

# Active Evaluation of General Agents: Problem Definition and Comparison of Baseline Algorithms

Marc Lanctot  
Google DeepMind  
Montreal, Canada  
lanctot@google.com

Ian Gemp  
Google DeepMind  
London, United Kingdom  
imgemp@google.com

Kate Larson  
Google DeepMind, University of Waterloo  
Waterloo, Canada  
katelarson@google.com

Michael Kaisers  
Google DeepMind  
Paris, France  
mkaisers@google.com

## ABSTRACT

As intelligent agents become more generally-capable, *i.e.*, able to master a wide variety of tasks, the complexity and cost of properly evaluating them rises significantly. Tasks that assess specific capabilities of the agents can be correlated and stochastic, requiring many samples for accurate comparisons, leading to added costs. In this paper, we propose a formal definition and a conceptual framework for *active evaluation* of agents across multiple tasks, which assesses the performance of ranking algorithms as a function of number of evaluation data samples. Rather than curating, filtering, or compressing existing data sets as a preprocessing step, we propose an online framing: on every iteration, the ranking algorithm chooses the task and agents to sample scores from. Then, evaluation algorithms report a ranking of agents on each iteration and their performance is assessed with respect to the ground truth ranking over time. Several baselines are compared under different experimental contexts, with synthetic generated data and simulated online access to real evaluation data from Atari game-playing agents. We find that the classical Elo rating system—while it suffers from well-known failure modes, in theory—is a consistently reliable choice for efficient reduction of ranking error in practice. A recently-proposed method, Soft Condorcet Optimization, shows comparable performance to Elo on synthetic data and significantly outperforms Elo on real Atari agent evaluation. When task variation from the ground truth is high, selecting tasks based on proportional representation leads to higher rate of ranking error reduction.

## KEYWORDS

general evaluation; multitask evaluation; ranking; active learning; game theory; social choice theory

### ACM Reference Format:

Marc Lanctot, Kate Larson, Ian Gemp, and Michael Kaisers. 2026. Active Evaluation of General Agents: Problem Definition and Comparison of Baseline Algorithms. In *Proc. of the 25th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2026)*, Paphos, Cyprus, May 25 – 29, 2026, IFAAMAS, 17 pages. <https://doi.org/10.65109/PUBN4692>



This work is licensed under a Creative Commons Attribution International 4.0 License.

*Proc. of the 25th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2026)*, C. Amato, L. Dennis, V. Mascardi, J. Thangarajah (eds.), May 25 – 29, 2026, Paphos, Cyprus. © 2026 International Foundation for Autonomous Agents and Multiagent Systems ([www.ifaamas.org](http://www.ifaamas.org)). <https://doi.org/10.65109/PUBN4692>

## 1 INTRODUCTION

Recent progress in AI research has led to intelligent agents which are increasingly more generally-capable. This has led to an increase in the number of multi-task benchmarks that assess different capabilities or axes of the agents. One classic multitask benchmark is the Arcade Learning Environment (ALE), which drove progress on general agent development, leading to the rise of deep reinforcement learning [7, 39]. The aim was not simply to assess agent performance on a single classic Atari video game, but across a suite of more than fifty different games, similar to general game-playing challenges [18, 41]. In recent years, large language models (LLMs) are being more widely adopted by the general public for a variety of different applications. LLMs are being evaluated for their many different capabilities: mathematical and programming ability, question-answering, general knowledge, abstract reasoning, etc.; however, reporting a complete evaluation of them on a public multi-task benchmark, Holistic Evaluation of Language Models (HELM), can cost thousands of dollars [30]. In addition, tasks are often correlated causing redundancy in evaluation data and wasted computational resources. Scores attributed to agents/models can also be noisy further increasing costs.

In addition to mounting costs, proper multi-task evaluation of general agent abilities is inherently challenging. There are many subtleties involved such as misconceptions in interpretation of results, overfitting to the domain (or evaluation context), and biases [12, 35, 48], statistical uncertainty [1], trade-offs between diversity and sensitivity [53], and non-transitivities as well as classical metrics that can be misleading or not robust [6, 27, 31]. While some of these issues are solvable with long-term commitments to better evaluation practices, some are possible to address with guarantees provided by evaluation methodologies rooted in social choice theory and game theory. For example, recent principled evaluation techniques have offer greater interpretability and robustness based on consistency and invariance properties that are inherited from their game-theoretic and social choice theoretic foundations [6, 27, 28, 31].

In this paper, we make the following contributions:

- We formally define the problem of *active evaluation* of general agents (*i.e.*, those evaluated on multi-task benchmarks); the goal is not only to rank agents but to do so *efficiently*.
- We define metrics that naturally balance between *identifying* the best agents and *ranking* them in the correct order.

- We extend recent offline algorithms for multitask evaluation (Voting-as-Evaluation [27], Nash averaging [6], and Metri-tocracy [43]) to the (online) active evaluation setting.
- We compare our extended algorithms and several baselines on synthetically-generated evaluation data and on an online simulation of Atari agent game-playing data.

Unlike traditional multitask evaluation where benchmarks and datasets are procured separately from evaluation, in active evaluation, algorithms are in full control of the data they get to see and choose the tasks and models to score *online*.

## 2 BACKGROUND AND TERMINOLOGY

In this section, we describe the basic terminology that will be used throughout the paper.

A multi-task evaluation problem for general agents consists of:

- A set of  $n$  **tasks** (or environments),  $V = \{v_1, v_2, \dots, v_n\}$ , and
- A set of  $m$  **agents** (or models),  $A = \{a_1, a_2, \dots, a_m\}$ .
- A mechanism for assessing the performance of any agent  $a_i \in A$  on task  $v_k \in V$ .

A common example in the field of deep reinforcement learning is the Arcade Learning Environment (ALE) [7]. In the ALE,  $n = 57$  and each task is one of the classic Atari 2600 games; the agents vary across papers but are commonly deep reinforcement learning algorithms, human agents, and/or reference agents such as random (for example,  $m = 8$  in [23]); the mechanism to assess the performance of an agent is its average in-game score on the Atari games after thousands of episodes (game plays). Similarly, multi-task benchmarks have been used for language models [30] and LLMs as agents [33] ( $n = 8, m = 27$ ).

The ultimate goal is to identify the “best” agent (or agents) or, most generally, to find a full ranking of agents. This is done by comparing their performance and aggregating them across tasks. A **ranking** of agents is a total order (permutation) over  $A$ . For example if  $m = 4$ , then one potential ranking,  $\succ$ , denoted

$$a_3 \succ a_0 \succ a_2 \succ a_1, \quad (1)$$

corresponds to  $a_3$  being the top-ranked (best) agent, followed by  $a_0$  (second best),  $a_2$  (third best), and  $a_1$  (worst). To quantify distances between rankings, such as the number of pairwise disagreements in order of agents, we use Kendall-tau distance.

*Definition 2.1.* Let  $A_1, A_2$  be finite sets of elements such that  $A_1 \subseteq A_2$ . Let  $\succ_1, \succ_2$  be rankings of agents in  $A_1$  and  $A_2$ , respectively. The **Kendall-tau distance** between two rankings is defined as

$$K_d(\succ_1, \succ_2) = \sum_{\{a_i, a_j\} \in A_1, i < j} \text{PAIRWISEORDER}_{i,j}(\succ_1, \succ_2), \quad (2)$$

where  $\text{PAIRWISEORDER}_{i,j}(\succ_1, \succ_2) = 0$  if the relative order of  $a_i$  and  $a_j$  agree across the rankings, *i.e.*, either  $(a_i \succ_1 a_j \text{ and } a_i \succ_2 a_j)$  or  $(a_j \succ_1 a_i \text{ and } a_j \succ_2 a_i)$ , or 1 otherwise (*i.e.*, they disagree).

Note that we use a slightly more general definition than the standard Kendall-tau distance (where  $A_2$  can be a larger set) to allow computing distances between partial rankings (over  $A_1 \subset A_2$ ) from the full ranking of agents  $A_2 = A$ . Since the maximum distance is  $\binom{|A_1|}{2} = \frac{|A_1|(|A_1|-1)}{2}$ ,  $K_d$  can be normalized to be in  $[0, 1]$ :

*Definition 2.2.* Let  $A_1, A_2$  be finite sets of elements such that  $A_1 \subseteq A_2$ . Let  $\succ_1, \succ_2$  be permutations over elements in  $A_1$  and  $A_2$ , respectively. The **normalized Kendall-tau distance**  $\succ_1$  and  $\succ_2$  is defined as

$$K_n(\pi_1, \pi_2) = \frac{2K_d(\succ_1, \succ_2)}{|A_1|(|A_1|-1)}. \quad (3)$$

We will use  $K_d$  and  $K_n$  ranking distances to the ground truth ranking as a measure of error.

### 2.1 Evaluation Systems for General Agents

Elo is a classic rating system that uses a simple logistic model learned from win/loss/draw outcomes [16]. A rating,  $\theta_i$ , is assigned to each player  $i$  such that the probability of player  $i$  beating player  $j$  is predicted as

$$\hat{p}_{ij} = \frac{1}{1 + 10^{(\theta_j - \theta_i)/400}}.$$

While Elo was designed specifically to rate players in the two-player zero-sum, perfect information game of Chess. Elo does not inherently model task information and its ratings are computed directly from head-to-head game outcomes. It is widely used including for the evaluation of LLMs in LMSYS Chatbot Arena [54].

Another way to evaluate general agents is to use social choice theory; two previous works are Vote N’Rank or Voting-as-Evaluation (VasE) [27, 44]. In VasE, the tasks are *votes* and there are no scores (necessarily); the main benefits are axiomatic properties and robustness via various forms of consistency based on the particular voting rule used. For example, when applying VasE to Atari game-playing, each task is an Atari game (57 tasks in total) and the average score of each agent in the game is sorted to get a complete vote (ranking over agents). Most recently, there have also been online mechanisms to compute rankings with various properties, such as Soft Condorcet Optimization (SCO) [28] and KemenyEl [19]. For example, SCO is based on gradient descent of a smooth Kendall-tau distance and the optimum of the SCO’s loss function is guaranteed to top-rank a Condorcet winner, if it exists.

Finally, there are game-theoretic methods that can also be used for general agent evaluation, such as Nash Averaging [6, 32, 38]. These methods derive ratings from equilibria of agents played in an evaluation meta-game; they offer interpretability and robustness via being invariant to **clones**: informally, a cloned agent is one that mimics the preferences of another agent, not changing any of the relative rankings between any other agents, across all tasks.

### 2.2 Active Learning

Active learning is a subfield of machine learning (also sometimes called *experimental design*) that addresses the challenge of expensive or time-consuming data labeling [40, Section 19.4]. Unlike traditional supervised learning, where a model passively trains on a large, fully labeled dataset, an active learning algorithm is “curious”: it starts by training on a small number of labeled examples and then intelligently queries a domain expert (or “oracle”) to label new data points that it identifies as being the most informative.

One approach to active learning is based on the information-theoretic principle of information gain that compares the entropy of the classifier on select data points to the expected entropy. This leads to the classical method maximum entropy sampling [47].

Maximum entropy sampling is part of a wider class of uncertainty sampling [46, 50] methods that sample preferentially on data points where the classifier is least certain. It was shown to be a good baseline approach for active logistic regression [51] which could be used as a basis for an active version of Elo.

### 2.3 Multi-armed Bandits

In the **stochastic multi-armed bandit** problem, an agent must choose one of several “arms” (or options) in each round, with the goal of maximizing the total reward over time. The core assumption is that the reward for pulling any given arm is drawn from a fixed, but unknown, probability distribution. The main challenge becomes the exploration-exploitation tradeoff: the agent must *explore* by trying different arms to learn their reward distributions, while also *exploiting* the arm that currently estimated best to maximize immediate gains. Performance is measured by regret: the cumulative difference between the agent’s reward and the reward it would have gotten by always picking the single best arm. A widely-used stochastic bandit algorithm is Upper Confidence Bounds (UCB) [3].

In the **adversarial bandit** (also called non-stochastic) setting [11], the problem becomes much harder as all statistical assumptions are dropped. Instead of rewards coming from a fixed distribution, they are chosen by an “adversary” in each round. This adversary can set the rewards for every arm arbitrarily, and in the strongest model, can even do so after observing the algorithm’s strategy. Regret is accumulated similarly, but comparing to the best single arm in hindsight. Two commonly used algorithms for adversarial bandits are Exp3 [4] and regret-matching [21].

There are well-known connections between two-player zero-sum games and adversarial bandits [8, 14, 55]: the average strategies of two adversarial bandits in self-play converge to a Nash equilibrium with high probability. This fact will be used to introduce online variants of one VasE algorithm and Nash averaging in Section 4.2.2.

## 3 ACTIVE EVALUATION

Active evaluation aims to minimize a metric of ranking efficiency. This section discusses several quality metrics for online ranking algorithms, and proposes Average Generalized Ranking Error (AGRE, definition 3.3) as the objective that measures efficiency.

At first the designer must choose an **experimental context**: a **data generator** and arguments for its parameters. The experimental context is the main component that ultimately determines the shape of the data. In this paper, we propose two models for synthetic data generation (see Section 3.2) and incremental access to existing data sets. A task, which may be stochastic, generates **scores**  $s(v, a) \sim S_{v,a}$ , where  $S_{v,a}$  is a random variable that depends on the task and the agent. Data takes the form of pairwise evaluations  $(s(v, a_i), s(v, a_j))$  between agents  $a_i, a_j \in A$  on a task  $v \in V$ .

The main experiment loop is summarized in Algorithm 1. There are  $T$  iterations. In each round  $t \in \{1, \dots, T\}$ , an algorithm chooses a task  $v^t$  and two agents  $a_i^t, a_j^t$  and receives data point  $s^t$ . The algorithm then uses  $s^t$  and outputs a new ranking  $\succ^t$ . A metric of performance  $f$  is used to assess the performance of the ranking  $\succ^t$ ; we describe several in the next subsection.

---

### Algorithm 1: Active Evaluation main loop

---

**Input:** A data generator  $\mathcal{D}$   
**Input:** A performance metric  $f$   
**Input:** An active evaluation algorithm  $\mathcal{A}$

- 1 **for** round  $t \in \{1, 2, \dots, T\}$  **do**
- 2      $(v^t, a_i^t, a_j^t) \sim \text{CHOOSETASKANDAGENTS}(\mathcal{A}, A, V, t)$
- 3     Sample scores  $s^t = (s(v^t, a_i^t), s(v^t, a_j^t)) \sim \mathcal{D}$
- 4     Update ranking of agents:  $\succ^t = \mathcal{A}(s^t, t)$
- 5     Report performance  $f(\succ^t)$  at round  $t$
- 6 **return** final performance  $f(\succ^T)$

---

### 3.1 Metrics

Many metrics can be used to assess the performance of an algorithm for active evaluation. Let us first assume the existence of a correct underlying **ground truth ranking**,  $\succ^*$  (which is available when generating data synthetically). In this case, the simplest metric might be to track the Kendall-tau distance to the ground truth over time,  $K_d(\succ^t, \succ^*)$ . To make the value less dependent on the number of agents, we could use normalized Kendall-tau distance,  $K_n$ .

However, it is common that the evaluator may only be interested in the top  $k$  agents (maybe first, second, and third, e.g.,  $k = 3$ ) or even more extremely only the top-performing agent ( $k = 1$ ). In these cases, the error metric should only take into account the performance of agents within these top ranks.

Denote by  $\succ_{[k]}$  the sub-ranking of  $\succ$  including only the top  $k$  agents, and  $A_{\succ_{[k]}}$  the set of agents among the top- $k$  ranked agents. For example, in the ranking depicted in equation (1),  $\succ_{[2]} = a_3 \succ a_0$  and  $A_{\succ_{[2]}} = \{a_0, a_3\}$ .

*Definition 3.1 (Top- $k$  Identification Error).* Let  $A$  be a set of agents,  $\succ^*$  be a ground truth ranking of agents in  $A$  with  $|A| = m$ , and  $k \in \{1, \dots, m\}$ . The **top- $k$  identification error** is the proportion of incorrectly identified top- $k$  agents:

$$\text{IDE}(A, \succ, \succ^*, k) = 1 - \frac{|A_{\succ_{[k]}} \cap A_{\succ^*_{[k]}}|}{k}.$$

Since the cardinality of the intersection varies from 0 to  $k$ ,  $0 \leq \text{IDE}(A, \succ, \succ^*, k) \leq 1$ . Denote the *sub-ranking*  $\succ_{k^*}$  as the relative order of the top- $k$  agents in  $\succ^*$  within  $\succ$ . Now we generalize the ranking error to include both ranking error and top- $k$  identification error:

*Definition 3.2 (Generalized Top- $k$  Ranking Error).* Let  $A$  be a set of agents and  $\succ^*$  be a ground truth ranking of agents in  $A$ , and  $k \in \{1, 2, \dots, m\}$ . The generalized top- $k$  ranking error is a combination of an error for identifying the set of top- $k$  agents and determining their correct rank  $\text{GRE}(A, \succ, \succ^*, k) =$


$$\alpha(k) \cdot \text{IDE}(A, \succ, \succ^*, k) + (1 - \alpha(k)) \cdot K_n(\succ_{k^*}, \succ^*),$$

where

$$\alpha(k) = \frac{m - k}{m - 1} \in [0, 1]$$

is the weight placed on the top- $k$  identification error.

Since  $\alpha(k) + (1 - \alpha(k)) = 1$ , and both  $\text{IDE}(A, \succ, \succ^*, k)$  and  $K_n$  are in  $[0, 1]$ , it follows that  $\text{GRE}(A, \succ, \succ^*, k) \in [0, 1]$ . When  $k = 1$ , all of the weight is placed on identification whereas when  $k = m$



**Figure 1: Example of synthetic data generation process with  $m = 4$  and  $n = 5$ . The task rankings  $(\succ_{v_1}^*, \succ_{v_2}^*, \dots, \succ_{v_5}^*)$  have Kendall-tau distances 4, 1, 4, 1, and 5, respectively, from  $\succ^*$ .**

all of the weight is placed on Kendall-tau distance, with linear weighting between the two error sources for other values of  $k$ . This linear weighting naturally places weight proportionally to favor identification error when  $k$  is small and Kendall-tau ranking error when  $k$  is large. Finally, we can also think of the *average* ranking error similarly to regret analysis of online learning [8, 14, 22]:

*Definition 3.3 (Average Generalized Top- $k$  Ranking Error).* Let  $A$  be a set of agents,  $\succ^*$  be a ground truth ranking of agents in  $A$ , and  $k \in \{1, 2, \dots, m\}$ . The *average generalized ranking error* (AGRE) of an algorithm that outputs rankings  $\succ^t$  on round  $t$  is:

$$\text{AGRE}(\succ^{1:T}, \succ^*, T) = \frac{\sum_{t=1}^T \text{GRE}(A, \succ^t, \succ^*, k)}{T}.$$

At any horizon  $T$ , the AGRE provides a metric of how efficiently the ranking algorithm identifies good rankings, as it cumulates all errors at horizons  $t < T$ .

### 3.2 Models for Synthetic Data

We now describe models for synthetic data generation. A key feature of using synthetic data is that the ground truth is known, hence the metrics described in Section 3.1 can be used directly.

*3.2.1 Mallows.* The Mallows model is a common way to model a number of noisy rankings from a central ranking,  $\succ^*$  over  $m$  alternatives [10, 36]. Given a *dispersion* parameter  $\phi \in [0, 1]$ , when  $\phi = 0$  only the central ranking is sampled; when  $\phi = 1$  all rankings are equally likely. The probability of sampling a ranking  $\succ$  is

$$\frac{1}{Z(\phi, m)} \phi^{K_d(\succ^*, \succ)}, \quad (4)$$

where  $Z(\phi, m)$  is a normalization constant.

A random permutation of  $A$  is generated to serve as to ground truth ranking  $\succ^*$ . We then use Equation 4 to generate **task rankings** over models,  $\succ_v^*$ , where each task ranks models. Intuitively, if  $\phi$  is close to 0, then we are more likely to draw task rankings which are close to  $\succ^*$  as measured by Kendall-tau distance. As  $\phi$  approaches 1, we are more likely to draw a more diverse set of task rankings since the effect of the Kendall-tau distance on the probability of any ranking being generated is smaller.

For each task ranking  $\succ_v^*$ ,  $m$  distributions of scores,  $S_{v,a}$  for  $a \in A$  mentioned in Section 3, are generated by sampling uniform random numbers in a bounded score interval and assigning the values as

the means of  $S_{v,a}$  ( $\mu_{v,a}$ ), such that they are sorted according to  $\succ_v^*$ . Each  $S_{v,a}$  is then normally-distributed according to  $\mathcal{N}(\mu_{s,a}, \sigma)$ . A summary of the process is depicted in Figure 1.

*3.2.2 Plackett-Luce.* We also try an alternative synthetic data model: Plackett-Luce. Results are consistent with those on Mallows; the description is similar, and is presented in Appendix C.3.

## 4 ALGORITHMS FOR ACTIVE EVALUATION

We now describe several algorithms which we will compare experimentally in the following section (summarized in Table 1). Algorithms for active evaluation are naturally decompose into three steps (with line references to Algorithm 1):

- (1) the **task selection** chooses a task  $v^t \in V$  (line 2),
- (2) the **agent selection** chooses two agents  $(a_i^t, a_j^t)$  (line 2), and
- (3) the process that updates tracked information based on the score sample  $s^t$  (line 4).

### 4.1 Baselines and Bandit Approaches

The most basic procedure for steps (1) and (2) is UNIFORMRANDOM which samples  $v \sim \text{UNIFORM}(V)$  uniformly at random and  $(a_i, a_j) \sim \text{UNIFORM}(A)$  without replacement. The simplest update rule tracks the cumulative average score  $\bar{S}_a$  aggregated across all tasks and outputs a ranking  $\succ^t$  by sorting  $\bar{S}_a$  for  $a \in A$ . Denote this simplest baseline algorithm UNIFORMAVERAGING.

A second baseline, BATCHELO is based on the Elo rating system: it also samples tasks and agents uniformly randomly, and maintains a historical record of win/loss outcomes where  $s(v^t, a_i^t) > s(v^t, a_j^t)$  counts as a win for  $a_i$  and loss for  $a_j$ . Then an efficient batch implementation of Elo [24] computes the best fit for ratings  $\theta$  and  $\succ^t$  is obtained by sorting the ratings in descending order. We also use ONLINEELO, the standard update typically used on competitive gaming platforms from single-game outcomes.

*4.1.1 Stochastic Bandit-based Approaches.* We suggest bandit-based approaches based on Upper Confidence Bounds (UCB) algorithm [3] as often used in Monte Carlo tree search [26].

BASICUCB is a single multi-armed bandit with one arm assigned to each agent (scores aggregated across tasks). Each agent  $a \in A$  is assigned a UCB score:

$$\text{UCBScore}(a, T) = \bar{s}_a + C\sqrt{\log(2T)/T_a},$$

where  $\bar{s}_a$  is the average score of agent  $a$ ,  $C$  is a constant set to  $\sqrt{2}\Delta_S$ ,  $2T$  is the total number of updates and,  $T_a$  is the number of updates to agent  $a$ . Tasks are selected uniformly at random. The final ranking is a sorting of visit counts per arm.

*4.1.2 Dueling Bandit Approaches.* When using a dueling bandit [52] approach feedback comes in the form of a preference between the two arms pulled. Dueling bandits have been used for ranking, for example Copeland dueling bandits [57]. We use a recent dueling bandit based on Kemeny ranking called KemenyEl [19, Algorithm 1]. Like UCB, KemenyEl uses statistics to estimate the noise between pairwise preferences, factoring it into its decision rule. Unlike UCB, the Kemeny rule is used for ranking.

We adapt KemenyEl to the online evaluation setting in the following way. The approximation accuracy of KemenyEl holds with

Name	Task Selection	Agent Selection	Update
UNIFORMAVERAGING	Uniform Random	Uniform Random	Track cumulative average $\bar{s}_a$ per agent (aggregated across tasks)
BASICUCB	Uniform Random	Agent-level UCB	Update means and visit counts for agent-level bandit
KEMENYEL	Uniform Random	Uniform Random	Update $Q$ matrix and statistics for pair of agents selected
BATCHELO	Uniform Random	Uniform Random	Track cumulative wins, losses and draws; rank by Elo ratings
BATCHCOPELAND	Uniform Random	Uniform Random	Growing-batch Copeland voting [15]
BATCHRANKEDPAIRS	Uniform Random	Uniform Random	Growing-batch Ranked Pairs voting [49]
BATCHMAXIMALLOTTERIES	Uniform Random	Agent Strategy	Growing-batch Iterative Maximal Lotteries voting [27]
BATCHSCO	Uniform Random	Uniform Random	Growing-batch version of SCO [28]
BATCHNASHAVERAGING	Task Strategy	Agent Strategy	Growing-batch version of Nash Averaging [6]
MEANMODELCOPELAND	Uniform Random	Uniform Random	Estimate task rankings via $\bar{s}_{v,a}$ and return VasE(Copeland)
MEANMODELRANKEDPAIRS	Uniform Random	Uniform Random	Estimate task rankings via $\bar{s}_{v,a}$ and return VasE(Ranked Pairs)
MEANMODELMAXLOT	Uniform Random	Uniform Random	Estimate task rankings via $\bar{s}_{v,a}$ and return VasE(Iterative Max Lot)
PROP REPRESENTATION	Greedy Prop. Rep.	Uniform Random	Online version of Greedy Metritocracy algorithm [43]
ONLINEELO	Uniform Random	Uniform Random	Track ratings only; update ratings from each game outcome
ONLINESCO	Uniform Random	Uniform Random	Track and update ratings only (analogously to online Elo) [28]
ONLINEMAXIMALLOTTERIES	Uniform Random	Agent Strategy	Fully-online Maximal Lotteries (based on adversarial bandits) [17]
ONLINENASHAVERAGING	Task Strategy	Agent Strategy	Fully-online Nash Averaging (based on adversarial bandits) [6]

**Table 1: A summary of algorithms for active evaluation. For methods based on zero-sum games, the Task Strategy and Agent Strategy are sampling distributions in the underlying game, i.e., with uniform exploration mixed in: e.g.,  $\sigma'(a) = \frac{\epsilon}{|A|} + (1 - \epsilon)\sigma$ .**

probability  $1 - \delta$ , where  $\delta$  is chosen by the user and determines the number of samples per pairwise matchup. In the online setting, we first choose an initial  $\delta_0 = 0.1$ ; once the number of samples per matchup is reached (i.e., one matchup per round of Alg. 1), a new epoch is started ( $k \leftarrow k + 1$ ) and we set  $\delta_k = \frac{\delta_{k-1}}{2}$ . Similarly, we set the acceptable Kendall-tau distance error,  $\rho_0$ , to be half of the total Kendall-tau distance and then halve it every epoch.

## 4.2 Growing Batch and Online Algorithms

Any algorithm that is normally run offline can be run online by simply collecting evaluation data over time and running the algorithm on the data set on iteration  $t$ . We call this the **growing-batch** version of an offline algorithm. VasE methods interpret evaluations as data sets similarly to how Elo does: for example, given two score samples for task  $v$  if,  $s(v^t, a_i^t) > s(v^t, a_j^t)$ , this counts as a vote  $a_i \succ a_j$ . Given the pairwise preference collection, we examine several Condorcet methods: BATCHCOPELAND, BATCHRANKEDPAIRS, and BATCHMAXIMALLOTTERIES corresponding to growing-batch versions of the Copeland, Ranked Pairs, and Maximal Lotteries voting schemes [15, 17, 27, 49]. We use an iterative variant of Maximal Lotteries [27] which provides a ranking over agents.

We also use a batch version of Soft Condorcet Optimization [28], which is a gradient descent algorithm run on a smooth Kendall-tau loss and its fully-online version which is analogous to Elo’s standard online update. We introduce a new fully-online version of Maximal Lotteries that we describe separately in Section 4.2.2.

A different way to apply VasE algorithms online is to use a **mean model**. Mean-model VasE algorithms maintain  $nm$  empirical averages of each score  $\bar{s}_{v,a}$ . Estimated per-task agent rankings  $\hat{\succ}_v$  are obtained by sorting  $\bar{s}_{v,a}$  for  $a \in A$ ; then, a voting rule returns a ranking from the  $n$  ranked-ballot votes. PROPORTIONALREPRESENTATION is a mean-model adaptation of the recently-proposed Greedy Metritocracy algorithm [43, Algorithm 1]. In particular, it

uses estimated task rankings  $\hat{\succ}_v$  as votes to determine an *reduced task set* based on proportional representation, recomputing this subset on each iteration based on updated estimates. The group size parameter ( $g = \lceil g_f \cdot n \rceil$ ) captures the trade-off between the granularity of representation and the size of subset needed to satisfy positional representation. Tasks are selected uniformly at random from this reduced task set, with a small probability (0.05) of choosing uniformly form the entire set of tasks to ensure exploration. The reported ranking is then obtained by running Ranked Pairs votes over the votes from estimated task rankings.

All of the mean-model and growing-batch algorithm use a **burn-in** phase: for the first  $nm$  iterations, the task and agent sampling one cycle through a uniform shuffling of the unique  $(v, a) \in V \times A$  pairs. This burn-in phase has been shown to benefit stochastic gradient descent algorithms [45]. Data is still collected and rankings reported in the same way during the burn-in period.

**4.2.1 Growing Batch Nash Averaging.** One way to define ratings of players is to analyze the one-shot meta-game played between agents and tasks [6]. This matrix game is a two-player zero-sum game where the row player chooses an agent  $a \in A$  and the column player chooses a task  $v \in V$ . The utility to the row player for the joint choice  $(a, v)$  is the score  $s_{v,a}$  and the utility to the column player is  $-s_{v,a}$ . In the original paper, these utilities were averaged over many samples, so utilities took the form e.g.,  $(\bar{s}_{v,a}, -\bar{s}_{v,a})$ . The rating of agent  $a$ ,  $\theta_a$ , is defined as the expected utility of strategy  $a$  versus the task player’s Nash equilibrium strategy, which can be computed efficiently (polynomial time) since the game is two-player zero-sum. A growing-batch version employs Nash averaging on the batch of data collected so far: utilities are averages, and it computes ratings as in the original Nash averaging. A ranking is then obtained by sorting the ratings.

Nash averaging can be also run fully online by using adversarial bandits. We elaborate on this in the next section.

**4.2.2 Online Max. Lotteries and Nash Averaging.** Similarly to Nash averaging, the Maximal Lotteries voting method also involves solving a two-player zero-sum “margin game” and using its solution to define the set of winners. The game is an two-player zero-sum agent-vs-agent game: the utility to the row player for strategy  $a_i$  versus the column player’s choice  $a_j$  is the voter margin  $\delta(a_i, a_j) = N(a_i, a_j) - N(a_j, a_i)$ , where  $N(a_i, a_j)$  is the number of votes where  $a_i \succ a_j$ . As above, the growing-batch version of Maximal Lotteries derives this game from the growing data set; like Nash averaging, it can also be run fully online.

We defer the details of solving two-player zero-sum games using adversarial bandits to Appendix A. We base our implementations on regret-matching [21]. In Nash averaging, when agent  $i$  is sampled against task  $v^t$  the utility is  $s(v^t, a_i)$  (obtained in Line 3 of Algorithm 1). The solution to the game (approximate Nash equilibrium) is extracted as a pair of distributions (stochastic strategies)  $\sigma = (\sigma_1, \sigma_2)$  for players 1 and 2. In the case of Nash averaging,  $\sigma_1 \in \Delta(A)$  and  $\sigma_2 \in \Delta(V)$ . In the case of Maximal Lotteries,  $\sigma_1, \sigma_2 \in \Delta(A)$  where  $\Delta(A)$  represents a simplex over elements in  $A$ . The utility for the row player choosing  $a_i$  playing against column player choosing  $a_j$  when task  $v^t$  was sampled (uniformly) is:  $\hat{u}_i(v^t, a_i, a_j) = 2 \cdot \mathbb{I}[s(v^t, a_i) > s(v^t, a_j)] - 1$ , where  $\mathbb{I}$  is the indicator function. In both cases (Nash averaging and Maximal Lotteries), to extract a ranking, we define the ratings as expected values against the equilibrium strategy; e.g., for  $a \in A$ :  $\text{RATING}(a) = \mathbb{E}_{b \sim \sigma_2} u_1(a, b)$ , where  $u_1(a, b)$  is the average utility to player 1 given choices  $a$  and  $b$  as described in the games above. The rankings are then extracted by sorting the ratings. This naturally extends Maximal Lotteries’ distribution over winner agents to a ranking among agents [37].

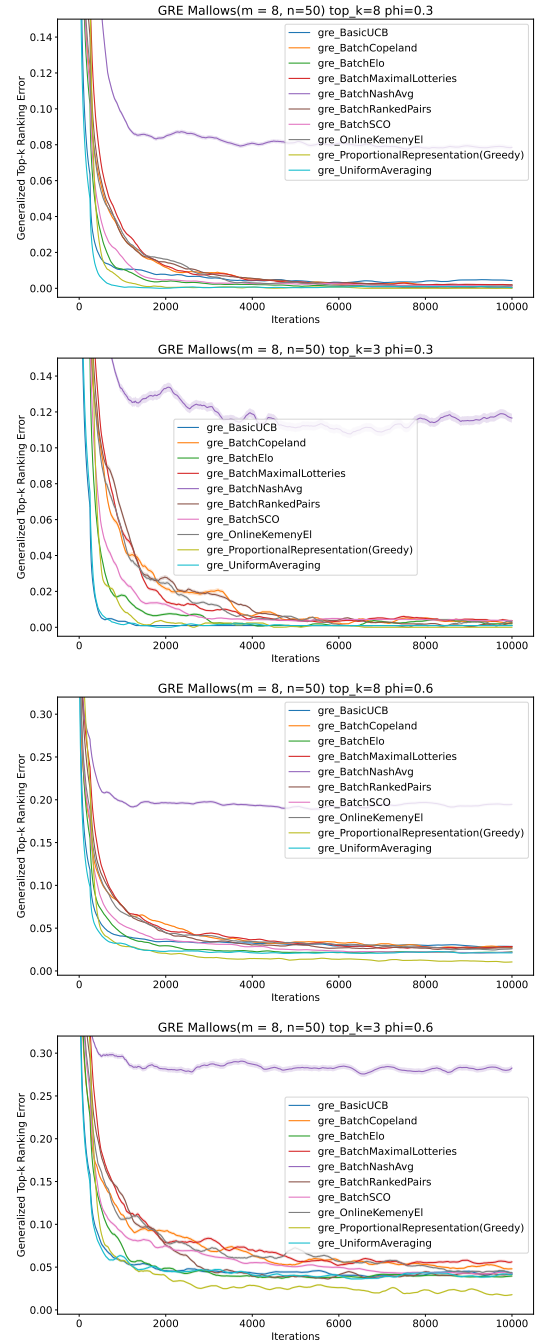
**4.3 Summary of Known Convergence Results**

Since the game online Nash Averaging plays is a sample-based version of the true underlying game, the average strategy used by adversarial bandits is guaranteed to converge to an approximate equilibrium with high probability, and the error decreases as  $\epsilon_t = O(1/\sqrt{t})$  where  $t$  is the number of iterations [4], and ratings will differ from the true ratings by at most  $\epsilon_t$ . Similarly, the solution computed by Batch Maximal Lotteries is obtained by solving a two-player zero-sum margin game derived directly from a preference profile. The error of the estimated expected utility of a cell in the margin game (defined in Section 4.2.2) decreases as  $O(1/\sqrt{t})$  and hence the Maximal Lottery solution we derive from it will be a  $O(1/\sqrt{t})$ -approximate solution.

Asymptotic guarantees on GRE requires knowing the ground truth and, possibly, also the shape of the data both of which are generally not known. When data is shaped as described in Sec 3.2.1, the mean model variants’ estimated means converge to the true means of the underlying score distributions at a rate of  $1/\sqrt{n}$  by the Monte Carlo square root law; consequently, the estimated task rankings will converge to the true task rankings and the overall ranking to a VasE-aggregated one from the task rankings.

**5 EXPERIMENTAL RESULTS**

We now compare the performance of each active evaluation algorithm, first on synthetic data then on real Atari agent data.



**Figure 2: Exp Condition 1: Generalized top-k ranking error (GRE) over iterations for several algorithms. Each line is a mean over the same 100 seeds with 95% confidence intervals in shaded area. Each data point is a sliding-window average of the most recent 250 values of GRE for the reported ranking  $\succ^t$  at iteration  $t$ . The top two graphs use dispersion  $\phi = 0.3$ , with  $k \in \{3, 8\}$ , where  $k = 8$  corresponds to normalized Kendall-tau distance. The bottom two use  $\phi = 0.6$  and vary  $k$  similarly.**

### 5.1 Experiments with Synthetic Data

For experiments using synthetic data using models described in Section 3.2. For score distributions, means are drawn from the interval  $\mu_{v,a} \in [0, 100]$  and use standard deviation of  $\sigma = 20$ .

**5.1.1 Experimental Condition 1: Ranking Error Reduction Rate.** A primary motivation of research on active evaluation is designing algorithms that use/sample data efficiently, so a main research question is how well algorithms can report a ranking that is close to the ground truth as a function of number of samples. We use a moderately sized problem with  $m = 8$  agents and  $n = 50$  tasks; each seed determines a new ground truth ranking, task rankings, and score distributions. Since the fully-online variants of algorithms are approximating the growing-batch versions, we restrict this first experiment to the growing batch algorithms with basic baselines and bandit approaches. Figure 2 shows the generalized top- $k$  ranking error (GRE, definition 3.2) as a function if iterations of Algorithm 1.


Our first observation is that, when  $k = 3$ , it is possible to drive the error to 0 within 2000 iterations and several algorithms achieve this. Our second observation, UNIFORMAVERAGING– the most basic of baselines– performs particularly well in all settings, and in three of four cases is reduces the error fastest in the first 1000 iterations, though BASICUCB has comparable performance when  $\phi = 0.6$ . In general, UNIFORMAVERAGING, BASICUCB, BATCHELO, BATCHSCO, and PROPORTIONALREPRESENTATION reduce the error most quickly at the beginning. In Nash averaging, we observe a phenomenon similar to “adversarial task selection” reported in [27]: since it is a zero-sum game of agents versus tasks, the task player prefers to choose tasks where the agents’ utilities are minimized, resulting in an equilibrium distribution skewed towards task rankings that may be arbitrarily far from the ground truth. Hence, the ranking overfits to these adversarially-chosen tasks. When  $\phi = 0.6$ , PROPORTIONALREPRESENTATION is a clear winner, deviating significantly from the other methods. We try group size proportions  $g_f \in \{0.1, 0.2, 0.25, 0.3, 0.35, 0.5, 0.6, 0.75\}$ , finding that  $g_f = 0.1$  to be the best value and use it by default. We show trade-offs between group sizes, reductions in task sets, and GRE in Appendix C.1.

In Table 2, we report on the longer-term performance in the form of average GRE at  $t = 10000$  for  $\phi = 0.3$  and  $k \in \{3, 8\}$ , which can be interpreted as an approximate average area-under-the-GRE-curve. We observe that BASICUCB and UNIFORMAVERAGING perform best when  $k = 3$  and UNIFORMAVERAGING performs best when  $k = 8$ . In both cases, BATCHELO also reduces the rank error at a fast rate.

**5.1.2 Experimental Condition 2: Batch versus Online.** Firstly, we see that there are several methods that reach 0 error in less than 10000 iterations: BATCHSCO, BATCHELO, ONLINEELO, PROPORTIONALREPRESENTATION, and all of the mean-model variants. Second, while the batch version of Maximal Lotteries reduces rank error faster than its online variant, in Nash averaging this is reversed, possibly due to the average strategy in the adversarial bandit changing more smoothly. Third, mean-model variants and PROPORTIONALREPRESENTATION reduce error rate faster than the batch versions of these algorithms, several reaching 0 error in less than 6000 iterations (MEANMODELMAXIMALLOTTERIES and MEANMODEL RANKEDPAIRS).

Algorithm	AGRE( $k, \phi$ )		
	(3, 0.3)	(8, 0.3)	(3, 0.6)
BASICUCB	<b>0.003</b>	0.007	0.046
UNIFORMAVERAGING	0.003	<b>0.002</b>	0.049
PROPORTIONALREPRESENTATION	0.009	0.005	<b>0.037</b>
BATCHELO	0.009	0.006	0.050
BATCHSCO	0.012	0.007	0.063
KEMENYEL	0.017	0.010	0.072
BATCHCOPELAND	0.018	0.010	0.072
BATCHMAXLOT	0.019	0.013	0.079
BATCHRANKEDPAIRS	0.020	0.012	0.063
BATCHNASHAVG	0.124	0.088	0.285

**Table 2: Exp. Condition 1: Average Generalized Top- $k$  ranking error,  $\text{AGRE}(\succ^{1:t}, \succ^*, t, k = \cdot)$ , at  $t = 10^4$  and various values of  $k$  and  $\phi$ . All values have 95% confidence intervals  $< 10^{-3}$ .**



**Figure 3: Exp Condition 2: Generalized top- $k$  ranking error (GRE) over iterations for several algorithms. Each data point is a sliding-window average of the most recent 250 values of GRE for the reported ranking  $\succ^t$  at iteration  $t$ . Both graphs use  $k = 3$ . Please note the log scale for both axes.**

### 5.2 Experiments with Real Data

In this section, we simulate the online sampling of evaluation data by presenting these data sets as data generators.

**5.2.1 Atari.** We use previous results of agents playing all 57 classic Atari games via the Arcade Learning Environment (ALE) [7]. In particular, we use the data from the “Agent57 paper” table in [5, Section H.4] which comprises  $m = 8$  agents and  $n = 57$  tasks (Atari games). The agents are varied among uniform random, human expert, and

Algorithm	AGRE( $k = 3$ )	AGRE( $k = 8$ )
BATCHSCO	<b>0.019031</b>	<b>0.008095</b>
ONLINESCO	0.025931	0.009034
KEMENYEL	0.027322	0.011631
BATCHRANKEDPAIRS	0.027873	0.012952
BATCHELO	0.032967	0.010907
BATCHCOPELAND	0.037158	0.013470
ONLINEELO	0.037259	0.012504
MEANMODELCOPELAND	0.042865	0.013617
BATCHMAXIMALLOTTERIES	0.043418	0.021718
MEANMODELMAXLOT	0.045105	0.013727
PROPORTIONALREPRESENTATION	0.045472	0.013558
MEANMODELRANKEDPAIRS	0.048279	0.014740
ONLINEMAXLOT	0.135168	0.057263
ONLINENASHAVG	0.240771	0.105709
BASICUCB	0.241976	0.074383
UNIFORMAVERAGING	0.242354	0.259529

**Table 3: Atari: Average Generalized Top- $k$  ranking error,  $\text{AGRE}(\succ^{1:t}, \succ^*, t, k = \cdot)$  at  $t = 10^4$  on simulated Atari data generator. 95% confidence intervals for each entry was  $< 10^{-3}$ .**

several deep reinforcement learning algorithms. In particular, unlike much of the previous work on deep reinforcement learning in ALE, this table reports both mean scores  $\mu_{v,a}$  and standard deviations  $\sigma_{v,a}$ . Hence, we simulate a data generator by setting the score distributions  $S_{v,a} = \mathcal{N}(\mu_{v,a}, \sigma_{v,a})$  where  $\mathcal{N}$  is a normal distribution, and otherwise identically to the synthetic data generators. Since the scale of rewards is different in each game, we also linearly normalize the scores per game from range  $[\min_{a \in A} \mu_{v,a}, \max_{a \in A} \mu_{v,a}]$  to within range  $[0, 100]$ . This score normalization does not affect the task rankings, is required by Nash averaging [6] and averaging methods, and helps maintain a consistent interval for exploration constant in bandit methods. In the absence of a proper ground truth we compute the ranking that minimizes the average Kendall-tau distance to the data:  $\succ^* = \min_{\succ'} \frac{1}{57} \sum_{v \in V} K_d(\succ', \succ_v)$ ; this ranking is a maximum likelihood estimator of the ground truth under a Mallows model [10, Chapter 8]. We also show in Appendix B.1 that it also recovers the ground truth ranking on synthetic data.

The average generalized top- $k$  ranking error is summarized in Table 3 (for full graphs of GRE over iterations, see Figure 9 in Appendix C). In this setting, BATCHSCO performed best followed by ONLINESCO and five of the top six were VasE methods. The order of the simple baselines BASICUCB and UNIFORMAVERAGING is completely reversed from the synthetic data setting: they perform worst in Atari; this indicates that in some domains one may need to be careful when aggregating scores across tasks. Nash averaging also performs poorly, which could also be due to the adversarially-chosen task distribution. We do note that there are  $\frac{15}{57}$  tasks whose rankings different significantly from the ground truth, *i.e.*, where  $K_d(\succ_v^*, \succ^*) \geq 8$ ; skewing a task distribution toward any of these high-distance tasks could significantly hinder the ability to find the ground truth ranking. We compare the task variation from the ground truth to samples from the synthetic data in Appendix B.

## 6 RELATED WORK

Most of the related previous work discussed in this paper (such as VasE and game-theoretic evaluation) are based on the standard offline case, *i.e.*, run on a batch of data. However, given the immense costs of evaluating language models, there are some recent works studying active evaluation for language models.

Li et al. [29] explicitly models dependencies in the test examples allowing predictions for remaining examples, and uses reinforcement learning to learn a policy to select subsets. Their method adaptively selects different subsets of prompts per model based on each model’s strengths and weaknesses. Angelopoulos et al. [2] propose a cost-optimal framework for how to balance obtaining evaluation data from strong-but-costly raters versus weak-but-cheap raters based on prediction-powered inference. Costs per evaluation are explicitly assumed and a main result is an optimal policy to actively sample prompts. In contrast, our conceptual framing is not specifically intended for (however, compatible with) language model evaluation and focuses on task and model selection rather than LM prompt subselection and trade-offs in rater quality.

There is work in computational social choice that examines the sample complexity required for noisy votes to reveal the truth, under what conditions, and algorithms which are guaranteed to find the ground truth in the limit of infinite samples [13].

## 7 CONCLUSION AND FUTURE WORK

In this paper, we formally described the problem of active evaluation for general (multi-task) agents and proposed data generation techniques. We found that across synthetic data and real data experiments, BATCHELO is a consistently reliable choice for obtaining low ranking error across samples, performing especially well in the low task variation ( $\phi = 0.3$ ) regime. However, BATCHSCO and ONLINESCO are comparable and come in as a close second in synthetic data, while outperforming BATCHELO two-fold in Atari agent evaluation. In the synthetic data experiments, BATCHELO was outperformed by UNIFORMAVERAGING and BASICUCB but they both performed poorly on the Atari experiment, and Nash averaging also performing poorly; evaluators may need to exercise caution when aggregating or comparing score values across tasks. BATCHELO was significantly outperformed by BATCHSCO, ONLINESCO, and BATCHRANKEDPAIRS, with SCO being the clearly best method on the Atari experiment. In Atari (unlike in the synthetic data), batch variants seemed to outperform their fully-online versions; however, online and mean-model variants are still viable choices that offer lower computational complexity.

Future work could derive formal properties and statistical approximation guarantees as a function of the number of rounds. Also, only a few of the algorithms we compared adaptively modified the task selection. We hope that future work investigates the benefit of adaptive sampling strategies for data collection that use confidence intervals, uncertainty sampling, expected model change, error reduction, or task/agent score correlations [19, 46, 51].

## ACKNOWLEDGMENTS

We would like to thank Manon Revel for the helpful feedback on a previous draft of the paper and Serena Wang for help with the Metritocracy (proportional representation) algorithm.

## REFERENCES

- [1] Rishabh Agarwal, Max Schwarzer, Pablo Samuel Castro, Aaron C. Courville, and Marc G. Bellemare. 2021. Deep Reinforcement Learning at the Edge of the Statistical Precipice. *CoRR abs/2108.13264* (2021). arXiv:2108.13264 <https://arxiv.org/abs/2108.13264>
- [2] Anastasios N. Angelopoulos, Jacob Eisenstein, Jonathan Berant, Alekh Agarwal, and Adam Fisch. 2025. Cost-Optimal Active AI Model Evaluation. arXiv:2506.07949 [cs.LG] <https://arxiv.org/abs/2506.07949>
- [3] P. Auer, N. Cesa-Bianchi, and P. Fischer. 2002. Finite-time Analysis of the Multi-armed Bandit Problem. *Machine Learning* 47 (2002), 235–256.
- [4] Peter Auer, Nicolò Cesa-Bianchi, and Robert Schapire. 1998. Gambling in a rigged casino: The adversarial multi-armed bandit problem. *Foundations of Computer Science, 1975., 16th Annual Symposium on* (07 1998). <https://doi.org/10.1109/SFCS.1995.492488>
- [5] Adrià Puigdomènech Badia, Bilal Piot, Steven Kapturovski, Pablo Sprechmann, Alex Vitvitskiy, Daniel Guo, and Charles Blundell. 2020. Agent57: Outperforming the Atari Human Benchmark. arXiv:2003.13350 [cs.LG] <https://arxiv.org/abs/2003.13350>
- [6] David Balduzzi, Karl Tuyls, Julien Perolat, and Thore Graepel. 2018. Re-evaluating evaluation. In *Advances in Neural Information Processing Systems*, S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett (Eds.), Vol. 31. Curran Associates, Inc.
- [7] M. G. Bellemare, Y. Naddaf, J. Veness, and M. Bowling. 2013. The Arcade Learning Environment: An Evaluation Platform for General Agents. *Journal of Artificial Intelligence Research* 47 (jun 2013), 253–279.
- [8] A. Blum and Y. Mansour. 2007. *Learning, Regret minimization, and Equilibria*. Cambridge University Press, Chapter 4.
- [9] Branislav Bošanský, Viliam Lisý, Marc Lanctot, Jiří Čermák, and Mark H.M. Winands. 2016. Algorithms for computing strategies in two-player simultaneous move games. *Artificial Intelligence* 237 (2016), 1–40.
- [10] Felix Brandt, Vincent Conitzer, Ulle Endriss, Jérôme Lang, and Ariel D. Procaccia (Eds.). 2016. *Handbook of Computational Social Choice*. Cambridge University Press.
- [11] Sébastien Bubeck and Nicolò Cesa-Bianchi. 2012. Regret Analysis of Stochastic and Nonstochastic Multi-armed Bandit Problems. *Foundations and Trends in Machine Learning* 5, 1 (2012), 1–122.
- [12] Ryan Burnell, Wout Schellaert, John Burden, Tomer D Ullman, Fernando Martinez-Plumed, Joshua B Tenenbaum, Danaja Rutar, Lucy G Cheke, Jascha Sohl-Dickstein, Melanie Mitchell, et al. 2023. Rethink reporting of evaluation results in AI. *Science* 380, 6641 (2023), 136–138.
- [13] Ioannis Caragiannis, Ariel D. Procaccia, and Nisarg Shah. 2016. When Do Noisy Votes Reveal the Truth? *ACM Trans. Econ. Comput.* 4, 3, Article 15 (March 2016), 30 pages. <https://doi.org/10.1145/2892565>
- [14] Nicolò Cesa-Bianchi and Gabor Lugosi. 2006. *Prediction, Learning, and Games*. Cambridge University Press.
- [15] Arthur Herbert Copeland. 1951. A ‘reasonable’ social welfare function. Seminar on Mathematics in Social Sciences, University of Michigan (unpublished).
- [16] Arpad E. Elo. 1978. *The Ratings of Chess Players, Past and Present* (2nd ed.). Arco Publishing, Inc.
- [17] P. C. Fishburn. 1984. Probabilistic social choice based on simple voting comparisons. *Review of Economic Studies* 51, 4 (1984), 683–692.
- [18] Michael Genesereth, Nathaniel Love, and Barney Pell. 2005. General Game Playing: Overview of the AAAI Competition. *AI Magazine* 26, 2 (June 2005), 62.
- [19] Anne-Marie George and Christos Dimitrakakis. 2024. Eliciting Kemeny Rankings. *Proceedings of the AAAI Conference on Artificial Intelligence* 38, 11 (Mar. 2024), 12164–12171. <https://doi.org/10.1609/aaai.v38i11.29105>
- [20] Jonathan Gray, Adam Lerer, Anton Bakhtin, and Noam Brown. 2020. Human-Level Performance in No-Press Diplomacy via Equilibrium Search. *CoRR abs/2010.02923* (2020). arXiv:2010.02923 <https://arxiv.org/abs/2010.02923>
- [21] Sergiu Hart and Andreu Mas-Colell. 2001. *A Reinforcement Procedure Leading to Correlated Equilibrium*. Springer Berlin Heidelberg, Berlin, Heidelberg, 181–200. [https://doi.org/10.1007/978-3-662-04623-4\\_12](https://doi.org/10.1007/978-3-662-04623-4_12)
- [22] Elad Hazan. 2015. *Introduction to online convex optimization*. Now Publishers Inc. Foundations and Trends in Optimization, Volume 2, Issue 3-4.
- [23] Matteo Hessel, Joseph Modayil, Hado van Hasselt, Tom Schaul, Georg Ostrovski, Will Dabney, Daniel Horgan, Bilal Piot, Mohammad Gheshlaghi Azar, and David Silver. 2017. Rainbow: Combining Improvements in Deep Reinforcement Learning. *CoRR abs/1710.02298* (2017). arXiv:1710.02298 <http://arxiv.org/abs/1710.02298>
- [24] David R. Hunter. 2004. MM Algorithms for Generalized Bradley-Terry Models. *The Annals of Statistics* 32, 1 (2004), 384–406.
- [25] J. Kemeny. 1959. Mathematics Without Numbers. *Daedalus* 88, 4 (1959), 577–591.
- [26] Levente Kocsis and Csaba Szepesvári. 2006. Bandit based monte-carlo planning. In *Proceedings of the 17th European Conference on Machine Learning* (Berlin, Germany) (ECML’06). Springer-Verlag, Berlin, Heidelberg, 282–293. [https://doi.org/10.1007/11871842\\_29](https://doi.org/10.1007/11871842_29)
- [27] Marc Lanctot, Kate Larson, Yoram Bachrach, Luke Marris, Zun Li, Avishkar Bhoopchand, Thomas Anthony, Brian Tanner, and Anna Koop. 2023. Evaluating Agents using Social Choice Theory. arXiv:2312.03121 [cs.AI]
- [28] Marc Lanctot, Kate Larson, Michael Kaisers, Quentin Berthet, Ian Gemp, Manfred Diaz, Roberto-Rafael Maura-Rivero, Yoram Bachrach, Anna Koop, and Doina Precup. 2025. Soft Condorcet Optimization for Ranking of General Agents. In *Proceedings of the 24th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2025)*. <https://arxiv.org/abs/2411.00119>
- [29] Yang Li, Jie Ma, Miguel Ballesteros, Yassine Benajiba, and Graham Horwood. 2025. Active Evaluation Acquisition for Efficient LLM Benchmarking. In *Proceedings of the Forty-Second International Conference on Machine Learning (ICML)*.
- [30] Percy Liang, Rishi Bommasani, Tony Lee, Dimitris Tsipras, Dilara Soylu, Michihiro Yasunaga, Yian Zhang, Deepak Narayanan, Yuhuai Wu, Ananya Kumar, Benjamin Newman, Binhang Yuan, Bobby Yan, Ce Zhang, Christian Cosgrove, Christopher D. Manning, Christopher Ré, Diana Acosta-Navas, Drew A. Hudson, Eric Zelikman, Esin Durmus, Faisal Ladhak, Frieda Rong, Hongyu Ren, Huaxiu Yao, Jue Wang, Keshav Santhanam, Laurel Orr, Lucia Zheng, Mengyue Kongul, Mirac Suzgun, Nathan Kim, Neel Guha, Niladri Chatterji, Omar Khattab, Peter Hendersong, Qian Huang, Ryan Chi, Sang Michael Xie, Shibani Santurkar, Surya Ganguli, Tatsunori Hashimoto, Thomas Icard, Tianyi Zhang, Vishrav Chaudhary, William Wang, Xuechen Li, Yifan Mai, Yuhui Zhang, and Yuta Koreeda. 2022. Holistic Evaluation of Language Models. *CoRR abs/2211.09110* (2022). arXiv:2211.09110 [cs.CL] See latest results at: <https://crfm.stanford.edu/helm/latest/>.
- [31] Siqi Liu, Ian Gemp, Luke Marris, Georgios Piliouras, Nicolas Heess, and Marc Lanctot. 2025. Re-evaluating Open-ended Evaluation of Large Language Models. arXiv:2502.20170 [cs.GT] <https://arxiv.org/abs/2502.20170>
- [32] Siqi Liu, Ian Gemp, Luke Marris, Georgios Piliouras, Nicolas Heess, and Marc Lanctot. 2025. Re-evaluating Open-ended Evaluation of Large Language Models. In *The Thirteenth International Conference on Learning Representations*. arXiv:2502.20170 [cs.GT] <https://arxiv.org/abs/2502.20170>
- [33] Xiao Liu, Hao Yu, Hanchen Zhang, Yifan Xu, Xuanyu Lei, Hanyu Lai, Yu Gu, Hangliang Ding, Kaiwen Men, Kejuan Yang, Shudan Zhang, Xiang Deng, Aohan Zeng, Zhengxiao Du, Chenhui Zhang, Sheng Shen, Tianjun Zhang, Yu Su, Huan Sun, Minlie Huang, Yuxiao Dong, and Jie Tang. 2023. AgentBench: Evaluating LLMs as Agents. arXiv:2308.03688 [cs.AI]
- [34] R. D Luce. 1959. *Individual Choice Behavior: A Theoretical Analysis*. Wiley.
- [35] Marlos C. Machado, Marc G. Bellemare, Erik Talvitie, Joel Veness, Matthew J. Hausknecht, and Michael Bowling. 2018. Revisiting the Arcade Learning Environment: Evaluation Protocols and Open Problems for General Agents. *Journal of Artificial Intelligence Research* 61 (2018), 523–562.
- [36] C. L. Mallows. 1957. Non-Null Ranking Models. *Biometrika* 44 (1957), 114–130.
- [37] Luke Marris, Marc Lanctot, Ian Gemp, Shayegan Omidshafiei, Stephen McAleer, Jerome Connor, Karl Tuyls, and Thore Graepel. 2022. Game Theoretic Rating in N-player general-sum games with Equilibria. arXiv:2210.02205 [cs.GT] <https://arxiv.org/abs/2210.02205>
- [38] Luke Marris, Siqi Liu, Ian Gemp, Georgios Piliouras, and Marc Lanctot. 2025. Deviation Ratings: A General, Clone-Invariant Rating Method. In *Proceedings of The 7th Games, Agents, and Incentives Workshop (GAIW)*. arXiv:2502.11645 [cs.GT] <https://arxiv.org/abs/2502.11645>
- [39] Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Andrei A. Rusu, Joel Veness, Marc G. Bellemare, Alex Graves, Martin Riedmiller, Andreas K. Fidjeland, Georg Ostrovski, Stig Petersen, Charles Beattie, Amir Sadik, Ioannis Antonoglou, Helen King, Dharmarajan Kumaran, Daan Wierstra, Shane Legg, and Demis Hassabis. 2015. Human-level control through deep reinforcement learning. *Nature* 518, 7540 (2015), 529–533. <https://doi.org/10.1038/nature14236>
- [40] Kevin P. Murphy. 2022. *Probabilistic Machine Learning: An introduction*. MIT Press. <http://probml.github.io/book1>
- [41] Diego Perez-Liebana, Spyridon Samothrakis, Julian Togelius, Tom Schaul, and Simon Lucas. 2016. General Video Game AI: Competition, Challenges and Opportunities. *Proceedings of the AAAI Conference on Artificial Intelligence* 30, 1 (Mar. 2016). <https://doi.org/10.1609/aaai.v30i1.9869>
- [42] R. L Plackett. 1975. The Analysis of Permutations. *Applied Statistics* 24, 2 (1975), 193–202.
- [43] Ariel Procaccia, Benjamin Schiffer, Serena Wang, and Shirley Zhang. 2025. Metriocracy: Representative Metrics for Lite Benchmarks. arXiv:2506.09813 [cs.LG] <https://arxiv.org/abs/2506.09813>
- [44] Mark Rofin, Vladislav Mikhailov, Mikhail Florinsky, Andrey Kravchenko, Tatiana Shavrina, Elena Tutubalina, Daniel Karabekyan, and Ekaterina Artemova. 2023. Vote ‘n Rank: Revision of Benchmarking with Social Choice Theory. In *Proceedings of the 17th Conference of the European Chapter of the Association for Computational Linguistics*, Andreas Vlachos and Isabelle Augenstein (Eds.). Association for Computational Linguistics, Dubrovnik, Croatia, 670–686. <https://doi.org/10.18653/v1/2023.eacl-main.48>
- [45] Nicolas Le Roux, Mark Schmidt, and Francis Bach. 2013. A Stochastic Gradient Method with an Exponential Convergence Rate for Finite Training Sets. arXiv:1202.6258 [math.OC] <https://arxiv.org/abs/1202.6258>
- [46] Burr Settles. 2009. *Active learning literature survey*. Computer Sciences Technical Report 1648. University of Wisconsin–Madison.

- [47] M. Shewry and H. Wynn. 1987. Maximum entropy sampling. *Journal of Applied Statistics* 14 (1987), 165–170.
- [48] Shivalika Singh, Yiyang Nan, Alex Wang, Daniel D’Souza, Sayash Kapoor, Ahmet Üstün, Sanmi Koyejo, Yuntian Deng, Shayne Longpre, Noah A. Smith, Beyza Ermis, Marzieh Fadaee, and Sara Hooker. 2025. The Leaderboard Illusion. arXiv:2504.20879 [cs.AI] <https://arxiv.org/abs/2504.20879>
- [49] N. Tideman. 1987. Independence of Clones as a Criterion for Voting Rules. *Social Choice and Welfare* 4, 3 (1987), 185–206.
- [50] Simon Tong and Daphne Koller. 2002. Support vector machine active learning with applications to text classification. *The Journal of Machine Learning Research* 2 (2002), 45–66.
- [51] Yazhou Yang and Marco Loog. 2018. A benchmark and comparison of active learning for logistic regression. *Pattern Recognition* 83 (2018), 401–415.
- [52] Y. Yue, J. Broder, R. Kleinberg, and T. Joachims. 2012. The  $K$ -armed dueling bandits problem. *J. Comput. System Sci.* 78 (2012), Issue 5.
- [53] Guanhua Zhang and Moritz Hardt. 2024. Inherent trade-offs between diversity and stability in multi-task benchmarks. In *Proceedings of the 41st International Conference on Machine Learning (Vienna, Austria) (ICML’24)*. JMLR.org, Article 2435, 19 pages.
- [54] Lianmin Zheng, Wei-Lin Chiang, Ying Sheng, Siyuan Zhuang, Zhanghao Wu, Yonghao Zhuang, Zi Lin, Zhuohan Li, Dacheng Li, Eric. P. Xing, Hao Zhang, Joseph E. Gonzalez, and Ion Stoica. 2023. Judging LLM-as-a-judge with MT-Bench and Chatbot Arena. arXiv:2306.05685 [cs.CL] See also <https://lmsys.org/blog/2023-05-03-arena/> and <https://huggingface.co/spaces/lmsys/chatbot-arena-leaderboard>. Accessed August 13th, 2023.
- [55] Martin Zinkevich. 2003. Online convex programming and generalized infinitesimal gradient ascent. In *Proceedings of the Twentieth International Conference on Machine Learning*.
- [56] Martin Zinkevich, Michael Johanson, Michael Bowling, and Carmelo Piccione. 2008. Regret Minimization in Games with Incomplete Information. In *Twentieth Conference on Neural Information Processing Systems*. 905–912.
- [57] Masrour Zoghi, Zohar Karnin, Shimon Whiteson, and Maarten de Rijke. 2015. Copeland Dueling Bandits. arXiv:1506.00312 [cs.LG] <https://arxiv.org/abs/1506.00312>