
MaxHedge: Maximising a Maximum Online

Stephen Pasteris

University College London
London, UK
s.pasteris@cs.ucl.ac.uk

Fabio Vitale

Sapienza University
Italy & INRIA Lille, France
fabio.vitale@inria.fr

Kevin Chan

Army Research Lab
Adelphi, MD, USA
kevin.s.chan.civ@mail.mil

Shiqiang Wang

IBM Research
Yorktown Heights, NY, USA
wangshiq@us.ibm.com

Mark Herbster

University College London
London, UK
m.herbster@cs.ucl.ac.uk

Abstract

We introduce a new online learning framework where, at each trial, the learner is required to select a subset of actions from a given known action set. Each action is associated with an energy value, a reward and a cost. The sum of the energies of the actions selected cannot exceed a given energy budget. The goal is to maximise the cumulative *profit*, where the profit obtained on a single trial is defined as the difference between the *maximum reward* among the selected actions and the *sum of their costs*. Action energy values and the budget are known and fixed. All rewards and costs associated with each action change over time and are revealed at each trial only after the learner’s selection of actions. Our framework encompasses several online learning problems where the environment changes over time; and the solution trades-off between minimising the costs and maximising the *maximum* reward of the selected subset of actions, while being constrained to an action energy budget. The algorithm that we propose is efficient and general that may be specialised to multiple natural online combinatorial problems.

1 Introduction

In this paper we propose a novel online framework where learning proceeds in a sequence of trials and the goal is to select, at each trial, a *subset of actions* maximising a profit while taking into account a certain constraint. More precisely, we are given a finite set of actions enumerated from 1 to n . Each action is associated with three values: (i) a *cost*, (ii) an amount of *energy* (both of which are required to perform the given action), as well as (iii) a *reward*. Both the cost and the reward associated with each action may change over time, are unknown at the beginning of each trial, and their values are all revealed after the learner’s selection. The energy associated with each action is instead known by the learner and does not change over time. At each trial, the learner is required to select a subset of actions such that the *sum* of their energies does not exceed a *fixed energy budget*. The goal of the learner is to maximise the cumulative *profit*, where the profit obtained on a single trial is defined as the difference between the *maximum reward* among the selected actions and the *sum of their costs*. We denote by T the total number of trials.

Our framework is general and flexible in the sense that it encompasses several online learning problems. In the general case, the main challenge lies in the fact that the rewards are not known at the beginning of each trial, and the learner’s profit depends only on the *maximum* reward among the selected subset of actions, instead of the *sum* of all their rewards. In particular, it is worth mentioning three different problems which can be seen as special cases of our online learning framework; these are variants of the Facility Location problem, the 0-1 Knapsack problem, and the Knapsack Median problem.

When all action energy values are equal to 0, we obtain

an online learning variant of the Facility Location problem (see, e.g., [4], [21], [18]). The goal of this specific problem may be viewed as selecting and opening a subset of facilities at each given trial, to service a sequence of users which arrive one at a time. At any given trial t , each action’s cost may be interpreted as the cost for opening a facility for the t -th user. The reward associated with a given facility represents what the user potentially gains when it is opened. More specifically, the rewards can be seen as quantities dependent on the distance between the user and the facilities in some metric space. In this context, it is reasonable to assume that the profit obtained on a single trial depends *only on the maximum reward* among the opened facilities taking into account the facility costs. This represents a natural setting for the Facility Location problem, because in practical scenarios users may arrive sequentially and each arrival requires a connection to an open facility. Online versions of the Facility Location problem have been studied in several contexts (see, e.g., [5], [20]). However, as far as we are aware, our work is the first study of this specific online setting for the Facility Location problem. Moreover, this dynamic model is natural and interesting within this context¹, because in practice the location of the next user, which in turn determines the reward associated with each facility, is often unknown to the learner. At each trial, after the user’s location is known, the rewards are then revealed, because disclosing the user’s location enables to compute all distances between the facilities and the user, which are previously unknown. In fact, this corresponds to assuming all rewards are revealed at the end of the trial.

When, instead, all rewards are equal to 0 and all costs are negative, the problem can be seen as an online learning variant of the 0-1 Knapsack problem (see, e.g., [6], [19]). In this case, each action corresponds to an item whose weight is equal to the associated energy, the energy budget represents the knapsack capacity, and the absolute value of each action cost can be viewed as the corresponding item value. Our formulation makes the problem challenging especially because the item values are revealed only after the learner’s selection.

Finally, when all costs are equal to 0, we obtain an online learning variant of the Knapsack Median problem ([1], [15]). In this problem (which is a generalization of the k -median problem – see, e.g., [2], [10]), we are given a set of clients and facilities, as well as a distance metric. The goal is to open a subset of facilities such that the total connection cost (distance to nearest open facility) of all clients is minimized, while the sum of the open facility weights is limited by a budget threshold.

¹Similar arguments hold also for motivating the other two special cases mentioned in this section.

In our framework, at each trial the rewards can express the closeness of each facility, and the action energy budget represents therefore the threshold of sum of the opened facility weights.

This problem has practical applications in multiple domains. In computer networks, network administrators may want to understand where to place network monitors or intrusion detection systems. Network packets or malicious attacks are related to the events playing a crucial role in this scenario, and a limited amount of network resources are available to detect or observe network behavior. Another class of application examples include municipal emergency services. A service center needs to deploy responders (police, paramedics, fire rescue), and with limited resources, personnel must be deployed sparingly. A further application is related to deploying program instances in a distributed computing environment (e.g. distributed cloud). These systems must respond to user requests and are subject to constraints. Both costs and rewards may be unknown in practical real-world scenarios at the beginning of each trial.

For our general problem we propose and rigorously analyse a very scalable algorithm called MAXHEDGE based on the complex interplay between satisfying the energy budget constraint and bounding the profit by a concave function, which in turn is related to the online gradient descent algorithm. Moreover, the total time required per trial by our learning strategy is quasi-linear in n . We measure the performance of proposed solution with respect to the difference between its cumulative profit and a discounted cumulative profit of the best fixed subset of actions.

In summary, our framework captures several real-world problems, where the environment changes over time and the solutions trade-off between minimising the costs and maximising the *maximum* reward among the selected subset of actions, while being constrained to an action energy budget. The framework is very general and the proposed algorithm is very efficient and may be specialised to several natural online combinatorial problems. Finally we provide a guarantee on the rewards achieved and costs incurred as compared to the best fixed subset of actions.

1.1 Related Work

The closest work to our online learning framework is perhaps addressed in [13], where the authors describe an online learning algorithm for structured concepts that are formed by components. Each component and each concept can be respectively seen as an action and a feasible subset of actions. Despite several similarities, the algorithm they proposed cannot be used when

the rewards are non-zero, because we focus on the *maximum* reward among the selected subsets of actions, whereas in [13] the profit corresponds to the *sum* of the rewards of all selected components/actions. When all rewards are instead equal to 0, then we have the online variant of the 0-1 Knapsack problem described above as one of the three special cases of our general problem. In Appendix C we prove that if the algorithm presented in [13] could handle the Knapsack problem, then the classical version of the Knapsack problem could be solved in polynomial time, which therefore implies that it cannot address this problem unless $P = NP$.

The Hedge Algorithm described in [7] obtains a regret linear in n . However, as we have exponentially many possible subsets of selected actions, a vanilla application of Hedge would require an exponential amount of time and space to solve our problem.

Another class of problems and algorithms that are not far from ours is represented by online decision problems where efficient strategies use, as a subroutine, an approximation algorithm for choosing a concept to maximise an inner product ([8], [12], [11]). Again, these learning strategies cannot handle the case of non-zero reward. They also have a significantly higher time complexity than MAXHEDGE in the case of zero reward.

In [3], [9], and [22], the authors address the problem of online maximisation of non-negative monotone submodular functions. Although our profit is submodular, it is not necessarily either non-negative or monotone. However, as we shall show in Appendix D, we could, for the facility location special case, combine much of the mechanics of our paper with the second algorithm of [3], essentially doing gradient ascent with the exact expected profit instead of an approximate expected profit (as is done in MAXHEDGE). Even though this new algorithm could be as efficient as the one presented in our paper, its theoretical guarantees are worse. Note that in the Knapsack and Knapsack Median special cases, the profit is indeed monotone and non-negative. For these special cases, the approach presented in [22] comes close to solving the problem, but only guarantees that the expectation of the total energy does not exceed the budget rather than the actual total energy. In addition, for the Knapsack Median problem, [22] uses quadratic time and space. As far as we are aware, there does not exist any trivial reduction to use [3] or [9] for these special cases.

Finally, in [16] the authors address a problem of online minimisation of submodular function. Our paper maximises, instead of minimising, a submodular function (the profit). This is a very different problem.

Some of the techniques behind the development of

MAXHEDGE were inspired by [25].

2 Preliminaries and Problem Setup

2.1 Preliminaries

Vectors in \mathbb{R}^n will be indicated in bold. Given a vector $\mathbf{v} \in \mathbb{R}^n$ we define v_i to be its i -th component. Given vectors $\mathbf{v}, \mathbf{x} \in \mathbb{R}^n$ define $\langle \mathbf{v}, \mathbf{x} \rangle := \sum_{i=1}^n v_i x_i$. Define \mathcal{P} to be the set of n -dimensional real vectors in which every component is non-negative. Given a closed convex set $C \subseteq \mathbb{R}^n$ and a vector $\mathbf{x} \in \mathbb{R}^n$, we define $\Pi_C(\mathbf{x})$ as the projection (under the Euclidean norm) of \mathbf{x} onto C . Let \mathbb{N} be the set of the positive integers. For $l \in \mathbb{N}$ define $[l] = \{1, 2, 3, \dots, l\}$. Given an event E let $\neg E$ be the event that E does *not* occur. Let $\mathbb{P}(E)$ be the probability that event E occurs. For a random variable Y , let $\mathbb{E}(Y)$ be the expected value of Y . Given a differentiable function $h : \mathbb{R}^n \rightarrow \mathbb{R}$ and a vector $\mathbf{x} \in \mathbb{R}^n$ let $\nabla h(\mathbf{x})$ be the derivative of h evaluated at \mathbf{x} . Let $\partial_i h(\mathbf{x})$ be the i -th component of $\nabla h(\mathbf{x})$.

2.2 Problem Setup

In this section we formally define our problem. We have a set of n actions enumerated from 1 to n . Each action i has an *energy* $z_i \in [0, \beta]$ for some $\beta < 1$, and on each trial t each action i has a *cost* $c_i^t \in \mathbb{R}$ (which can be negative) and a *reward* $r_i^t \in \mathbb{R}^+$. The learner knows \mathbf{z} , but \mathbf{c}^t and \mathbf{r}^t are revealed to the learner only at the end of trial t . On each trial t the learner has to select a set $X^t \subseteq [n]$ of actions, such that the total energy $\sum_{i \in X^t} z_i$ of the selected actions is no greater than 1. In selecting the set X^t , the learner pays a cost equal to $\sum_{i \in X^t} c_i^t$. On each trial t the learner then receives the maximum reward $\max_{i \in X^t} r_i^t$, over all actions selected (defined as equal to zero if X^t is empty). Hence, the *profit* obtained by the learner on trial t is equal to $\max_{i \in X^t} r_i^t - \sum_{i \in X^t} c_i^t$.

Formally, this online problem can be defined as follows: We have a vector $\mathbf{z} \in \mathcal{P}$ known to the learner. On trial t :

1. Nature selects vectors $\mathbf{c}^t \in \mathbb{R}^n$ and $\mathbf{r}^t \in \mathcal{P}$ (but does not reveal these vectors to learner)
2. Learner selects a set $X^t \subseteq [n]$ with $\sum_{i \in X^t} z_i \leq 1$
3. Learner obtains profit:

$$\mu^t(X^t) := \max_{i \in X^t} r_i^t - \sum_{i \in X^t} c_i^t$$

4. \mathbf{c}^t and \mathbf{r}^t are revealed to the learner.

In this paper we write, for a trial t , the cost vector \mathbf{c}^t as the sum $\mathbf{c}^{t+} + \mathbf{c}^{t-}$ where $c_i^{t+} := \max\{0, c_i^t\}$ and $c_i^{t-} := \min\{0, c_i^t\}$, i.e. \mathbf{c}^{t+} and \mathbf{c}^{t-} are the positive and negative parts of the cost vector, respectively.

In order to bound the cumulative profit of our algorithm we, for some $\alpha, \delta \in [0, 1]$, some set $S \subseteq [n]$ and some trial t , define the (α, δ) -discounted profit $\hat{\mu}_{\alpha, \delta}^t(S)$ as:

$$\hat{\mu}_{\alpha, \delta}^t(S) := \alpha \max_{i \in S} r_i^t - \alpha \sum_{i \in S} c_i^{t-} - \delta \sum_{i \in S} c_i^{t+}$$

which would be the profit obtained on trial t if we selected the subset of actions S , and all the rewards and negative costs were multiplied by α and all positive costs were multiplied by δ .

In this paper we provide a randomised quasi-linear time (per trial) algorithm MAXHEDGE that, for any set $S \subseteq [n]$ with $\sum_{i \in S} z_i \leq 1$, obtains an expected cumulative profit bounded below by:

$$\mathbb{E} \left(\sum_{t=1}^T \mu^t(X^t) \right) \geq \sum_{t=1}^T \hat{\mu}_{\alpha, \delta}^t(S) - n\sqrt{2T}\delta(\hat{r} + \hat{c})$$

where $\delta := (1 - \sqrt{\beta})^2$, $\alpha := 1 - \exp(- (1 - \sqrt{\beta})^2)$, $\hat{r} := \max_{t \in [T], i \in [n]} r_i^t$ and $\hat{c} := \max_{t \in [T], i \in [n]} |c_i^t|$.

This paper is structured as follows. In Section 3 we introduce the algorithms that define MAXHEDGE. In Section 4 we prove that the sets of actions selected by MAXHEDGE are feasible. In Section 5 we prove the above bound on the cumulative profit. In Section 6 we give special cases of the general problem.

3 Algorithms

We now present our learning strategy MAXHEDGE, describing the two subroutines ‘‘Algorithm 1’’ and ‘‘Algorithm 2’’ (see the pseudocode below). MAXHEDGE maintains a vector $\boldsymbol{\omega} \in C$ where $C := \{\mathbf{x} \in [0, 1]^n : \langle \mathbf{x}, \mathbf{z} \rangle \leq 1\}$. We define $\boldsymbol{\omega}^t$ to be the vector $\boldsymbol{\omega}$ at the start of trial t . We initialise $\boldsymbol{\omega}^1 \leftarrow \mathbf{0}$. On trial t MAXHEDGE (randomly) constructs X^t from $\boldsymbol{\omega}^t$ using Algorithm 1. After receiving \mathbf{r}^t and \mathbf{c}^t MAXHEDGE updates $\boldsymbol{\omega}$ (from $\boldsymbol{\omega}^t$ to $\boldsymbol{\omega}^{t+1}$) using Algorithm 2. Algorithm 2 also uses a ‘‘learning rate’’ $\hat{\eta}^t$ which is defined from $\hat{\eta}^{t-1}$. We define $\hat{\eta}^0 := \infty$.

Algorithm 1 operates with a partition of all possible actions and, for each set in this partition, Algorithm 1 draws a certain subset of actions from it. Given a set in the partition, the number of actions drawn from it and the probability distribution governing the draws depends on β and $\boldsymbol{\omega}^t$. The subset, X^t , of actions selected by Algorithm 1 satisfies the following three crucial properties (proved in sections 4 and 5):

- The total energy of all actions selected is no greater than 1.
- Given an arbitrary set $Z \subseteq [n]$, the probability that X^t and Z intersect is lower bounded by $1 - \exp(-\delta \sum_{i \in Z} \omega_i^t)$.
- Given an action i , the probability that it is selected on trial t is upper bounded by $\delta \omega_i^t$.

In the analysis we shall construct, from \mathbf{r}^t and \mathbf{c}^t , a convex function $h^t : C \rightarrow \mathbb{R}$. Using the second and third properties (given above) of Algorithm 1, we show that $h^t(\boldsymbol{\omega}^t)$ is an upper bound on the negative of the expected profit on trial t . Algorithm 2 computes the gradient $\mathbf{g}^t := \nabla h^t(\boldsymbol{\omega}^t)$ and updates $\boldsymbol{\omega}$ using online gradient descent on C .

The last line of Algorithm 2 requires us to project (with Euclidean distance) the vector \mathbf{y}^t onto the set C , i.e. we must compute the \mathbf{x} that minimises the value $\|\mathbf{x} - \mathbf{y}^t\|$ subject to $\mathbf{x} \in C$. Note that minimising $\|\mathbf{x} - \mathbf{y}^t\|$ is equivalent to minimising $\|\mathbf{x} - \mathbf{y}^t\|^2 = \langle \mathbf{x}, \mathbf{x} \rangle - 2\langle \mathbf{y}^t, \mathbf{x} \rangle + \langle \mathbf{y}^t, \mathbf{y}^t \rangle$, which is in turn equivalent to minimising $\langle \mathbf{x}, \mathbf{x} \rangle - 2\langle \mathbf{y}^t, \mathbf{x} \rangle$. The constraints defining the set C then imply that this projection is a case of the continuous bounded quadratic knapsack problem which can be solved in linear time (see, e.g., [14]).

The bottleneck of the algorithms is hence the ordering step in Algorithm 2 which takes a time of $\mathcal{O}(n \log(n))$

Algorithm 1 Constructing X^t

- 1: $C \leftarrow \{\mathbf{x} \in [0, 1]^n : \langle \mathbf{x}, \mathbf{z} \rangle \leq 1\}$
- 2: $\beta \leftarrow \max_{i \in [n]} z_i$
- 3: $\tau \leftarrow 1 - \sqrt{\beta}$
- 4: $\delta \leftarrow (1 - \sqrt{\beta})^2$
- 5: $\Gamma \leftarrow \{q \in \mathbb{N} : \exists i \in [n] \text{ with } \tau^q \beta < z_i \leq \tau^{q-1} \beta\}$
- 6: For all $q \in \Gamma$ set $\Omega_q \leftarrow \{i \in [n] : \tau^q \beta < z_i \leq \tau^{q-1} \beta\}$.

7: **Input:** $\boldsymbol{\omega}^t \in C$

- 8: For all $q \in \Gamma$ set $\pi_q^t \leftarrow \sum_{i \in \Omega_q} \omega_i^t$
- 9: For all $q \in \Gamma$ set $\zeta_q^t \leftarrow \lfloor \delta \pi_q^t \rfloor$
- 10: For all $q \in \Gamma$ and for all $k \leq \zeta_q^t$ draw $\xi_{q,k}^t$ randomly from Ω_q such that $\xi_{q,k}^t \leftarrow i$ with probability ω_i^t / π_q^t
- 11: For all $q \in \Gamma$ and for $k := \zeta_q^t + 1$ draw $\xi_{q,k}^t$ randomly from $\Omega_q \cup \{0\}$ such that, for $i \in \Omega_q$ we have $\xi_{q,k}^t \leftarrow i$ with probability $(\delta \pi_q^t - \lfloor \delta \pi_q^t \rfloor) \omega_i^t / \pi_q^t$, and we have $\xi_{q,k}^t \leftarrow 0$ with probability $\lfloor \delta \pi_q^t \rfloor + 1 - \delta \pi_q^t$.
NB: In the case that $\pi_q^t = 0$ we define $0/0 = 0$

12: **Output:** $X^t \leftarrow \{\xi_{q,k}^t : q \in \Gamma, k \leq \zeta_q^t + 1\} \setminus \{0\}$

Algorithm 2 Computing ω^{t+1}

- 1: $C \leftarrow \{\mathbf{x} \in [0, 1]^n : \langle \mathbf{x}, \mathbf{z} \rangle \leq 1\}$
- 2: $\beta \leftarrow \max_{i \in [n]} z_i$
- 3: $\delta \leftarrow (1 - \sqrt{\beta})^2$

- 4: **Input:** $\omega^t \in C$, $\mathbf{c}^t \in \mathbb{R}^n$, $\mathbf{r}^t \in \mathcal{P}$

- 5: Order $[n]$ as $[n] = \{\sigma(t, 1), \sigma(t, 2), \dots, \sigma(t, n)\}$ such that $r_{\sigma(t,j)}^t \geq r_{\sigma(t,j+1)}^t$ for all $j \in [n-1]$
- 6: For all $j \in [n]$ set $\epsilon_j^t \leftarrow \exp\left(-\delta \sum_{k=1}^j \omega_{\sigma(t,k)}^t\right)$
- 7: For all $j \in [n]$ set:

$$\lambda_j^t \leftarrow r_{\sigma(t,n)}^t \epsilon_n^t + \sum_{k=j}^{n-1} \left(r_{\sigma(t,k)}^t - r_{\sigma(t,k+1)}^t\right) \epsilon_k^t$$
- 8: For all $j \in [n]$ set:

$$g_{\sigma(t,j)}^t \leftarrow \delta \left(c_{\sigma(t,j)}^{t+} + c_{\sigma(t,j)}^{t-}\right) \exp\left(-\delta \omega_{\sigma(t,j)}^t\right) - \lambda_j^t$$
- 9: $\hat{\eta}^t \leftarrow \min\{\hat{\eta}^{t-1}, \sqrt{n}/\|g^t\|\}$
- 10: $\eta^t \leftarrow \hat{\eta}^t / \sqrt{2t}$
- 11: $\mathbf{y}^t \leftarrow \omega^t - \eta^t g^t$

- 12: **Output:** $\omega^{t+1} \leftarrow \Pi_C(\mathbf{y}^t)$

4 The Feasibility of X^t

In this section we show that the the total energy of the actions selected by Algorithm 1 is no greater than 1, as required in our problem definition (see Section 2.2). We first introduce the sets and quantities used in the selection of the actions.

Definition 4.1. We define the following:

- $\tau := 1 - \sqrt{\beta}$ and $\delta := (1 - \sqrt{\beta})^2$
- For all $q \in \mathbb{N}$ we define $\Omega_q := \{i \in [n] : \tau^q \beta < z_i \leq \tau^{q-1} \beta\}$
- $\Gamma := \{q \in \mathbb{N} : \Omega_q \neq \emptyset\}$. Note that $\{\Omega_q : q \in \Gamma\}$ is a partition of $[n]$
- On trial t , for all $q \in \Gamma$ define $\pi_q^t := \sum_{i \in \Omega_q} \omega_i^t$ and $\zeta_q^t := \lfloor \delta \pi_q^t \rfloor$.

Algorithm 1 works by drawing actions $\{\xi_{q,k}^t : q \in \Gamma, k \in [\zeta_q^t + 1]\}$ randomly, where the number of actions (including a “null action” 0) drawn (with replacement) from Ω_q is equal to $\zeta_q^t + 1$ and the probability distribution of the draws is dependent on ω^t .

The following theorem ensures that the choice of X^t made by our method satisfies our problem’s energy constraint.

Theorem 4.2. On trial t we have $\sum_{i \in X^t} z_i \leq 1$.

Proof. See Appendix A □

5 Bounding the Cumulative Profit

In this section we bound the cumulative profit of MAXHEDGE.

5.1 The Probability of Intersection

In this subsection we first bound below the probability that, on a trial t , an arbitrary set Z intersects with X^t . We start with the following lemma:

Lemma 5.1. Given a set D of independent draws from $|D|$ -many probability distributions, such that the probability of an event E happening on draw $d \in D$ is ρ_d , then the probability of event E happening on either of the draws is lower bounded by:

$$1 - \exp\left(-\sum_{d \in D} \rho_d\right)$$

Proof. See Appendix A □

We now bound the probability of intersection:

Theorem 5.2. For any trial t and any subset $Z \subseteq [n]$ we have $\mathbb{P}(Z \cap X^t \neq \emptyset) \geq 1 - \exp(-\delta \sum_{i \in Z} \omega_i^t)$.

Proof. See Appendix A □

We now bound the probability that some arbitrary action is selected on trial t :

Theorem 5.3. Given some action $i \in [n]$ and some trial $t \in [T]$ we have $1 - \exp(-\delta \omega_i^t) \leq \mathbb{P}(i \in X^t) \leq \delta \omega_i^t$.

Proof. See Appendix A □

5.2 Approximating the Expected Profit

In this subsection we define a convex function $h^t : C \rightarrow \mathbb{R}$ and show that the expected profit on trial t is bounded below by $-h^t(\omega^t)$.

Definition 5.4. For each trial t we order $[n]$ as $[n] = \{\sigma(t, 1), \sigma(t, 2), \dots, \sigma(t, n)\}$ where $r_{\sigma(t,j)}^t \geq r_{\sigma(t,j+1)}^t$ for all $j \in [n-1]$. We also define $\sigma(t, n+1) := 0$ and $r_0^t := 0$.

Definition 5.5. Given a trial t and a number $j \in [n]$ we define the function $f_j^t : \mathcal{P} \rightarrow \mathbb{R}$ by:

$$f_j^t(\gamma) := \left(r_{\sigma(t,j)}^t - r_{\sigma(t,j+1)}^t\right) \left(1 - \exp\left(-\delta \sum_{k=1}^j \gamma_{\sigma(t,k)}\right)\right)$$

We also define the function $h^t : \mathcal{P} \rightarrow \mathbb{R}$ as:

$$h^t(\gamma) = \langle \mathbf{c}^{t+}, \gamma \rangle + \sum_{i=1}^n c_i^{t-} (1 - \exp(-\delta \gamma_i)) - \sum_{j=1}^n f_j^t(\gamma)$$

Theorem 5.6. For all $t \in [T]$, the function h^t is convex.

Proof. See Appendix A \square

The rest of this subsection proves that the expected profit on trial t is bounded below by $-h^t(\omega^t)$.

Lemma 5.7. On trial t we have $\max_{i \in X^t} r_i^t = \sum_{j=1}^n (r_{\sigma(t,j)}^t - r_{\sigma(t,j+1)}^t) \mathcal{I}(\exists k \leq j : \sigma(t,k) \in X^t)$.

Proof. See Appendix A \square

Lemma 5.8. On trial t we have $\mathbb{E}(\max_{i \in X^t} r_i^t) = \sum_{j=1}^n (r_{\sigma(t,j)}^t - r_{\sigma(t,j+1)}^t) \mathbb{P}(\exists k \leq j : \sigma(t,k) \in X^t)$.

Proof. Direct from lemma 5.7 using linearity of expectation. \square

Theorem 5.9. On trial t we have $\mathbb{E}(\mu^t(X^t)) \geq -h^t(\omega^t)$.

Proof. See Appendix A \square

5.3 The Gradient

In this subsection we show how to construct the gradient of h^t and bound its magnitude. We start with the following definitions.

Definition 5.10. On any trial t and for any $j \in [n]$ we define:

- $\epsilon_j^t := \exp\left(-\delta \sum_{k=1}^j \omega_{\sigma(t,k)}^t\right)$
- $\lambda_j^t := \sum_{k=j}^n (r_{\sigma(t,k)}^t - r_{\sigma(t,k+1)}^t) \epsilon_k^t$
- $g_{\sigma(t,j)}^t := \delta \left(c_{\sigma(t,j)}^{t+} + c_{\sigma(t,j)}^{t-} \exp\left(-\delta \omega_{\sigma(t,j)}^t\right) - \lambda_j^t \right)$

We first show that \mathbf{g}^t is the gradient of h^t evaluated at ω^t .

Theorem 5.11. On any trial t we have $\mathbf{g}^t = \nabla h^t(\omega^t)$.

Proof. See Appendix A \square

We now bound the magnitude of the gradient.

Lemma 5.12. For any trial t we have $\|\mathbf{g}^t\|^2 \leq n\delta^2(\hat{r} + \hat{c})^2$.

Proof. Since $\omega_{\sigma(t,k)}^t \geq 0$ for all $k \in [n]$ we have $\epsilon_j^t \in [0, 1]$ for all $j \in [n]$. This gives us, for all $j \in [n]$, that $\delta \lambda_j^t \leq \delta r_{\sigma(t,n)}^t + \delta \sum_{k=j}^{n-1} (r_{\sigma(t,k)}^t - r_{\sigma(t,k+1)}^t) = \delta r_{\sigma(t,j)}^t \leq \delta \hat{r}$. Since also $\lambda_j^t \geq 0$ this implies that $-g_i^t \leq \delta \hat{r} + \delta \hat{c}$ and that $g_i^t \leq \delta \hat{c}$ so $(g_i^t)^2 \leq (\delta \hat{r} + \delta \hat{c})^2$. This then implies the result. \square

5.4 Online Gradient Descent

In this subsection we show that Algorithm 2 corresponds to the use of online gradient descent over C with convex functions $\{h^t : t \in [T]\}$ and we use the standard analysis of online gradient descent to derive a lower bound on the cumulative profit.

From here on we compare the performance of our algorithm against any fixed set S of actions such that $\sum_{i \in S} z_i \leq 1$. We define ϕ as the vector in \mathbb{R}^n such that, for all $i \in [n]$, we have $\phi_i := 0$ if $i \notin S$ and $\phi_i := 1$ if $i \in S$. It is clear that $\phi \in C$.

Definition 5.13. Our learning rates are defined as follows:

- $\hat{\eta}^t := \min_{t' \leq t} (\sqrt{n}/\|\mathbf{g}^{t'}\|)$
- $\eta^t := \hat{\eta}^t/\sqrt{2t}$

The next result follows from the standard analysis of online gradient descent.

Theorem 5.14. We have:

$$\sum_{t=1}^T (h^t(\omega^t) - h^t(\phi)) \leq \frac{R^2}{2\eta^T} + \frac{1}{2} \sum_{t=1}^T \eta^t \|\mathbf{g}^t\|^2$$

where $R := \max_{x,y \in C} \|x - y\|$.

Proof. For all trials t : From Theorem 5.6 we have that h^t is a convex function. From Theorem 5.11 we have that $\mathbf{g}^t = \nabla h^t(\omega^t)$. We also have that $\eta^{t+1} \leq \eta^t$ so since, by Algorithm 2, we have $\omega^{t+1} = \Pi_C(\omega^t - \eta^t \mathbf{g}^t)$ and $\phi \in C$, the standard analysis of online gradient descent (see, e.g., [24]) gives us the result. \square

We now bound the right hand side of the equation in Theorem 5.14.

Definition 5.15. Define $s := \max_{t \in [T]} \|\mathbf{g}^t\|^2/n$.

Lemma 5.16. For any trial t , $\eta^t \|\mathbf{g}^t\|^2 \leq n\sqrt{s/(2t)}$.

Proof. We have $\hat{\eta}^t = \min_{t' \leq t} (\sqrt{n}/\|\mathbf{g}^{t'}\|) \leq \sqrt{n}/\|\mathbf{g}^t\|$. This implies that $\eta^t = \hat{\eta}^t/\sqrt{2t} \leq \sqrt{n}/(\|\mathbf{g}^t\|\sqrt{2t})$ so $\eta^t \|\mathbf{g}^t\|^2 \leq \sqrt{n}\|\mathbf{g}^t\|/\sqrt{2t}$ which, by definition of s , is bounded above by $n\sqrt{s/(2t)}$. \square

Lemma 5.17. $R^2/\eta^T \leq n\sqrt{2sT}$

Proof. We have $\hat{\eta}^T = \min_{t \in [T]} \sqrt{n}/\|\mathbf{g}^t\| = 1/\sqrt{s}$ so $\eta^T := \hat{\eta}^T/\sqrt{2T} \leq 1/\sqrt{2sT}$. Since $C \subseteq [0, 1]^n$, if $\mathbf{x}, \mathbf{y} \in C$ then each component of $\mathbf{x} - \mathbf{y}$ has magnitude bounded above by 1 which implies that $R^2 \leq n$. Putting together gives the result. \square

Theorem 5.18. *We have:*

$$\sum_{t=1}^T (h^t(\omega^t) - h^t(\phi)) \leq n\sqrt{2T}\delta(\hat{r} + \hat{c})$$

Proof. See Appendix A \square

The next result bounds $h^t(\phi)$.

Lemma 5.19. *On trial t we have $\hat{\mu}_{\alpha,\delta}^t(S) \leq -h^t(\phi)$ where $\alpha := 1 - e^{-\delta}$.*

Proof. Let $j' = \operatorname{argmax}_{j \in [n]: \sigma(t,j) \in S} r_{\sigma(t,j)}^t$ which is equal to the minimum j such that $\sigma(t,j) \in S$. Note that for all $j \geq j'$ we have $\sum_{k=1}^j \phi_{\sigma(t,k)} \geq 1$ and hence $f_j^t(\phi) \geq (r_{\sigma(t,j)}^t - r_{\sigma(t,j+1)}^t)(1 - e^{-\delta})$ so $\sum_{j=1}^n f_j^t(\phi) \geq \sum_{j=j'}^n f_j^t(\phi) \geq \sum_{j=j'}^n (r_{\sigma(t,j)}^t - r_{\sigma(t,j+1)}^t)(1 - e^{-\delta}) = r_{\sigma(t,j')}^t(1 - e^{-\delta}) = (1 - e^{-\delta}) \max_{i \in S} r_i^t$.

Also note that $\delta \langle c^{t+}, \phi \rangle = \delta \sum_{i \in S} c_i^{t+}$ and that $\sum_{i=1}^n c_i^{t-}(1 - \exp(-\delta\phi_i)) = \sum_{i \in S} c_i^{t-}(1 - e^{-\delta})$. Combining with the above gives us the result. \square

Putting together we obtain the main result:

Theorem 5.20. *We have:*

$$\mathbb{E} \left(\sum_{t=1}^T \mu^t(X^t) \right) \geq \sum_{t=1}^T \hat{\mu}_{\alpha,\delta}^t(S) - n\sqrt{2T}\delta(\hat{r} + \hat{c})$$

where $\delta := (1 - \sqrt{\beta})^2$, $\alpha := 1 - \exp(-(1 - \sqrt{\beta})^2)$, $\hat{r} := \max_{t \in [T], i \in [n]} r_i^t$ and $\hat{c} := \max_{t \in [T], i \in [n]} |c_i^t|$.

Proof. Let $\alpha := 1 - e^{-\delta}$. By Theorem 5.9 we have, for all $t \in [T]$, that $\mathbb{E}(\mu^t(X^t)) \geq -h^t(\omega^t)$. By Lemma 5.19 we have, for all $t \in [T]$, that $\hat{\mu}_{\alpha,\delta}^t(S) \leq -h^t(\phi)$. Hence we have that $\hat{\mu}_{\alpha,\delta}^t(S) - \mathbb{E}(\mu^t(X^t)) \leq h^t(\omega^t) - h^t(\phi)$. By Theorem 5.18 we then have:

$$\begin{aligned} & \sum_{t=1}^T (\hat{\mu}_{\alpha,\delta}^t(S) - \mathbb{E}(\mu^t(X^t))) \\ & \leq \sum_{t=1}^T (h^t(\omega^t) - h^t(\phi)) \\ & \leq n\sqrt{2T}\delta(\hat{r} + \hat{c}) \end{aligned}$$

Rearranging gives us:

$$\sum_{t=1}^T \mathbb{E}(\mu^t(X^t)) \geq \sum_{t=1}^T \hat{\mu}_{\alpha,\delta}^t(S) - n\sqrt{2T}\delta(\hat{r} + \hat{c})$$

\square

6 Special Cases

The following online variants of classic computer science problems are special cases of the general problem.

6.1 Facility Location Problem

The (inverted) facility location problem is defined by a vector $c \in \mathcal{P}$ and vectors $r^1, r^2, \dots, r^T \in \mathcal{P}$. A feasible solution is any $X \subseteq [n]$. The aim is to maximise the objective function:

$$\sum_{t=1}^T \max_{i \in X} r_i^t - \sum_{i \in X} c_i$$

An example of the problem is as follows. There are n sites and T users, all located in some metric space. We have to choose a set X of sites to open a facility on. Opening a facility on site i costs us c_i . Each user pays us a reward based on how near it is to the closest open facility. If the nearest open facility to user t is at site i then user t rewards us r_i^t . The objective is to maximise the total profit.

In our online variant of the (inverted) facility location problem, learning proceeds in trials. On trial t :

1. For all sites i , the cost, c_i^t of opening a facility on site i is revealed to the learner.
2. The learner chooses a set X^t of sites in which to open facilities on.
3. User t requests the use of a facility, revealing r^t to the learner.
4. Learner incurs profit: $\max_{i \in X^t} r_i^t - \sum_{i \in X^t} c_i^t$

The objective is to maximise the cumulative profit. Note that this is the special case of our problem when, for all $i \in [n]$ and $t \in [T]$ we have $z_i = 0$ and $c_i^t \geq 0$. Given some set S the expected cumulative profit of MAXHEDGE is then bounded below by:

$$\sum_{t=1}^T \left((1 - 1/e) \max_{i \in S} r_i^t - \sum_{i \in S} c_i^t \right) - n\sqrt{2T}(\hat{r} + \hat{c})$$

6.2 Knapsack Median Problem

The (inverted) knapsack median problem is defined by a vector $z \in \mathcal{P}$ and vectors $r^1, r^2, \dots, r^T \in \mathcal{P}$. A feasible solution is any $X \subseteq [n]$ with $\sum_{i \in X} z_i \leq 1$. The aim is to maximise the objective function $\sum_{t=1}^T \max_{i \in X} r_i^t$.

An example of the problem is as follows. There are n sites and T users, all located in some metric space. We

have to choose a set X of sites to open a facility on. Opening a facility on site i has a fee of z_i and we have a budget of 1 to spend on opening facilities. Each user pays us a reward based on how near it is to the closest open facility. If the nearest open facility to user t is at site i then user t rewards us r_i^t . The objective is to maximise the total reward.

In our online variant of the (inverted) knapsack median problem, learning proceeds in trials. The learner has knowledge of the fee z_i for every site i . On trial t :

1. The learner chooses a set X^t of sites in which to open facilities on. The total fee, $\sum_{i \in X^t} z_i$, can't exceed 1.
2. User t requests the use of a facility, revealing \mathbf{r}^t to the learner.
3. Learner incurs reward: $\max_{i \in X^t} r_i^t$

The objective is to maximise the cumulative reward. Note that this is the special case of our problem when, for all $i \in [n]$ and $t \in [T]$ we have $c_i^t = 0$. Given some set S with $\sum_{i \in S} z_i \leq 1$ the expected cumulative reward of MAXHEDGE is then bounded below by:

$$(1 - \exp(-\delta)) \sum_{t=1}^T \max_{i \in S} r_i^t - \delta n \sqrt{2T} \hat{r}$$

where $\delta := (1 - \sqrt{\max_{i \in [n]} z_i})^2$

6.3 0-1 Knapsack Problem

The knapsack problem is defined by a vector $\mathbf{z} \in \mathcal{P}$ and a vector $\mathbf{v} \in \mathcal{P}$. A feasible solution is any $X \subseteq [n]$ with $\sum_{i \in X} z_i \leq 1$. The aim is to maximise the objective function $\sum_{i \in X} v_i$.

An example of the problem is as follows. We have n items. Item i has a value v_i and a weight z_i . The objective is to place a set $X \subseteq [n]$ of items in the knapsack that maximises the total value of all items in the knapsack subject to their total weight being no greater than 1.

In our online variant of the knapsack problem, learning proceeds in trials. The learner has knowledge of the weight z_i for every item i . On trial t :

1. The learner chooses a set X^t of items to place in the knapsack. The total weight, $\sum_{i \in X^t} z_i$, can't exceed 1.
2. For each item i , the value v_i^t , of item i on this trial is revealed to the learner
3. Learner incurs profit: $\sum_{i \in X^t} v_i^t$

The objective is to maximise the cumulative profit. Note that this is the special case of our problem when, for all $i \in [n]$ and $t \in [T]$ we have $r_i^t = 0$ and $c_i^t \leq 0$ (noting that $c_i^t = -v_i^t$). Given some set S with $\sum_{i \in S} z_i \leq 1$ the expected cumulative profit of MAXHEDGE is then bounded below by:

$$(1 - \exp(-\delta)) \sum_{t=1}^T \sum_{i \in S} v_i^t - \delta n \sqrt{2T} \hat{c}$$

where $\delta := (1 - \sqrt{\max_{i \in [n]} z_i})^2$

7 Conclusions and Ongoing Research

We presented and investigated in depth a novel online framework, capable of encompassing several online learning problems and capturing many practical problems in the real-world. The main challenge of the general version of this problem lies in the fact that the learner's profit depends on the maximum reward of the selected actions, instead of the sum of all their rewards. We proposed and rigorously analysed a very scalable and efficient learning strategy MAXHEDGE.

Current ongoing research includes:

- Deriving a lower bound on the achievable profit.
- Complementing our results with a set of experiments on synthetic and real-world datasets.
- Several real systems usually have a switching cost for turning on/off services, which translates in our framework to the cost incurred whenever an action selected at any given trial is not selected in the preceding one. This represents an interesting direction for further research, which is certainly motivated by practical problems.

Acknowledgements. This research was sponsored by the U.S. Army Research Laboratory and the U.K. Ministry of Defence under Agreement Number W911NF-16-3-0001. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the U.S. Army Research Laboratory, the U.S. Government, the U.K. Ministry of Defence or the U.K. Government. The U.S. and U.K. Governments are authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation hereon. Fabio Vitale acknowledges support from the ERC Starting Grant "DMAP 680153", the Google Focused Award "ALL4AI", and grant "Dipartimenti di Eccellenza 2018-2022", awarded to the Department of Computer Science of Sapienza University.

References

- [1] M. Charikar, S. Guha, E. Tardos, D.B. Shmoys. A constant-factor approximation algorithm for the k-median problem. In *ACM Symposium on Theory of Computing (STOC)*, pp. 1–10, ACM (1999).
- [2] M. Charikar and S. Guha. Improved combinatorial algorithms for the facility location and k-median problems. In *IEEE Foundations of Computer Science*, pages 378–388, 1999.
- [3] L. Chen, H. Hassani, A. Karbasi. Online Continuous Submodular Maximization. In *International Conference on Artificial Intelligence and Statistics, AISTATS 2018*.
- [4] G. Cornuejols, G. L. Nemhauser, and L. A. Wolsey. The uncapacitated facility location problem. In *Pitu B. Mirchandani and Richard L. Francis, editors, Discrete Location Theory*, pages 119–171. John Wiley and Son, Inc., New York, 1990.
- [5] M. Cygan, A. Czumaj, M. Mucha, P. Sankowski. Online Facility Location with Deletions. In *Annual European Symposium on Algorithms, ESA 2018*.
- [6] G.B. Dantzig. Discrete variable extremum problems, In *Operations Research* 5 (1957) 266–277.
- [7] Y. Freund and R. E. Schapire. A decision-theoretic generalization of on-line learning and an application to boosting. In *Journal of Computer and System Sciences*, 55:119–139, 1997.
- [8] T. Fujita, K. Hatano, E Takimoto Combinatorial Online Prediction via Metarounding. In *Algorithmic Learning Theory* (2013), 68-82,
- [9] D. Golovin, A. Krause, M. Streeter. Online Submodular Maximization under a Matroid Constraint with Application to Learning Assignments In *Technical report, arXiv, 2014*.
- [10] K. Jain and V. Vazirani. Approximation algorithms for metric facility location and k-median problems using the primal-dual schema and Lagrangian relaxation. In *J. ACM*, 48(2):274–296, 2001.
- [11] A. Kalai and S. Vempala. Efficient algorithms for online decision problems, In *Journal of Computer and System Sciences*, vol. 71, no. 3, pp. 291–307, 2005.
- [12] S. Kakade, A. Kalai, and K. Ligett. Playing games with approximation algorithms. In *ACM Symposium on the Theory of Computing*, pages 546-555, STOC 2007.
- [13] W.M. Koolen, M.K. Warmuth, J. Kivinen. Hedging structured concepts. In *Conference on Learning Theory*, Omnipress, 2010, pp. 239-254.
- [14] K. C. Kiwiel. Breakpoint searching algorithms for the continuous quadratic knapsack problem. In *Math. Program*, 112(2): 473-491 (2008).
- [15] A. Kumar. Constant factor approximation algorithm for the knapsack median problem. In *ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pp. 824–832, SIAM (2012).
- [16] E. Hazan, and S. Kale. Online submodular minimization. In *Journal of Machine Learning Research (JMLR)*, 2012.
- [17] D. P. Helmbold and S. V. N. Vishwanathan. Online Learning of Combinatorial Objects via Extended Formulation. In *International Conference on Algorithmic Learning Theory, ALT 2018*.
- [18] N. Laoutaris, G. Smaragdakis, K. Oikonomou, I. Stavrakakis, A. Bestavros. Distributed Placement of Service Facilities in Large-Scale Networks. In *IEEE INFOCOM*, pages 2144–2152, 2007.
- [19] S. Martello, P. Toth. An upper bound for the zero-one knapsack problem and a branch and bound algorithm, In *European Journal of Operational Research* 1 (1977) 169–175.
- [20] A. Meyerson. Online Facility Location. In *IEEE Symposium on Foundations of Computer Science, FOCS 2001*.
- [21] D. Shmoys, E. Tardos and K. Aardal. Approximation algorithms for facility location problems. In *ACM Symposium on Theory of Computing (STOC 1997)*, pages 265–274, ACM Press, 1997.
- [22] M. Streeter and D. Golovin. An Online Algorithm for Maximizing Submodular Functions. In *Conference on Neural Information Processing Systems, NIPS 2008*.
- [23] H. Yu, M. J. Neely, X. Wei. Online Convex Optimization with Stochastic Constraints. In *Conference on Neural Information Processing Systems, NIPS 2017*.
- [24] M. Zinkevich. Online convex programming and generalized infinitesimal gradient ascent. In *International Conference on Machine Learning*, 2003.
- [25] S. Pasteris, S. Wang, M. Herbster, T. He. Service Placement with Provable Guarantees in Heterogeneous Edge Computing Systems. To appear in *IEEE INFOCOM*, 2019.

A Missing Proofs

Proof of Theorem 4.2

Proof. Define $z_0 := 0$. We have $X^t = \{\xi_{q,k}^t : q \in \Gamma, k \leq \zeta_q^t + 1\} \setminus \{0\}$ so:

$$\begin{aligned}
 \sum_{i \in X^t} z_i &\leq \sum_{q \in \Gamma, k \leq \zeta_q^t + 1} z_{\xi_{q,k}^t} \\
 &= \sum_{q \in \Gamma} \sum_{k \leq \zeta_q^t + 1} z_{\xi_{q,k}^t} \\
 &\leq \sum_{q \in \Gamma} \sum_{k \leq \zeta_q^t + 1} \tau^{q-1} \beta \\
 &= \sum_{q \in \Gamma} (\zeta_q^t + 1) \tau^{q-1} \beta \\
 &\leq \sum_{q \in \Gamma} (\delta \pi_q^t + 1) \tau^{q-1} \beta \\
 &= \beta \sum_{q \in \Gamma} \tau^{q-1} + \beta \delta \sum_{q \in \Gamma} \pi_q^t \tau^{q-1} \\
 &\leq \beta \sum_{q=1}^{\infty} \tau^{q-1} + \beta \delta \sum_{q \in \Gamma} \pi_q^t \tau^{q-1} \\
 &= \frac{\beta}{1-\tau} + \beta \delta \sum_{q \in \Gamma} \pi_q^t \tau^{q-1} \\
 &= \frac{\beta}{1-\tau} + \beta \delta \sum_{q \in \Gamma} \sum_{i \in \Omega_q} \omega_i^t \tau^{q-1} \\
 &= \frac{\beta}{1-\tau} + \delta \sum_{q \in \Gamma} \sum_{i \in \Omega_q} \omega_i^t \tau^{-1} (\tau^q \beta) \\
 &\leq \frac{\beta}{1-\tau} + \delta \sum_{q \in \Gamma} \sum_{i \in \Omega_q} \omega_i^t \tau^{-1} z_i \\
 &\leq \frac{\beta}{1-\tau} + \frac{\delta}{\tau} \sum_{q \in \Gamma} \sum_{i \in \Omega_q} \omega_i^t z_i \\
 &\leq \frac{\beta}{1-\tau} + \frac{\delta}{\tau} \sum_{i \in [n]} \omega_i^t z_i \\
 &\leq \frac{\beta}{1-\tau} + \frac{\delta}{\tau} \\
 &= \sqrt{\beta} + (1 - \sqrt{\beta}) \\
 &= 1
 \end{aligned}$$

□

Proof of Lemma 5.1

Proof. Given some $d \in D$ let E_d be the event that E happens on draw d . Given some set $D' \subseteq D$ let $E_{D'}$ be the probability that E happens on either of the draws in D' . We prove the result by induction on $|D|$. The inductive hypothesis clearly holds for

$|D| = 0$ as then $\mathbb{P}(E_D) = 0 = 1 - 1 = 1 - \exp(0) = 1 - \exp(-\sum_{d \in D} \rho_d)$.

Now suppose the inductive hypothesis holds for $|D| = l$ for some l . We now show that it holds for $|D| = l + 1$ which will complete the proof. To show this choose some $d \in D$ and set $D' = D \setminus \{d\}$. We have:

$$\begin{aligned}
 \mathbb{P}(\neg E_D) &= \mathbb{P}(\neg E_{D'} \wedge \neg E_d) \\
 &= \mathbb{P}(\neg E_{D'}) \mathbb{P}(\neg E_d) \\
 &= \mathbb{P}(\neg E_{D'}) (1 - \mathbb{P}(E_d)) \\
 &= \mathbb{P}(\neg E_{D'}) (1 - \rho_d)
 \end{aligned} \tag{1}$$

$$\begin{aligned}
 &\leq \mathbb{P}(\neg E_{D'}) \exp(-\rho_d) \\
 &\leq (1 - \mathbb{P}(E_{D'})) \exp(-\rho_d) \\
 &\leq \exp\left(-\sum_{d \in D'} \rho_d\right) \exp(-\rho_d)
 \end{aligned} \tag{2}$$

$$= \exp\left(-\sum_{d \in D} \rho_d\right)$$

Where Equation 1 is due to the independence of the draws and Equation 2 is from the inductive hypothesis (noting $|D'| = l$). We now have

$$\mathbb{P}(E_D) = 1 - \mathbb{P}(\neg E_D) \tag{3}$$

$$\geq 1 - \exp\left(-\sum_{d \in D} \rho_d\right) \tag{4}$$

which proves the inductive hypothesis. □

Proof of Theorem 5.2

Proof. Note that the event $Z \cap X^t \neq \emptyset$ happens if either of $\{\xi_{q,k}^t : q \in \mathbb{N}, k \leq \zeta_q^t\}$ are in Z . For every $q \in \mathbb{N}$ and $k \in [\zeta_q^t]$ we have:

$$\begin{aligned}
 \mathbb{P}(\xi_{q,k}^t \in Z) &= \sum_{i \in Z \cap \Omega_q} \mathbb{P}(\xi_{q,k}^t = i) \\
 &= \sum_{i \in Z \cap \Omega_q} \frac{\omega_i^t}{\pi_q^t}
 \end{aligned}$$

and for $k = \zeta_q^t + 1$ we similarly have:

$$\mathbb{P}(\xi_{q,k}^t \in Z) = (\delta \pi_q^t - \lfloor \delta \pi_q^t \rfloor) \sum_{i \in Z \cap \Omega_q} \frac{\omega_i^t}{\pi_q^t}$$

So letting $\rho_{(q,k)} := \mathbb{P}(\xi_{q,k}^t \in Z)$ we have:

$$\begin{aligned}
 &\sum_{k=1}^{\zeta_q^t + 1} \rho_{(q,k)} \\
 &= \sum_{k=1}^{\zeta_q^t} \sum_{i \in Z \cap \Omega_q} \frac{\omega_i^t}{\pi_q^t} + (\delta \pi_q^t - \lfloor \delta \pi_q^t \rfloor) \sum_{i \in Z \cap \Omega_q} \frac{\omega_i^t}{\pi_q^t}
 \end{aligned}$$

$$\begin{aligned}
 &= \zeta_q^t \sum_{i \in Z \cap \Omega_q} \frac{\omega_i^t}{\pi_q^t} + (\delta \pi_q^t - \lfloor \delta \pi_q^t \rfloor) \sum_{i \in Z \cap \Omega_q} \frac{\omega_i^t}{\pi_q^t} \\
 &= (\zeta_q^t + \delta \pi_q^t - \lfloor \delta \pi_q^t \rfloor) \sum_{i \in Z \cap \Omega_q} \frac{\omega_i^t}{\pi_q^t} \\
 &= \delta \pi_q^t \sum_{i \in Z \cap \Omega_q} \frac{\omega_i^t}{\pi_q^t} \\
 &= \delta \sum_{i \in Z \cap \Omega_q} \omega_i^t
 \end{aligned}$$

So, plugging into Lemma 5.1 with $D := \{(q, k) : q \in \mathbb{N}, k \in [\zeta_q^t + 1]\}$, we get:

$$\begin{aligned}
 \mathbb{P}(Z \cap X^t \neq \emptyset) &= \mathbb{P}(\exists i \in Z : i \in X^t) \\
 &\geq 1 - \exp\left(-\sum_{d \in D} \rho_d\right) \\
 &= 1 - \exp\left(-\sum_{q=1}^{\infty} \sum_{k=1}^{\zeta_q^t+1} \rho_{(q,k)}\right) \\
 &= 1 - \exp\left(-\delta \sum_{q=1}^{\infty} \sum_{i \in Z \cap \Omega_q} \omega_i^t\right) \\
 &= 1 - \exp\left(-\delta \sum_{i \in Z} \omega_i^t\right)
 \end{aligned}$$

□

Proof of Theorem 5.3

Proof. From Theorem 5.2 we have $\mathbb{P}(i \in X^t) = \mathbb{P}(\{i\} \cap X^t \neq \emptyset) \geq 1 - \exp(-\delta \omega_i^t)$. Choosing $q \in \mathbb{N}$ such that $i \in \Omega_q$ we also have:

$$\begin{aligned}
 \mathbb{P}(i \in X^t) &\leq \sum_{k=1}^{\zeta_q^t+1} \mathbb{P}(\xi_{q,k}^t = i) \\
 &= \zeta_q^t \omega_i^t / \pi_q^t + (\delta \pi_q^t - \lfloor \delta \pi_q^t \rfloor) \omega_i^t / \pi_q^t \\
 &= \delta \pi_q^t \omega_i^t / \pi_q^t \\
 &= \delta \omega_i^t
 \end{aligned}$$

□

Proof of Theorem 5.7

Proof. For all $j \in [n-1]$ the function $\delta \sum_{k=1}^j \gamma_{\sigma(t,k)}$ is concave and the function $(1 - \exp(-x))$ is concave and monotonic increasing which implies their combination, $(1 - \exp(-\delta \sum_{k=1}^j \gamma_{\sigma(t,k)}))$, is concave. Hence, as $(r_{\sigma(t,j)}^t - r_{\sigma(t,j+1)}^t) \geq 0$, f_j^t is concave. Similarly, as all components of \mathbf{c}^{t-} are negative, we have that

$\sum_{i=1}^n c_i^{t-} (1 - \exp(-\delta \gamma_i))$ is convex. Since also the function $\mathbf{c}^{t+} \cdot \boldsymbol{\gamma}$ is convex we then have that h^t is a positive sum of convex functions and is therefore convex. □

Proof of Lemma 5.6

Proof. Let $l := \min\{j \in [n] : \sigma(t, j) \in X^t\}$. Note that $r_{\sigma(t,l)}^t = \max_{i \in X^t} r_i^t$.

From the definition of l we have:

- For all $j < l$, $\mathcal{I}(\exists k \leq j : \sigma(t, k) \in X^t) = 0$
- For all $j \geq l$, $\mathcal{I}(\exists k \leq j : \sigma(t, k) \in X^t) = 1$

This implies that:

$$\begin{aligned}
 &\sum_{j=1}^n (r_{\sigma(t,j)}^t - r_{\sigma(t,j+1)}^t) \mathcal{I}(\exists k \leq j : \sigma(t, k) \in X^t) \\
 &= \sum_{j=l}^n (r_{\sigma(t,j)}^t - r_{\sigma(t,j+1)}^t) \\
 &= r_{\sigma(t,l)}^t \\
 &= \max_{i \in X^t} r_i^t
 \end{aligned}$$

□

Proof of Theorem 5.9

Proof. For all $j \in [n]$, Theorem 5.2 with $Z := \{\sigma(t, k) : k \leq j\}$ implies that:

$$\begin{aligned}
 &\mathbb{P}(\exists k \leq j : \sigma(t, k) \in X^t) \\
 &= \mathbb{P}(\{\sigma(t, k) : k \leq j\} \cap X^t \neq \emptyset) \\
 &\geq 1 - \exp\left(-\delta \sum_{k=1}^j \omega_{\sigma(t,k)}^t\right)
 \end{aligned}$$

Lemma 5.8 then gives us:

$$\begin{aligned}
 &\mathbb{E}\left(\max_{i \in X^t} r_i^t\right) \\
 &\geq \sum_{j=1}^n (r_{\sigma(t,j)}^t - r_{\sigma(t,j+1)}^t) \left(1 - \exp\left(-\delta \sum_{k=1}^j \omega_{\sigma(t,k)}^t\right)\right) \\
 &= \sum_{j=1}^n f_j^t(\boldsymbol{\omega}^t)
 \end{aligned}$$

By Theorem 5.3 we also have:

$$\begin{aligned}
 &\sum_{i \in [n]: c_i^t < 0} c_i^t \mathbb{P}(i \in X^t) \\
 &\leq \sum_{i \in [n]: c_i^t < 0} c_i^t (1 - \exp(-\delta \omega_i^t))
 \end{aligned}$$

$$= \sum_{i \in [n]} c_i^{t-} (1 - \exp(-\delta \omega_i^t))$$

and:

$$\begin{aligned} & \sum_{i \in [n]: c_i^t > 0} c_i^t \mathbb{P}(i \in X^t) \\ & \leq \sum_{i \in [n]: c_i^t > 0} c_i^t \delta \omega_i^t \\ & = \sum_{i \in [n]} c_i^{t+} \delta \omega_i^t \\ & = \delta \langle \mathbf{c}^{t+}, \boldsymbol{\omega}^t \rangle \end{aligned}$$

so:

$$\sum_{i=1}^n c_i^t \mathbb{P}(i \in X^t) \leq \delta \langle \mathbf{c}^{t+}, \boldsymbol{\omega}^t \rangle + \sum_{i \in [n]} c_i^{t-} (1 - \exp(-\delta \omega_i^t))$$

Hence, we have:

$$\begin{aligned} & \mathbb{E}(\mu^t(X^t)) \\ & = \mathbb{E}\left(\max_{i \in X^t} r_i^t - \sum_{i \in X^t} c_i^t\right) \\ & = \mathbb{E}\left(\max_{i \in X^t} r_i^t - \sum_{i=1}^n c_i^t \mathcal{I}(i \in X^t)\right) \\ & = \mathbb{E}\left(\max_{i \in X^t} r_i^t\right) - \sum_{i=1}^n c_i^t \mathbb{E}(\mathcal{I}(i \in X^t)) \quad (5) \\ & = \mathbb{E}\left(\max_{i \in X^t} r_i^t\right) - \sum_{i=1}^n c_i^t \mathbb{P}(i \in X^t) \\ & \geq \sum_{j=1}^n f_j^t(\boldsymbol{\omega}^t) - \sum_{i=1}^n c_i^t \mathbb{P}(i \in X^t) \\ & \geq \sum_{j=1}^n f_j^t(\boldsymbol{\omega}^t) - \delta \langle \mathbf{c}^{t+}, \boldsymbol{\omega}^t \rangle - \sum_{i \in [n]} c_i^{t-} (1 - \exp(-\delta \omega_i^t)) \\ & = -h^t(\boldsymbol{\omega}^t) \end{aligned}$$

where Equation 5 is due to linearity of expectation. \square

Proof of Theorem 5.11

Proof. Suppose we have some $l \in [n]$. For $j < l$ we have $\partial_{\sigma(t,l)} f_j^t(\boldsymbol{\omega}^t) = 0$ and for $j \geq l$ we have

$$\begin{aligned} & \partial_{\sigma(t,l)} f_j^t(\boldsymbol{\omega}^t) \\ & = \delta \left(r_{\sigma(t,j)}^t - r_{\sigma(t,j+1)}^t \right) \exp\left(-\delta \sum_{k=1}^j \omega_{\sigma(t,k)}^t\right) \\ & = \delta \left(r_{\sigma(t,j)}^t - r_{\sigma(t,j+1)}^t \right) \epsilon_j^t \end{aligned}$$

This implies that:

$$\begin{aligned} \partial_{\sigma(t,l)} \sum_{j=1}^n f_j^t(\boldsymbol{\omega}^t) & = \delta \sum_{j=l}^n \left(r_{\sigma(t,j)}^t - r_{\sigma(t,j+1)}^t \right) \epsilon_j^t \\ & = \delta \lambda_l^t \end{aligned}$$

which gives us

$$\begin{aligned} & \partial_{\sigma(t,l)} h^t(\boldsymbol{\omega}^t) \\ & = \delta c_{\sigma(t,j)}^{t+} + \delta c_{\sigma(t,j)}^{t-} \exp(-\delta \omega_{\sigma(t,j)}^t) - \delta \lambda_l^t \\ & = g_{\sigma(t,l)}^t \end{aligned}$$

\square

Proof of Theorem 5.18

Proof. From Lemma 5.16 we have:

$$\begin{aligned} \sum_{t=1}^T \eta^t \|\mathbf{g}^t\|^2 & \leq \sum_{t=1}^T n \sqrt{s/(2t)} \\ & = n \sqrt{s/2} \sum_{t=1}^T \frac{1}{\sqrt{t}} \\ & \leq n \sqrt{s/2} \left(1 + \int_{t=1}^T \frac{1}{\sqrt{t}} \right) \\ & = n \sqrt{s/2} (2\sqrt{T} - 1) \\ & \leq 2n \sqrt{sT/2} \end{aligned}$$

So using Lemma 5.17 and plugging into Theorem 5.14 we have:

$$\begin{aligned} \sum_{t=1}^T (h^t(\boldsymbol{\omega}^t) - h^t(\boldsymbol{\phi})) & \leq \frac{1}{2} n \sqrt{2sT} + n \sqrt{sT/2} \\ & = 2n \sqrt{sT/2} \\ & = n \sqrt{2sT} \end{aligned}$$

By Definition 5.15 and Lemma 5.12 we have $s \leq \delta^2(\hat{r} + \hat{c})^2$. Plugging this into the above gives us the result. \square

B When $\max_{i \in [n]} z_i$ is large

When $\max_{i \in [n]} z_i$ is large we construct X^t from ω_i^t in the following way:

1. Define $A = \{i \in [n] : z_i \in [1/2, 1]\}$
2. Define $a^t := \sum_{i \in A} \omega_i^t / 4$
3. Flip a biased coin with probability of heads equal to a^t
4. If heads sample i from A with probability $\omega_i^t / (4a_t)$ and set $X^t = \{i\}$
5. If tails define $\beta := 1/2$ and run Algorithm 1.

We defer the analysis.

C On the Knapsack Problem and Component Hedge

In this section we show that Component Hedge cannot solve the online knapsack problem special case in polynomial time unless $P = NP$.

The Component Hedge algorithm described in [13] requires a set of concepts $\mathcal{X} \subseteq \{0, 1\}^n$. In the knapsack problem, \mathcal{X} is the subset of all \mathbf{x} that fit in the knapsack, i.e. those $\mathbf{x} \in \{0, 1\}^d$ with $\langle \mathbf{x}, \mathbf{z} \rangle \leq 1$. For Component Hedge, the convex hull \mathcal{H} of \mathcal{X} must have a number of constraints polynomial in n and that any $\mathbf{y} \in \mathcal{H}$ can, in polynomial time, be decomposed into a convex combination of $n + 1$ concepts. We now show that if this is true then we can solve the knapsack problem in polynomial time:

Let \mathbf{c} be the vector defining the objective function of the knapsack problem (i.e. we seek the $\mathbf{x} \in \mathcal{X}$ that maximises $\langle \mathbf{c}, \mathbf{x} \rangle$). Now, since \mathcal{H} has polynomial constraints we can efficiently choose the maximiser of $\langle \mathbf{c}, \mathbf{y} \rangle$ (for $\mathbf{y} \in \mathcal{H}$) via linear programming. Let this maximiser be \mathbf{y} . Now decompose \mathbf{y} into a convex combination, $\sum_{i=1}^{n+1} (m_i \mathbf{x}^i)$, of $n + 1$ concepts \mathbf{x}^i (where $\sum_{i=1}^{n+1} m_i = 1$). Now suppose, for contradiction, that $\langle \mathbf{c}, \mathbf{x}^j \rangle < \langle \mathbf{c}, \mathbf{y} \rangle$ for some j with $m_j > 0$. Without loss of generality let $j = n + 1$.

Then, since for all $\mathbf{x} \in \mathcal{X}$ with have $\langle \mathbf{c}, \mathbf{x} \rangle \leq \langle \mathbf{c}, \mathbf{y} \rangle$ (as $\mathcal{X} \subset \mathcal{H}$) we have: $\langle \mathbf{c}, \mathbf{y} \rangle = \langle \mathbf{c}, (\sum_{i=1}^{n+1} m_i \mathbf{x}^i) \rangle = \sum_{i=1}^{n+1} m_i \langle \mathbf{c}, \mathbf{x}^i \rangle = \sum_{i=1}^n m_i \langle \mathbf{c}, \mathbf{x}^i \rangle + m_{n+1} \langle \mathbf{c}, \mathbf{x}_{n+1} \rangle < \sum_{i=1}^n m_i \langle \mathbf{c}, \mathbf{x}^i \rangle + m_{n+1} \langle \mathbf{c}, \mathbf{y} \rangle \leq \sum_{i=1}^n m_i \langle \mathbf{c}, \mathbf{y} \rangle + m_{n+1} \langle \mathbf{c}, \mathbf{y} \rangle = \sum_{i=1}^{n+1} m_i \langle \mathbf{c}, \mathbf{y} \rangle = \langle \mathbf{c}, \mathbf{y} \rangle$ which is a contradiction.

Hence, for all j with $m_j > 0$ we have $\langle \mathbf{c}, \mathbf{x}^j \rangle = \langle \mathbf{c}, \mathbf{y} \rangle$ so the maximising \mathbf{x} is found in polynomial time. This contradicts the assumption that $P \neq NP$.

D Using Submodular Maximisation

We consider the facility location special case (i.e. $\mathbf{z} := \mathbf{0}$ and all costs are positive). In this case it is possible to obtain our bound by selecting, on a trial t , each action i with probability ω_i^t . We can then use some definitions in our paper to define the following functions on a trial t :

- For all $j \in [n]$ define $\hat{f}_j^t(\gamma) := \left(r_{\sigma(t,j)}^t - r_{\sigma(t,j+1)}^t \right) \left(1 - \prod_{k=1}^j (1 - \gamma_{\sigma(t,k)}) \right)$
- Define $\hat{h}^t(\gamma) = \sum_{j=1}^n \hat{f}_j^t(\gamma) - \langle \mathbf{c}^t, \gamma \rangle$

We can then use the arguments in our paper to write the expected profit as: $\mathbb{E}(\mu^t(X^t)) = \hat{h}^t(\omega^t)$ which is

continuous submodular. We can use a slight modification of our method of computing the gradient of h^t to compute the gradient of \hat{h}^t in quasi-linear time. With this in hand we can then plug the gradient into one of the algorithms of [3]:

- The first algorithm requires the submodular function to be monotone and non-negative whereas ours, in general, is neither. Even if our expected profit was monotone and non-negative, for their algorithm to have a regret linear in \sqrt{T} their per-trial time complexity bound becomes $\Omega(n\sqrt{T})$ which is much larger than ours.
- In the gradient ascent based second algorithm, the submodular function also needs to be monotone and non-negative but the cost vector can be separated from the reward vector in the analysis allowing us to use the expected maximum reward which is monotone non-negative. The result of this analysis is similar to ours expect that for them $\alpha = 1/2$ instead of the better $\alpha = 1 - e^{-1}$ that we have (NB: α is the discount on the comparator selection S).