

Test-Time Collective Action: Proxy-Based Perturbations for Correcting Algorithmic Harms

Meghana Bhange^{1,2}, Ulrich Aïvodji^{1,2}, Elliot Creager^{3,4}

¹ ÉTS Montréal

² Mila - Quebec AI Institute

³ University of Waterloo

⁴ Vector Institute

meghana-shashikant.bhange.1@ens.etsmtl.ca, ulrich.aivodji@etsmtl.ca, creager@uwaterloo.ca

Abstract

When machine learning systems under-perform for particular subgroups, affected users typically have no way to correct these disparities without relying on platform-level fixes. Existing approaches to algorithmic fairness rely on provider-centric approaches to correct these failures, leaving users with no external lever when faced with harm. Recent work in Algorithmic Collective Action shows that coordinated users can steer an algorithmic system toward a collective goal, but the existing mechanisms require the provider to retrain on the collective’s modified data which users may not have control over. We propose *Test-Time Collective Action (TTCA)*, a framework through which a group of users who share query access to the platform, can correct disparities affecting under-served subgroup without participating in the platform’s training loop. We implement this through a proxy-based mechanism where the collective pools query access to a black-box API to extract a proxy of the platform, then optimizes a per-class universal perturbation against the proxy. Each member applies this perturbation to their own inputs at submission time, requiring no cooperation from the platform. We empirically evaluate the mechanism on CIFAR-10, CIFAR-100, and FairFace, showing that modestly-sized collectives close most of the subgroup accuracy gap, transfer across architectures (a small proxy can attack a larger platform), and improve worst-group accuracy, equal-opportunity gap, and disparate impact. A query-budget analysis comparing a per-user black-box attack baseline shows that pooling is cheaper than each subgroup member attacking alone. Test-time collective action thus offers corrective intervention to users when platform-side remediation is unavailable or delayed.

- **Code -**
<https://github.com/tisl-lab/TestTimeCollectiveAction>
- **Extended version with appendix -**
<https://arxiv.org/abs/2605.27689>

1 Introduction

Machine learning systems can exhibit performance disparities across different user groups, with documented examples in computer vision (Buolamwini and Gebru 2018), safety features of language models (Chehbouni et al. 2024) and automated speech recognition (Koenecke et al. 2020). Shelby et al. (2023) describe such failures as quality-of-service

(QoS) harms. When users encounter these harms, provider-centric fairness interventions such as pre-processing, in-processing, and post-processing methods are commonly used to address them (Caton and Haas 2024). This could be considered reasonable if platforms are willing to act towards reducing these harms; however, when remediation is unavailable, it leaves affected users with no lever of their own to address the situation.

Recent work in Protective Optimization Technologies (Kulynych et al. 2020) and Algorithmic Collective Action (Hardt et al. 2023) identifies promising avenues for user-side interventions that allow affected users to regain agency in machine learning systems. Algorithmic Collective Action (Hardt et al. 2023) shows how even a small, organized collective can have an out-sized effect in steering a platform’s outputs in the collective’s favor. However, these mechanisms require the platform to re-train on user contributions, a condition that users rarely have control over in real-world settings. Radiya-Dixit et al. (2022) also discusses this limitation by showing that data poisoning-style interventions, as user-side protections, do not always work because providers can train on data collected before an attack was discovered, or simply wait for a more robust model before processing poisoned data. Hence, it is important for users to also have methods for taking corrective action at test time. In this work, we explore a test-time method that allows a group of users to pool their data contributions to generate a perturbation that mitigates the algorithmic harms of a frozen black-box model using only API query access. The high-level pipeline is shown in Figure 1. which involves extracting a collective proxy from pooled queries, then applying the resulting perturbation at test-time.

Shelby et al. (2023) define algorithmic harms as adverse experiences arising from the interaction between technical system components and societal power dynamics (we refer the reader to their survey for a full taxonomy on algorithmic harm). This work focuses on **quality-of-service (QoS) harms**, in which a system disproportionately under-performs for particular groups while maintaining strong performance for others. We focus on two regimes which could be the source of disparity: (i) a **balanced (or entire) platform**, where the platform sees the full training distribution and the under-served subgroup is whichever groups the model happens to perform worst on, and (ii) an **unbalanced**




Figure 1: *Test-Time Collective Action (TTCA)*. *Top*: members pool their per-user query budgets to query the platform with a public image pool $\mathcal{D}_{\text{pub}}$, then train a collective proxy on the resulting query–response pool $\mathcal{D}_{\text{proxy}}$, where the proxy boundary tries to approximate platform boundary. The collective optimizes a per-class perturbation δ_y against it. *Bottom*: at test-time, each member of the under-served subgroup applies δ_y to their own input x , sending $x + \delta_y$ to the platform in place of x .

platform, where a controlled fraction of groups is deliberately under-sampled before training so that the under-served subgroup is the one the platform was systematically under-trained on. In both cases, the collective’s goal is to close the accuracy gap on the affected subgroup without degrading performance for other groups.

In both settings, we assume that affected users are non-adversarial with respect to the platform’s task objective and their goal is to attain the same quality of service that the rest of the population already receives. We assume scenarios in which the model is broadly useful to users but fails to deliver comparable performance or utility to a particular subgroup due to data imbalance or even deployment choices. Our framing thus treats this collective intervention as a form of a *corrective action* for unequal outcomes when platform-side remediation is unavailable or delayed.

As an example of this setting, consider a user-driven photo-sharing platform deploys an image classifier that auto-tags every uploaded image and feeds those tags into its search and recommendation systems. The tool is broadly useful for users: most users’ photos surface for the searches they should, reach an audience whose interests match. But for some users, the classifier produces systematically more errors than it does for the rest of the user base. The affected users do not want to break the search for the rest of the platform; they want their content surfaced as fairly as everyone else’s.

We study this setting empirically on standard image-classification benchmarks throughout this paper and evaluate the feasibility of such an intervention. Figure 2 shows the intervention in action on CIFAR-100 across both balanced and unbalanced platform training where a single per-class perturbation considerably improves the accuracy of the worst-off subgroup without disturbing the rest of the label space.

In this work, we present Test-Time Collective Action as a mechanism to correct quality-of-service harms. We evaluate its performance on CIFAR-10, CIFAR-100, and FairFace

under both balanced and unbalanced platform training. We show the following:

1. Subgroup accuracy improves even at the smallest pool size and once a proxy is faithful enough to the predictions of the platform model, additional queries do not significantly increase subgroup accuracy. Indicating that a collective can generate an effective perturbation without exhausting their query budget or needing to label the full public dataset available to them.
2. Affected users attacking alone with per-image perturbations require up to 308× more total platform queries to reach comparable subgroup accuracy than a collective that pools queries once to train a proxy and generate reusable perturbations at $B=700$, narrowing to about 105× at the collective’s accuracy ceiling ($B=2500$ and above). Pooling can become cheaper than individual intervention with as few as five to ten contributing members.
3. Lastly, the gains transfer to standard group-fairness metrics such as worst-group accuracy and equal-opportunity gap, which positions test-time collective action as an effective user-side fairness intervention.

2 Related Works

When users face the unwanted consequences of algorithmic systems and anticipate that platform-level remedies will not be available, the remaining lever is what users can do outside the system. Kulynych et al. (2020) propose *Protective Optimization Technologies* (POTs) as an alternative to internal fixes by platform. POTs enable individuals or collectives to act from outside optimization systems rather than relying on platform cooperation. *Data leverage* (Vincent et al. 2021) refers to the power that individuals or groups have operating in data-driven technologies, dependent on their contributions. By withholding, redirecting, or altering data, users can affect the model’s performance and outcomes. DeVrio, Eslami, and Holstein (2024) extend this lens to *responses*




Figure 2: *Per-class accuracy before and after the perturbation*: Test-time collective action helps the affected subgroup (defined here as bottom-10% under-performing classes, but can alternatively include user demographics where available) for both balanced and unbalanced platform training regimes, without participating in the platform’s training loop (subgroup selection and dataset regimes are defined in §4.3 and §4.4 respectively). The plot shows CIFAR-100 classes plotted in balanced-platform subgroup accuracy, unbalanced-platform subgroup accuracy coordinates. *Left*: bottom-10% classes on the balanced platform (blue circles) move rightward after applying δ_y , improving subgroup accuracy for balanced-platform regime. *Right*: deliberately under-represented classes on the unbalanced platform regime (red squares) move upward after δ_y . Other classes (gray) stay near the $y = x$ diagonal in both panels, so the perturbation acts as corrective recourse for the targeted subgroup rather than as a system-wide attack.

from below, examining how people collectively build and use tools to resist or redirect algorithmic influence.

Algorithmic Collective Action (ACA) (Hardt et al. 2023) formalizes how even a small, coordinated group of users can steer the algorithmic system in their favor by planting signals or erasing them. Ben-Dov et al. (2025) extend this to group fairness by implementing the erasure strategy, in which a minority-only collective relabels its own data to approximate counterfactual fairness. Recent work considers multiple simultaneous collectives (Karan et al. 2025) and robustness under differential privacy (Solanki et al. 2025).

The current work in ACA assumes that the platform trains on the modified data of the collective; however, in practice, users may not always have control over the training loops of the platform. Radiya-Dixit et al. (2022) discuss that once a perturbation is scraped, it cannot be revised, so the model trainer can simply wait for newer architectures that incidentally resist the perturbation, or adaptively train against the public attack tool. Test-time collective action alongside train-time action may give users more axes of agency in the system. An example of a deployment-time strategy explored in ACA literature is *Decline Now*. Sigg, Hardt, and Mendler-Dünnér (2025) model the #DeclineNow campaign, in which DoorDash drivers collectively reject low-

paying orders to trigger algorithmic price increases on the deployed allocation system.

Outside the ACA literature, evasion attacks (Biggio et al. 2013) form the test-time counterpart to data-poisoning interventions. Instead of altering the data the platform trains on, the user perturbs the inputs at test-time to control the model’s behavior. We re-purpose these mechanisms for correction purposes in this work (see the Adversarial Perturbations in §3 for details), showing that a small group of coordinated users can *improve* model performance and fairness without requiring model retraining, or even knowledge of its architecture.

Most directly related to our work is fairness-aware adversarial perturbation (FAAP) (Wang et al. 2022), which also uses a universal adversarial perturbation at test time to mitigate fairness disparities of a deployed model. The works overlap on the broad strategy but differ in their mechanisms and framing. FAAP trains a generator-discriminator pair on a deployed model to hide protected attribute information at test time. The motivation is that if the model cannot detect a sample’s protected attribute, its predictions cannot correlate with that attribute. Our method takes a different approach, aiming to correct misclassifications of under-served groups using a universal perturbation that pushes the sample toward its true class, rather than censoring protected attributes. Our setting differs in both agency and access assumptions. FAAP studies fairness-aware perturbations as a bias-mitigation mechanism for deployed models, whereas TTCA frames perturbation generation as a user-side collective action problem under black-box API access. Affected users do not assume platform cooperation or internal model access; instead, they pool limited queries, extract a proxy, and generate reusable class-specific perturbations that can be applied at test time. Thus, TTCA contributes a collective, black-box, query-budgeted formulation of corrective perturbations rather than a new platform-side fairness intervention.

3 Background

Black-box Model Extraction. Model extraction attacks exploit query access provided by prediction APIs to approximate, or “steal”, a deployed model (Tramèr et al. 2016). In a typical extraction attack, the adversary, who has black-box access, queries the target model on auxiliary inputs, collects the resulting query-response pairs, and uses them to train a surrogate model. Following prior work Jagielski et al. (2020), the quality of the extraction can be measured using various metrics, such as accuracy (performance on the underlying learning task) and fidelity (matching the target model’s predictions on a reference set). In this work, we use *Knockoff Nets* (Orekondu, Schiele, and Fritz 2019) to construct the collective proxy, since it assumes no knowledge of the train/test data used by the target model. *Knockoff Nets* offers two query strategies: a random strategy that uniformly samples from a public data pool, and an adaptive strategy that learns a sampling policy from auxiliary labels in that pool. We use the former to construct the collective’s proxy model.

Adversarial Perturbations. Adversarial examples add small perturbations to an input, which can cause a model to change its output (Szegedy et al. 2014; Goodfellow, Shlens, and Szegedy 2015). Such perturbations belong to the class of evasion attacks (Biggio et al. 2013), in which the adversary manipulates inputs at test time to induce erroneous predictions, without interfering with the model’s training process. Moosavi-Dezfooli et al. (2017) introduced *Universal Adversarial Perturbation* (UAP), which aims to find a single, universal perturbation that can misclassify images with high probability. Targeted variants of UAP, such as *localized patches* and *iterative tFGSM aggregation*, were subsequently developed by Brown et al. (2017) and Hirano and Takemoto (2020), respectively. In the context of per-image perturbation, query-only black-box attacks such as *SimBA* (Guo et al. 2019) fit a fresh perturbation to each image using only the platform’s confidence scores. We use SimBA later as the individual baseline against which we compare the collective intervention.

Targeted UAPs are especially interesting in the context of collective user-side interventions because they can enable users to steer an algorithmic system toward a specific decision without accessing its training pipeline. Targeted UAP methods optimize a perturbation that moves inputs from non-target classes across the model’s decision boundary and into the target class. However, when the perturbation is optimized against a proxy of the deployed model rather than the model itself, the proxy’s decision boundary is an approximation of the platform’s, and a perturbation tuned to cross the proxy’s boundary may fail to cross the platform’s. Zeng et al. (2023) identify a closely related failure mode in the context of backdoor attacks, where surrogate models trained on out-of-distribution data, and propose instead optimizing the perturbation toward the target class itself rather than across its boundary. Our mechanism adopts this target-directed perturbation, but inverts both the threat model and the objective. The perturbation is applied at test time rather than during training (as in clean-label poisoning), and the intended predicted class is the input’s ground-truth class. Thus, the goal is not to cause misclassifications but to steer the deployed platform model toward the correct decision.

Group fairness Group fairness has been formalized along several axes to assess whether a classifier’s outcomes are distributed equitably across groups. *Disparate impact* (Feldman et al. 2015) measures the rate at which an unprivileged group receives the favorable outcome compared to the rate for the privileged group. *Equal opportunity* (Hardt, Price, and Srebro 2016) shifts the focus from outcome rates to error rates, requiring that true-positive rates be equalized across sensitive groups conditional on the advantaged outcome, so groups that genuinely qualify for a favorable label are recognized at equal rates. *Worst-group accuracy* (WGA) (Sagawa et al. 2020) measures the minimum classification accuracy across all groups, regardless of which group is “privileged.” Each metric captures a distinct fairness concern (rate of favor, error parity, worst-case service quality). The setting we study uses these metrics to determine whether the intervention improves group fairness in a corrective direction. The

metrics are used to check if a crafted perturbation can reduce the existing disparity at inference time.

4 Problem and Experimental Setup

4.1 Threat Model

A platform operates a classifier $f_{\text{plat}} : \mathcal{X} \rightarrow [0, 1]^{|\mathcal{Y}|}$ as a black-box API: users submit images $x \in \mathcal{X}$ and receive the platform’s probability vector $f_{\text{plat}}(x)$ over the label space, whose $\text{argmax } \hat{y} = \text{argmax}_y f_{\text{plat}}(x)_y$ is the predicted class. Weights, architecture, training data, and intermediate representations are not exposed. The collective therefore observes only input–output behavior through the API and has no direct access to the model’s internal parameters, structure, training data, or hidden representations.

A *collective* C of n users coordinates to correct misclassifications affecting an under-served subgroup $\mathcal{G}$. A collective C may have members from $\mathcal{G}$ as well as other members in the system although may not necessarily overlap. Each member is constrained by a per-user query limit of B requests, so the collective can pool a total budget of $Q = nB$ queries. Members share three resources: query access to f_{plat} within their individual budgets; a public, unlabeled images pool $\mathcal{D}_{\text{pub}}$ that is disjoint from the platform’s private training data; and the infrastructure to aggregate query results and distribute the generated perturbations. In our experiments, we vary the pooled budget Q directly and report subgroup accuracy and fairness as functions of this collective budget. The decomposition $Q = nB$ is used only when we compare the collective to an individual user acting alone, where we hold B fixed at each point of the sweep and vary n , so an individual at budget B compares directly to a collective of n users at the same per-user cost (total budget nB). Fixing B rather than Q isolates the effect of collective coordination as each user faces the same individual cost, so any improvement from cooperation is attributable to pooling rather than to one user having more queries available.

The collective’s goal is to improve the platform’s accuracy on the affected subgroup $\mathcal{G}$ at test time, without requiring the platform to retrain on modified inputs. The collective’s lever is to add a small, bounded perturbation to the inputs of subgroup members before they are submitted to the API for prediction. Thus, this intervention repurposes the logic of evasion attacks toward a corrective end. Concretely, for each class y represented in $\mathcal{G}$, the collective constructs a single perturbation δ_y with $\|\delta_y\|_{\infty} \leq \varepsilon$ so that all subgroup members of class y apply to their own inputs at test time. The perturbed input is $x' = \text{clip}(x + \delta_y, 0, 1)$, where $\text{clip}(z, 0, 1) = \min(\max(z, 0), 1)$ is applied element-wise to keep pixels in the valid range. The collective seeks to maximize subgroup accuracy:

$$\text{SubgroupAcc} = \Pr[f_{\text{plat}}(x') = y \mid y \in \mathcal{G}], \quad (1)$$

The same δ_y has to work across every subgroup member of class y , which is what makes the perturbation *universal* (where a single per-class perturbation, distributed once, is applied independently by each affected user).

4.2 Mechanism

Crafting δ_y would normally require access to gradients of the platform classifier, which are unavailable through a black-box API. The collective’s workaround is to first train a surrogate classifier from the platform’s query responses, and then optimize the perturbation against this surrogate.

(1) Query pooling. First, each member uses their query budget to label a subset of $\mathcal{D}_{\text{pub}}$ with predictions from f_{plat} . These are aggregated into a shared pool $\mathcal{D}_{\text{proxy}} = \{(x, f_{\text{plat}}(x))\}$ of Q pairs, each linking an input x to f_{plat} ’s output rather than to a ground-truth label or one supplied by the collective. So the proxy trained on $\mathcal{D}_{\text{proxy}}$ imitates f_{plat} ’s predictions and even inherent biases rather than producing an accurate classifier of its own. This is the only stage that consumes platform queries; all subsequent steps are performed offline using $\mathcal{D}_{\text{proxy}}$.

(2) Proxy extraction. The collective then trains a surrogate f_{proxy} on $\mathcal{D}_{\text{proxy}}$ via KnockoffNets random sampling (Orekony, Schiele, and Fritz 2019), fitting a network with cross-entropy against the platform’s soft-probability outputs. Because gradients on a single proxy can be noisy, we additionally train $K = 5$ proxies with independent random seeds and ensemble their outputs as $f_{\text{ens}} = \frac{1}{K} \sum_{k=1}^K f_{\text{proxy}}^{(k)}$. Previous work has shown that ensembling multiple surrogate models can improve the transferability of crafted perturbations to target models (Yao et al. 2025).

(3) Perturbation optimization. For each class y in $\mathcal{G}$, the collective optimizes a perturbation δ_y against the proxy ensemble under an ℓ_∞ budget:

$$\delta_y = \arg \min_{\|\delta\|_\infty \leq \varepsilon} \mathbb{E}_{x \sim \mathcal{D}_{\text{proxy}}^y} [\text{CE}(f_{\text{ens}}(\tilde{x}), y)], \quad (2)$$

where $\mathcal{D}_{\text{proxy}}^y$ is the subset of queries the proxy labelled as y and $\tilde{x} = \text{clip}(x + \delta_y, 0, 1)$ is the perturbed input. A small total-variation regularizer (10^{-4}) and a Gaussian blur on δ at each step bias the perturbation toward low-frequency components. This design choice is motivated by prior findings on black-box transferability (Guo, Frank, and Weinberger 2019). This formalization can be seen as a user-side counterpart to boosting-style fairness corrections such as multi-accuracy (Kim, Ghorbani, and Zou 2019) where the corrective signal is optimized against a proxy ensemble from the under-served group’s side and applied to the input at inference, rather than added by the platform to its own classifier during post-processing.

(4) Test-time application. The collective distributes the set $\{\delta_y : y \in \mathcal{G}\}$ to its members. A member with input x of class y checks whether $y \in \mathcal{G}$; if so, they submit $\text{clip}(x + \delta_y, 0, 1)$ instead of x ; otherwise the input is sent unchanged. The pool $\mathcal{D}_{\text{proxy}}$ is used for dual purposes. It is used for generating a training signal for *Proxy extraction* in stage (2) and optimization set for *Perturbation optimization* in stage (3), so no platform queries are spent beyond the initial Q .

4.3 Datasets, Partitions, and Subgroups

We evaluate our method on different image-classification datasets. CIFAR-10 and CIFAR-100 (Krizhevsky and Hinton 2009) which are standard image-classification benchmarks of 32×32 color images with 10 and 100 classes respectively. For CIFAR-100 we use the fine-grained labels rather than the 20 coarse super-class labels. In addition to the CIFAR datasets we also use FairFace (Kärkkäinen and Joo 2021) which is a face-attribute dataset of 108,501 images, balanced across 7 race groups (White, Black, Indian, East Asian, Southeast Asian, Middle Eastern, Latino) and labelled with binary gender, race, and 9 age buckets. FairFace is used here because it gives us a real-world benchmark where the target label (gender) interacts with sensitive attributes (race, age), allowing us to evaluate the mechanism on demographic fairness rather than just per-class accuracy disparities. As the dataset is balanced across race/gender groups and unbalanced across age groups it also allows us to model the balanced and unbalanced platform regimes. We follow the dataset’s binary gender annotation. However, gender recognition from faces is itself a contested practice (Keyes 2018; Scheuerman, Paul, and Brubaker 2019), and we return to this scope question in §7.

Disjoint platform/collective partitions. For each dataset, the official train split is shuffled with a fixed seed and divided 50:50 (stratified by label) into a *platform* partition $\mathcal{D}_P$ and a *collective* partition $\mathcal{D}_C$. This is motivated by the assumption that the collective has no access to the platform’s training set, but can rely on auxiliary samples from the same underlying data distribution. This assumption reflects the fact that collective members are not external attackers probing an unrelated system; they are platform users whose inputs belong to the same operational environment as the model’s test-time inputs. Thus, the collective’s advantage comes from pooling user-side data and query access, not from privileged access to the platform’s training set or internal model state. Within each partition, a further 90:10 split yields train/val pools. The official test split is left intact, and neither party trains on it; it is only used for evaluation.

Groups and subgroup definition. A *group* is the unit at which fairness metrics are computed. On CIFAR datasets, the group is defined by the class label. On the FairFace dataset, the group is a (gender, sensitive attribute) pair, where the sensitive attribute is one of the 7 race buckets or one of the 9 age buckets tagged as attributes in the dataset. The under-served subgroup $\mathcal{G}$ is the bottom 10% of groups by per-group accuracy on the shared test set.

4.4 Balanced (or Entire) vs. Unbalanced Platform Training

We study QoS harm under two platform-training cases that differ in *why* the under-served subgroup $\mathcal{G}$ is under-served. This distinction is important because the collective sees only the platform’s outputs and not the reason behind them. An effective intervention should thus work whether the reason for the disparity comes from an incidental model weakness or from platform under-training on a subgroup due to a lack

of data. The corrective objective from Equation 1 is identical across both and only the construction of $\mathcal{G}$ differs.

Balanced (or entire) platform. The platform is trained on the full partition $\mathcal{D}_P^{\text{train}}$ with no deliberate sub-sampling. $\mathcal{G}$ is defined *post hoc* as the bottom-10% groups by per-group accuracy on the shared test set. The collective’s pool and the platform’s data come from the same distribution, so the under-service arises from whatever combination of model weakness and natural data imbalance the dataset itself exhibits, rather than from any sub-sampling we have introduced.

Unbalanced platform. In this case, the platform is instead trained on a deliberately skewed variant of $\mathcal{D}_P^{\text{train}}$. 10% of groups are selected at random as under-represented and, within each, only a small fraction of examples is retained while the remaining groups are kept whole. The retention rate is set per dataset based on how aggressively the per-group pool can be cut without pushing the platform into underfitting: 10% for CIFAR-10 and FairFace, and 30% for CIFAR-100, where the per-class platform pool is already small. Under-represented groups are picked uniformly at random for CIFAR-10/100 and FairFace (sensitive attribute: race); for FairFace (sensitive attribute: age) as the data is only balanced across race groups and not age buckets there is already a natural imbalance in the dataset. Thus, we instead pick the smallest age buckets to mirror the dataset’s natural imbalance.

4.5 Models and Training Parameters

Platform. The platform backbones are ResNet-18 (IMAGENET1K_V1), ResNet-50 (IMAGENET1K_V2) (He et al. 2016) and RegNet-Y-16GF (IMAGENET1K_V2) (Xu et al. 2023), pretrained on ImageNet, with only the final fc head replaced by a linear layer of size $|\mathcal{Y}|$. Both backbones are fine-tuned end-to-end on $\mathcal{D}_P^{\text{train}}$ with AdamW ($\eta = 10^{-3}$, weight decay 10^{-4}), a cosine schedule over 30 epochs, batch size 128, and cross-entropy loss. We select the checkpoint with the highest validation accuracy.

Proxy. The KnockoffNets surrogate uses the same backbone family as the platform by default. We use the implementation from the Adversarial Robustness Toolbox (Nicolae et al. 2018) where the implementation passes `use_probability=True` to KnockoffNets, so the proxy distills the full output distribution rather than only the argmax label). We also test a cross-architecture setting in which a smaller ResNet-18 proxy is trained against a ResNet-50 and RegNet-Y-16GF platform, modeling the case where users do not know the platform’s architecture. Proxy hyper-parameters (AdamW, label smoothing 0.1, 30 epochs, batch 128) are held fixed across configurations. We sweep the total query budget Q at six log-spaced points from 500 to $0.6|\mathcal{D}_C^{\text{train}}|$, and run five seeds per configuration so that subgroup-accuracy curves report mean $\pm$ standard deviation.

Perturbation. Perturbation is optimized according to Stage (3) of mechanism mentioned in §4.2. Hyperparameter used for optimization are Adam ($\eta = 10^{-2}$, mini-batch

64, 200 iterations) with ℓ_∞ projection after each step. We use $\varepsilon = 0.031$ for CIFAR and $\varepsilon = 0.12$ for FairFace.

5 Experiments

Subgroup Accuracy and Fairness Metrics The primary outcome we measure is the subgroup accuracy from Equation 1 which is the platform’s mean accuracy on the under-served subgroup $\mathcal{G}$ after members apply δ_y to their inputs at test time. We measure this as a function of the pooled query budget Q across both platform-training regimes (§4.4) and across platform backbones (ResNet-18, ResNet-50, ResNet-50 platform model $\rightarrow$ ResNet-18 proxy and RegNet-Y-16GF platform model $\rightarrow$ ResNet-18 proxy), with five seeds per configuration. To measure if the intervention improves group fairness for $\mathcal{G}$ we also report the following metrics (§3): worst-group accuracy, equal-opportunity gap, and disparate impact ratio. The collective’s intervention is corrective if subgroup accuracy rises while these metrics improve or hold.

Individual baseline. A natural alternative to collective pooling is for each affected user to attack the platform independently. In this baseline, each user runs a per-image targeted black-box attack on their own input, paying only for queries associated with that image and stopping once the platform predicts the correct class. This models an uncoordinated scenario in which users attempt to correct their own misclassifications without sharing data, queries, or perturbations with others.

We instantiate this baseline with targeted SimBA using the pixel variant (Guo et al. 2019). SimBA is appropriate for our setting because it requires only query access to the platform, the confidence score assigned to the target class, and the predicted label, which aligns with the type of API feedback assumed for the collective intervention. For each user, we run SimBA until the first successful correction or until the per-user query cap B is reached. The realized query cost q_i is therefore the actual number of queries spent by user i , capped at B , rather than simply the configured budget.

In the experiments, we sweep B , the per-user query budget that may be imposed by the platform through rate limiting. Both the individual and collective modes are evaluated on the same 1,000-user under-served subgroup selected in Section 4.4, using CIFAR-100 across five seeds. We focus on CIFAR-100 because it has the largest under-served subgroup among our benchmarks, making aggregate-cost scaling most visible.

For the individual baseline, each affected user attacks separately. Let $q_i = \min\{\tau_i, B\}$, where τ_i is the first query at which the individual attack succeeds for user i , and B is the per-user query budget. If the attack does not succeed within budget, then $q_i = B$. The total query cost of the individual baseline is therefore $\sum_i q_i$. For the collective intervention, users spend N total queries to build the proxy model and generate the class-wise perturbations. Once constructed, the same perturbation δ_y can be applied by all 1,000 affected users in class y at inference time without additional platform queries. To measure how many users are needed to build such a pool, we report the *break-even collective size*




Figure 3: *Subgroup accuracy after perturbation as a function of pooled query budget Q (log):* For balanced and unbalanced platform regimes even at the smallest query budget, the subgroup accuracy is above the pre-perturbation baseline (dashed gray line) and the cross-architecture trajectories stay close to the same-architecture curves as Q grows.

$K^* = \lceil N/B \rceil$, alongside the realized solo cost $\sum_i q_i$. Here, K^* is conservative: it is the minimum number of contributors required if each contributes their full budget B to obtain a pool of size N that matches the individual SimBA baseline at the same budget. If more than K^* users participate, each contributor’s share falls below B while the pool size remains N , reducing the per-user query cost. Alternatively, if additional members each contribute up to B , the collective can build a larger pool and potentially improve accuracy beyond the individual baseline. The remaining $1,000 - K^*$ subgroup users can still benefit from δ_y without contributing any queries. Because the perturbation is reusable for the current deployed model, future affected users in class y who encounter the same failure mode can also apply it directly rather than generating a new fix from scratch.

6 Results


The experiments show that test-time collective action can be an effective corrective intervention. The result is consistent across CIFAR-10, CIFAR-100, and FairFace; under both balanced and deliberately skewed platform-training regimes; and even when the collective’s proxy is smaller than the platform it queries.

6.1 Subgroup accuracy improves at low query budgets

Subgroup Accuracy improvements Figures 3a and 3b show that the mechanism works at even the smallest query budgets across all three benchmarks and both platform-training regimes. We report subgroup accuracy after applying the collective perturbation as a function of the pooled query budget Q , distinguishing between platforms trained

on a balanced corpus (Figure 3a) and those trained on deliberately skewed data (Figure 3b). In most panels, the post-perturbation curves exceed the pre-perturbation baseline (dashed gray line) even at the smallest query budgets. This indicates that relatively small query pools are sufficient to improve subgroup accuracy. Appendix A.3 further probes what δ_y encodes in isolation by applying the perturbation to a blank white image. The results reveal a class-specific signal already at $N = 500$, suggesting that the collective can produce an effective perturbation even with a modest query budget. These findings show that the intervention benefits affected users without requiring the collective to exhaust its available queries. Higher pooled query budgets improve subgroup accuracy across all datasets until performance approaches its ceiling. Platforms trained on deliberately skewed data exhibit substantially higher seed-to-seed variance than those trained on balanced data. This variance is likely driven by the random selection of underrepresented groups across seeds. Consistent with this explanation, variance is much lower in the deterministic FairFace’s by-age setting, where we select the smallest naturally occurring age buckets rather than sampling the affected groups at random.

Cross-architecture transfer The collective may rely on a smaller proxy model when the platform architecture is unknown or when users face computational constraints. Figure 3 also evaluates cross-architecture transfer by training the proxy as a ResNet-18 while the platform uses a ResNet-50 or RegNet-Y-16GF. Across all datasets, the cross-architecture trajectory closely follows the same-architecture setting, indicating that perturbations crafted on a smaller ResNet-18 proxy transfer effectively to the larger ResNet-50 or RegNet-Y-16GF platform.



● cifar10 ● fairface (race) ● balanced teacher
● cifar100 ● fairface (age) ▲ unbalanced teacher

Figure 4: *Proxy fidelity vs subgroup accuracy after δ_y* : For CIFAR-10 and FairFace, subgroup accuracy saturates well below highest fidelity; However, CIFAR-100 remains in the climbing region and continues to benefit from higher Q . This indicates a faithful enough proxy suffices for most datasets, so users do not always need to exhaust their entire query budgets.

Proxy fidelity and Subgroup Accuracy Effective perturbations do not always require a high-fidelity proxy. Figure 4 plots the proxy’s fidelity against the subgroup accuracy the resulting perturbation achieves. Markers cluster tightly along an upward trend across every dataset. Hence, a larger Q leads to a more faithful proxy, and a more faithful proxy results in a perturbation that transfers back to the platform more reliably. On CIFAR-10 and FairFace, where the proxy reaches 60-90% fidelity within the budgets we test, the accuracy curve flattens near ceiling. CIFAR-100, however it never exceeds $\sim 42\%$ fidelity in our sweep and is still climbing, so additional queries continue to pay off. Here we can see once marginal fidelity stops translating into marginal accuracy, members can stop pooling rather than exhausting their per-user budgets or labeling the full public pool.

6.2 Individual vs Collective Action

Subgroup-accuracy break-even occurs with only a few members. We sweep the per-user query budget B and compare the individual SimBA baseline to the collective intervention on CIFAR-100 across seeds. Table 1 reports the individual baseline’s mean subgroup accuracy, the collective accuracy achieved by the smallest matching query pool, the corresponding pool size N , the break-even collective size $K^* = \lceil N/B \rceil$, and the realized cumulative solo cost $\sum_i q_i$. At $B = 100$, individual SimBA reaches only 55.1% subgroup accuracy, with an average cost of approximately 92 queries per user. This average reflects both users whose predictions are corrected quickly and users who exhaust the 100-query cap without success. In contrast, the collective reaches higher subgroup accuracy, 57.1%, with only

Query Budget (B)	Individual acc. (%)	Collective acc. (%)	Total Collective Query Pool (N)	Collective Members (K^*)	Avg. Individual queries ($\bar{q}$)	Total Individual queries ($\sum q_i$)
100	55.1 $\pm$ 1.5	57.1 $\pm$ 2.4	500 $\pm$ 0	5	92 $\pm$ 3	91,817 $\pm$ 2,719
200	56.8 $\pm$ 1.4	59.2 $\pm$ 1.2	780 $\pm$ 256	4	180 $\pm$ 6	180,050 $\pm$ 5,703
300	58.3 $\pm$ 1.1	60.2 $\pm$ 2.2	874 $\pm$ 209	4	265 $\pm$ 8	264,973 $\pm$ 8,326
400	59.9 $\pm$ 1.3	65.9 $\pm$ 4.5	1,415 $\pm$ 643	4	347 $\pm$ 10	346,875 $\pm$ 10,471
500	61.4 $\pm$ 0.9	66.4 $\pm$ 3.7	1,508 $\pm$ 494	3	425 $\pm$ 12	425,414 $\pm$ 12,460
600	62.5 $\pm$ 0.6	67.9 $\pm$ 3.3	1,689 $\pm$ 403	4	502 $\pm$ 14	501,548 $\pm$ 13,812
700	64.0 $\pm$ 0.5	68.9 $\pm$ 2.2	1,869 $\pm$ 0	3	575 $\pm$ 15	574,971 $\pm$ 14,542
1,000	67.8 $\pm$ 1.2	72.1 $\pm$ 3.6	2,566 $\pm$ 955	3	779 $\pm$ 18	778,518 $\pm$ 17,620
1,500	73.4 $\pm$ 1.0	77.5 $\pm$ 1.1	4,286 $\pm$ 1,508	3	1,072 $\pm$ 25	1,071,639 $\pm$ 25,140
2,000	78.4 $\pm$ 0.9	80.2 $\pm$ 2.3	6,309 $\pm$ 1,508	4	1,313 $\pm$ 33	1,313,186 $\pm$ 33,334
2,500	82.2 $\pm$ 0.8	82.3 $\pm$ 1.5	14,293 $\pm$ 7,522	6	1,509 $\pm$ 42	1,508,935 $\pm$ 42,121
4,000	90.1 $\pm$ 0.5	83.0 $\pm$ 2.6	22,000 $\pm$ 0	6	1,915 $\pm$ 57	1,914,676 $\pm$ 56,880
6,000	95.6 $\pm$ 0.6	83.0 $\pm$ 2.6	22,000 $\pm$ 0	4	2,184 $\pm$ 69	2,183,551 $\pm$ 69,278
8,000	97.9 $\pm$ 0.2	83.0 $\pm$ 2.6	22,000 $\pm$ 0	3	2,304 $\pm$ 80	2,303,648 $\pm$ 80,329

Table 1. *Query-budget analysis on CIFAR-100*: At every per-user budget B below the collective’s $\sim 83\%$ accuracy ceiling, pooling matches the individual SimBA baseline with far fewer platform queries, and the margin grows as B is higher. The highlighted row is the largest B at which a collective of K^* members still matches individual; the rows below it are budgets where individual surpasses the collective ceiling (by using two orders of magnitude more queries), and no pool size, including the full $\mathcal{D}_{\text{pub}}$, closes the gap.

$N = 500$ total query-response pairs. This corresponds to $K^* = 5$, meaning that five users each contributing their full 100-query budget are sufficient to build a pool that outperforms the individual baseline. Across all budgets up to the collective’s accuracy ceiling, K^* remains in the single digits (ranging from three to six), indicating that only a small number of contributors is needed for pooling to become competitive in subgroup accuracy.

Pooling also reduces the average per-contributor cost. The cost-sharing calculation also favors pooling once the collective is slightly larger than K^* . With K contributors, the accounting cost per contributor is N/K . This falls below the average solo cost once

$$K \geq \left\lceil \frac{N}{\bar{q}} \right\rceil, \quad \text{where } \bar{q} = \frac{1}{|\mathcal{G}|} \sum_{i \in \mathcal{G}} q_i.$$

At $B = 100$, this threshold is $\lceil 500/92 \rceil = 6$, only one more than the subgroup-accuracy break-even. This is a cost-sharing calculation rather than an additional experiment over all six-user subsets: it means that if six contributors contribute the same $N = 500$ query pool, their average contribution is lower than the average number of queries spent by a user attacking alone. Near the collective’s accuracy ceiling, the threshold remains modest. At $B = 2,500$, the collective requires approximately $N = 14,293$ queries, while the average solo cost is about 1,509 queries, so roughly ten contributors are enough for the average per-contributor share to fall below the average solo cost.

This comparison highlights the main advantage of pooling. A targeted black-box attack can correct an individual user’s prediction when that user has enough query budget, but the cost is not amortized: the next user facing the same failure must run a new per-image attack and pay a separate query cost. By contrast, the collective pays the proxy-extraction and perturbation-generation cost once, amortizes it across contributors, and produces a reusable perturbation δ_y that can be applied by non-contributing affected users at test time. This is especially valuable when many users face the same failure mode, or when individual users cannot sustain the solo attack loop because of account-level rate limits or other query constraints. The benefit persists for future affected users as long as the deployed model remains unchanged.

Well-resourced individuals can exceed the collective subgroup accuracy. The comparison also reveals a trade-off between individual and collective intervention, as shown in Figure 5. The individual baseline reaches a higher subgroup accuracy on CIFAR-100, 97.9%, than the collective, which saturates around 83.0%. This is expected: SimBA optimizes a separate perturbation for each input, using the platform’s own scores on that specific image, and can continue querying until the prediction crosses the decision boundary or the budget is exhausted. By contrast, the collective learns a universal perturbation δ_{y^*} that must work for all affected subgroup members of class y^* . Above the collective’s accuracy ceiling, no break-even pool size is defined because the present universal-perturbation mechanism cannot reach the individual baseline’s accuracy, even when the available public pool is exhausted. However, this higher individual accuracy comes at a much larger aggregate query cost. The collective reaches its ceiling with roughly 23,000 total platform queries paid once and reused across users, whereas bringing the 1,000 subgroup users to the individual ceiling requires approximately 2.3 million cumulative queries. Thus, well-resourced individuals can obtain higher accuracy by attacking alone, but collective pooling remains far more query-efficient at the subgroup level.

6.3 Fairness

Subgroup-accuracy gains translate into improvements on group-fairness metrics. Figure 6 compares worst-group accuracy (WGA), equal-opportunity gap (EOD), and disparate impact ratio (DI) before and after applying the collective perturbation. Each marker corresponds to one experimental setting, defined by a dataset, sensitive attribute, and platform-training regime, evaluated at the smallest pool size that reaches the 75% subgroup-accuracy target. We use this 75% target as a common reference point: it is high enough to indicate that the perturbation is effective, but low enough to be reached across all settings. This avoids confounding the fairness comparison with differences in accuracy ceilings across datasets and regimes. Most settings move in the improving direction for all three metrics: above the $y=x$ diagonal for WGA and DI, and below it for EOD. WGA shows the largest and most consistent improvement, which is expected because it is the metric most directly tied to the under-served




Figure 5: *Subgroup Accuracy vs Total Platform queries (log)*: Test-time collective action requires far fewer platform queries than individuals attacking alone however has a lower accuracy ceiling. Individual SimBA (gray) attains 97.9% after ≈ 2.3 M cumulative queries from the 1,000-user subgroup. In contrast collective, at accuracy ceiling, spends 22,000 queries once to train a shared proxy, then derives a single universal perturbation δ that every member applies at inference, so the platform-query requests stays flat at $N = 22,000$ regardless of how many of the 1,000 users participate; However, even when the collective consumes the entire available public pool, accuracy ceiling is at 83.02%, which is below the individual ceiling.

subgroup targeted by the perturbation. In particular, settings that initially have WGA in the 0.3–0.4 range typically improve beyond 0.4, and several CIFAR-10 settings reach 0.85 or higher after the intervention. This translation from subgroup accuracy to group fairness is partly by construction. In Stage (4) of §4.2, each member applies δ_y only to inputs whose true class is y , so improvements on the targeted subgroup naturally affect group-level fairness metrics. To check that the perturbation is not equally effective as a generic misclassification tool, Appendix A.1 evaluates cross-target use, where the same δ_y is applied to inputs whose true class is not y . The corrective rate substantially exceeds the cross-target rate at every pool size, indicating that the perturbation is a stronger lever for correction than for misuse.

7 Discussion

Platform-side defenses. In this work, we assume that platform model is a static black-box classifier and the collective optimizes their perturbation on the query-response pair from this model. However, a platform may also deploy adversarial defenses such as Adaptive Randomized Smoothing (Lyu et al. 2024) or Nasty Teacher (Ma et al. 2021) to reduce the effectiveness of universal perturbations on the target model or make training of proxy models more difficult. Furthermore, the intervention discussed in




Figure 6: *Fairness metrics before vs. after the collective perturbation*: Each marker is one (dataset, sensitive attribute, platform-training regime) configuration at the smallest pool size that reaches our 75% target accuracy; the shaded half-plane in each panel is the improving direction. Most markers fall on the improving side across all three metrics.

this work focuses on users encountering quality-of-service harms. A platform has limited incentive to deploy defenses against this kind of intervention because users’ motivation is corrective rather than adversarial, and the platform would presumably prefer corrected predictions. However, the incentives may differ in the case of allocative harms, where the disparity arises from the distribution of a limited resource. Appendix A.2 shows that platform-side defenses can be effective in preventing TTCA but also cost the platform in terms of accuracy degradation. Whether platform-side robustness disproportionately blocks corrective recourse compared to malicious use is an open empirical question.

Conservative individual baseline. We compare collective action against targeted SimBA (pixel variant), a simple query-only attack chosen to represent the individual baseline. More query-efficient per-image black-box (Feng et al. 2022; Andriushchenko et al. 2020; Chen, Jordan, and Wainwright 2020) could reach individual success in fewer queries, which would reduce the cost gap we report. We expect the collective to remain cheaper in aggregate because it pays a fixed proxy-extraction and perturbation-generation cost once, amortizes that cost across the subgroup, and allows non-contributing users to benefit from the resulting perturbation. However, stronger individual attacks would shift the break-even K^* upward and tighten the per-contributor cost threshold $\lceil N/\bar{q} \rceil$. Appendix A.4. shows experiments for Square Attack (Andriushchenko et al. 2020) where the break-even point is higher than SimBA baseline. The Individual vs. Collective comparison in §6.1 should therefore be interpreted as a conservative estimate of the collective’s cost advantage, not as its expected magnitude under the strongest possible individual attack.

Stronger adversarial mechanisms. We propose one specific way of implementing a test-time collective action framework where we use a KnockoffNets ensemble with a per-class optimized perturbation. Other adversarial methods could be substituted at any stage of the pipeline, both in how

the proxy is extracted from the platform and in how the perturbation is optimized against it. Evaluating these systematically is a natural extension and would sharpen the practical guidance the framework can offer a real collective.

Beyond quality-of-service framings. Our framing treats subgroup-accuracy disparities as quality-of-service harms that affected users should be able to correct. This assumption holds only when the underlying task objective is broadly useful to users. However, this may not work when the task itself is the cause of the harm. For instance, face-attribute classification, particularly in automated gender recognition, has been criticized in various works. Keyes (2018) argues that by treating gender as a binary and externally visible attribute, the harm of automated gender recognition is structural rather than a matter of performance, leading to systems erasing trans identities. Scheuerman, Paul, and Brubaker (2019) also in their evaluations find that facial analysis services underperform on transgender individuals and are unable to classify non-binary genders. Shelby et al. (2023) formalize such effects as *representational harms*, including erasure of social groups, alienation of those whose membership the system fails to acknowledge, denial of the opportunity to self-identify, and reification of essentialist categories. Keyes (2018) recommends that designers of gendered systems should ask whether the system needs to be gendered at all, and if so, how to gender it in a way that recognizes a wide range of people. Hence, although our experiments show that subgroup-accuracy gaps can be reduced technically in a gender classifier, as in our FairFace results, they do not address the harm of deploying such a classifier in the first place. When the task itself is harmful, user-side action may need to move beyond the corrective framing developed in this paper and instead take forms such as refusal, opt-out, or contestation (Zong and Matias 2024).

8 Conclusion

We proposed and evaluated test-time collective action: a mechanism through which a coordinated group of users, relying only on standard API access to a deployed classifier, can improve subgroup accuracy for an under-served group without participating in, or relying on, the platform’s training pipeline. We frame this as a form of *corrective intervention*: the platform is assumed to be broadly useful, but uneven in the quality of service it provides, and the central question is whether affected users can close this gap from outside the training loop when platform-side remediation is unavailable or delayed. We answer this by using a proxy-based mechanism where the collective generate a per-class perturbations. This not only addresses the issue by closing the subgroup-accuracy gaps but also translates into fairness metrics and is orders of magnitude fewer API queries than uncoordinated individual interventions. This work contributes to the ongoing discussion on external, user-driven mechanisms to correct harms. TTCA also acts as a complementary approach to the ACA framework in the test-time setting, showing that collective action remains a viable option even when users cannot influence the training loop or know whether their data was used to train the model.

Acknowledgements

The resources used in preparing this research were provided, in part, by the Province of Ontario, the Government of Canada through CIFAR, and companies sponsoring the Vector Institute (www.vectorinstitute.ai/partnerships/). The authors thank the Digital Research Alliance of Canada for computing resources. Ulrich Aïvodji is supported by an NSERC Discovery grant (RGPIN-2022-04006) and an IVADO's Canada First Research Excellence Fund to develop Robust, Reasoning and Responsible Artificial Intelligence (R³AI) grant (RG-2024-290714). Elliot Creager is supported by an NSERC Discovery grant (RGPIN-2024-05116).

References

- Andriushchenko, M.; Croce, F.; Flammarion, N.; and Hein, M. 2020. Square Attack: A Query-Efficient Black-Box Adversarial Attack via Random Search. In *ECCV (23)*, Lecture Notes in Computer Science, 484–501. Springer.
- Ben-Dov, O.; Samadi, S.; Sanyal, A.; and Tifrea, A. 2025. Fairness for the People, by the People: Minority Collective Action. *arXiv preprint arXiv:2508.15374*.
- Biggio, B.; Corona, I.; Maiorca, D.; Nelson, B.; Srndic, N.; Laskov, P.; Giacinto, G.; and Roli, F. 2013. Evasion Attacks against Machine Learning at Test Time. In *ECML/PKDD (3)*, Lecture Notes in Computer Science, 387–402. Springer.
- Brown, T. B.; Mané, D.; Roy, A.; Abadi, M.; and Gilmer, J. 2017. Adversarial patch. *arXiv preprint arXiv:1712.09665*.
- Buolamwini, J.; and Gebru, T. 2018. Gender Shades: Intersectional Accuracy Disparities in Commercial Gender Classification. In *FAT*, Proceedings of Machine Learning Research, 77–91. PMLR.
- Caton, S.; and Haas, C. 2024. Fairness in Machine Learning: A Survey. *ACM Comput. Surv.*, 56(7).
- Chehbouni, K.; Roshan, M.; Ma, E.; Wei, F. A.; Taïk, A.; Cheung, J. C. K.; and Farnadi, G. 2024. From Representational Harms to Quality-of-Service Harms: A Case Study on Llama 2 Safety Safeguards. In *ACL (Findings)*, Findings of ACL, 15694–15710. Association for Computational Linguistics.
- Chen, J.; Jordan, M. I.; and Wainwright, M. J. 2020. Hop-SkipJumpAttack: A Query-Efficient Decision-Based Attack. In *SP*, 1277–1294. IEEE.
- DeVrio, A.; Eslami, M.; and Holstein, K. 2024. Building, Shifting, & Employing Power: A Taxonomy of Responses From Below to Algorithmic Harm. In *FAccT*, 1093–1106. ACM.
- Feldman, M.; Friedler, S. A.; Moeller, J.; Scheidegger, C.; and Venkatasubramanian, S. 2015. Certifying and Removing Disparate Impact. In *KDD*, 259–268. ACM.
- Feng, R.; Mangaokar, N.; Chen, J.; Fernandes, E.; Jha, S.; and Prakash, A. 2022. GRAPHITE: Generating Automatic Physical Examples for Machine-Learning Attacks on Computer Vision Systems. In *EuroS&P*, 664–683. IEEE.
- Goodfellow, I. J.; Shlens, J.; and Szegedy, C. 2015. Explaining and Harnessing Adversarial Examples. In *ICLR (Poster)*.
- Guo, C.; Frank, J. S.; and Weinberger, K. Q. 2019. Low Frequency Adversarial Perturbation. In *UAI*, Proceedings of Machine Learning Research, 1127–1137. AUAI Press.
- Guo, C.; Gardner, J. R.; You, Y.; Wilson, A. G.; and Weinberger, K. Q. 2019. Simple Black-box Adversarial Attacks. In *ICML*, Proceedings of Machine Learning Research, 2484–2493. PMLR.
- Hardt, M.; Mazumdar, E.; Mendler-Dünner, C.; and Zrnic, T. 2023. Algorithmic Collective Action in Machine Learning. In *ICML*, Proceedings of Machine Learning Research, 12570–12586. PMLR.
- Hardt, M.; Price, E.; and Srebro, N. 2016. Equality of Opportunity in Supervised Learning. In *NIPS*, 3315–3323.
- He, K.; Zhang, X.; Ren, S.; and Sun, J. 2016. Deep Residual Learning for Image Recognition. In *CVPR*, 770–778. IEEE Computer Society.
- Hirano, H.; and Takemoto, K. 2020. Simple iterative method for generating targeted universal adversarial perturbations. *Algorithms*, 13(11): 268.
- Jagielski, M.; Carlini, N.; Berthelot, D.; Kurakin, A.; and Papernot, K. 2020. High Accuracy and High Fidelity Extraction of Neural Networks. In *USENIX Security Symposium*, 1345–1362. USENIX Association.
- Karan, A.; Vincent, N.; Karahalios, K.; and Sundaram, H. 2025. Algorithmic Collective Action with Two Collectives. In *FAccT*, 1468–1483. ACM.
- Kärkkäinen, K.; and Joo, J. 2021. FairFace: Face Attribute Dataset for Balanced Race, Gender, and Age for Bias Measurement and Mitigation. In *WACV*, 1547–1557. IEEE.
- Keyes, O. 2018. The Misgendering Machines: Trans/HCI Implications of Automatic Gender Recognition. *Proc. ACM Hum. Comput. Interact.*, 2(CSCW): 88:1–88:22.
- Kim, M. P.; Ghorbani, A.; and Zou, J. Y. 2019. Multiaccuracy: Black-Box Post-Processing for Fairness in Classification. In *AIES*, 247–254. ACM.
- Koenecke, A.; Nam, A.; Lake, E.; Nudell, J.; Quartey, M.; Mengesha, Z.; Toups, C.; Rickford, J. R.; Jurafsky, D.; and Goel, S. 2020. Racial disparities in automated speech recognition. *Proc. Natl. Acad. Sci. USA*, 117(14): 7684–7689.
- Krizhevsky, A.; and Hinton, G. 2009. Learning multiple layers of features from tiny images. Technical Report 0, University of Toronto, Toronto, Ontario.
- Kulynych, B.; Overdorf, R.; Troncoso, C.; and Gürses, S. F. 2020. POTs: protective optimization technologies. In *FAT**, 177–188. ACM.
- Lyu, S.; Shaikh, S.; Shpilevskiy, F.; Shelhamer, E.; and Lécuyer, M. 2024. Adaptive Randomized Smoothing: Certified Adversarial Robustness for Multi-Step Defences. In *NeurIPS*.
- Ma, H.; Chen, T.; Hu, T.; You, C.; Xie, X.; and Wang, Z. 2021. Undistillable: Making A Nasty Teacher That CAN-NOT teach students. In *ICLR*. OpenReview.net.
- Moosavi-Dezfooli, S.; Fawzi, A.; Fawzi, O.; and Frossard, P. 2017. Universal Adversarial Perturbations. In *CVPR*, 86–94. IEEE Computer Society.

- Nicolae, M.-I.; Sinn, M.; Tran, M. N.; Buesser, B.; Rawat, A.; Wistuba, M.; Zantedeschi, V.; Baracaldo, N.; Chen, B.; Ludwig, H.; et al. 2018. Adversarial Robustness Toolbox v1.0.0.
- Orekondy, T.; Schiele, B.; and Fritz, M. 2019. Knock-off Nets: Stealing Functionality of Black-Box Models. In *CVPR*, 4954–4963. Computer Vision Foundation / IEEE.
- Radiya-Dixit, E.; Hong, S.; Carlini, N.; and Tramèr, F. 2022. Data Poisoning Won’t Save You From Facial Recognition. In *ICLR*. OpenReview.net.
- Sagawa, S.; Koh, P. W.; Hashimoto, T. B.; and Liang, P. 2020. Distributionally Robust Neural Networks. In *ICLR*. OpenReview.net.
- Scheuerman, M. K.; Paul, J. M.; and Brubaker, J. R. 2019. How Computers See Gender: An Evaluation of Gender Classification in Commercial Facial Analysis Services. *Proc. ACM Hum. Comput. Interact.*, 3(CSCW): 144:1–144:33.
- Shelby, R.; Rismani, S.; Henne, K.; Moon, A.; Ros-tamzadeh, N.; Nicholas, P.; Yilla-Akbari, N.; Gallegos, J.; Smart, A.; García, E.; and Virk, G. 2023. Sociotechnical Harms of Algorithmic Systems: Scoping a Taxonomy for Harm Reduction. In *AIES*, 723–741. ACM.
- Sigg, D.; Hardt, M.; and Mendler-Dünner, C. 2025. Decline Now: A Combinatorial Model for Algorithmic Collective Action. In *CHI*, 912:1–912:17. ACM.
- Solanki, R.; Bhange, M.; Aïvodji, U.; and Creager, E. 2025. Crowding out the noise: Algorithmic collective action under differential privacy.
- Szegedy, C.; Zaremba, W.; Sutskever, I.; Bruna, J.; Erhan, D.; Goodfellow, I. J.; and Fergus, R. 2014. Intriguing properties of neural networks. In *ICLR (Poster)*.
- Tramèr, F.; Zhang, F.; Juels, A.; Reiter, M. K.; and Risten-part, T. 2016. Stealing Machine Learning Models via Prediction APIs. In *USENIX Security Symposium*, 601–618. USENIX Association.
- Vincent, N.; Li, H.; Tilly, N.; Chancellor, S.; and Hecht, B. J. 2021. Data Leverage: A Framework for Empowering the Public in its Relationship with Technology Companies. In *FAccT*, 215–227. ACM.
- Wang, Z.; Dong, X.; Xue, H.; Zhang, Z.; Chiu, W.; Wei, T.; and Ren, K. 2022. Fairness-aware Adversarial Perturbation Towards Bias Mitigation for Deployed Deep Models. In *CVPR*, 10369–10378. IEEE.
- Xu, J.; Pan, Y.; Pan, X.; Hoi, S. C. H.; Yi, Z.; and Xu, Z. 2023. RegNet: Self-Regulated Network for Image Classification. *IEEE Trans. Neural Networks Learn. Syst.*, 34(11): 9562–9567.
- Yao, W.; Zhang, Z.; Tang, H.; and Liu, Y. 2025. Understanding Model Ensemble in Transferable Adversarial Attack. In *ICML*, Proceedings of Machine Learning Research. PMLR / OpenReview.net.
- Zeng, Y.; Pan, M.; Just, H. A.; Lyu, L.; Qiu, M.; and Jia, R. 2023. Narcissus: A Practical Clean-Label Backdoor Attack with Limited Information. In *CCS*, 771–785. ACM.
- Zong, J.; and Matias, J. N. 2024. Data Refusal from Below: A Framework for Understanding, Evaluating, and Envisioning Refusal as Design. *ACM J. Responsib. Comput.*, 1(1).