# Constant Approximation for Stochastic Orienteering Problem with ( $1+\epsilon$ )-Budget Relaxiation 

Yiyao Jiang<br>Institute for Interdisciplinary Information Sciences, Tsinghua University

## 1 Introduction

In the competitive practice of orienteering, participants are given a map with marked locations. Starting from a given point, they must try to visit as many locations as possible within a given time limit. Each point has a given reward, and the goal for competitors is to maximize the total reward within the time limit. This problem was first studied in [8] from an algorithmic perspective; see [13] for a more recent survey.

A similar problem to Orienteering Problem (OP) is the Prize Collecting Traveling Salesman Problem (PCTSP) [5]. In the latter, the salesman's costs are the time spent traveling between points, plus a point-dependent penalty that is incurred for each point that he fails to visit. PCTSP's goal is to minimize the traveling time plus the penalties of the points. Comparing with OP, PCTSP is not limited in time, whereas OP is strictly bounded with time limit.

### 1.1 The (Simple) Orienteering Problem

An instance of the Orienteering Problem (OP) consists of a metric space ( $V, d$ ) with $|V|=n$ points and distances $d(u, v) \in \mathbb{R}^{+}$for each $(u, v) \in V \times V$; moreover, the instance specifies a starting point $\rho \in V$, a total time budget $B$, and a reward $r_{v} \in \mathbb{R}^{+}$for each $v \in V$. The object is to chart a path that maximizes the sum of rewards $r_{v}$ while staying within the time budget $B$, where the distances $d(u, v)$ are interpreted as time costs. In this setting, and as all quantities are real numbers, we can assume without loss of generality that $B=1$ after normalization.

In 1998 Arkin et al. [1] obtained a 2-approximation for OP with planar Euclidean distances. Furthermore, Chen and Har-Peled [7] obtained a polynomialtime approximation scheme (PTAS) for OP in fixed-dimensional Euclidean space.

For general metric spaces $(V, d)$, Blum et al. [2] obtained the first constantfactor approximation (with factor 4) in 2007. The best result until now is a $(2+\epsilon)$-approximation algorithm by Chekuri et al. [6].

The Orienteering Problem finds many natural applications. For example, a traveling salesman with limited time [12] (who must maximize his total profit in one day).

### 1.2 Stochastic Orienteering Problem

In the Stochastic Orienteering Problem (SOP) [9] a job is associated to each vertex $v \in V$, which must be completed at $v$ in order to obtain the reward $r_{v}$. The time necessary to complete the job, moreover, is randomly distributed. For example, a salesman may need to wait for his customer to come downstairs, or a tourist need time to enjoy the scenery, and so on.

Generally, one is not allowed to give up a job until it is finished. In certain models one is allowed to give up a job and its reward after starting the job, while being forbidden from re-trying to complete the job later, once it has been given up.

The time required by job $v$ is a random variable $S_{v}$, where $S_{v}$ has a known discrete distribution $\pi_{v}$ for each $v \in V$, where $\pi_{v}$ modeled as a function from $\mathbb{R}^{+}$ into $[0,1]$; the distributions $\left\{\pi_{v}: v \in V\right\}$ are thus part of an instance of SOP.

We note that SOP can be considered as a kind of combination of OP and of the Stochastic Knapsack Problem (SKP) [11]. The Stochastic Knapsack Problem, indeed, corresponds to an instance of SOP in which $d(u, v)=0$ for all $(u, v) \in V \times V$. For SKP, moreover, a $(2+\epsilon)$-approximation algorithm is known [4].

For a non-adaptive variant of SOP in which the path is chosen in advance and in which the distances $d(u, v)$ are restricted to integers, Gupta et al. [9] have shown a constant factor approximation algorithm. In that setting (and due to the integer restriction on distances) the time budget $B$ is not normalized to 1 , and remains a parameter of the problem. In fact, Gupta et al. show that their algorithm is an $O(\log \log B)$-approximation algorithm for SOP with arbitrary adaptive policies. In the same formal setting, moreover, Bansal and Nagarajan [3] have shown that a $\Omega\left((\log \log B)^{1 / 2}\right)$ gap between adaptive and non-adaptive policies is unavoidable. This contrasts with the case of SKP, for which nonadaptive policies can approximate adaptive policies to within a constant factor [4]. (In OP the distinction between adaptive and non-adaptive policies is moot.)

In this paper we present an $O\left(\epsilon^{-14}\right)$-approximation for the (original, adaptive) SOP with time budget $(1+\epsilon) B=1+\epsilon, 0<\epsilon<1$. More precisely, we show a non-adaptive policy of time budget $1+\epsilon$ such that the expected reward of the optimal adaptive policy at time budget 1 is at most $O\left(\epsilon^{-14}\right)$ times the expected reward of our [non-adpative, time $1+\epsilon$ ] policy. Our algorithm therefore achieves an $O(1)$-approximation for any constant $0<\epsilon<1$.

Theorem 1. There is a polynomial time, non-adaptive $O\left(\epsilon^{-14}\right)$-approximation for $S O P$ with respect to a relaxed time budget of $(1+\epsilon) B, 0<\epsilon<1$.

## 2 Definitions and Notations

We write $\mathbb{R}^{+}$for the nonnegative real numbers $[0, \infty)$. In this paper all quantities are real-valued.

### 2.1 Orienteering

As introduced in Section 1, an instance of the Orienteering Problem (OP) is defined on an underlying metric space $(V, d)$ with $|V|=n$ points and distances $d(u, v) \in \mathbb{R}^{+}$for each $(u, v) \in V \times V$. We are given a starting "root" point $\rho \in V$ and a total time budget $B$. An instance of OP thus corresponds to a tuple $I_{\mathrm{O}}=\left(V, d,\left\{r_{v}\right\}_{v \in V}, B, \rho\right)$.

### 2.2 Stochastic Orienteering

In the Stochastic Orienteering Problem (SOP) each point $v \in V$ is related to a unique stochastic job. The job associated to $v$ has a fixed reward $r_{v} \geq 0$, and a random variable of processing time size $S_{v}$ with a known but arbitrary discrete probability distribution $\pi_{v}: U \rightarrow[0,1]$, where $U \subseteq \mathbb{R}^{+}$is a countable set, such that $\sum_{t \in U} \pi_{v}(t)=1$ for all $v \in V$. (Time size is the time that we have to wait at the point before receiving the reward for job $v$. Note that we can take the same underlying time set $U$ for all $v \in V$, wlog.) An instance of SOP thus corresponds to a tuple $I_{\text {SO }}=\left(\pi, V, d,\left\{r_{v}\right\}_{v \in V},\left\{S_{v}\right\}_{v \in V}, B, \rho\right)$.

At each step we travel to a point $v$ from a current point $u$ (at time cost $d(u, v))$ and then process the job $v$, at time $\operatorname{cost} S_{v}$, and once a job $v$ is selected, one must wait for $v$ to complete. If we are still within the time budget $B$ when the job is completed we receive the reward $r_{v}$, and choose a new point as the next destination.

For generality, $\forall t \in U$ we assume $t \in[0, B]$. If one time size $t>B$, we can just truncated it to $B$, since we will always stop at time budget $B$ before completing the job and never get the reward. Also, we can assume wlog that $d(u, \rho) \leq B$ for all $u \in V$, which also implies that $d(u, v) \in[0,2 B]$ for all $u, v \in V$.

### 2.3 Task Orienteering Problem

Our analysis refers to a simplified version of the Stochastic Orienteering Problem known as the "Task Orienteering Problem" (TOP); TOP corresponds to an instance of SOP in which each time size $S_{v}$ is deterministic. An instance of TOP thus corresponds to a tuple $I_{\mathrm{TO}}=\left(V, d,\left\{r_{v}\right\}_{v \in V},\left\{s_{v}\right\}_{v \in V}, B, \rho\right)$ with $s_{v} \in \mathbb{R}^{+}$ for each $v \in V$.

## 3 Algorithm

We begin with an analysis of the Task Orienteering Problem described in section 2.3 , to be used as a component in our main analysis.

### 3.1 An Algorithm for TOP

We first provide a polynomial time algorithm AlgTO (Algorithm 1) and then prove that AlgTO gives an $O(1)$-approximation to TOP.

Definition 1. (Valid OP Instance) Given an instance $I_{\mathrm{TO}}=\left(V, d,\left\{r_{v}\right\}_{v \in V},\left\{s_{v}\right\}_{v \in V}, B, \rho\right)$ of TOP, we define the following valid OP instance $I_{\mathrm{O}}\left(I_{\mathrm{TO}}\right):=\left(V^{\prime}, d^{\prime},\left\{r_{v}\right\}_{v \in V^{\prime}}, B, \rho\right)$, where

1. $V^{\prime}:=\left\{v \in V: d(\rho, v)+s_{v} \leq B\right\}$;
2. $d^{\prime}(u, v):=d(u, v)+s_{v} / 2+s_{u} / 2$ for all $u, v \in V^{\prime}, u \neq v$;
3. $d^{\prime}(u, u):=0$ for all $u \in V$;
4. $r_{v}^{\prime}:=r_{v}$ for all $v \in V^{\prime}$.

Lemma 1. The function $d^{\prime}: V^{\prime} \times V^{\prime} \rightarrow \mathbb{R}^{+}$constructed in Definition 1 is a metric.

The proof is easy.

| Algorithm | 1 | AlgTO | for | Task | Orienteering | on | input |
| :--- | :--- | :--- | :--- | :--- | :--- | :--- | :--- |
| $I_{\mathrm{TO}}$ | $=$ |  |  |  |  |  |  | $\frac{\left(V, d,\left\{r_{v}\right\}_{v \in V},\left\{s_{v}\right\}_{v \in V}, B, \rho\right)}{\text { 1. let } I_{O}\left(I_{\mathrm{TO}}\right):=\left(V^{\prime}, d^{\prime},\left\{r_{v}^{\prime}\right\}_{v \in V^{\prime}}, B, \rho\right) \text {; }}$

2. run AlgOrient on the (Simple) Orienteering Problem $I_{O}$,
get the path $P$, the reward $\mathrm{OPT}_{\mathrm{O}}^{\prime}=\sum_{v \in P} r_{v}^{\prime}$, and the ending point $\rho^{\prime}$;
3. compare $\sum_{v \in P \backslash\left\{\rho^{\prime}\right\}} r_{v}^{\prime}$ and $r_{\rho^{\prime}}^{\prime}$;
4. output the path

$$
P^{\prime}:= \begin{cases}\rho \longrightarrow \rho^{\prime}, & \text { if } r_{\rho^{\prime}}^{\prime} \geq \sum_{v \in P \backslash\left\{\rho^{\prime}\right\}} r_{v}^{\prime}, \\ P \backslash\left\{\rho^{\prime}\right\}, & \text { otherwise }\end{cases}
$$

6. and the reward $\mathrm{OPT}_{\mathrm{TO}}^{\prime}:=\sum_{v \in P^{\prime}} r_{v}^{\prime}$.

Theorem 2. Algorithm AlgTO (Algorithm 1) is a polynomial time $O(1)$-approximation for the Task Orienteering Problem.

Proof. Assume the optimal result for Task Orienteering problem $I_{\mathrm{TO}}$ is $\mathrm{OPT}_{\mathrm{TO}}$. Now we prove that $\mathrm{OPT}_{\mathrm{TO}}^{\prime}$ returning by the polynomial time algorithm AlgTO (Algorithm 1) is an $\Omega(1)$-approximation to $\mathrm{OPT}_{\mathrm{TO}}$.

In Algorithm 1 line 1, as in Definition 1 item 1, we delete all points $v \in V$ which are obviously not in the optimal path of $I_{\mathrm{TO}}$, because even we only reach to one point $v$ from the starting point $\rho$, we still do not have enough time to reach the point $v$ and complete the job with time $d(\rho, v)+s_{v}>B$.

In Definition 1 item 2, for each point $v \in V^{\prime}$, we divide the fixed job time $s_{v}$ by 2 , and add $s_{v} / 2$ to the lengths of all edges adjacent to point $v$. (Note that, although we add $s_{v} / 2$ to the lengths of all edges adjacent to point $v$, in real algorithm path, we only count twice of $s_{v} / 2$ as doing the job in time $s_{v}$ : the edge into point $v$, and the edge out of point $v$ ). Define the new distance as $d^{\prime}(u, v)$.

In Lemma 1, we prove that $\left(V^{\prime}, d^{\prime}\right)$ is a metric, thus we create an Orienteering Problem instance $I_{\mathrm{O}}=\left(V^{\prime}, d^{\prime},\left\{r_{v}\right\}_{v \in V}, B, \rho\right)$ successfully in Algorithm 1 step 1.

In Algorithm 1 line 2, we use the polynomial time algorithm AlgOrient (by Blum et al. [2] mentioned in Section 2.1) on OP instance $I_{\mathrm{O}}$. And in line 3, the path $P$ of reward $\mathrm{OPT}_{\mathrm{O}}^{\prime}$ is a constant approximation of the optimal solution for $I_{\mathrm{O}}$ by algorithm AlgOrient. Note that any path in $I_{\mathrm{TO}}$ including the optimal path of $I_{\mathrm{TO}}$ is surely a feasible path in $I_{\mathrm{O}}$ (since the only change is the distances in $I_{\mathrm{TO}}$ are longer), so $\mathrm{OPT}_{\mathrm{O}}^{\prime}=\Omega\left(\mathrm{OPT}_{\mathrm{TO}}\right)$.

Note that in path $P$, we always get in and out of a point $v$ except for the starting point $\rho$ and ending point $\rho^{\prime}$. So the $s_{v} / 2$ are always counted twice, and in $I_{\text {TO }}$ we do complete all the jobs on the path except for $\rho$ and $\rho^{\prime}$. Thus the only incidence we cannot complete path $P$ in $I_{\mathrm{TO}}$ is the jobs on $\rho$ and $\rho^{\prime}$ in path $P$.

At the starting point $\rho$ in $I_{\mathrm{TO}}$, we can always create a fake "starting point" $\rho^{\star}$ which is the same $d^{\prime}($,$) as \rho$ but $d^{\prime}\left(\rho, \rho^{\star}\right)=d^{\prime}\left(\rho^{\star}, \rho\right)=d^{\prime}\left(\rho^{\star}, \rho^{\star}\right)=0, s_{\rho^{\star}}=0$ and $r_{\rho^{\star}}=0$, and we set $\rho$ as a normal point. Run AlgTO (Algorithm 1) on $I_{\mathrm{TO}}$ with $\rho^{\star}$. If $\rho$ is still in path, we just start from $\rho$ and use $s_{\rho}$ time to get the reward $r_{\rho}$ then go through path $P$; else, $\rho$ is not in path, we also start from $\rho$, but give up the job and reward on $\rho$, then go through path $P$. All these changes do not change the $\mathrm{OPT}_{\mathrm{O}}^{\prime}$, and do not effect on $\mathrm{OPT}_{\mathrm{TO}}$. Thus, we now only need to concern about $\rho^{\prime}$.

In Algorithm 1 line 4, we compare the reward $r_{\rho^{\prime}}^{\prime}$ at the ending point $\rho^{\prime}$ (note that $d\left(\rho, \rho^{\prime}\right)+s_{\rho^{\prime}} \leq B$ in Definition 1 indicates that path $\rho \longrightarrow \rho^{\prime}$ is feasible in $I_{\mathrm{TO}}$ ), and the rewards $\sum_{v \in P \backslash\left\{\rho^{\prime}\right\}} r_{v}^{\prime}$ of all the points in path $\widehat{P}$ except for $\rho^{\prime}$. Since $\mathrm{OPT}_{\mathrm{O}}^{\prime}=\sum_{v \in P} r_{v}^{\prime}=r_{\rho^{\prime}}^{\prime}+\sum_{v \in P \backslash\left\{\rho^{\prime}\right\}} r_{v}^{\prime}, \max \left\{r_{\rho^{\prime}}^{\prime}, \sum_{v \in P \backslash\left\{\rho^{\prime}\right\}} r_{v}^{\prime}\right\} \geq$ $\mathrm{OPT}_{\mathrm{O}}^{\prime} / 2=\Omega\left(\mathrm{OPT}_{\mathrm{TO}}\right)$.

In Algorithm 1 line 5 and 6, we thus choose the larger one between $r_{\rho^{\prime}}^{\prime}$ and $\sum_{v \in P \backslash\left\{\rho^{\prime}\right\}} r_{v}^{\prime}$ as $\mathrm{OPT}_{\mathrm{TO}}^{\prime}$. If $r_{\rho^{\prime}}$ is larger, then we just go to a single point $\rho^{\prime}$ and complete the job; if the other is larger, we just ignore $\rho^{\prime}$, go and do the jobs through the path $P$ except for $\rho^{\prime}$.

Thus, we have a polynomial time $O(1)$-approximation algorithm AlgTO for Task Orienteering problem.

### 3.2 The Algorithm for Stochastic Orienteering

Definition 2. (Valid TOP Instance) Given an instance $I_{\text {SO }}=\left(\pi, V, d,\left\{r_{v}\right\}_{v \in V},\left\{S_{v}\right\}_{v \in V}, B, \rho\right)$ of SOP, and a value $\epsilon>0$, we define the following valid TOP instance with parameter $\epsilon$ to be $I_{\mathrm{TO}}\left(\epsilon, I_{\mathrm{SO}}\right):=\left(\widehat{V}, \widehat{d},\left\{\widehat{r}_{u}\right\}_{\forall u \in \widehat{V}},\left\{\widehat{s}_{u}\right\}_{\forall u \in \widehat{V}}, \widehat{B}, \widehat{\rho}\right)$, where we

1. define $\widehat{V}$ : let $\widehat{V}:=V$,
for all $u \in V$, if $d(\rho, u)>B, \widehat{V}:=\widehat{V} \backslash\{u\}$,
create a virtual point $\widehat{\rho}, \widehat{V}:=V \cup\{\widehat{\rho}\} ;$
2. for all $u, v \in \widehat{V}$, define distances $\widehat{d}(u, v):$ if $u, v \in V, \widehat{d}(u, v):=d(u, v)$,
for all $u \in V$, set $\widehat{d}(u, \widehat{\rho})=\widehat{d}(\widehat{\rho}, u)=B$,
set $d(\widehat{\rho}, \widehat{\rho})=0$;
3. for all $u \in \widehat{V}$, define rewards $\widehat{r}_{u}$ : if $u \in V, \widehat{r}_{u}=r_{u}$, set $\widehat{r}_{\widehat{\rho}}=0$;
4. for all $u \in \widehat{V}$, define deterministic job times $\widehat{s}_{u}$ : if $u \in V$, $\widehat{s}_{u}=E\left[S_{u}^{\prime}\right]$, where $S_{u}^{\prime}=\min \left\{S_{u}, B\right\}$,
set $\widehat{s}_{\widehat{\rho}}=0$;
5. define time bound $\widehat{B}$ : let $\widehat{B}:=\left(1+\epsilon^{13}\right) B$;
6. define the starting point: set $\widehat{\rho}$ to be the new starting point.

Lemma 2. The $(\widehat{V}, \widehat{d})$ of the $T O P$ instance $I_{\mathrm{TO}}$ in Definition 2 is a metric space.

The proof is easy.

```
Algorithm 2 Algorithm AlgSO for SOP on input \(I_{S O}=\)
\(\left(\pi, V, d,\left\{r_{v}\right\}_{v \in V},\left\{S_{v}\right\}_{v \in V}, B, \rho\right)\) and parameter \(0<\epsilon<1\)
    for all \(v \in V\) do
        let \(R_{v}:=r_{v} \cdot \operatorname{Pr}_{S_{v} \sim \pi_{v}}\left[S_{v} \leq(B-d(\rho, v))\right]\) be the expected reward of the single-
    node tour from \(\rho\) to \(v\);
    w.p. \(1 / 2\), just visit the point \(\widehat{v}\) with the highest \(R_{\widehat{v}}\) and exit.
    let \(I_{\mathrm{TO}}\left(\epsilon, I_{\mathrm{SO}}\right):=\left(\widehat{V}, \widehat{d},\left\{\widehat{r}_{u}\right\}_{\forall u \in \widehat{V}},\left\{\widehat{s}_{u}\right\}_{\forall u \in \widehat{V}}, \widehat{B}, \widehat{\rho}\right)\);
    run AlgTO (Algorithm 1) on the valid TOP instance \(I_{\mathrm{TO}}\left(\epsilon, I_{\mathrm{SO}}\right)\),
        get the path \(P\), the reward \(\widehat{\mathrm{OPT}_{\mathrm{TO}}}\);
    replace the starting point \(\widehat{\rho}\) in path \(P\) of \(I_{\mathrm{TO}}\left(\epsilon, I_{\mathrm{SO}}\right)\) by \(\rho\) and output path \(\widehat{P}\);
    go through path \(\widehat{P}\) in \(I_{\text {SO }}\) with time budget \((1+\epsilon) B\) and output the total reward
        \(\widehat{\mathrm{OPT}_{\mathrm{SO}}}\).
```

We want to prove that Algorithm 2 is a polynomial time $O\left(\epsilon^{-14}\right)$-approximation for Stochastic Orienteering problem with respect to a relaxed time budget of $(1+\epsilon) B, 0<\epsilon<1$.

Define the expected reward on optimal adaptive policy on $I_{\text {SO }}$ as OPT, and define the expected reward on optimal policy on $I_{\mathrm{TO}}\left(\epsilon, I_{\mathrm{SO}}\right)$ as $\mathrm{OPT}_{\mathrm{TO}}$.

We design to have either a single-point tour $\rho \longrightarrow v$ of expected reward $R_{v}=$ $\Omega(\mathrm{OPT})$ with $50 \%$ chance, or the path $\widehat{P}$ on $I_{\mathrm{SO}}$ of reward $\widehat{\mathrm{OPT}_{\mathrm{SO}}}=\Omega(\mathrm{OPT})$ with $50 \%$ chance. For this purpose, we want the following theorem.

Theorem 3. Given an instance $I_{\text {SO }}$ for which an optimal adaptive strategy has an expected reward of OPT, either there is a single-point tour with expected reward $\Omega(\mathrm{OPT})$, or the valid Task Orienteering instance $I_{\mathrm{TO}}\left(\epsilon, I_{\mathrm{SO}}\right)$ has reward $\mathrm{OPT}_{\mathrm{TO}}=\Omega(\mathrm{OPT})$, with constant approximation parameter $\epsilon^{14}$.

We will prove Theorem 3 later in Section 4. For now, let us assume that Theorem 3 is correct, and use it to complete the proof of Theorem 1. The following proof is for Theorem 1.

### 3.3 Proof for Theorem 1

Suppose we enter Algorithm 2 line 4, and in line 6, AlgTO (Algorithm 1) finds a path $P=\left(\widehat{\rho}, v_{1}, v_{2}, \ldots, v_{k}\right)$ with reward $\widehat{\mathrm{OPT}_{\mathrm{TO}}}$. And then we get the path $\widehat{P}=\left(\rho, v_{1}, v_{2}, \ldots, v_{k}\right)$.

Lemma 3. In $I_{\mathrm{SO}}=\left(\pi, V, d,\left\{r_{v}\right\}_{v \in V},\left\{S_{v}\right\}_{v \in V}, B, \rho\right)$, for any point $v_{i} \in \widehat{P}$, the probability of successfully reaching to $v$ and finishing the job at $v$ before violating the budget $(1+\epsilon) B$ is at least $1-\epsilon^{13}$.

Proof. Note that now we relax the time budget on $I_{\text {SO }}$ to $(1+\epsilon) B$, but we still have time budget $\left(1+\epsilon^{13}\right) B$ on $I_{\mathrm{TO}}\left(\epsilon, I_{\mathrm{SO}}\right)$ as in Definition 2.

In Definition 2, for $I_{\mathrm{TO}}\left(\epsilon, I_{\mathrm{SO}}\right)$, $\sum_{i=1}^{k}\left(\widehat{s}_{v_{i}}+\widehat{d}\left(v_{i-1}, v_{i}\right)\right) \leq\left(1+\epsilon^{13}\right) B, S_{v_{i}}^{\prime}=$ $\min \left\{S_{v_{i}}, B\right\}$, and $\widehat{s}_{v_{i}}=E\left[S_{v_{i}}^{\prime}\right]$. Since $\widehat{d}\left(\widehat{\rho}, v_{1}\right)=B, E\left[S_{v_{1}}^{\prime}\right]+\sum_{i=2}^{k}\left(E\left[S_{v_{i}}^{\prime}\right]+\right.$ $\left.\widehat{d}\left(v_{i-1}, v_{i}\right)\right) \leq \epsilon^{13} B$.

Define $S_{v_{i}}^{\star}= \begin{cases}S_{v_{1}}^{\prime} & \text { if } i=1, \\ S_{v_{i}}^{\prime}+\widehat{d}\left(v_{i-1}, v_{i}\right) & \text { if } 2 \leq i \leq k .\end{cases}$
Then $\sum_{i=1}^{k} E\left[S_{v_{i}}^{\star}\right] \leq \epsilon^{13} B$. By using Markov's inequality, for each $j \leq k$, we have

$$
\operatorname{Pr}\left[\sum_{i=1}^{j} S_{v_{i}}^{\star} \leq \epsilon B\right] \geq 1-\epsilon^{13}
$$

The change from $S_{u}$ to $S_{u}^{\prime}$ does not effect on optimal adaptive policy on $I_{\text {SO }}$, the distance $d\left(\rho, v_{1}\right) \leq B$, and $d\left(v_{i-1}, v_{i}\right)=\widehat{d}\left(v_{i-1}, v_{i}\right)$. So in path $\widehat{P}$ of $I_{\text {SO }}$ with time budget $(1+\epsilon) B$, for each point $v_{j}(1 \leq j \leq k)$, the probability we successfully reach to point $v_{j}$, and complete the job on point $v_{j}$ is not less than $\operatorname{Pr}\left[\sum_{i=1}^{j} S_{v_{i}}^{\star} \leq \epsilon B\right]$. Thus the probability of successfully reaching to $v$ and finishing the job at $v$ before violating the budget $(1+\epsilon) B$ is at least $1-\epsilon^{13}$. And this complete the proof for Lemma 3.

Proof (The proof of Theorem 1). We prove that AlgSO (Algorithm 2) is a polynomial time $O\left(\epsilon^{-14}\right)$-approximation for Stochastic Orienteering problem with respect to a relaxed time budget of $(1+\epsilon) B, 0<\epsilon<1$.

We assume that Theorem 3 is correct. Then either there is a single-point tour with expected reward $\Omega(\mathrm{OPT})$, or the valid Task Orienteering instance $I_{\mathrm{TO}}\left(\epsilon, I_{\mathrm{SO}}\right)$ has reward $\mathrm{OPT}_{\mathrm{TO}}=\Omega(\mathrm{OPT})$, with a constant approximation parameter $\epsilon^{14}$.

1. There is a single-point tour $\rho \longrightarrow v$ satisfies $R_{v}=\Omega\left(\epsilon^{14} \mathrm{OPT}\right)$.

We have $50 \%$ chance to obtain the highest $R_{\widehat{v}}$. And $R_{\widehat{v}} \geq R_{v}=\Omega\left(\epsilon^{14} \mathrm{OPT}\right)$.
2. The valid Task Orienteering instance $I_{\mathrm{TO}}\left(\epsilon, I_{\mathrm{SO}}\right)$ has reward $\mathrm{OPT}_{\mathrm{TO}}=$ $\Omega(\mathrm{OPT})$.
We have $50 \%$ chance to get $\widehat{\mathrm{OPT}_{\mathrm{SO}}}$.
In Algorithm 2 line 6 , by applying Theorem 2, AlgTO (Algorithm 1) will find a path $P=\left(\widehat{\rho}, v_{1}, v_{2}, \ldots, v_{k}\right)$ with reward $\widehat{\mathrm{OPT}_{\mathrm{TO}}}=\Omega\left(\mathrm{OPT}_{\mathrm{TO}}\right)$.

By assuming Theorem 3 is correct, we have $\mathrm{OPT}_{\mathrm{TO}}=\Omega\left(\epsilon^{14} \mathrm{OPT}\right)$. Then $\widehat{\mathrm{OPT}_{\mathrm{TO}}}=\Omega\left(\mathrm{OPT}_{\mathrm{TO}}\right)=\Omega\left(\epsilon^{14} \mathrm{OPT}\right)$.
Now applying Lemma 3 on path $\widehat{P}$ that we get in Algorithm 2 line 7, and that gives us an expected reward of at least $\widehat{\mathrm{OPT}_{\mathrm{SO}}}=\widehat{\mathrm{OPT}_{\mathrm{TO}}} \cdot\left(1-\epsilon^{13}\right)=$ $\Omega\left(\epsilon^{14} \mathrm{OPT}\right)$.

So we always have at least $50 \%$ probability to get $\Omega\left(\epsilon^{14} \mathrm{OPT}\right)$ by Algorithm 2 , and thus we complete the proof of Theorem 1.

## 4 Optimal Policy for Stochastic Orienteering

The main method to deal with the huge optimal decision tree defined in Section 4.1 is Discretization and Block Policy [10]. Since all quantities are all real-valued, by scaling, we can assume $B=1$ and the size of each item is distributed between 0 and $B$. The relaxed time budget $B+O(\varepsilon)$ should be less than $2 B$.

### 4.1 Policy and Decision Tree

A policy $\sigma$ for an SOP instance $(\pi, V, d, B)$ can be represented as a decision tree $T_{\sigma}(\pi, V, d, B)$.


Fig. 1. Decision tree $T_{\sigma}$

In Figure 1, each node of $T_{\sigma}$ is labeled by a vertex $v \in V$, with the root of $T_{\sigma}$ being labeled by $\rho$. The vertices on a path from the root to a node are distinct, hence $T_{\sigma}$ has depth at most $n=|V|$. A node of $T_{\sigma}$ has countably many children, or one for each element of $U$, where $U=\left\{t \in \mathbb{R}^{+}: \pi_{v}(t)>0\right.$ for some $\left.v \in V\right\}$ is the (countable) support set of the time size distributions. Let $z \in T_{\sigma}$ be a node of label $v \in V$, and let $e$ be the $t$-th edge emanating from $z, t \in U$; then $e$ has probability $\pi_{e}:=\pi_{v}(t)=\operatorname{Pr}\left[S_{v}=t\right]$; moreover if $z$ 's parent node is node $y \in T_{\sigma}$ of label $u \in V$, then $e$ has weight $w_{e}:=d(u, v)+t$.

We will sometimes refer to nodes in $T_{\sigma}$ by their label. (With little chance of confusion.) For a node $v \in T_{\sigma}$, we define $P(v)$ to be the path of edges of leading
from the root to $v$. We define $W(v):=\sum_{e \in P(v)} w_{e}$ and $\Phi(v):=\prod_{e \in P(v)} \pi_{e}$. Thus $W(v)$ is the total time elapsed before reaching $v$ and $\Phi(v)$ is the probability of reaching $v$.

We note that a given policy $\sigma$ will be "compatible" with any tuple ( $\pi^{\prime}, V^{\prime}$, $\left.d^{\prime}, B^{\prime}\right)$ such that $V^{\prime}=V$ and such that $\pi^{\prime}$ has the same support $U$ as $\pi$. Namely, the modified $\pi^{\prime}$ and $d^{\prime}$ will induce modified edge weights $w_{e}^{\prime}$ and modified edge probabilities $p_{e}^{\prime}$. We emphasize this by including $(\pi, V, d, B)$ as arguments to $T_{\sigma}$ (though we will always use the same $V$ for a given policy $\sigma$ ).

We let

$$
R(\sigma, \pi, V, d, B)=\sum_{v \in T_{\sigma}, e=(v, u): W(v)+w_{e} \leq B} \pi_{e} \cdot r_{v} \Phi(v)
$$

denote the expected reward that policy $\sigma$ achieves with respect to the metric $d$, time size distributions $\pi$, and time budget $B$. We use OPT to denote the expected reward of the optimal adaptive policy.

For the following lemma, we also consider a modified node label $X_{v}:=S_{v}+$ $d(u, v)$, where $u$ is the parent and $v$ is the child. Thus, while $w_{e}=t+d(u, v)$ is a number, $X_{v}$ is a random variable. (Moreover, $X_{v}$ "ignores" the fact that edge $e$ is the $t$-th edge of $v$, i.e., that edge $e$ is associated to a particular outcome of $S_{v}$.)

Lemma 4. (part of Lemma 2.4 in [4]) For any policy $\sigma$ on instance ( $\pi, V, d, B$ ), there exists a policy $\sigma^{\prime}$ such that

$$
R\left(\sigma^{\prime}, \pi, V, d, B\right)=(1-O(\epsilon)) R(\sigma, \pi, V, d, B)
$$

and for any realization path $P$ in $T_{\sigma^{\prime}}(\pi, V, d, B), \sum_{v \in P} E\left[X_{v}\right]=O(B / \epsilon)$.
Here we mention that policy $\sigma^{\prime}$ is just cutting some branches on policy $\sigma$ to ensure $\sum_{v \in P} E\left[X_{v}\right]=O(B / \epsilon)$. Note that policy $\sigma$ is all designed by us, we surely may have a node $v$ that $\sum_{v \in P} E\left[X_{v}\right]>O(B / \epsilon)$ though $v$ has no contribution to reward.

### 4.2 Discretization

We present how to discretize the decision tree in Section 4.1 with given $\epsilon$. We use the discretization method similar in paper [10] Section 2.1. We split all jobs into two parts: small size vertices and big size vertices, and discretize them separately.

Denote $\widetilde{S}_{v}$ for value of vertex $v$ after discretization, and the new distribution $\widetilde{\pi}_{v}$. And for each node $v$ in the tree, define $\widetilde{X}_{v}:=\widetilde{S}_{v}+d(u, v)$ as a random variable which is similar in the definition of $X_{v}$ in Section 4.1.

1. Small size region if $S_{v} \leq \epsilon^{4}$.

There exists a value $0 \leq h \leq \epsilon^{4}$, such that $\operatorname{Pr}\left[S_{v} \geq h \mid S_{v} \leq \epsilon^{4}\right] \cdot \epsilon^{4}=$

$$
E\left[S_{v} \mid S_{v} \leq \epsilon^{4}\right] \text {. Then set: } \widetilde{S}_{v}= \begin{cases}0, & 0 \leq S_{v}<h \\ \epsilon^{4}, & h \leq S_{v} \leq \epsilon^{4} \\ S_{v}, & S_{v}>\epsilon^{4}\end{cases}
$$

2. Large size region if $S_{v}>\epsilon^{4}$.

We simply discretize it as $\dot{\widetilde{S}}_{v}=\left\lfloor\frac{S_{v}}{\epsilon^{5}}\right\rfloor \epsilon^{5}$.
We denote the set of the discretized size by $S=\left\{s_{0}, s_{1}, \cdots, s_{z-1}\right\}$ where $s_{0}=0, s_{1}=\epsilon^{5}, s_{2}=2 \epsilon^{5}, \ldots, s_{z-1}$. Note that $s_{1}=\epsilon^{5}, s_{2}=2 \epsilon^{5}, \ldots, s_{1 / \epsilon-1}=$ $\epsilon^{4}-\epsilon^{5}$ are also in $S$ though their probability is 0 . The total number of values of time size is $|S|=z=O\left(B / \epsilon^{5}\right)$ which is a constant.

### 4.3 Canonical Policies

We need to use canonical policies introduced in [4]. A policy $\widetilde{\sigma}$ is a canonical policy if it makes decisions based on the discretized sized of vertices, but not their actual sizes. Under the canonical policy, we will keep trying new vertices if the total discretized size budget does not exceed, even the actual budget overflows. But no reward will get from these vertices which make actual budget exceeding.

Lemma 5. (Lemma 4.2 in [10])) Let $\pi$ be the distribution of vertice size and $\widetilde{\pi}$ be the discretized version of $\pi$. Then, the following statements hold:

1. For any policy $\sigma$, there exists a canonical policy $\tilde{\sigma}$ such that

$$
R(\widetilde{\sigma}, \widetilde{\pi}, V, d,(1+4 \epsilon) B)=(1-O(\epsilon)) R(\sigma, \pi, V, d, B)
$$

2. For any canonical policy $\tilde{\sigma}$,

$$
R(\widetilde{\sigma}, \pi, V, d,(1+4 \epsilon) B)=(1-O(\epsilon)) R(\widetilde{\sigma}, \widetilde{\pi}, V, d, B)
$$

Proof sketch: for the first result, we prove that there is randomized canonical policy $\sigma_{r}$ such that $R\left(\sigma_{r}, \widetilde{\pi}, V, \widetilde{d},(1+4 \epsilon) B\right)=(1-O(\epsilon)) R(\sigma, \pi, V, d, B)$. Thus such a deterministic policy $\widetilde{\sigma}$ exists. The randomized policy $\sigma_{r}$ is derived from $\sigma$ as follows. $T_{\sigma_{r}}$ keeps the same tree structure as $T_{\sigma}$. If $\sigma_{r}$ visits a vertex $v$ and observes a discretized size $s \in S$, it chooses a random branch in $\sigma$ among those sizes that are mapped to $s$. Let $t_{1}, t_{2}, \ldots, t_{k}$ are possible size realization of $s$ as per size distribution $\pi$ and let $\pi_{v}\left(t_{1}\right), \pi_{v}\left(t_{2}\right), \ldots, \pi_{v}\left(t_{k}\right)$ be the corresponding probabilities. Hence we choose one of the branches corresponding to size $t_{1}, t_{2}, \ldots, t_{k}$ w.p. $\pi_{v}\left(t_{1}\right), \pi_{v}\left(t_{2}\right), \ldots, \pi_{v}\left(t_{k}\right)$ (normalized) respectively. For more detail, see Lemma 4.2 in [10].

From Lemma 5, we conclude that the reward will not loss too much if we using canonical policies to applying the policies with $T_{\widetilde{\sigma}}$ described in section 4.2, but not $T$.

### 4.4 Block Tree

Now we partition the decision tree $T_{\widetilde{\sigma}}$ created in Section 4.3 into blocks.
For any node $v$ in the decision tree $T_{\widetilde{\sigma}}$, we define the leftmost path of $v$ to be the realization path which starts at $v$,ends at a leaf, and consists of only edges corresponding to size zero. We define the block starting with node $v$ (denote
as $\operatorname{seg}(v))$ in $T_{\widetilde{\sigma}}$ as the maximal prefix of the leftmost path of $v$ such that: If $E[\tilde{X}(v)]>\epsilon^{13}, \operatorname{seg}(v)$ is the singleton node $\{v\}$. Otherwise, $E[\widetilde{X}(\operatorname{seg}(v))] \leq \epsilon^{13}$.

We partition $T_{\widetilde{\sigma}}$ into blocks as follows. We say a node $v$ is a starting node if $v$ is a root or $v$ corresponds to a non-zero size realization of its parent. For each starting node $v$, we greedily partition the leftmost path of $v$ into blocks, i.e., delete $\operatorname{seg}(v)$ and recurse on the remaining part.

Lemma 6. A canonical policy $\widetilde{\sigma}$ with expected reward $(1-O(\epsilon))$ OPT can be partitioned into several blocks that satisfy the following properties:

1. There are at most $O\left(\epsilon^{-14}\right)$ blocks on any root-leaf path in the decision tree.
2. There are $|S|=O\left(B / \epsilon^{5}\right)$ children for each block.
3. Each block $M$ with more than one node satisfies that $\sum_{b \in M} E\left[\widetilde{X_{b}}\right] \leq \epsilon^{13}$.

Proof. 1. Fix a particular root-to-leaf path $R$. Let us bound the number of blocks on $R$.
We have $\sum_{v \in R} E[\widetilde{X}(v)]=O(B / \epsilon)=O(1 / \epsilon)=O\left(\epsilon^{-1}\right)$ by Lemma 4 . By definition of Block Policies, any single-point block $v$ in $T_{\widetilde{\sigma}}$ satisfies $E[\widetilde{X}(v)]>$ $\epsilon^{13}$. Then there are at most $O\left(\epsilon^{-1}\right) / \epsilon^{13}=O\left(\epsilon^{-14}\right)$ single-point in path $R$.
By definition of Discretization, any $\widetilde{X}(v)$ with non-zero size after discretization, is no less than $\epsilon^{4}$. Then there are at most $O\left(1 / \epsilon^{4}\right)=O\left(\epsilon^{-4}\right)$ nodes $w$ corresponding to a non-zero size realization.
This gives a bound $O\left(\epsilon^{-14}+\epsilon^{-4}\right)=O\left(\epsilon^{-14}\right)$ blocks on any root-leaf path in the decision tree.
2. After Dicretization in Section 4.2, we have only $O\left(B / \epsilon^{5}\right)$ number of values which means there are $|S|=O\left(B / \epsilon^{5}\right)$ children for each block.
3 . This is exactly what we define our blocks.
And then we prove Theorem 3.
Proof (The proof of Theorem 3). By Properties 1,2 in Lemma 6 above, there are only constant number of blocks in the decision tree $T_{\widetilde{\sigma}}$ with block policies. Since there are at most $O\left(\epsilon^{-14}\right)$ blocks in the optimal policy path in decision tree $T_{\widetilde{\sigma}}$ of reward $(1-O(\epsilon)) \mathrm{OPT}$, there exists a block $M$ with $\Omega\left(\epsilon^{14} \mathrm{OPT}\right)$ reward.

We consider two situations in block $M$ with different number of nodes.
Block $M$ is a single-node block of node $v$.
And the node $v$ has a huge reward. Then the single-node tour from the starting point $\rho$ to $v$ has an expected reward $\Omega\left(\epsilon^{14} \mathrm{OPT}\right)$.

Otherwise, block $M$ has at least two nodes.
And block $M$ has an expected reward $\Omega\left(\epsilon^{14} \mathrm{OPT}\right)$. By Property 3 in Lemma 6 , we have $\sum_{b \in M} E\left[\widetilde{X_{b}}\right] \leq \epsilon^{13}$. Since the distance from the starting point $\widetilde{\rho}$ to the first node in $M$ is at most 1 , the path $P^{\star}$ which connects $\widetilde{\rho}$ and the block $M$ in order has length at most $1+\epsilon^{13}$. Thus $P^{\star}$ is a feasible path for $I_{\mathrm{TO}}\left(\epsilon, I_{\mathrm{SO}}\right)$ as defined in Definition 2, and it has an expected reward $\Omega\left(\epsilon^{14} \mathrm{OPT}\right) \leq \mathrm{OPT}_{\mathrm{TO}}$.

Thus, either there is a single-point tour with expected reward $\Omega(\mathrm{OPT})$, or the valid Task Orienteering instance $I_{\mathrm{TO}}\left(\epsilon, I_{\mathrm{SO}}\right)$ has reward $\mathrm{OPT}_{\mathrm{TO}}=\Omega(\mathrm{OPT})$, with constant approximation parameter $\epsilon^{14}$. And this completes the proof for Theorem 3.

## 5 Conclusion

We describe an $O(1)$-approximation algorithm with time budget $(1+\epsilon) B$. As the expected reward of the optimal policy at time budget $B$ is at most $O\left(\epsilon^{-14}\right)$ times the expected reward of our [time $(1+\epsilon) B$ ] policy, we get an $O(1)$-approximation for any constant $0<\epsilon<1$.

## References

1. Esther M. Arkin, Joseph S. B. Mitchell, and Giri Narasimhan. Resourceconstrained geometric network optimization. In Proceedings of the fourteenth annual symposium on Computational geometry, pages 307-316, 1998.
2. Blum Avrim, Chawla Shuchi, Karger David, R., Lane Terran, Meyerson Adam, and Minkoff Maria. Approximation algorithms for orienteering and discounted-reward tsp. SIAM Journal on Computing, 37(2):653, 2007.
3. Nikhil Bansal and Viswanath Nagarajan. On the Adaptivity Gap of Stochastic Orienteering. Springer International Publishing, 2014.
4. Anand Bhalgat, Ashish Goel, and Sanjeev Khanna. Improved approximation results for stochastic knapsack problems. Proceedings of the Twenty Second Annual Acm Siam Symposium on Discrete Algorithms, 27(2):1647-1665, 2011.
5. Daniel Bienstock, Michel X. Goemans, David Simchi-Levi, and David Williamson. A note on the prize collecting traveling salesman problem. Mathematical Programming, 59(3):413-420, 1993.
6. Chandra Chekuri, Nitish Korula, and Martin Pal. Improved algorithms for orienteering and related problems. Acm Transactions on Algorithms, 8(3):661-670, 2008.
7. Ke Chen and Sariel Har-Peled. The orienteering problem in the plane revisited. Proc.symp.on Computational Geometry, pages 247 - 254, 2006.
8. Bruce L. Golden, Larry Levy, and Rakesh Vohra. The orienteering problem. Naval Research Logistics, 34(34):307-318, 1987.
9. Anupam Gupta, Ravishankar Krishnaswamy, Viswanath Nagarajan, and R. Ravi. Approximation algorithms for stochastic orienteering. In Proceedings of the twentythird annual ACM-SIAM symposium on Discrete Algorithms, pages 1522-1538, 2012.
10. Jian Li and Wen Yuan. Stochastic combinatorial optimization via poisson approximation. Proceedings of the Annual Acm Symposium on Theory of Computing, pages 971-980, 2012.
11. Keith W. Ross and Danny H K Tsang. Stochastic knapsack problem. Communications IEEE Transactions on, 37(7):740-747, 1989.
12. T. Tsiligirides. Heuristic methods applied to orienteering. Journal of the Operational Research Society, 35(9):797-809, 1984.
13. Pieter Vansteenwegen, Wouter Souffriau, and Dirk Van Oudheusden. The orienteering problem: A survey. Mitteilungen Klosterneuburg Rebe Und Wein Obstbau Und Fruchteverwertung, 209(209):1-10, 2011.
