# FED-MOZOO: FEDERATED MULTI-OBJECTIVE ONLINE ZEROTH-ORDER OPTIMIZATION WITH INTERMITTENT COMMUNICATION

**Anonymous authors**
Paper under double-blind review

## ABSTRACT

Federated online learning provides a promising framework for sequential decision-making over distributed data, but its effectiveness is challenged by communication constraints, heterogeneous local data, and dynamically changing objectives. When multiple objectives must be optimized simultaneously, existing federated online methods mainly rely on first-order gradient information, which may be unavailable in black-box or gradient-inaccessible environments. This motivates the study of federated multi-objective online optimization with limited function-value feedback and intermittent communication. To this end, we propose FED-MOZOO, a two-point zeroth-order method that performs local preference-based scalarization and projected updates, followed by periodic model aggregation. For convex $G$-Lipschitz losses over a radius-$B$ domain, we establish a dimension- and communication-dependent regret guarantee for the virtual global model: $\mathcal{O}(BG(\sqrt{d/T} + d^{1/4}/\sqrt{R} + 1/T))$, where $T$ is the online horizon and $R$ is the number of communication rounds. The bound characterizes the effects of zeroth-order feedback and intermittent communication, and recovers the standard $\mathcal{O}(BG\sqrt{d/T})$ scaling when communication occurs at every round. Experiments on synthetic and convex multi-objective learning tasks evaluate regret, communication–accuracy trade-offs, preference tracking, Pareto-front quality, and the scaling behavior with dimension and communication frequency.

## 1 INTRODUCTION

Federated learning (FL) has emerged as an effective paradigm for distributed machine learning by enabling multiple clients to collaboratively train a shared model without transferring their private local data to a central server. The seminal federated averaging framework (McMahan et al., 2017) demonstrated that performing local updates followed by periodic model aggregation can substantially reduce communication overhead. This has motivated extensive research on federated optimization under data heterogeneity and limited communication resources (Kairouz & McMahan, 2021; Li et al., 2020). However, most existing federated learning methods are primarily designed for static optimization problems, where the underlying objective functions remain fixed throughout the entire training process. In many practical applications, such as real-time edge computing, adaptive recommendations, and sequential decision-making systems (Zinkevich, 2003; Shalev-Shwartz, 2012), this stationary assumption easily fails. In these environments, streaming data arrives sequentially, causing the local data distributions and the associated loss functions to change continuously over time. When objectives are dynamic, traditional static convergence criteria no longer reflect the system's performance, which naturally leads to the study of federated online learning.

Federated online learning combines the distributed structure of FL with the online convex optimization (OCO) framework, where algorithms are evaluated through cumulative regret over time rather than asymptotic convergence to a fixed point (Zinkevich, 2003; Hazan, 2016). Recent studies have investigated federated online optimization under intermittent communication and limited feedback, showing that the communication interval and the number of local update steps directly affect the regret bound (Woodworth et al., 2021; Kwon et al., 2023; Liu et al., 2024). In particular, Patel et al. (2023) developed federated online and bandit convex optimization frameworks and analyzed the trade-off between client collaboration and communication frequency under both first-order and



Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

zeroth-order feedback. These results demonstrate that federated online learning must simultaneously handle temporal variation in loss functions and model drift caused by infrequent communication. Nevertheless, almost all existing federated online optimization methods focus exclusively on a single objective, assuming that all clients work together to minimize the same scalar loss function.

This single-objective formulation is overly restrictive when real-world systems must balance multiple potentially conflicting performance metrics. For example, an edge system often needs to minimize prediction error while simultaneously reducing operational latency and power consumption. Multi-objective optimization (MOO) provides a principled framework for addressing such trade-offs by finding Pareto-optimal solutions (Désidéri, 2012; Sener & Koltun, 2018). Recent works have extended this formulation to federated settings by optimizing multiple objectives across heterogeneous clients (Yang et al., 2023; Hartmann et al., 2025). When these objectives also change dynamically over time, the task becomes a federated multi-objective online learning problem, where distributed clients must make sequential decisions while jointly tracking time-varying Pareto trade-offs under intermittent communication. Jiang et al. (2023) proposed a theoretical framework for multi-objective online convex optimization and introduced sequence-wise Pareto suboptimality metrics to measure online performance under changing objectives. However, existing methods in both federated multi-objective learning and multi-objective online optimization rely strictly on first-order gradient information, assuming that exact or stochastic analytical gradients are always available.

This reliance on first-order information is a major bottleneck when objectives are accessible only through function evaluations, as in black-box simulators, physical experiments, proprietary modules, or non-differentiable system metrics. Zeroth-order (ZO) optimization provides a practical alternative by constructing gradient estimates purely from function-value queries without requiring explicit gradient expressions (Nesterov & Spokoiny, 2017). While zeroth-order techniques have been studied in federated learning under data heterogeneity (Fang et al., 2022; Li et al., 2025), and in federated online learning under communication constraints (Patel et al., 2023), these methods remain confined to single-objective settings. Conversely, current federated multi-objective online methods still assume first-order feedback. In this black-box online setting, the estimation error of multi-objective zeroth-order queries interacts with time-varying losses and client drift under intermittent communication. This technical challenge leads to a central question: *How can distributed clients collaboratively track time-varying Pareto optimality when only zeroth-order function-value feedback is available under intermittent communication?*

To answer this question, we study *federated multi-objective online learning with zeroth-order feedback*, where multiple clients collaboratively optimize a sequence of potentially conflicting dynamic objectives while accessing only function values. We propose **Fed-MOZOO**, a federated multi-objective online zeroth-order optimization method that combines two-point zeroth-order estimation, preference-based scalarization, and periodic model aggregation. Each client performs multiple local updates between communication rounds using only function evaluations, which significantly cuts communication frequency without requiring analytical gradients. To establish theoretical guarantees, we first analyze an intermediate stochastic first-order federated tracking problem, and then characterize the additional errors caused by zeroth-order smoothing and directional estimation. The resulting regret analysis explicitly captures the combined effects of the number of clients, local update frequency, communication intervals, and zeroth-order parameters.

Our main contributions are summarized as follows:

- We formulate a federated multi-objective online optimization problem under zeroth-order feedback, which extends federated online optimization to the setting of multiple conflicting dynamic objectives with only function-value evaluations.
- We propose **Fed-MOZOO**, an algorithm that integrates two-point zeroth-order estimation, preference-based scalarization, and periodic federated aggregation, enabling communication-efficient online optimization without analytical gradient access.
- We establish a regret analysis by first analyzing an intermediate stochastic first-order federated problem and then bounding the zeroth-order estimation and smoothing errors. The resulting bounds explicitly characterize the effects of client numbers, local update frequency, communication rounds, and zeroth-order parameters.

- We empirically evaluate Fed-MOZOO on both convex and multi-objective online learning tasks, demonstrating its effectiveness under limited feedback and reduced communication frequency.

## 2 Related Work

### 2.1 Federated Multi-Objective Optimization

Multi-objective optimization has been increasingly studied in federated learning (FL) to address competing goals such as predictive performance, fairness, robustness, and resource efficiency. Classical multi-objective methods such as multiple-gradient descent (Désidéri, 2012) and multi-objective formulations for multi-task learning (Sener & Koltun, 2018) provide the foundation for balancing conflicting objectives. Building on these ideas, Yang et al. (2023) introduced a federated multi-objective learning framework and developed federated multi-gradient descent methods to handle objective and data heterogeneity across clients. Subsequent studies have considered specific trade-offs in federated systems, including fairness and model performance (Pan et al., 2023; 2024; Badar et al., 2024; Shen et al., 2025), as well as communication efficiency (Askin et al., 2025). A recent survey further provides a taxonomy of multi-objective methods in FL and identifies objective conflicts, client heterogeneity, and resource constraints as important challenges (Hartmann et al., 2025). Although these methods demonstrate the benefits of Pareto-oriented optimization in federated systems, they mainly focus on static optimization and assume access to first-order gradient information. In particular, they do not explicitly address sequentially changing objective functions or the limited feedback setting in which gradients are unavailable.

### 2.2 Online and Zeroth-Order Federated Optimization

Online optimization provides a natural framework for sequential decision-making under dynamically changing environments, where performance is characterized by cumulative regret (Zinkevich, 2003; Hazan, 2016). Jiang et al. (2023) extended this framework to multi-objective online convex optimization and introduced the sequence-wise Pareto suboptimality gap together with a doubly regularized online mirror multiple descent method, establishing a foundation for regret analysis in centralized multi-objective online learning. In parallel, federated online optimization has been studied under limited feedback and intermittent communication. Patel et al. (2023) considered federated online and bandit convex optimization and established regret guarantees for both first-order and zeroth-order feedback, while recent work has further investigated bandit online optimization and distributed zeroth-order online optimization under communication constraints and delays (Fokkema et al., 2024; Inoue et al., 2025). Zeroth-order optimization has also become an important approach for federated learning because it eliminates the need for explicit gradients and can reduce communication overhead. FedZO (Fang et al., 2022) introduced stochastic zeroth-order optimization for federated learning, followed by recent studies on zeroth-order adaptive optimization (Xie et al., 2024), historical-trajectory-assisted optimization (Wu et al., 2024), large-scale model fine-tuning (Qin et al., 2024), dimension-free communication (Li et al., 2025), and Hessian-informed zeroth-order federated optimization (Li et al., 2026). Despite these advances, existing online zeroth-order federated methods primarily consider single-objective problems, whereas existing multi-objective online methods are mainly centralized and first-order. Our work bridges these two lines by studying federated multi-objective online learning with zeroth-order feedback, where objective conflicts, temporal variation, client drift, and communication frequency are jointly characterized.

## 3 Problem Setting

We consider $M$ clients, $P$ objectives, and an online horizon of $T = KR$ local rounds, where $K$ is the number of local updates between consecutive communications and $R$ is the number of communication rounds. Let $[M] = \{1, \dots, M\}$ and $[P] = \{1, \dots, P\}$. The decision set is the Euclidean ball

$$\mathcal{X} = \mathbb{B}_2(B) := \{x \in \mathbb{R}^d : \|x\|_2 \le B\}. \tag{3.1}$$

At round $t$, client $m$ has a local decision $x_t^m \in \mathcal{X}$ and incurs a vector-valued loss

$$\mathbf{F}_t^m(x) = \left[f_{t,1}^m(x), \dots, f_{t,P}^m(x)\right]^\top. \tag{3.2}$$

For the regret theorem we use the standard oblivious-sequence model: the loss functions and preference sequence are fixed independently of the algorithm's zeroth-order randomization. This condition is needed because the comparator in Eq. (3.4) is selected after the full sequence is realized; extending the same comparator guarantee to fully adaptive adversaries requires additional care. The global vector loss is

$$\mathbf{F}_t(x) = \frac{1}{M}\sum_{m=1}^{M} \mathbf{F}_t^m(x). \tag{3.3}$$

**Assumption 1** (Convex Lipschitz objectives)**.** *For every $m \in [M]$, $t \in [T]$, and $v \in [P]$, $f_{t,v}^m$ is convex and $G$-Lipschitz on a $\delta$-neighborhood of $\mathcal{X}$ containing all query points $x \pm \delta u$ with $x \in \mathcal{X}$ and $\|u\|_2 = 1$.*

At the beginning of round $t$, a preference vector $\lambda_t \in \Delta_P$ is available, where $\Delta_P := \left\{\lambda \in \mathbb{R}_+^P : \sum_{v=1}^P \lambda_v = 1\right\}$. The local and global scalarized losses are $\phi_t^m(x) = \langle \lambda_t, \mathbf{F}_t^m(x)\rangle, \phi_t(x) = \frac{1}{M}\sum_{m=1}^M \phi_t^m(x)$. Because $\lambda_t \in \Delta_P$, the scalarized losses remain convex and $G$-Lipschitz.

Let

$$x_\lambda^\star \in \arg\min_{x\in\mathcal{X}} \sum_{t=1}^{T} \phi_t(x) \tag{3.4}$$

be the best static comparator for the realized preference sequence. We use two regret notions.

**Definition 1** (Local-play regret)**.** *The average local-play regret is*

$$\overline{\mathrm{Reg}}_{T,\mathrm{loc}}^{\lambda} := \frac{1}{MT}\mathbb{E}\left[\sum_{t=1}^{T}\sum_{m=1}^{M}\left(\phi_t^m(x_t^m) - \phi_t^m(x_\lambda^\star)\right)\right]. \tag{3.5}$$

Define the virtual network average $x_t := \frac{1}{M}\sum_{m=1}^M x_t^m$. Since $\mathcal{X}$ is convex, $\bar{x}_t \in \mathcal{X}$. At synchronization rounds, $\bar{x}_t$ coincides with the shared server model.

**Definition 2** (Global-model regret)**.** *The average global-model regret is*

$$\overline{\mathrm{Reg}}_{T}^{\lambda} := \frac{1}{T}\mathbb{E}\left[\sum_{t=1}^{T}\left(\phi_t(\bar{x}_t) - \phi_t(x_\lambda^\star)\right)\right]. \tag{3.6}$$

The corresponding cumulative regrets are obtained by multiplying the average regrets by $T$. The local-play quantity evaluates the actual local iterates, whereas the global-model quantity measures the quality of a coherent shared model and is therefore sensitive to consensus drift.

For a fixed preference $\lambda \in \mathrm{int}(\Delta_P)$, minimizing the cumulative scalarized objective is consistent with Pareto optimality of the cumulative vector objective.

**Proposition 1** (Pareto consistency for fixed preferences)**.** *If $\lambda_t \equiv \lambda \in \mathrm{int}(\Delta_P)$ and $x_\lambda^\star$ minimizes $\sum_t \langle \lambda, \mathbf{F}_t(x)\rangle$ over $\mathcal{X}$, then $x_\lambda^\star$ is Pareto optimal for the cumulative vector objective $\sum_t \mathbf{F}_t(x)$.*

The proof is immediate by contradiction and is included in Appendix C for completeness. For time-varying $\lambda_t$, the regret in Eq. (3.6) instead measures performance under the prescribed preference schedule and does not imply objective-wise regret guarantees for every component separately.

# 4 FED-MOZOO

## 4.1 PREFERENCE-WEIGHTED TWO-POINT ESTIMATOR

At round $t$, client $m$ samples $u_t^m \sim \mathrm{Unif}(\mathbb{S}^{d-1})$ and queries the vector-valued oracle at

$$x_t^{m,+} = x_t^m + \delta u_t^m, \qquad x_t^{m,-} = x_t^m - \delta u_t^m. \tag{4.1}$$

The objective-wise estimator is

$$g_{t,v}^m = \frac{d}{2\delta}\left[f_{t,v}^m(x_t^{m,+}) - f_{t,v}^m(x_t^{m,-})\right] u_t^m, \tag{4.2}$$

**Algorithm 1** FED-MOZOO

**Require:** stepsize $\eta$, smoothing radius $\delta$, communication interval $K$, horizon $T = KR$

Initialize $x_1^m = 0$ for all $m \in [M]$
**for** $t = 1, \dots, T$ **do**
  Receive preference $\lambda_t \in \Delta_P$
  **for** each client $m \in [M]$ **in parallel do**
    Sample $u_t^m \sim \mathrm{Unif}(\mathbb{S}^{d-1})$
    Query $\mathbf{F}_t^m(x_t^m + \delta u_t^m)$ and $\mathbf{F}_t^m(x_t^m - \delta u_t^m)$
    Form $g_t^m$ using (4.3)
    $\widetilde{x}_{t+1}^m \leftarrow \Pi_{\mathcal{X}}(x_t^m - \eta g_t^m)$
  **end for**
  **if** $t \bmod K = 0$ **then**
    $x_{t+1} \leftarrow M^{-1} \sum_{m=1}^M \widetilde{x}_{t+1}^m$ and broadcast $x_{t+1}^m \leftarrow x_{t+1}$
  **else**
    $x_{t+1}^m \leftarrow \widetilde{x}_{t+1}^m$ for all $m$
  **end if**
**end for**

and the preference-weighted estimator is

$$g_t^m = \sum_{v=1}^P \lambda_{t,v} g_{t,v}^m = \frac{d}{2\delta} \left[\phi_t^m(x_t^{m,+}) - \phi_t^m(x_t^{m,-})\right] u_t^m. \tag{4.3}$$

Thus a single random direction is shared across objectives and only two vector-valued oracle calls are needed per local update.

For the analysis, define the *ball-smoothed* scalarized loss

$$\widehat{\phi}_{t,\delta}^m(x) := \mathbb{E}_{v \sim \mathrm{Unif}(\mathbb{B}_2(1))} \left[\phi_t^m(x + \delta v)\right]. \tag{4.4}$$

The use of ball smoothing is important: the standard sphere-query identity implies that the symmetric estimator in Eq. (4.3) is unbiased for the gradient of Eq. (4.4).

**Lemma 1** (Two-point estimator)**.** *Under Assumption 1, for every $m$ and $t$,*

$$\mathbb{E}_t[g_t^m] = \nabla \widehat{\phi}_{t,\delta}^m(x_t^m), \tag{4.5}$$

$$\mathbb{E}_t \|g_t^m\|_2^2 \le 2dG^2, \tag{4.6}$$

$$\sup_{x \in \mathcal{X}} \left|\widehat{\phi}_{t,\delta}^m(x) - \phi_t^m(x)\right| \le G\delta. \tag{4.7}$$

*Consequently,* $\mathbb{E}_t \|g_t^m - \mathbb{E}_t g_t^m\|_2^2 \le 2dG^2$.

The proof is given in Appendix A. The dimension-linear second moment follows from the symmetry of the two-point estimator and spherical concentration, in the spirit of the sharp two-point analysis of (Shamir, 2017); it is sharper than the naive $d^2G^2$ bound from pointwise Lipschitz continuity.

### 4.2 Local updates and periodic aggregation

Each client performs the projected update $x_{t+1}^m = \Pi_{\mathcal{X}}\left(x_t^m - \eta g_t^m\right)$. If $t$ is a multiple of $K$, the server averages the tentative local iterates and broadcasts the result, $x_{t+1}^m = \frac{1}{M} \sum_{i=1}^M \widetilde{x}_{t+1}^i, \qquad m \in [M]$. Otherwise, $x_{t+1}^m = \widetilde{x}_{t+1}^m$. Algorithm 1 summarizes the proposed method.

Each client uses $2T$ vector-valued function queries. The system performs $R = T/K$ synchronization rounds, transmitting $\mathcal{O}(MRd)$ model scalars up to constant factors. Increasing $K$ therefore reduces communication linearly, but the theory below shows that it increases consensus drift.

## 5 Theoretical Analysis

We now derive the regret guarantee of FED-MOZOOin a form that makes the origin of every error term explicit. The analysis is organized around four steps. First, the two-point estimator is interpreted

as an unbiased stochastic gradient of a ball-smoothed scalarized loss. Second, a projected-descent potential inequality controls the regret accumulated by the local client iterates. Third, we quantify how far the clients can drift from their network average during a block of $K$ local updates. Finally, a local-to-global transfer inequality converts local progress into regret of the coherent network-average model. This decomposition separates zeroth-order estimation, smoothing, and intermittent communication rather than hiding them in a single big-$\mathcal{O}$ term.

Let $\mathcal{F}_t$ denote the sigma-field generated by all iterates and zeroth-order directions before the directions $\{u_t^m\}_{m=1}^M$ are sampled at round $t$. Under the oblivious-sequence model in Section 3, the losses $\{\phi_t^m\}_{m=1}^M$, the preference vector $\lambda_t$, and the iterates $\{x_t^m\}_{m=1}^M$ are $\mathcal{F}_t$-measurable. We write $\mathbb{E}_t[\cdot] := \mathbb{E}[\cdot \mid \mathcal{F}_t]$. For compactness, define

$$\widehat{\phi}_t^m := \widehat{\phi}_{t,\delta}^m, \qquad \widehat{\phi}_t(x) := \frac{1}{M}\sum_{m=1}^M \widehat{\phi}_t^m(x), \qquad Q := G\sqrt{2d}. \tag{5.1}$$

Lemma 1 implies

$$\mathbb{E}_t[g_t^m] = \nabla\widehat{\phi}_t^m(x_t^m), \qquad \mathbb{E}_t\|g_t^m\|_2^2 \le Q^2, \qquad |\widehat{\phi}_t^m(x) - \phi_t^m(x)| \le G\delta. \tag{5.2}$$

Moreover, averaging preserves convexity and the $G$-Lipschitz property, so every $\widehat{\phi}_t^m$ and $\widehat{\phi}_t$ is convex and $G$-Lipschitz.

### 5.1 Potential decrease for the local iterates

For any comparator $x \in \mathcal{X}$, define the mean squared distance

The key point is that the same potential works both on ordinary local-update rounds and on synchronization rounds. Indeed, Euclidean projection is non-expansive, while model averaging is contractive for the mean squared distance to a common comparator.

**Lemma 2** (One-step smoothed progress)**.** *For any $x \in \mathcal{X}$ and any round $t$,*

$$\frac{1}{M}\sum_{m=1}^M \mathbb{E}_t\Big[\widehat{\phi}_t^m(x_t^m) - \widehat{\phi}_t^m(x)\Big] \le \frac{V_t(x) - \mathbb{E}_t[V_{t+1}(x)]}{2\eta} + \frac{\eta Q^2}{2}. \tag{5.3}$$

**Derivation.** For a tentative local iterate $\widetilde{x}_{t+1}^m = \Pi_{\mathcal{X}}(x_t^m - \eta g_t^m)$, we know that

$$\|\widetilde{x}_{t+1}^m - x\|_2^2 \le \|x_t^m - \eta g_t^m - x\|_2^2 = \|x_t^m - x\|_2^2 - 2\eta\langle g_t^m, x_t^m - x\rangle + \eta^2\|g_t^m\|_2^2. \tag{5.4}$$

If no communication occurs, $x_{t+1}^m = \widetilde{x}_{t+1}^m$. If communication occurs, all clients are replaced by their average and Jensen's inequality yields $\Big\|\frac{1}{M}\sum_{m=1}^M \widetilde{x}_{t+1}^m - x\Big\|_2^2 \le \frac{1}{M}\sum_{m=1}^M \|\widetilde{x}_{t+1}^m - x\|_2^2$.

Thus averaging never increases $V_{t+1}(x)$. Taking $\mathbb{E}_t$, using Eq. (5.2), and applying convexity, $\widehat{\phi}_t^m(x_t^m) - \widehat{\phi}_t^m(x) \le \Big\langle \nabla\widehat{\phi}_t^m(x_t^m), x_t^m - x \Big\rangle$, produces Eq. (5.3). The complete argument is given in Appendix B.

Taking $x = x_\lambda^\star$, summing Eq. (5.3) over $t$, and using $x_1^m = 0$ and $\|x_\lambda^\star\|_2 \le B$ gives the smoothed local-play bound

$$\frac{1}{MT}\sum_{t=1}^T\sum_{m=1}^M \mathbb{E}\Big[\widehat{\phi}_t^m(x_t^m) - \widehat{\phi}_t^m(x_\lambda^\star)\Big] \le \frac{B^2}{2\eta T} + \frac{\eta Q^2}{2}. \tag{5.5}$$

Eq. (5.5) is important conceptually: periodic model averaging introduces no additional penalty for the local-play metric because the averaging operation contracts the comparator potential.

### 5.2 Consensus drift within one communication block

The global-model regret is different because it evaluates the network-average model $\bar{x}_t$. We therefore need to control the discrepancy between $x_t^m$ and $\bar{x}_t$.

**Lemma 3** (Pointwise and average consensus drift)**.** *Let $s_r = (r-1)K + 1$ be the first index of communication block $r$, and write $t = s_r + q$ with $q \in \{0, \dots, K-1\}$. Then*

$$\frac{1}{M}\sum_{m=1}^{M} \mathbb{E}\|x_t^m - \bar{x}_t\|_2 \le 2\eta q Q, \tag{5.6}$$

$$\frac{1}{M}\sum_{m=1}^{M} \mathbb{E}\|x_t^m - \bar{x}_t\|_2^2 \le 4\eta^2 q^2 Q^2. \tag{5.7}$$

*Consequently,*

$$\frac{1}{T}\sum_{t=1}^{T}\frac{1}{M}\sum_{m=1}^{M} \mathbb{E}\|x_t^m - \bar{x}_t\|_2 \le \eta Q(K-1). \tag{5.8}$$

The proof uses only synchronization at the beginning of each block and the bound on the expected update magnitude. Since $\|x_{j+1}^m - x_j^m\|_2 \le \eta\|g_j^m\|_2$ between two communication events, a client can move by at most order $q\eta Q$ after $q$ local steps. The network average can move by the same order, yielding the factor 2 in Eq. (5.6). Averaging $2q\eta Q$ over $q = 0, \dots, K-1$ gives Eq. (5.8). Appendix B.3 provides the complete first- and second-moment derivation.

### 5.3 From local progress to the global model

The remaining step relates the loss at the network average to the losses at the local iterates.

**Lemma 4** (Local-to-global transfer)**.** *For every round $t$ and comparator $x \in \mathcal{X}$,*

$$\mathbb{E}\Big[\widehat{\phi}_t(\bar{x}_t) - \widehat{\phi}_t(x)\Big] \le \frac{1}{M}\sum_{m=1}^{M} \mathbb{E}\Big[\widehat{\phi}_t^m(x_t^m) - \widehat{\phi}_t^m(x)\Big] + \frac{G}{M}\sum_{m=1}^{M} \mathbb{E}\|x_t^m - \bar{x}_t\|_2. \tag{5.9}$$

This result follows directly from $G$-Lipschitz continuity: $\widehat{\phi}_t^m(\bar{x}_t) \le \widehat{\phi}_t^m(x_t^m) + G\|x_t^m - \bar{x}_t\|_2$. Combining Lemmas 2–4 with the smoothing inequality in Eq. (5.2) yields the main generic bound.

**Theorem 1** (Generic local-play and global-model regret)**.** *Under Assumption 1 and the oblivious-sequence model stated in Section 3, for any $\eta > 0$ and $\delta > 0$,*

$$\overline{\mathrm{Reg}}_{T,\mathrm{loc}}^{\lambda} \le \frac{B^2}{2\eta T} + \frac{\eta Q^2}{2} + 2G\delta, \tag{5.10}$$

$$\overline{\mathrm{Reg}}_{T}^{\lambda} \le \frac{B^2}{2\eta T} + \frac{\eta Q^2}{2} + \eta G Q(K-1) + 2G\delta. \tag{5.11}$$

This bound separates the effects of the initial potential, zeroth-order updates, consensus drift, and smoothing bias. In particular, the communication-dependent term $\eta GQ(K-1)$ vanishes when $K = 1$. An $M^{-1/2}$ speedup is not asserted here because the potential argument controls the average of the individual estimator second moments rather than the variance of an averaged gradient estimator. An explicit $M$-dependent collaboration gain would require additional structure or a different update/oracle-sharing mechanism.

**Corollary 1** (Communication-dependent global-model regret)**.** *Let $A_{d,K} := 2d + 2(K-1)\sqrt{2d}$, $\eta = \frac{B}{G\sqrt{TA_{d,K}}}$. Then, for every $\delta > 0$,*

$$\overline{\mathrm{Reg}}_{T}^{\lambda} \le BG\sqrt{\frac{A_{d,K}}{T}} + 2G\delta. \tag{5.12}$$

*In the exact two-point oracle model considered in this paper, the second-moment bound in Lemma 1 is independent of $1/\delta$. Hence no noisy-oracle bias–variance balancing is required. Choosing, for example, $\delta = \frac{B}{T}$ produces*

$$\overline{\mathrm{Reg}}_{T}^{\lambda} \le BG\left(\sqrt{\frac{A_{d,K}}{T}} + \frac{2}{T}\right) \tag{5.13}$$

$$\le BG\left(\sqrt{\frac{2d}{T}} + \frac{2^{3/4}d^{1/4}}{\sqrt{R}} + \frac{2}{T}\right), \qquad T = KR. \tag{5.14}$$

**Interpretation and sanity checks.** When $K = 1$, the clients synchronize after every update and the consensus term vanishes exactly; Eq. (5.13) reduces to the standard derivative-free online scale $\mathcal{O}(BG\sqrt{d/T})$ up to the lower-order smoothing term. For fixed $K$, $R = T/K$ and the average regret decays as $\mathcal{O}(T^{-1/2})$. If $R$ is held fixed while $T$ grows, the communication term in Eq. (5.14) does not vanish, reflecting the fact that an increasingly long sequence of local updates cannot maintain a uniformly accurate coherent global model with only finitely many synchronizations. Finally, the $d^{1/4}/\sqrt{R}$ factor is not a separate zeroth-order variance term; it appears after optimizing the stepsize in the presence of the drift contribution $\eta GQ(K-1)$ with $Q = G\sqrt{2d}$.

**Scope of the guarantee.** The result assumes convex Lipschitz objectives, exact two-point function-value feedback, and an oblivious loss/preference sequence. It is a preference-conditioned scalarized regret guarantee. For a fixed strictly positive preference, the cumulative scalarized comparator is Pareto consistent as stated in Proposition 1; for time-varying preferences, the theorem does not imply simultaneous sublinear regret for every objective component. The corresponding precise Pareto interpretation is given in Appendix C.

# 6 EXPERIMENTS

We test three questions: **(Q1)** online regret under two-point feedback, **(Q2)** the effect of intermittent communication on the shared model, and **(Q3)** preference-conditioned trade-offs and adaptation to changes in $\lambda_t$. Theorem-audit ablations are reported in Appendix D. Unless stated otherwise, report at least five independent seeds with 95% confidence intervals.

## 6.1 SETUP AND BASELINES

**Benchmarks.** We use (i) bounded conflicting quadratic streams with controllable dimension, client heterogeneity, and temporal drift; (ii) linear logistic regression on Adult Dua & Graff (2019) with two group-conditioned risks; and (iii) linear softmax regression on MNIST Lecun et al. (1998) with two class-group losses. The synthetic benchmark admits an accurate static comparator, while the two real-data tasks remain convex and therefore aligned with the theory. Non-IID client partitions, streaming protocol, and comparator construction are detailed in Appendix D.

**Baselines.** The primary references are **Fed-MOGD** (same federation and exact gradients), **Central-MOZO** (two-point feedback with full communication), and **Local-MOZO** (no aggregation). Fixed-preference Pareto plots additionally include FMGDA and FSMGDA Yang et al. (2023) and the centralized first-order online MOO reference Jiang et al. (2023). First-order methods are oracle references, not like-for-like black-box competitors. Zeroth-order query budgets are matched across methods, and communication rounds are always reported.

## 6.2 MAIN COMPARISONS

Figure 1 summarizes the main empirical questions. Panels (a–b) report running average global-model regret on synthetic and convex classification streams, comparing FED-MOZOOwith Fed-MOGD, Central-MOZO, and Local-MOZO. Panels (c–d) fix $T$ and vary $K$ (equivalently $R = T/K$), reporting final global-model regret together with local-play regret or consensus error to isolate the communication-sensitive term in Theorem 1. Panels (e–f) evaluate multi-objective behavior for $P = 2$: a fixed-preference sweep $\lambda = (\alpha, 1-\alpha)$ traces objective pairs, while a piecewise-constant $\lambda_t$ schedule examines the response to preference switches. The primary metric is average global-model regret in Eq. (3.6); we additionally report local-play regret, consensus error, per-objective losses, Pareto hypervolume, communication rounds $R$, transmitted model scalars, and vector-valued oracle calls. Dimension, communication, and smoothing ablations, together with the detailed plotting protocol, are provided in Appendix D.

# 7 CONCLUSION

We introduced FED-MOZOOfor federated multi-objective online optimization with function-value feedback and intermittent communication. Its preference-weighted symmetric two-point estimator

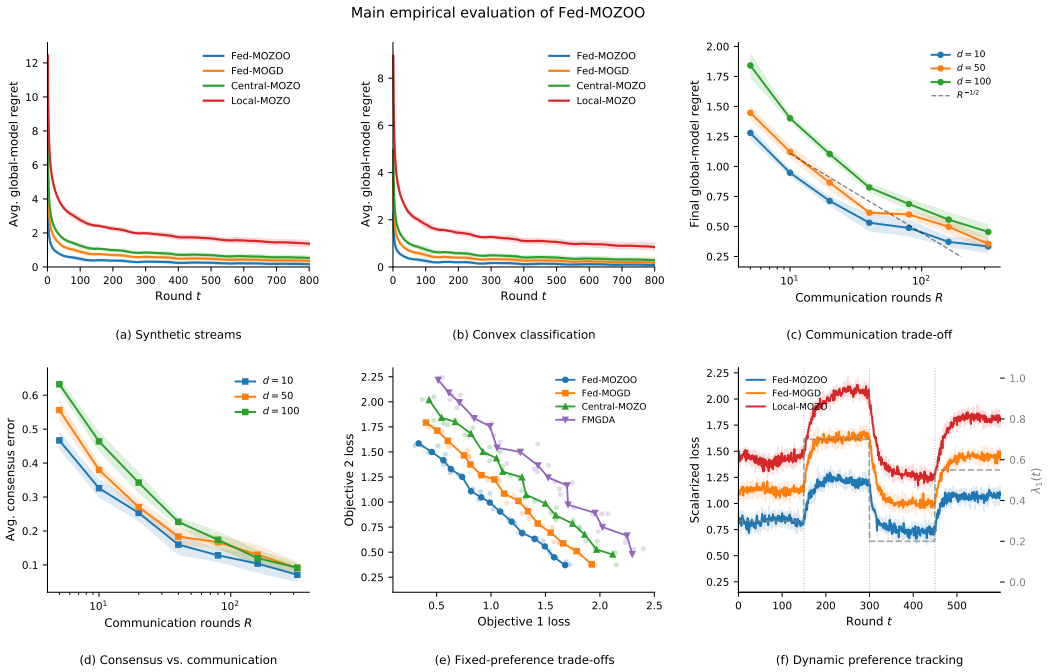


Figure 1: **Main empirical evaluation.** (a–b) Online global-model regret. (c–d) Communication trade-off and its consensus/local-play diagnostic. (e) Fixed-preference objective trade-offs. (f) Dynamic preference tracking. Lower regret is better; in the objective plane the lower-left direction is preferred.

supports local projected updates and periodic model averaging without requiring gradients. By separating local-play regret from global-model regret, the analysis identifies consensus drift as the explicit price of maintaining a coherent shared model. For convex Lipschitz losses, the average global-model regret is $\mathcal{O}(BG(\sqrt{d/T} + d^{1/4}/\sqrt{R} + 1/T))$. The resulting decomposition links black-box estimation, online adaptation, and communication frequency in a single guarantee and motivates the regret, communication, preference, and scaling experiments. Natural extensions include collaboration gains with explicit $M$-dependence, noisy or delayed feedback, non-convex models, and learned rather than prescribed preference sequences.

## AI Use Statement

The authors did not use generative AI tools in the research, implementation, analysis, or writing of this submission.

## Ethics Statement

The method does not require raw client data at the server. Public attributes, when used, serve only to study multi-objective trade-offs and are not claimed as fairness guarantees. No new human-subject data were collected.

## Reproducibility Statement

Algorithm 1 fully specifies the method; assumptions and proofs are in Sections 3–5 and the appendix. Experimental protocols appear in Appendix D. Code will be made available upon request.

## References

Baris Askin, Pranay Sharma, Gauri Joshi, and Carlee Joe-Wong. Federated communication-efficient multi-objective optimization. arXiv preprint arXiv:2410.16398, 2025.

Maryam Badar, Sandipan Sikdar, Wolfgang Nejdl, and Marco Fisichella. FairTrade: Achieving pareto-optimal trade-offs between balanced accuracy and fairness in federated learning. In Proceedings of the AAAI Conference on Artificial Intelligence, pp. 10962–10970, 2024.

Jean-Antoine Désidéri. Multiple-gradient descent algorithm (MGDA) for multiobjective optimization. Comptes Rendus Mathématique, 350(5–6):313–318, 2012.

Dheeru Dua and Casey Graff. UCI machine learning repository, 2019. University of California, Irvine, School of Information and Computer Sciences.

Wenzhi Fang, Ziyi Yu, Yuning Jiang, Yuanming Shi, Colin N. Jones, and Yong Zhou. Communication-efficient stochastic zeroth-order optimization for federated learning. IEEE Transactions on Signal Processing, 70:5058–5073, 2022.

Hidde Fokkema, Dirk Van der Hoeven, Tor Lattimore, and Jack J. Mayo. Online newton method for bandit convex optimisation extended abstract. In Proceedings of the 37th Conference on Learning Theory (COLT), pp. 1713–1714, 2024.

Maria Hartmann, Grégoire Danoy, and Pascal Bouvry. Multi-objective methods in federated learning: A survey and taxonomy. arXiv preprint arXiv:2502.03108, 2025.

Elad Hazan. Introduction to online convex optimization. Foundations and Trends in Optimization, 2 (3–4):157–325, 2016.

Keito Inoue, Naoki Hayashi, and Shigemasa Takai. Distributed zeroth-order online optimization with communication delays. IET Control Theory & Applications, 19(1):e12759, 2025.

Jiyan Jiang, Wenpeng Zhang, Shiji Zhou, Lihong Gu, Xiaodong Zeng, and Wenwu Zhu. Multi-objective online learning. In International Conference on Learning Representations (ICLR), 2023.

Peter Kairouz and H. Brendan McMahan. Advances and open problems in federated learning. Foundations and Trends in Machine Learning, 14(1–2):1–210, 2021.

Dohyeok Kwon, Jonghwan Park, and Songnam Hong. Tighter regret analysis and optimization of online federated learning. IEEE Transactions on Pattern Analysis and Machine Intelligence, 45 (12):15772–15789, 2023.

Y. Lecun, L. Bottou, Y. Bengio, and P. Haffner. Gradient-based learning applied to document recognition. Proceedings of the IEEE, 86(11):2278–2324, 1998.

Tian Li, Anit Kumar Sahu, Ameet Talwalkar, and Virginia Smith. Federated learning: Challenges, methods, and future directions. IEEE Signal Processing Magazine, 37(3):50–60, 2020.

Zhe Li, Bicheng Ying, Zidong Liu, Chaosheng Dong, and Haibo Yang. Achieving dimension-free communication in federated learning via zeroth-order optimization. In International Conference on Learning Representations (ICLR), 2025.

Zhe Li, Bicheng Ying, Zidong Liu, Chaosheng Dong, and Haibo Yang. Converge faster, talk less: Hessian-informed federated zeroth-order optimization. In International Conference on Learning Representations (ICLR), 2026.

Jiandong Liu, Lan Zhang, Fengxiang He, Chi Zhang, Shanyang Jiang, and Xiang-Yang Li. Communication-efficient regret-optimal distributed online convex optimization. IEEE Transactions on Parallel and Distributed Systems, 35(11):2270–2284, 2024.

Brendan McMahan, Eider Moore, Daniel Ramage, Seth Hampson, and Blaise Aguera y Arcas. Communication-efficient learning of deep networks from decentralized data. In Proceedings of the 20th International Conference on Artificial Intelligence and Statistics (AISTATS), pp. 1273–1282, 2017.

Yurii Nesterov and Vladimir Spokoiny. Random gradient-free minimization of convex functions. Foundations of Computational Mathematics, 17(2):527–566, 2017.

Zibin Pan, Shuyi Wang, Chi Li, Haijin Wang, Xiaoying Tang, and Junhua Zhao. FedMDFG: Federated learning with multi-gradient descent and fair guidance. In Proceedings of the AAAI Conference on Artificial Intelligence, pp. 9364–9371, 2023.

Zibin Pan, Chi Li, Fangchen Yu, Shuyi Wang, Haijin Wang, Xiaoying Tang, and Junhua Zhao. FedLF: Layer-wise fair federated learning. In Proceedings of the AAAI Conference on Artificial Intelligence, pp. 14527–14535, 2024.

Kumar Kshitij Patel, Lingxiao Wang, Aadirupa Saha, and Nathan Srebro. Federated online and bandit convex optimization. In Proceedings of the 40th International Conference on Machine Learning (ICML), pp. 27439–27460, 2023.

Zhen Qin, Daoyuan Chen, Bingchen Qian, Bolin Ding, Yaliang Li, and Shuiguang Deng. Federated full-parameter tuning of billion-sized language models with communication cost under 18 kilobytes. In Proceedings of the 41st International Conference on Machine Learning (ICML), pp. 41473–41497, 2024.

Ozan Sener and Vladlen Koltun. Multi-task learning as multi-objective optimization. In Advances in Neural Information Processing Systems (NeurIPS), pp. 525–536, 2018.

Shai Shalev-Shwartz. Online learning and online convex optimization. Foundations and Trends in Machine Learning, 4(2):107–194, 2012.

Ohad Shamir. An optimal algorithm for bandit and zero-order convex optimization with two-point feedback. Journal of Machine Learning Research, 18(52):1–11, 2017.

Yuhao Shen, Wei Xi, Yunyun Cai, Yuwei Fan, He Yang, and Jizhong Zhao. Multi-objective federated learning: Balancing global performance and individual fairness. Future Generation Computer Systems, 162:107468, 2025.

Blake E. Woodworth, Brian Bullins, Ohad Shamir, and Nathan Srebro. The min-max complexity of distributed stochastic convex optimization with intermittent communication. In Proceedings of the 34th Conference on Learning Theory (COLT), pp. 4386–4437, 2021.

Chenlin Wu, Xiaoyu He, Zike Li, Jing Gong, and Zibin Zheng. A historical trajectory assisted optimization method for zeroth-order federated learning. arXiv preprint arXiv:2409.15955, 2024.

Ping Xie, Xiangrui Gao, Fan Li, Ling Xing, Yu Zhang, and Hanxiao Sun. Communication-efficient zeroth-order adaptive optimization for federated learning. Mathematics, 12(8):1148, 2024.

Haibo Yang, Zhuqing Liu, Jia Liu, Chaosheng Dong, and Michinari Momma. Federated multi-objective learning. In Advances in Neural Information Processing Systems (NeurIPS), pp. 39602–39625, 2023.

Martin Zinkevich. Online convex programming and generalized infinitesimal gradient ascent. In Proceedings of the 20th International Conference on Machine Learning (ICML), pp. 928–936, 2003.