forall m \in \mathcal{M}$ according to (22);
- 9: Draw the discrete task offloading decision c_{t+1}^m randomly according to $\mathbf{q}_t^m, \forall m \in \mathcal{M}$;
- 10: Acquire the analog-amplitude resource allocation decisions $\mathbf{x}_{t+1}$ by solving (35);
- 11: Deploy decisions $(c_{t+1}, \mathbf{x}_{t+1})$ to the MEC system to observe y_{t+1} ;
- 12: $\bar{\mathcal{D}}_{t+1} = \bar{\mathcal{D}}_t \cup \{(\bar{\mathbf{z}}_{t+1}, y_{t+1})\}$ and update $w_{t+1}^m(k)$ via (25), $\forall k \in \{0, 1, \dots, N\}, m \in \mathcal{M}$;
- 13: **end for**

milder fading effect. The total average channel gain follows the free-space path loss model $|\bar{h}_t^{m,n}|^2 = A_d (\frac{3 \times 10^8}{4\pi\phi d_{m,n}})^{PL}, \forall t$, where $A_d = 4.11$ denotes the antenna gain, $\phi = 915$ MHz is the carrier frequency, $d_{m,n}$ represents the distance (measured by meters) between WD m and BS n , and $PL = 3$ signifies the pass loss exponent.

In addition, the means of time-varying edge CPU frequencies $\{f_{c,t}^n\}_{n,t}$, task computational workloads $\{L_t^m\}_{m,t}$, and task input data sizes $\{I_t^m\}_{m,t}$ are 26 GHz, 125 Mcycles, and 1250 KBytes, respectively [3], [4], [7], [8]. Specifically, the generation rules are as follows

$$f_{c,t}^n = (\bar{f}_c + \tilde{f}_{c,t}^n) \times 10^9 \text{ Hz} \quad \forall n, t \quad (37a)$$

$$L_t^m = (\bar{L} + \tilde{L}_t^m) \times 10^6 \text{ Cycles} \quad \forall m, t \quad (37b)$$

$$I_t^m = (\bar{I} + \tilde{I}_t^m) \times 10^4 \text{ Bytes} \quad \forall m, t \quad (37c)$$

where $\bar{f}_c = 26$, $\bar{L} = 125$, $\bar{I} = 125$, and the dynamic components $\tilde{f}_{c,t}^n$, $\tilde{L}_t^m$, and $\tilde{I}_t^m$ are evolved based on the following first-order Markovian processes

$$\tilde{f}_{c,1}^n = e_{f,1}^n \quad \tilde{L}_1^m = e_{L,1}^m \quad \tilde{I}_1^m = e_{I,1}^m \quad (38a)$$

$$\tilde{f}_{c,t+1}^n = \sqrt{1-\eta} \tilde{f}_{c,t}^n + \sqrt{\eta} e_{f,t+1}^n, \quad e_{f,t+1}^n \sim \mathcal{N}(0, 3) \quad (38b)$$

$$\tilde{L}_{t+1}^m = \sqrt{1-\eta} \tilde{L}_t^m + \sqrt{\eta} e_{L,t+1}^m, \quad e_{L,t+1}^m \sim \mathcal{N}(0, 3) \quad (38c)$$

$$\tilde{I}_{t+1}^m = \sqrt{1-\eta} \tilde{I}_t^m + \sqrt{\eta} e_{I,t+1}^m, \quad e_{I,t+1}^m \sim \mathcal{N}(0, 3) \quad (38d)$$

where the process noises $e_{f,t}^n$, $e_{L,t}^m$, and $e_{I,t}^m$ are i.i.d., and $\eta \in [0, 1]$ is the parameter adjusting the level of temporal dynamics in these system state variables. In particular, $\eta = 0$ represents the time-invariant scenario, while $\eta = 1$ indicates the independent system dynamics across time slots [37].

Besides, the peak transmit power P_{peak} and computational frequency f_{peak} of each WD are equal to 100 mW and 10^8 Hz, respectively. To be aligned with commercial practise, the




Fig. 2. Average regret under different temporal kernel hyperparameters ρ for the time-varying BO approach.

computing efficiency coefficient ξ of the WDs in (3) is chosen as $\xi = 10^{-26}$ [44]. We set the channel additive white Gaussian noise power $\sigma^2 = 10^{-10}$ W, and the bandwidth $W = 2$ MHz. The prior weights of the time delay and energy consumption cost of the WDs in (10) are set as $\beta_d = \beta_e = 0.5$.

For the proposed (contextual) time-varying BO approaches, the *Matérn* kernel (15) with parameter $\nu = 5/2$ is adopted for the kernel κ_x over continuous variables. The weight λ regarding the sum and product kernel compositions in (17) is set to 0.5. The coefficients $\zeta_t = \hat{\zeta}_t = 2, \forall t$, in UCB-based acquisition rules (23) and (35). Unless otherwise stated, the other kernel hyperparameters are optimized by maximizing the log marginal likelihood every $\delta = 10$ slots via multi-started gradient descent. The performance measure of the competing methods is given by the notion of regret. By denoting the maximizer of φ_t as $(\mathbf{c}_t^*, \mathbf{x}_t^*)$, the instantaneous regret per slot t is

$$g_t := \varphi_t(\mathbf{c}_t^*, \mathbf{x}_t^*) - \varphi_t(\mathbf{c}_t, \mathbf{x}_t), \quad (39)$$

based on which the cumulative and average regrets are denoted as $G_T := \sum_{t=1}^T g_t$ and $\bar{G}_T := G_T/T$, respectively. It is worth mentioning that $(\mathbf{c}_t^*, \mathbf{x}_t^*)$ are obtained by relying on explicit cost function in (P2) with known system state information. All the methods are run for 200 time slots and the average performances over 100 random repetitions are reported.

A. Effect of Kernel Hyperparameters

To study the effect of temporal and contextual kernel hyperparameters on the performance of the proposed BO approaches, a 2-BS MEC system with $M = 2$ WDs is first considered, where the distances from the WDs to BSs are $[d_{1,1}, d_{1,2}, d_{2,1}, d_{2,2}] = [20, 13, 15, 18]$ meters, the Rician factor in (36) used to generate the channel gain is $K = 4$, and η in (38) is set to 0.2. Fig. 2 depicts the average regret of time-varying BO as a function of the time slot under different values of the temporal kernel hyperparameter ρ in (18). It can be readily observed that the regret performance improves and then deteriorates as the value of ρ increases. Specifically,




Fig. 3. Average regret under different contextual and temporal kernel hyperparameters for the contextual time-varying BO approach.

$\rho = 0.048$ achieves the lowest average regret by best capturing the temporal variation in the black-box objective function.

Further considering the contextual time-varying BO where a *Matérn* kernel with $\nu = 5/2$ in (15) is adopted for the contextual κ_s , the curves of the average regret for various contextual and temporal kernel hyperparameters are presented in Fig. 3, where it is evident that the best-performing hyperparameter set is given by $\rho = 0.02$ and $l = 0.2$ in the temporal and contextual kernel, respectively. Notice that the best-performing hyperparameter ρ of the temporal kernel in the contextual time-varying BO is smaller than that in the time-varying BO. To put it equivalently, the temporal kernel in the latter captures more dynamics in the objective function than that in the former. This phenomenon can be explained by that the observed contextual state information including time-varying task computational workload L_t^m and input data size I_t^m accounts for a portion of the overall dynamics, yielding lower degree of dynamics to be represented by the temporal kernel in the contextual time-varying BO.

B. Performance Comparison

For performance comparison, four existing schemes are employed as baselines, namely, the MAB [38], bandit convex optimization (BCO) [15], the conventional time-invariant BO approach [23], and the random scheme. Since MAB can only cope with discrete decision variables, we discretized the analog-amplitude resource allocation variables into 5 levels and then adopted the multi-agent EXP3 method [38] for learning. In BCO, the analog-amplitude resource allocation variables are obtained by constructing gradient estimates using evaluated function values, while the discrete offloading variables are still sought based on MAB as in the proposed BO approaches. Besides, time-invariant BO method neglects both temporal and contextual information in MEC systems. We additionally include a random server selection scheme with resource allocation variables being half of their corresponding peak values following [37].

With properly selected temporal and contextual kernel hyperparameters, the average regret curves of all the




Fig. 4. Comparison of average regret under the 2-BS and 2-WD MEC system with Rician factor $K = 4$ and $\eta = 0.2$.




Fig. 5. Comparison of average regret under the 2-BS and 2-WD MEC system with Rician factor $K = 9$ and $\eta = 0.02$.

competing approaches are presented in Fig. 4 for the 2-BS and 2-WD MEC system with $[d_{1,1}, d_{1,2}, d_{2,1}, d_{2,2}] = [20, 13, 15, 18]$, $K = 4$ and $\eta = 0.2$. Specifically, the temporal kernel hyperparameter in the time-varying BO approach is chosen as $\rho = 0.048$. As for contextual time-varying BO algorithm, the temporal kernel hyperparameter ρ and the lengthscale l of the contextual kernel are set to 0.02 and 0.2, respectively. As shown in Fig. 4, it is evident that all the bandit-based methods outperform the random scheme. In addition, our proposed time-varying BO approach outperforms the three benchmarks, namely, time-invariant BO, MAB, and BCO, by around 1.21%, 8.51% and 25.72% in average regret after 200 time slots. This suggests the benefits of adapting temporal information-aided Bayesian approach to the black-box optimization with both categorical (i.e., task offloading) and analog-amplitude (i.e., resource allocation) variables. By further utilizing the observed context information (i.e., the characteristics of computational tasks) via the contextual kernel, the novel contextual time-varying BO method achieves 1.81% and 3% lower average regret than time-varying BO and traditional BO after 200 slots.




Fig. 6. Comparison of average regret under the 2-BS and 5-WD MEC system with Rician factor $K = 5.67$ and $\eta = 0.2$.

Further, the performances of the proposed BO approaches are investigated in the 2-BS and 2-WD MEC system with a smaller scale of system dynamics, that is given by the Rician factor $K = 9$ in (36) and the temporal variation parameter $\eta = 0.02$ in (38). The temporal kernel parameter ρ is set to 0.011 in the time-varying BO approach, while $\rho = 0.0045$ and contextual kernel lengthscale $l = 0.2$ are chosen in contextual time-varying BO. Here, the values of ρ in both cases are smaller than the counterparts in Fig. 4, what is in accordance with the degree of the underlying temporal dynamics. Compared with the alternative time-invariant BO, MAB, BCO, and random schemes, the proposed (contextual) time-varying BO methods reduce the average regret by approximately 1.49%, 34.77%, 47.75% and 47.89% after 200 slots as showcased in Fig. 5. In addition, the performance of the time-invariant BO method is close to the proposed time-varying BO alternatives due to such small-scale system dynamics.

C. Effect of Network Size

Lastly, the performances of all the schemes are assessed as the number of WDs and BSs varies. Consider first a 2-BS MEC system with a larger number $M = 5$ of WDs, where the time-varying system state is generated using the Rician factor $K = 5.67$ in (36) and temporal variation factor $\eta = 0.2$ in (38). In this case, $\rho = 0.018$ in time-varying BO approach, while $\rho = 0.006$ and $l = 0.5$ in contextual time-varying BO strategy. Still, the proposed (contextual) time-varying BO methods outperform the other four alternatives by leveraging temporal and contextual information as shown in Fig. 6. In addition, we observe that the random scheme attains lower average regret than the BCO methods. It is because the considered random scheme leverages the fixed half-peak resource allocation strategy, while the whole continuous action space needs to be explored in the BCO method for the resource allocation.

Moreover, fixing the number M of WDs as 2, the average EDC over slots is plotted as a function of the number N of BSs for all the competing methods in Fig. 7. Here, the




Fig. 7. Impact of MEC network size on average energy-delay cost.




Fig. 8. Impact of mean of task input data size on average energy-delay cost under real-world 2-BS and 2-WD MEC system.

Rician factor in (36) and value of η in (38) are set to $K = 4$ and $\eta = 0.2$ respectively. Apparently, the two proposed BO approaches achieve lower average EDC than the other four baselines. Additionally, the average EDC of all the methods decreases as the network size grows by better exploiting the diverse computing capacities and channel conditions of the edge servers.

In Fig. 8, we further illustrate the impact of the mean of task input data size on the average EDC over slots under a real-world 2-BS and 2-WD MEC system with Rician factor $K = 4$ and time variation parameter $\eta = 0.2$. Specifically, we consider a widely-used public dataset containing the geographical information of real-world edge servers and anonymous mobile users in Melbourne CBD area in Australia [45], [46]. In this experiment, the locations of $N = 2$ edge servers and $M = 2$ WDs are extracted from the dataset to simulate the MEC system. It is observed that the proposed BO approaches attain lower average EDC than the other baselines under different input data size settings in the real-world MEC system.

VI. CONCLUSION

BO for dynamic MEC management was studied in this paper. Different from prior works in time-varying MEC systems, the focus was online joint optimization of discrete task offloading decisions and analog-amplitude resource allocation strategies by minimizing the EDC using only bandit observations at queried points. Specifically, by exploiting both temporal and contextual information, we developed two novel BO approaches that incorporate the strength of the MAB framework. Numerical tests under different MEC network sizes demonstrated the effectiveness of the proposed BO approaches.

REFERENCES

- [1] J. Yan, Q. Lu, and G. B. Giannakis, "Bayesian optimization for task offloading and resource allocation in mobile edge computing," in *Proc. 56th Asilomar Conf. Signals, Syst., Comput.*, Oct. 2022, pp. 1086–1090.
- [2] Y. Mao, C. You, J. Zhang, K. Huang, and K. B. Letaief, "A survey on mobile edge computing: The communication perspective," *IEEE Commun. Surveys Tuts.*, vol. 19, no. 4, pp. 2322–2358, 4th Quart., 2017.
- [3] C. You, K. Huang, and H. Chae, "Energy efficient mobile cloud computing powered by wireless energy transfer," *IEEE J. Sel. Areas Commun.*, vol. 34, no. 5, pp. 1757–1771, May 2016.
- [4] C. You, K. Huang, H. Chae, and B.-H. Kim, "Energy-efficient resource allocation for mobile-edge computation offloading," *IEEE Trans. Wireless Commun.*, vol. 16, no. 3, pp. 1397–1411, Mar. 2017.
- [5] Y. Wang, M. Sheng, X. Wang, L. Wang, and J. Li, "Mobile-edge computing: Partial computation offloading using dynamic voltage scaling," *IEEE Trans. Commun.*, vol. 64, no. 10, pp. 4268–4282, Oct. 2016.
- [6] T. Q. Dinh, J. Tang, Q. D. La, and T. Q. S. Quek, "Offloading in mobile edge computing: Task allocation and computational frequency scaling," *IEEE Trans. Commun.*, vol. 65, no. 8, pp. 3571–3584, Aug. 2017.
- [7] S. Bi and Y. J. Zhang, "Computation rate maximization for wireless powered mobile-edge computing with binary computation offloading," *IEEE Trans. Wireless Commun.*, vol. 17, no. 6, pp. 4177–4190, Jun. 2018.
- [8] J. Yan, S. Bi, Y. J. Zhang, and M. Tao, "Optimal task offloading and resource allocation in mobile-edge computing with inter-user task dependency," *IEEE Trans. Wireless Commun.*, vol. 19, no. 1, pp. 235–250, Jan. 2020.
- [9] Y. Mao, J. Zhang, and K. B. Letaief, "Dynamic computation offloading for mobile-edge computing with energy harvesting devices," *IEEE J. Sel. Areas Commun.*, vol. 34, no. 12, pp. 3590–3605, Dec. 2016.
- [10] Y. Mao, J. Zhang, S. H. Song, and K. B. Letaief, "Stochastic joint radio and computational resource management for multi-user mobile-edge computing systems," *IEEE Trans. Wireless Commun.*, vol. 16, no. 9, pp. 5994–6009, Sep. 2017.
- [11] Z. Yang, S. Bi, and Y. A. Zhang, "Dynamic offloading and trajectory control for UAV-enabled mobile edge computing system with energy harvesting devices," *IEEE Trans. Wireless Commun.*, vol. 21, no. 12, pp. 10515–10528, Dec. 2022.
- [12] T. Chen, Q. Ling, Y. Shen, and G. B. Giannakis, "Heterogeneous online learning for 'thing-adaptive' fog computing in IoT," *IEEE Internet Things J.*, vol. 5, no. 6, pp. 4328–4341, Dec. 2018.
- [13] T. Chen, Q. Ling, and G. B. Giannakis, "An online convex optimization approach to proactive network resource allocation," *IEEE Trans. Signal Process.*, vol. 65, no. 24, pp. 6350–6364, Dec. 2017.
- [14] E. C. Hall and R. M. Willett, "Online convex optimization in dynamic environments," *IEEE J. Sel. Topics Signal Process.*, vol. 9, no. 4, pp. 647–662, Jun. 2015.
- [15] A. D. Flaxman, A. T. Kalai, and H. B. McMahan, "Online convex optimization in the bandit setting: Gradient descent without a gradient," in *Proc. ACM SODA*, Vancouver, BC, Canada, Jan. 2005, pp. 385–394.
- [16] Y. Ding and J. Lavaei, "Structured projection-free online convex optimization with multi-point bandit feedback," in *Proc. 60th IEEE Conf. Decis. Control (CDC)*, Dec. 2021, pp. 28–40.
- [17] O. Shamir, "An optimal algorithm for bandit and zero-order convex optimization with two-point feedback," *J. Mach. Learn. Res.*, vol. 18, no. 1, pp. 1703–1713, Jan. 2017.
- [18] T. Chen and G. B. Giannakis, "Bandit convex optimization for scalable and dynamic IoT management," *IEEE Internet Things J.*, vol. 6, no. 1, pp. 1276–1286, Feb. 2019.
- [19] B. Wu, T. Chen, W. Ni, and X. Wang, "Multi-agent multi-armed bandit learning for online management of edge-assisted computing," *IEEE Trans. Commun.*, vol. 69, no. 12, pp. 8188–8199, Dec. 2021.
- [20] B. Li, T. Chen, and G. B. Giannakis, "Secure mobile edge computing in IoT via collaborative online learning," *IEEE Trans. Signal Process.*, vol. 67, no. 23, pp. 5922–5935, Dec. 2019.
- [21] Y. Sun et al., "Adaptive learning-based task offloading for vehicular edge computing systems," *IEEE Trans. Veh. Technol.*, vol. 68, no. 4, pp. 3061–3074, Apr. 2019.
- [22] Y. Sun, S. Zhou, and J. Xu, "EMM: Energy-aware mobility management for mobile edge computing in ultra dense networks," *IEEE J. Sel. Areas Commun.*, vol. 35, no. 11, pp. 2637–2646, Nov. 2017.
- [23] P. I. Frazier, "A tutorial on Bayesian optimization," 2018, *arXiv:1807.02811*.
- [24] Q. Lu, K. D. Polyzos, B. Li, and G. B. Giannakis, "Surrogate modeling for Bayesian optimization beyond a single Gaussian process," 2022, *arXiv:2205.14090*.
- [25] C. E. Rasmussen and C. K. Williams, *Gaussian Processes for Machine Learning*. Cambridge, MA, USA: MIT Press, 2006.
- [26] Q. Lu, G. Karanikolas, Y. Shen, and G. B. Giannakis, "Ensemble Gaussian processes with spectral features for online interactive learning with scalability," in *Proc. Int. Conf. Artif. Intell. Statist.*, Jun. 2020, pp. 1910–1920.
- [27] Q. Lu, G. V. Karanikolas, and G. B. Giannakis, "Incremental ensemble Gaussian processes," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 45, no. 2, pp. 1876–1893, Feb. 2023.
- [28] K. D. Polyzos, Q. Lu, and G. B. Giannakis, "Ensemble Gaussian processes for online learning over graphs with adaptivity and scalability," *IEEE Trans. Signal Process.*, vol. 70, pp. 17–30, 2022.
- [29] J. Snoek, H. Larochelle, and R. P. Adams, "Practical Bayesian optimization of machine learning algorithms," in *Proc. Adv. Neural Inf. Process. Syst.*, vol. 25, Dec. 2012, pp. 1–9.
- [30] K. Korovina et al., "ChemBO: Bayesian optimization of small organic molecules with synthesizable recommendations," in *Proc. Int. Conf. Artif. Intell. Statist.*, Jun. 2020, pp. 3393–3403.
- [31] A. Cully, J. Clune, D. Tarapore, and J.-B. Mouret, "Robots that can adapt like animals," *Nature*, vol. 521, no. 7553, pp. 503–507, May 2015.
- [32] L. Maggi, A. Valcarce, and J. Hoyer, "Bayesian optimization for radio resource management: Open loop power control," *IEEE J. Sel. Areas Commun.*, vol. 39, no. 7, pp. 1858–1871, Jul. 2021.
- [33] R. M. Dreifuerst et al., "Optimizing coverage and capacity in cellular networks using machine learning," in *Proc. IEEE Int. Conf. Acoust., Speech Signal Process. (ICASSP)*, Jun. 2021, pp. 8138–8142.
- [34] S. Yang, B. Liu, Z. Hong, and Z. Zhang, "Bayesian optimization-based beam alignment for mmWave MIMO communication systems," 2022, *arXiv:2207.14174*.
- [35] H. Liu and G. Cao, "Deep learning video analytics through online learning based edge computing," *IEEE Trans. Wireless Commun.*, vol. 21, no. 10, pp. 8193–8204, Oct. 2022.
- [36] B. Ru, A. S. Alvi, V. Nguyen, M. A. Osborne, and S. J. Roberts, "Bayesian optimisation over multiple continuous and categorical inputs," in *Proc. Int. Conf. Mach. Learn.*, Apr. 2020, pp. 8276–8285.
- [37] I. Bogunovic, J. Scarlett, and V. Cevher, "Time-varying Gaussian process bandit optimization," in *Proc. Int. Conf. Artif. Intell. Statist.*, May 2016, pp. 314–323.
- [38] P. Auer, N. Cesa-Bianchi, Y. Freund, and R. E. Schapire, "The nonstochastic multiarmed bandit problem," *SIAM J. Comput.*, vol. 32, no. 1, pp. 48–77, Jan. 2002.
- [39] N. Srinivas, A. Krause, S. M. Kakade, and M. W. Seeger, "Information-theoretic regret bounds for Gaussian process optimization in the bandit setting," *IEEE Trans. Inf. Theory*, vol. 58, no. 5, pp. 3250–3265, May 2012.
- [40] V. Mnih et al., "Human-level control through deep reinforcement learning," *Nature*, vol. 518, pp. 529–533, Feb. 2015.
- [41] M. Titsias, "Variational learning of inducing variables in sparse Gaussian processes," in *Proc. Int. Conf. Artif. Intell. Statist.*, Apr. 2009, pp. 567–574.
- [42] M. Lázaro-Gredilla, J. Quiñero-Candela, C. E. Rasmussen, and A. R. Figueiras-Vidal, "Sparse spectrum Gaussian process regression," *J. Mach. Learn. Research*, vol. 11, no. 1, pp. 1865–1881, Aug. 2010.

- [43] A. Krause and C. Ong, "Contextual Gaussian process bandit optimization," in *Proc. Adv. Neural Inf. Process. Syst.*, vol. 24, Dec. 2011, pp. 1–9.
- [44] A. P. Miettinen and J. K. Nurminen, "Energy efficiency of mobile clients in cloud computing," in *Proc. 2nd USENIX Workshop Hot Topics Cloud Comput.*, Jun. 2010, pp. 1–7.
- [45] P. Lai et al., "Optimal edge user allocation in edge computing with variable sized vector bin packing," in *Proc. Int. Conf. Service-Oriented Comput.* Hangzhou, China: Springer, 2018, pp. 230–245.
- [46] G. Cui et al., "OL-EUA: Online user allocation for NOMA-based mobile edge computing," *IEEE Trans. Mobile Comput.*, vol. 22, no. 4, pp. 2295–2306, Apr. 2023.



Jia Yan (Member, IEEE) received the B.Eng. degree from the School of Electronic and Information Engineering, South China University of Technology, Guangzhou, China, in 2017, and the Ph.D. degree in information engineering from The Chinese University of Hong Kong in 2021. He was a Post-Doctoral Associate with the Department of Electrical and Computer Engineering, University of Minnesota, Twin Cities, Minneapolis, USA, from 2021 to 2023. He is currently an Assistant Professor with The Hong Kong University of Science and Technology (Guangzhou). His research interests include optimization and machine learning techniques in wireless communications and networking, particularly in mobile edge computing and edge intelligence.



Qin Lu (Member, IEEE) received the B.S. degree in electrical engineering from the University of Electronic Science and Technology of China in 2013 and the Ph.D. degree in electrical engineering from the University of Connecticut (UConn) in 2018.

She is currently an Assistant Professor with the School of Electrical and Computer Engineering, College of Engineering, University of Georgia. Previously, she was a Post-Doctoral Research Associate with the University of Minnesota, Twin Cities. Her research interests include signal processing, machine learning, data science, and communications, with special focus on Gaussian processes, Bayesian optimization, spatio-temporal inference over graphs, and data association for multi-object tracking. She received the National Scholarship from China twice. She was awarded Summer Fellowship and Doctoral Dissertation Fellowship from UConn. She was also a recipient of the Women of Innovation Award in Collegian Innovation and Leadership by Connecticut Technology Council in March 2018.



Georgios B. Giannakis (Life Fellow, IEEE) received the Diploma degree in electrical engineering (EE) from the National Technical University of Athens, Greece, in 1981, and the M.Sc. degree in EE, the M.Sc. degree in mathematics, and the Ph.D. degree in EE from the University of Southern California (USC) in 1983, 1986, and 1986, respectively.

From 1982 to 1986, he was with USC. He was with the University of Virginia from 1987 to 1998. Since 1999, he has been with the University of Minnesota (UMN), where he held an Endowed Chair of Telecommunications. He served as the Director of the Digital Technology Center from 2008 to 2021. Since 2016, he has been a UMN Presidential Chair of ECE. His research interests include statistical learning, communications, and networking—subjects on which he has published more than 490 journal articles, 800 conference papers, 26 book chapters, two edited books, and two research monographs. His current research interests include data science with applications to IoT and power networks with renewables. He is the (co-) inventor of 36 issued patents and the (co-)recipient of ten best journal article awards from the IEEE Signal Processing (SP) and Communications Societies, including the G. Marconi Prize. He also received the IEEE-SPS "Nobert Wiener" Society Award in 2019, the EURASIP's "A. Papoulis" Society Award in 2020, the Technical Achievement Awards from the IEEE-SPS in 2000 and EURASIP in 2005, the IEEE ComSoc Education Award in 2019, and the IEEE Fourier Technical Field Award in 2015. He is a member of the Academia European, the Academy of Athens, Greece, and the Royal Academy of Engineering (U.K.), and a fellow of the U.S. National Academy of Inventors, the European Academy of Sciences, and EURASIP. He has served the IEEE in a number of posts, including that of a Distinguished Lecturer for the IEEE-SPS.