

1 Online Time-Windows TSP with Predictions

2 Shuchi Chawla ✉ 

3 University of Texas at Austin, United States

4 Dimitris Christou ✉ 

5 University of Texas at Austin, United States

6 — Abstract —

7 In the *Time-Windows TSP* (TW-TSP) we are given requests at different locations on a network;
8 each request is endowed with a reward and an interval of time; the goal is to find a tour that
9 visits as much reward as possible during the corresponding time window. For the online version of
10 this problem, where each request is revealed at the start of its time window, no finite competitive
11 ratio can be obtained. We consider a version of the problem where the algorithm is presented with
12 predictions of where and when the online requests will appear, without any knowledge of the quality
13 of this side information.

14 Vehicle routing problems such as the TW-TSP can be very sensitive to errors or changes in the
15 input due to the hard time-window constraints, and it is unclear whether imperfect predictions can
16 be used to obtain a finite competitive ratio. We show that good performance can be achieved by
17 explicitly building slack into the solution. Our main result is an online algorithm that achieves a
18 competitive ratio logarithmic in the diameter of the underlying network, matching the performance of
19 the best offline algorithm to within factors that depend on the quality of the provided predictions. The
20 competitive ratio degrades smoothly as a function of the quality and we show that this dependence
21 is tight within constant factors.

22 **2012 ACM Subject Classification** Theory of computation → Online algorithms

23 **Keywords and phrases** Travelling Salesman Problem, Predictions, Learning-Augmented Algorithms,
24 Approximation

25 **Digital Object Identifier** 10.4230/LIPIcs.APPROX/RANDOM.2024.2

26 **Related Version** Due to space limitations, some proofs are omitted to the full version of this paper.

27 *Full Version:* <https://arxiv.org/abs/2304.01958> [13]

28 **Funding** *Shuchi Chawla:* This work was funded in part by NSF award CCF-2217069.



© Shuchi Chawla and Dimitris Christou;

licensed under Creative Commons License CC-BY 4.0

Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques
(APPROX/RANDOM 2024).

Editors: Amit Kumar and Noga Ron-Zewi; Article No. 2; pp. 2:1–2:22



Leibniz International Proceedings in Informatics

Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

29 **1** Introduction

30 Many optimization problems exhibit a large gap in how well they can be optimized offline
 31 versus when their input arrives in online fashion. In order to obtain meaningful algorithmic
 32 results in the online setting, a natural direction of investigation is to consider "beyond
 33 worst case" models that either limit the power of the adversary or increase the power of the
 34 algorithm. A recent line of work in this direction has considered the use of *predictions* in
 35 bridging the offline versus online gap. Predictions in this context are simply side information
 36 about the input that an online algorithm can use for its decision making; the true input is
 37 still adversarially chosen and arrives online. The goal is to show that on the one hand, if
 38 the predictions are aligned with the input, the algorithm performs nearly as well as in the
 39 offline setting (a property known as *consistency*); and on the other hand, if the predictions
 40 are completely unrelated to the input, the algorithm nevertheless performs nearly as well as
 41 the best online algorithm (a.k.a. *robustness*). *Put simply, good predictions should help, but*
 42 *bad predictions should not hurt, and ideally we should reap the benefits without any upfront*
 43 *knowledge about the quality of the predictions.*

44 Predictions have been shown to effectively bypass lower bounds for a variety of different
 45 online decision-making problems including, for example, caching [30, 31, 25], scheduling [12,
 46 24, 3], online graph algorithms [4], load balancing [27, 28], online set cover [6], matching
 47 problems [17], *k*-means [19], secretary problems [1, 18], network design [20, 33, 9] and more.¹

48 In this paper, we consider a problem whose objective function value is highly sensitive
 49 to changes in the input, presenting a significant challenge for the predictions setting. In
 50 the Traveling Salesman Problem with Time Windows (TW-TSP for short), we are given
 51 a sequence of service requests at different locations on a weighted undirected graph. Each
 52 request is endowed with a reward as well as a time window within which it should be serviced.
 53 The goal of the algorithm is to produce a path that maximizes the total reward of the requests
 54 visited within their respective time windows. In the online setting, the requests arrive one at
 55 a time at the start of their respective time windows, and the algorithm must construct a
 56 path incrementally without knowing the locations or time windows of future requests.

57 Vehicle routing problems such as the TW-TSP that involve hard constraints on the lengths
 58 of subpaths (e.g. the time at which a location is visited) are generally more challenging than
 59 their length-minimization counterparts. In particular, a small bad decision at the beginning
 60 of the algorithm, such as taking a slightly suboptimal path to the first request, can completely
 61 obliterate the performance of the algorithm by forcing it to miss out on all future reward.
 62 In the offline setting, this means that the approximation algorithm has to carefully counter
 63 any routing inefficiency in some segments by intentionally skipping reachable value in other
 64 segments. In the online setting, this means that no sublinear competitive ratio is possible.

65 *Given the sensitivity of the TW-TSP objective to small routing inefficiencies, is it possible*
to design meaningful online algorithms for this problem using imperfect predictions?

66 We consider a model where the algorithm is provided with a predicted sequence of requests
 67 at the beginning, each equipped with a predicted location and a predicted time window. The
 68 true sequence of requests is revealed over time as before. Of course if the predicted sequence
 69 is identical to the true request sequence, the algorithm can match the performance of the
 70 best offline algorithm. But what if the predictions are slightly off? Could these small errors
 71 cause large losses for the online algorithm? Can the algorithm tolerate large deviations?

¹ A comprehensive compendium of literature on the topic can be found at [29].

72 Our main result is an online algorithm for the TW-TSP based on predictions whose
 73 performance degrades smoothly as a function of the errors in prediction. We obtain this
 74 result by explicitly building slack into our solution and benchmark. In a slight departure
 75 from previous work on TW-TSP, we require the server to spend one unit of idle "service time"
 76 at each served request. We show that this is necessary to obtain a sublinear approximation
 77 even with predictions (Theorem 9). (However, in the absence of predictions, the setting
 78 with service times continues to admit a linear lower bound on the competitive ratio; See
 79 Theorem 8.) We then use service times judiciously in planning a route and accounting for
 80 delays caused by prediction errors.

81 There are two primary sources of error in predictions: (1) the predicted locations of
 82 requests may be far from the true locations; and, (2) the predicted time windows may be
 83 different from the true time windows. The competitive ratio of our algorithm depends linearly
 84 on each of these components, taking the maximum error over each predicted request and
 85 normalizing appropriately.² This dependence is tight to within constant factors. Besides this
 86 dependence on the prediction error, the competitive ratio depends logarithmically on the
 87 diameter of the underlying network, matching the performance of the best known offline
 88 algorithm for TW-TSP.

89 Although our competitive ratio is stated in terms of the maximum location or time
 90 window errors, where the maximum is taken over all requests in the instance, our algorithm
 91 performs well even when some of the errors are large and most errors are small. In particular,
 92 our algorithm's performance is *simultaneously* competitive against the maximum achievable
 93 reward over *any* subset of requests, scaled down by the maximum prediction error over that
 94 subset. (See "Extensions" in Section 3 for a formal statement.) In this respect, our guarantees
 95 fall within the framework of *metric error with outliers* proposed by [4]. On the other hand,
 96 when all or most requests are predicted poorly, our algorithm also inevitably performs poorly
 97 as it inherits lower bounds from fully online instances.

98 Importantly, our algorithm requires little to no information about how the predictions
 99 match up against the true requests. For the purpose of analysis, we measure the error
 100 in predictions with respect to some underlying matching between the predicted and true
 101 requests – the error parameters are then defined in terms of the maximum mismatch between
 102 any predicted request and its matched true request. This matching is never revealed to
 103 the algorithm and in fact the performance of the algorithm depends on the quality of the
 104 *best* possible matching between the predicted and true requests. The only information the
 105 algorithm requires about the quality of the predictions is the location error – the maximum
 106 distance between any prediction and its matched true request. Even for this parameter, an
 107 upper bound suffices (and at a small further loss, a guess suffices).

108 Our overall approach has several components. The first of these is to construct an instance
 109 of the TW-TSP over predicted requests that requires the server to spend some idle time at
 110 each request as a "service delay". We then extend offline TW-TSP algorithms to this service
 111 delay setting, obtaining a logarithmic in diameter approximation. We then follow and adapt
 112 this offline solution in the online setting. Every time the offline solution services a predicted
 113 request, we match this request to a previously revealed true request, take a detour from the
 114 computed path to visit and service the true request, and then resume the precomputed path.
 115 Altogether this provides the desired competitive ratio.

² Formally, the competitive ratio depends linearly on the ratio of the maximum location error of any predicted request to the minimum service time, as well as the ratio of the maximum time window error to the minimum time window length.

116 Our results further generalize to a setting where predictions are coarse in that each single
117 predicted location captures multiple potential true requests that are nearby. We show that
118 with prediction errors defined appropriately, we can again achieve a competitive ratio for
119 this "many to one matching" setting that is logarithmic in the diameter of the graph and
120 polynomial in the prediction error.

121 Finally, our algorithm and analysis incorporates error in estimating rewards of requests
122 in a straightforward manner achieving the optimal dependence on this third source of error.

123 Further related work

124 Using predictions in the context of online algorithm design was first proposed by [30] for
125 the well-studied caching problem. Since that work, the literature on online algorithm design
126 with predictions has grown rapidly. We point the interested reader to a compendium at [29]
127 for further references.

128 Metric error with outliers.

129 Azar et al. [4] initiated the study of predictions in the context of online graph optimization
130 problems, and proposed a framework for quantifying errors in predictions, called *metric error*
131 *with outliers*, that we adapt. The idea behind this framework is to capture two sources of
132 error: (1) Some true requests may not be captured by predictions and some predictions may
133 not correspond to any true requests; (2) For requests that are captured by predictions, the
134 predictions may not be fully faithful or accurate. The key observation is that it is possible
135 to design algorithms with performance that depends on these two sources of error *without*
136 *explicit upfront knowledge of the (partial) correspondence between predicted and true requests.*

137 In this work, we focus mostly on the second source of error, which we further subdivided
138 into three kinds of error in order to obtain a finer understanding of the relationship between
139 the competitive ratio and different kinds of error. As in the work of Azar et al. [4], we assume
140 that the correspondence between predicted and true requests is never explicitly revealed to
141 the algorithm. The performance of the algorithm nevertheless depends on the error of the
142 *best* matching between predicted and true requests. In Section 3 we describe how the first
143 source of error in Azar et al.'s framework can also be incorporated into our bounds.

144 TSP with predictions.

145 Recently a few papers [10, 23, 22] have considered the online TSP and related routing
146 problems with predictions. The input to the online TSP is similar to ours: requests arrive
147 over time in a graph, and a tour must visit each request after its arrival time. However,
148 the objective is different. In our setting, requests also have deadlines, and the algorithm
149 cannot necessarily visit all requests. The goal therefore is to visit as many as possible. In
150 the online TSP, there are no deadlines, and so the objective is to visit *all* requests as quickly
151 as possible, or in other words to minimize the makespan. This makespan minimization
152 objective is typically much easier than the deadline setting, as evidenced by constant factor
153 approximations for it in the offline, online, and predictions settings, as opposed to logarithmic
154 or worse approximations for the latter problem.

155 We mention that the algorithmic idea of precomputing an offline path based on the
156 predictions and then adapting it to the online input has also been used in [10, 23]. The main
157 challenge in the setting that our works considers, is that due to the existence of deadlines,
158 our algorithm needs to be careful on how it adapts its path, as taking a large detour could
159 result in entirely missing the time-windows of future (unrevealed) requests. We circumvent

160 this issue by introducing appropriately large idle times on the predicted requests that our
 161 offline solution visits.

162 TW-TSP *without* predictions.

163 The (offline) TW-TSP problem has a rich literature and has been studied for over 20 years.
 164 The problem is known to be NP-hard even for special cases, e.g. on the line [32], and when
 165 all requests have the same release times and deadlines (a.k.a. Orienteering) [11]. Orienteering
 166 admits constant factor approximations [11, 7, 14], and even a PTAS when requests lie in
 167 a fixed dimensional Euclidean space [2, 16]. For general time windows, constant factor
 168 approximations are only known for certain special cases: e.g. constant number of distinct
 169 time windows [15]; and on line graphs [32, 26, 8, 21]. For general graphs and time-windows,
 170 the best approximations known are logarithmic in input parameters [7, 14].

171 To the best of our knowledge, the online setting for TW-TSP has only been considered
 172 by Azar and Vardi [5]. Azar and Vardi assume that service times are non-zero and present
 173 competitive algorithms under the assumption that the smallest time window length L_{\min}
 174 is comparable to the diameter D of the graph, as no sublinear competitive ratio can be
 175 achieved if $L_{\min} < D/2$. We are able to beat this lower bound by relying on predictions.

176 Organization of the paper

177 We formally define the problem and our error model in Section 2. Section 3 describes our
 178 results and provides an outline for our analysis. All of our main results are covered in that
 179 section. Proofs of these results can be found in subsequent sections. In particular, Sections 4,
 180 5, and 6 fill in the details of our upper bound, and Section 7 proves the stated lower bounds.
 181 Proofs omitted from the main body of this paper can be found in the appendix of the full
 182 version [13].

183 2 Definitions

184 2.1 The Traveling Salesman Problem with Time-Windows

185 An instance of the TW-TSP consists of a network G and a (finite) sequence of service requests
 186 I . Here, $G = (V, E, \ell)$ is an undirected network with edge lengths $\{\ell_e\}_{e \in E}$. Extending the
 187 notion of distance to all vertex pairs in G , we define $\ell(u, v)$ for $u, v \in V$ to be the length of
 188 the shortest path from u to v . We assume without loss of generality that G is connected
 189 and that the edge lengths ℓ_e are integers. A service request $\sigma = (v_\sigma, r_\sigma, d_\sigma, \pi_\sigma)$ consists of a
 190 vertex $v_\sigma \in V$ at which the request arrives, a release time $r_\sigma \in \mathbb{Z}^+$, a deadline $d_\sigma \in \mathbb{Z}^+$ with
 191 $d_\sigma > r_\sigma$, and a reward $\pi_\sigma \in \mathbb{Z}^+$. We use $\Sigma \subseteq V \times \mathbb{Z}^+ \times \mathbb{Z}^+ \times \mathbb{Z}^+$ to denote the set of all
 192 possible client requests and $I \subset \Sigma$ to denote the set of requests received by the algorithm.

193 The solution to TW-TSP is a continuous *walk* on G that is allowed to remain idle on the
 194 vertices of the graph.³ Formally, the walk starts from some vertex at time $t = 0$; at every
 195 time-step that it occupies a vertex $u \in V$, it can either remain idle on u for some number
 196 of time-steps *or* it can move to some $v \in V$ by spending time $t = \ell_{(u,v)}$; we comment that
 197 while the path is mid-transition, no more decisions can be made. Notice that this creates

³ To keep our exposition simple, we do not specify a starting location for the walk. However, all of our algorithms can be adapted without loss to the case where a starting location is fixed, as described towards the end of Section 3.

198 a notion of a discrete time-horizon that will be important towards formalizing the online
199 variant of the problem.

200 We use $\mathcal{W}(G)$ to denote the set of all walks on G . Given a request $\sigma \in \Sigma$, we say that
201 a walk W *covers* it if W remains idle on vertex v_σ for at least one time-step,⁴ starting on
202 some step $\tau \in [r_\sigma, d_\sigma - 1]$. For a sequence of requests $I \subset \Sigma$, we use $\text{Cov}(W, I) \subseteq I$ to denote
203 the set of requests in I that are covered by W . Then, the reward obtained by walk W is
204 denoted by $\text{Rew}(W, I) := \sum_{\sigma \in \text{Cov}(W, I)} \pi_\sigma$. The objective of TW-TSP is to compute a walk
205 $W \in \mathcal{W}(G)$ of maximum reward. We denote this by $\text{OPT}(G, I) := \max_{W \in \mathcal{W}(G)} [\text{Rew}(W, I)]$.

206 2.2 The offline, online, and predictions settings

207 We assume that the network G is known to the algorithm upfront in all of the settings we
208 consider. In the **offline** version of the problem, the sequence of requests I is given to the
209 algorithm in advance. In the **online** version, requests $\sigma \in I$ arrive in an online fashion;
210 specifically, each request $\sigma \in I$ is revealed to the algorithm at its release time r_σ .

211 In the **predictions** setting, the true sequence of requests I arrives online, as in the online
212 setting. However, the algorithm is also provided with a predicted sequence $I' \subset \Sigma$ in advance,
213 where every request $\sigma' \in I'$ is endowed with a location, a time window, and a reward. The
214 quality of predictions is expressed in terms of their closeness to true requests. To this end,
215 we define three notions of mismatch or error. For a true request σ and predicted request σ' ,
216 the location error, time windows error, and reward error are defined as:

$$\begin{aligned} 217 \quad \text{LocErr}(\sigma, \sigma') &:= \ell(v_\sigma, v_{\sigma'}) \\ 218 \quad \text{TWErr}(\sigma, \sigma') &:= \max\{|r_\sigma - r_{\sigma'}|, |d_\sigma - d_{\sigma'}|\} \\ 219 \quad \text{RewErr}(\sigma, \sigma') &:= \max\{\pi_\sigma/\pi_{\sigma'}, \pi_{\sigma'}/\pi_\sigma\} \end{aligned}$$

220 We extend these definitions to the entire sequences I and I' through an underlying (but
221 unknown to the algorithm) matching between the requests in the two lists:

222 ► **Definition 1.** *Given two request sequences $I, I' \subset \Sigma$ with $|I| = |I'|$ and a perfect matching*
223 *$M : I \mapsto I'$, we define the location, time window, and reward errors for the matching M as:*

$$\begin{aligned} 224 \quad \Lambda_M &:= \max_{\sigma \in I} \text{LocErr}(\sigma, M(\sigma)) \\ 225 \quad \tau_M &:= \max_{\sigma \in I} \text{TWErr}(\sigma, M(\sigma)) \\ 226 \quad \rho_M &:= \max_{\sigma \in I} \text{RewErr}(\sigma, M(\sigma)) \end{aligned}$$

227 We use $n = |V|$ to denote the number of vertices in G , D to denote the diameter of
228 the graph, and L_{\min} and L_{\max} to denote the size of the smallest and largest time windows
229 respectively (of a true or predicted request) in the given instance; that is, we denote
230 $L_{\min} = \min_{\sigma \in I \cup I'} |d_\sigma - r_\sigma|$ and $L_{\max} = \max_{\sigma \in I \cup I'} |d_\sigma - r_\sigma|$. The competitive ratios of the
231 algorithms we develop depend on these parameters.

232 Knowns and unknowns.

233 We denote an instance of the TW-TSP with predictions by (G, I, I', M) . All components of
234 the instance are chosen adversarially. As mentioned earlier, the network G and the predicted

⁴ As we mentioned in the introduction, this requirement of a minimal one-unit service time is necessary in order to achieve any sublinear approximation for the online TW-TSP even with predictions. See Theorem 9 in Section 7.

235 sequence I' are provided to the algorithm at the start. The sequence I arrives online. We
 236 assume that the algorithm receives no direct information about the matching M , but is
 237 provided with an upper bound on the error Λ_M . We will also assume that the algorithm
 238 knows the parameter L_{\min} , although this is without loss of generality as the parameter can
 239 be inferred within constant factor accuracy from the predictions.⁵

240 2.3 The TW-TSP with service times

241 At a high level our algorithm has two components: an offline component that computes a
 242 high-reward walk over the predicted locations of requests, and an online component that
 243 largely follows this walk but takes "detours" to cover the arriving true sequence of requests.
 244 In particular, as the algorithm follows the offline walk, for each predicted location it visits
 245 where a "close by" true request is available, the algorithm takes a "detour" to this true
 246 request, returns back to the predicted location, and resumes the remainder of the walk. In
 247 order to incorporate the time spent taking these detours in our computation of the offline
 248 walk, we require the walk to spend some "service time" at each predicted location it covers.
 249 Accordingly, we define a generalization of the TW-TSP:

250 ► **Definition 2.** *The TW-TSP with Service Times (TW-TSP-S) takes as input a network G ,*
 251 *a sequence of service requests I , and a service time $S \in \mathbb{Z}^+$, and returns a walk $W \in \mathcal{W}(G)$.*
 252 *We say that W covers a request $\sigma \in I$, denoted $\sigma \in \text{Cov}(W, I, S)$, if it remains idle on vertex*
 253 *v_σ for at least S time steps, starting at some step $t \in [r_\sigma, d_\sigma - S]$. We define the reward of*
 254 *W as $\text{Rew}(W, I, S) := \sum_{\sigma \in \text{Cov}(W, I, S)} \pi_\sigma$. The optimal value of the instance is given by:*

$$255 \quad \text{OPT}(G, I, S) := \max_{W \in \mathcal{W}(G)} [\text{Rew}(W, I, S)].$$

256 Note that the original version of TW-TSP as defined previously simply corresponds
 257 to the special case of TW-TSP-S with service time $S = 1$, and in particular, we have
 258 $\text{Rew}(W, I) = \text{Rew}(W, I, 1)$, and $\text{OPT}(G, I) = \text{OPT}(G, I, 1)$.

259 3 Our results and an outline of our approach

260 Our main result is as follows.

261 ► **Theorem 3.** *Given any instance (G, I, I', M) of the TW-TSP with predictions whose errors*
 262 *satisfy $\tau_M \leq L_{\min}/2$ and $\Lambda_M \leq (L_{\min} - 1)/4$, there exists a polynomial-time online algorithm*
 263 *that takes the tuple (G, I', Λ_M) as offline input and I as online input, and constructs a walk*
 264 *$W \in \mathcal{W}(G)$ such that*

$$265 \quad \mathbb{E}[\text{Rew}(W, I)] \geq \frac{1}{O(\Lambda_M \cdot \rho_M^2 \cdot \log \min(D, L_{\max}))} \cdot \text{OPT}(G, I).$$

266 As mentioned previously, our algorithm consists of two components. The offline component
 267 constructs a potential walk in the network with the help of the predicted requests. Then an
 268 online component adapts this walk to cover true requests that arrive one at a time. We break
 269 up the design and analysis of our algorithm into four steps. The first two steps relate the
 270 offline instance we solve to the hindsight optimal solution for the online instance. The third

⁵ In particular, assuming $\tau_M \leq L_{\min}/2$, which is necessary for our results to hold, the time window of any true request can be no shorter than half the smallest time window of any predicted request.

271 step then applies an offline approximation to the predicted instance with appropriate service
 272 times. The final step deals with the online adaptation of the walk to the arriving requests.

273 The following four lemmas capture the four steps. First, we show (Section 4) that
 274 introducing a service time of S hurts the optimal value by at most a factor of $2S - 1$. As
 275 we prove in Lemma 14 of Section 4, this dependency on S is tight. Observe that we require
 276 $S \leq L_{\min}$, as for any tour to feasibly cover a request, the service time for that request must
 277 fit within its time window.

278 ► **Lemma 4.** *For any instance (G, I) of the TW-TSP with service times, and any integer*
 279 *$S \leq L_{\min}$, we have*

$$280 \quad \text{OPT}(G, I, S) \geq \frac{1}{2S - 1} \cdot \text{OPT}(G, I, 1).$$

281 Our second step (also in Section 4) relates the value of the optimal solution over the true
 282 requests I to the optimum over the predicted sequence I' . In both cases, we impose some
 283 service time requirements. Note that this argument needs to account for the discrepancy in
 284 locations, time windows, as well as the rewards of the true and predicted requests.

285 ► **Lemma 5.** *Let (G, I, I', M) be an instance of the TW-TSP with predictions, where Λ_M ,*
 286 *ρ_M , and τ_M denote the maximum location, reward, and time window errors of the instance*
 287 *respectively. Define $S := 4\Lambda_M + 1$ and $S' := 2\Lambda_M + 1$. Then, if $\tau_M \leq L_{\min}/2$ and*
 288 *$\Lambda_M \leq (L_{\min} - 1)/4$, we have*

$$289 \quad \text{OPT}(G, I', S') \geq \frac{1}{3\rho_M} \cdot \text{OPT}(G, I, S).$$

290 Our third step (Section 5) captures the offline component of our algorithm: computing
 291 an approximately optimal walk over the predicted requests with the specified service times.
 292 For this we leverage previous work on the TW-TSP without service times and show how to
 293 adapt it to capture the service time requirement.

294 ► **Lemma 6.** *Given any instance (G, I', S') of the TW-TSP with service times, there exists*
 295 *a polynomial time algorithm that returns a walk $W \in \mathcal{W}(G)$ with reward*

$$296 \quad \text{Rew}(W, I', S') = \frac{1}{\mathcal{O}(\log \min(D, L_{\max}))} \cdot \text{OPT}(G, I', S').$$

297 Finally, the fourth component (Section 6) addresses the online part of our algorithm.
 298 Given a walk computed over the predicted request sequence, it solves an appropriate online
 299 matching problem to construct detours to capture true requests. As in Lemma 5, this part
 300 again needs to account for the discrepancy in locations, time windows, as well as the rewards
 301 of the true and predicted requests.

302 ► **Lemma 7.** *Given an instance (G, I, I', M) of the TW-TSP with predictions satisfying*
 303 *$\tau_M \leq L_{\min}/2$ and $\Lambda_M \leq (L_{\min} - 1)/4$; a walk $W' \in \mathcal{W}(G)$; and any integer $S' \geq 2\Lambda_M + 1$,*
 304 *there exists an online algorithm (Algorithm 1) that returns a walk $W \in \mathcal{W}(G)$ with expected*
 305 *reward*

$$306 \quad \mathbb{E}[\text{Rew}(W, I, 1)] \geq \frac{1}{6\rho_M} \cdot \text{Rew}(W', I', S').$$

307 Theorem 3 follows immediately by putting Lemmas 4, 5, 6 and 7 together.

308 Lower bounds and tightness of our results.

309 We show that the online TW-TSP does not admit sublinear competitive algorithms in the
 310 absence of predictions if $L_{\min} < D$, even with non-zero service times. Furthermore, if the
 311 service times are all 0, no sublinear competitive ratio is possible even using predictions that
 312 are accurate in all respects except the request location. Therefore, in order to achieve a
 313 nontrivial competitive ratio, it is necessary to use predictions as well as to impose non-zero
 314 service times on the optimum. The proofs are presented in Section 7.

315 ► **Theorem 8.** *The competitive ratio of any randomized online algorithm for Online TW-TSP*
 316 *on instances with $L_{\min} \leq D$ and all service times equal to 1 is at most $1/n$.*

317 ► **Theorem 9.** *For any $S > 0$, there exists an instance (G, I, I', M) of the TW-TSP with*
 318 *predictions and service times 0, satisfying $\tau_M = 0$, $\rho_M = 1$, and $\Lambda_M = S$, such that any*
 319 *randomized online algorithm taking the tuple (G, I', Λ_M) as offline input and I as online*
 320 *input achieves a reward no larger than $O(1/n) \cdot \text{OPT}(G, I, 0)$. Here n is the number of*
 321 *vertices in G .*

322 As mentioned earlier, the best known approximation factor for the offline TW-TSP
 323 is $O(\log L_{\max})$ (which we show can be improved slightly to $O(\log \min(D, L_{\max}))$). We
 324 inherit this logarithmic dependence on D and L_{\max} in the predictions setting. Furthermore,
 325 any improvements to the offline approximation would immediately carry through into our
 326 competitive ratio as well. In particular, given an offline TW-TSP algorithm that achieves a
 327 competitive ratio of $\alpha(D, L_{\max})$, we obtain an online algorithm that achieves a competitive
 328 ratio of $O(\Lambda_M \cdot \rho_M^2 \cdot \alpha(D, L_{\max}))$.

329 The dependence of our bound on ρ_M can easily be seen to be tight – consider a star
 330 graph with requests on leaves, and edge lengths and time windows defined in such a manner
 331 that any feasible walk can cover at most one request. Then an uncertainty of a factor of
 332 ρ_M in the predicted rewards can force any online algorithm to obtain an $\Omega(\rho_M^2)$ competitive
 333 ratio even if the predictions are otherwise perfect. Finally, we show in Section 7.2 that the
 334 dependence of our competitive ratio on Λ_M is also tight:

335 ► **Theorem 10.** *For any $S > 0$, there exists an instance (G, I, I', M) of the TW-TSP with*
 336 *predictions satisfying $\tau_M = 0$, $\rho_M = 1$, and $\Lambda_M = S$ such that the competitive ratio of any*
 337 *randomized online algorithm taking the tuple (G, I', Λ_M) as offline input and I as online*
 338 *input asymptotically approaches $1/(S + 1)$.*

339 Extensions and generalizations.

340 We now describe some ways in which we can weaken the assumptions in Theorem 3 while
 341 maintaining its competitive ratio guarantee:

- 342 ■ **Lack of knowledge of Λ_M .** Our algorithm continues to work as intended if it is provided
 343 with an upper bound on Λ_M rather than the exact value of the parameter, with the
 344 performance of the algorithm degrading linearly with the upper bound, as in the theorem
 345 above. One such upper bound is simply $L_{\min}/4$. Moreover, by guessing Λ_M within a
 346 factor of 2 in the range $[0, L_{\min}/4]$, we can obtain the claimed approximation with a
 347 further loss of $O(\log L_{\min})$. Thus, our algorithm can achieve non-trivial guarantees that
 348 scale with the location error even in settings where no information is given about any of
 349 the prediction errors $\Lambda_M, \tau_M, \rho_M$.

350 ■ **Assumptions on τ_M and Λ_M .** It is easy to see that it is necessary to assume $\Lambda_M \leq L_{\min}$
 351 to obtain a nontrivial competitive ratio, as predictions with a location error larger than
 352 the time window size are of no value to the online algorithm. On the other hand, assuming
 353 $\tau_M \leq L_{\min}$ is not necessary. We can accommodate larger time window errors by following
 354 one out of roughly τ_M/L_{\min} different time shifts of the offline walk. This worsens our
 355 approximation factor by an additional factor of τ_M/L_{\min} . In particular, this algorithm
 356 achieves a competitive ratio of $O(\Lambda_M \cdot \rho_M^2 \cdot \tau_M/L_{\min} \cdot \log \min(D, L_{\max}))$.

357 ■ **Random rewards.** Our results also hold in the case of random rewards. Specifically,
 358 consider a setting where the rewards $\{\pi_\sigma\}_{\sigma \in I}$ are drawn from some joint (not necessarily
 359 product) distribution D over \mathbb{R}_+^I . In that case, we define $\text{Rew}(W, I) := \sum_{\sigma \in \text{Cov}(W, I)} \mathbb{E}[\pi_\sigma]$,
 360 and $\text{OPT}(G, I)$ as the maximum reward obtained by any walk $W \in \mathcal{W}(G)$.⁶ Finally, we
 361 define $\text{RewErr}(\sigma, \sigma')$ as the mismatch between π'_σ and $\mathbb{E}[\pi_\sigma]$. Our analysis provides the
 362 same approximation as before in this setting. See the full version for a formal proof.

363 ► **Corollary 11.** *Given an instance (G, I, I', M) of the TW-TSP with predictions where
 364 requests have randomly drawn rewards, and predictions errors satisfy that $\tau_M \leq L_{\min}/2$
 365 and also $\Lambda_M \leq (L_{\min} - 1)/4$, there exists a polynomial-time online algorithm that takes the
 366 tuple (G, I', Λ_M) as offline input and I as online input, and constructs a walk $W \in \mathcal{W}(G)$
 367 such that*

$$368 \quad \mathbb{E}[\text{Rew}(W, I)] \geq \frac{1}{O(\Lambda_M \cdot \rho_M^2 \cdot \log \min(D, L_{\max}))} \cdot \text{OPT}(G, I)$$

369 ■ **Rooted instances.** Next, we consider the case where a starting vertex v_0 is also specified,
 370 and the solution space $\mathcal{W}(G)$ includes all walks on G that *start* on vertex v_0 at $t = 0$.
 371 We can easily see that this setting is essentially equivalent to its unrooted counterpart,
 372 under the extra assumption that each request $\sigma = (v_\sigma, r_\sigma, d_\sigma, \pi_\sigma)$ satisfies the conditions
 373 $\ell(v_0, v_\sigma) \leq r_\sigma$. This is a reasonable assumption as no algorithm can visit a request σ
 374 before time $\ell(v_0, v_\sigma)$ anyway. Clearly, for any rooted instance (G, I, v_0) , the unrooted
 375 optimal $\text{OPT}(G, I)$ is an upper bound on the rooted optimal $\text{OPT}(G, I, v_0)$. On the other
 376 hand, the unrooted path computed by our algorithm can be transformed to a path of
 377 same reward rooted at v_0 by going directly from v_0 to the predicted request serviced first,
 378 as this distance is at most equal to the request's release time.

379 ■ **Partial matching.** Next we consider the case where not all true requests are captured
 380 by the predicted requests and, on the flip side, where some predicted requests do not
 381 correspond to true requests at all. Following the framework of [4], we consider partial
 382 matchings between I and I' , and define Δ_1^M to be the total reward of all true requests
 383 that are unmatched, and Δ_2^M to be the total predicted reward of predicted requests that
 384 are unmatched. Then, it is easy to see that our analysis goes through for the subsets of I
 385 and I' that are matched to each other, costing us an additive amount of no more than
 386 $\Delta_1^M + \Delta_2^M$. See the full version for a formal proof.

387 ► **Corollary 12.** *Given an instance (G, I, I') of the TW-TSP with predictions, let M be any
 388 (incomplete) matching between I and I' , and let the error parameters $\Lambda_M, \rho_M, \tau_M, \Delta_1^M$,
 389 and Δ_2^M be defined as above. Then, there exists an online algorithm that takes (G, I', Λ_M)*

⁶ Note that we do not allow the optimal walk to adapt to instantiations of rewards. Adaptive walks cannot be competed against in an online setting even with predictions.

390 as offline input and I as online input, and returns a walk $W \in \mathcal{W}(G)$ such that

$$391 \quad \mathbb{E}[\text{Rew}(W, I)] \geq \Omega\left(\frac{1}{\Lambda_M \cdot \rho_M^2 \cdot \log \min(D, L_{\max})}\right) \cdot (\text{OPT}(G, I) - \Delta_1^M) - \frac{\Delta_2^M}{\rho_M}.$$

392 ■ **Many to one matching.** Consider a setting where predictions are coarse in that each
 393 single predicted location captures multiple potential true requests. We can model such a
 394 setting within our predictions framework and obtain almost the same guarantee as in
 395 Theorem 3. In particular, for this setting, let M be a many-to-one matching from I to I' .
 396 We define the location error of a predicted request $\sigma' \in I'$ as the length of the shortest
 397 path that starts at σ' , visits all of the locations of the true requests that are preimages of
 398 σ' in M , spending one unit of time at each, and returns back to σ' . Observe that this
 399 location error is the length of the optimal solution to an orienteering problem rooted at σ' .
 400 Correspondingly, we want the reward associated with σ' to capture the total reward of all
 401 the true requests matched to σ' , and define its reward error accordingly. Finally, the time
 402 window error is defined as before, as a maximum over all pairs σ and σ' that are matched
 403 to each other. Our algorithm for the setting of Theorem 3 constructs a matching between
 404 I' and I in an online fashion. For this one to many setting, we solve instances of the
 405 orienteering problem rooted at each predicted request we visit. The performance of the
 406 algorithm accordingly worsens by a small constant factor and we achieve a competitive
 407 ratio of $O(\Lambda_M \rho_M^2 \log \min(D, L_{\max}))$ as before. Due to space limitations, the details of
 408 the proof are omitted from this version. See Section 8 of the full version [13] for further
 409 details.

410 ► **Theorem 13.** *Given an instance (G, I, I', M) of the TW-TSP with predictions where
 411 M is a many-to-one matching with errors as defined above, and satisfying $\tau_M \leq L_{\min}/2$
 412 and $\Lambda_M \leq L_{\min}/2$, there exists a polynomial-time online algorithm that takes the tuple
 413 (G, I', Λ_M) as offline input and I as online input, and constructs a walk $W \in \mathcal{W}(G)$ such
 414 that*

$$415 \quad \mathbb{E}[\text{Rew}(W, I)] \geq \frac{1}{O(\Lambda_M \cdot \rho_M^2 \cdot \log \min(D, L_{\max}))} \cdot \text{OPT}(G, I).$$

416 4 Relating the Optima

417 In this section we provide the proofs of Lemmas 4 and 5 that relate the optima over the
 418 true and the predicted request sequences, using service times as a mechanism to capture the
 419 prediction errors. We begin by proving that a service time of S can hurt the optimal by at
 420 most a factor of $2S - 1$.

421 ► **Lemma 4.** *For any instance (G, I) of the TW-TSP with service times, and any integer
 422 $S \leq L_{\min}$, we have*

$$423 \quad \text{OPT}(G, I, S) \geq \frac{1}{2S - 1} \cdot \text{OPT}(G, I, 1).$$

424 **Proof.** Let $W \in \mathcal{W}(G)$ be the walk that achieves the optimum $\text{OPT}(G, I, 1)$, and let the
 425 requests in I that are covered by W be denoted as $\sigma_i = (v_i, r_i, d_i, \pi_i)$ and ordered in the
 426 sequence in which they are covered by W . The lemma follows directly from the simple
 427 observation that if we don't service the $(S - 1)$ -requests *prior* and *after* some request σ_i ,
 428 then we can save enough time to service σ_i for S time-steps within its time window.

429 Formally, if $t_i \in [r_i, d_i - 1]$ is the step at which W begins servicing request σ_i , then by
 430 skipping the idle times on the $(S - 1)$ -previous and next requests we can remain idle on v_i
 431 from step $t_i - (S - 1)$ until step $t_i + S$ (since W already remained idle on v_i for 1 step) while
 432 still being able to keep up with walk W . Since $S \leq L_{\min}$, it is easy to verify that at least S
 433 of these time-steps are going to fall in the time-window $[r_i, d_i]$.

434 We now partition the requests $\sigma_i = (v_i, r_i, d_i, \pi_i)$ into $2S - 1$ sub-sequences, each of which
 435 starts at some request $i \in [S]$, and covers the requests $\sigma_i, \sigma_{i+(2S-1)}, \sigma_{i+2(2S-1)}$, and so forth.
 436 Each such sequence can be covered with a walk, with idle times built in as above, so as to be
 437 feasible for the instance (G, I, S) . Clearly, one of these walks obtains a reward of at least
 438 $\text{OPT}(G, I, 1)/(2S - 1)$, completing the proof. \blacktriangleleft

439 In the appendix of the full version, we show that the above lemma obtains a tight gap
 440 between the optima at different service times.

441 **► Lemma 14.** *For any pair of integers (L, S) such that $L \geq 2S - 2 \geq 1$, there exists a rooted*
 442 *instance (G, I) of the TW-TSP with service costs such that $L_{\min} = L$ and*

$$443 \quad \text{OPT}(G, I, S) = \frac{1}{2S - 1} \cdot \text{OPT}(G, I, 1).$$

444 Next, we provide the proof of Lemma 5 that relates the optima between the predicted
 445 and true request sequences, by appropriately addressing all three possible types of prediction
 446 errors.

447 **► Lemma 5.** *Let (G, I, I', M) be an instance of the TW-TSP with predictions, where $\Lambda_M,$
 448 $\rho_M,$ and τ_M denote the maximum location, reward, and time window errors of the instance
 449 respectively. Define $S := 4\Lambda_M + 1$ and $S' := 2\Lambda_M + 1$. Then, if $\tau_M \leq L_{\min}/2$ and
 450 $\Lambda_M \leq (L_{\min} - 1)/4$, we have*

$$451 \quad \text{OPT}(G, I', S') \geq \frac{1}{3\rho_M} \cdot \text{OPT}(G, I, S).$$

452 **Proof.** Let W be the walk that achieves the optimum $\text{OPT}(G, I, S)$, and let the requests
 453 in I covered by W be denoted as $\sigma_i = (v_i, r_i, d_i, \pi_i)$ and ordered in the sequence in which
 454 they are visited by W . Let $\sigma'_i = (v'_i, r'_i, d'_i, \pi'_i)$ denote the predicted request matched to σ_i ,
 455 that is, $\sigma'_i = M(\sigma_i)$. Observe that the total reward of all requests $\{\sigma'_i\}$ corresponding to
 456 $\sigma_i \in \text{Cov}(W, I, S)$ is at least $\text{Rew}(W, I, S)/\rho_M$.

457 We will consider a walk W' in G defined as follows. The walk W' follows W , visiting
 458 the requests σ_i in sequence. As soon as W starts servicing σ_i , W' takes a detour to
 459 visit σ'_i ; remains idle at σ'_i for S' time steps; returns back to σ_i ; remains idle at σ_i for
 460 $S - 2\ell(v_i, v'_i) - S' \geq 0$ time steps; and then resumes the walk W . Observe that W' is identical
 461 to W outside of the detours it takes to visit the σ'_i 's.

462 Our goal is to feasibly capture all of the reward contained in the σ'_i 's. The problem is
 463 that the walk W' may miss some of this reward due to the mismatch in the time windows of
 464 the true and predicted requests. To this end, we will consider two variations of the walk W' .
 465 Let $K := L_{\min}/2 \geq \tau_M$. The walk W'_1 is identical to W' except that it starts K steps after
 466 W' starts, and accordingly visits every location exactly K steps after W' visits it. The walk
 467 W'_2 is identical to W' except that it starts K steps before W' starts,⁷ and accordingly visits
 468 every location exactly K steps before W' visits it.

⁷ To be precise, this walk starts at the location where W' is at at step K .

469 Now consider some σ'_i corresponding to a request σ_i covered by W in the instance (G, I, S) .
 470 We claim that at least one of the walks W' , W'_1 , and W'_2 covers σ'_i in (G, I', S') . Let t be
 471 the time at which W' arrives at v'_i ; recall that W' remains at the node until at least $t + S'$.
 472 Note that $t \geq r_i$ and $t + S' \leq d_i$ due to $\sigma_i \in \text{Cov}(W, I, S)$.

473 First, suppose that $r'_i \leq t$ and $d'_i \geq t + S'$, then σ'_i is covered by W' in (G, I', S') . Next
 474 suppose that $r'_i > t$. Then, W'_1 arrives at v'_i at time $t + K \geq r_i + K \geq r_i + \tau_M \geq r'_i$. On
 475 the other hand, it remains at v'_i until time $t + K + S' < r'_i + K + S' \leq r'_i + L_{\min} \leq d'_i$.
 476 Therefore, σ'_i is covered by W'_1 . Finally, suppose that $d'_i < t + S'$. Then, W'_2 arrives at v'_i at
 477 time $t - K > d'_i - S' - K \geq d'_i - L_{\min} \geq r'_i$. On the other hand, it remains at v'_i until time
 478 $t - K + S' \leq d_i - K \leq d_i - \tau_M \leq d'_i$. Therefore, σ'_i is covered by W'_2 .

479 We get that at least one of W' , W'_1 , or W'_2 obtains at least a $1/3\rho_M$ fraction of
 480 $\text{OPT}(G, I, S)$, where the factor of ρ_M is lost due to the mismatch in the predicted rewards.
 481 The lemma follows directly from this. \blacktriangleleft

482 5 The offline approximation

483 In this section, we design an $O(\log \min(D, L_{\max}))$ deterministic and polynomial-time ap-
 484 proximation algorithm for the TW-TSP with service times, providing the proof of Lemma 6.
 485 Our proof relies on a series of reductions between different offline problems, applications of
 486 existing algorithms as well as the design of novel algorithmic components. We break up our
 487 argument into a series of lemmas. Due to space limitations, all the proofs are moved to the
 488 appendix of the full version [13].

489 1. First, we designing an $O(\log \min(D, L_{\max}))$ approximation algorithm for TW-TSP
 490 (without service times). Since the work of [14] already provides a $O(\log L_{\max})$ ap-
 491 proximation for the setting with integer time-windows (see Lemma 5.3 of [14]), it suffices
 492 to prove the following:

493 \blacktriangleright **Lemma 15.** *Given an instance of the TW-TSP (without service times) with $L_{\min} \geq 4D$,
 494 there exists a polynomial time algorithm that achieves an $O(1)$ approximation.*

495 Our proof relies on the observation that when time-windows are sufficiently large compared
 496 to the diameter of the graph, the problem essentially reduces to an instance of the
 497 well-studied Orienteering problem, for which constant approximation algorithms are
 498 known. We comment that similar ideas have been used in [5]. Then, it is straight-
 499 forward to combine this algorithm together with the algorithm of [14] to acquire an
 500 $O(\log \min(L_{\max}, D))$ approximation of TW-TSP.

501 \blacktriangleright **Lemma 16.** *There exists an $O(\log \min(D, L_{\max}))$ approximation algorithm for the
 502 TW-TSP problem.*

503 2. Next, we design a simple approximation-preserving reduction from TW-TSP with service
 504 times to TW-TSP (without service times). The main idea behind this reduction is to
 505 treat service times as edge lengths in an augmented graph whose diameter is roughly
 506 $D + S$. For instances with $S \leq D$, this increase becomes negligible and thus by combining
 507 our reduction with Lemma 16, we immediately get the following:

508 \blacktriangleright **Lemma 17.** *Given an instance (G, I, S) of the TW-TSP with service times such that
 509 $S \leq D$, there exists a polynomial time algorithm that achieves an $O(\log \min(D, L_{\max}))$
 510 approximation.*

511 3. Finally, we handle the case of large service times, specifically $S \geq D$. In that case, it
 512 turns out that we can reduce the instance to one over a uniform complete graph. Then,

513 the TW-TSP-S essentially becomes equivalent to the well-studied *Job Scheduling* problem,
514 for which constant approximations are known.

515 ▶ **Lemma 18.** *Given an instance (G, I, S) of the TW-TSP with service costs such that*
516 *$S \leq D$, there exists a polynomial time algorithm that achieves an $O(1)$ approximation.*

517 The proof of Lemma 6 follows immediately from Lemma 17 and Lemma 18. We comment that
518 any improvement in the best known approximation algorithm for TW-TSP will immediately
519 imply an improvement for all the results that this work presents. Lemma 18 essentially
520 enables us to assume that in all instances of interest, $S \leq D$. Under this assumption,
521 our reduction used in the proof of Lemma 17 essentially states that TW-TSP-S becomes
522 equivalent to TW-TSP in graphs of diameter $\mathcal{O}(D)$ and maximum window size $\mathcal{O}(L_{\max})$.
523 As an immediate corollary, given an offline TW-TSP algorithm that achieves a competitive
524 ratio of $\alpha(D, L_{\max})$, we immediately obtain an offline $O(\alpha(D, L_{\max}))$ approximation for
525 TW-TSP-S, that can be used in order to substitute Lemma 6 in our analysis and improve
526 the competitive ratio of Theorem 3.

527 6 The online algorithm

528 In this section we present an online algorithm that takes as input a pre-computed walk over
529 the predicted request sequence and solves an appropriate online matching problem in order
530 to construct detours that capture true requests, while taking into account the possible errors
531 in the predictions. The formal guarantee of our algorithm is given in Lemma 7, which we
532 restate for the reader's convenience:

533 ▶ **Lemma 7.** *Given an instance (G, I, I', M) of the TW-TSP with predictions satisfying*
534 *$\tau_M \leq L_{\min}/2$ and $\Lambda_M \leq (L_{\min} - 1)/4$; a walk $W' \in \mathcal{W}(G)$; and any integer $S' \geq 2\Lambda_M + 1$,*
535 *there exists an online algorithm (Algorithm 1) that returns a walk $W \in \mathcal{W}(G)$ with expected*
536 *reward*

$$537 \quad \mathbb{E}[\text{Rew}(W, I, 1)] \geq \frac{1}{6\rho_M} \cdot \text{Rew}(W', I', S').$$

538 We begin by establishing some notation. Let $W' \in \mathcal{W}(G)$ be any walk that services some
539 predicted requests in I' with a service time of S' . We use $\sigma'_i = (v'_i, r'_i, d'_i, \pi'_i)$ to denote the
540 predicted requests in I' that are covered by W' , ordered in the sequence in which they are
541 visited by W' . Likewise, we use $\sigma_i = (v_i, r_i, d_i, \pi_i) \in I$ to denote the true request matched
542 to the prediction σ'_i , that is, $\sigma'_i = M(\sigma_i)$.

543 At a high level, our algorithm follows the walk W' , but when it reaches a predicted
544 request σ'_i , it considers taking a detour to service a true request that is available at that
545 point of time. To this end, we define the set of "reachable" true requests as follows.

546 ▶ **Definition 19.** *Given a partial walk W that is at request $\sigma'_i \in I'$ at time t , we define the*
547 *set of reachable requests $R_i(W, t)$ to be the set of all $\sigma \in I$ such that:*

- 548 1. $r_\sigma \leq t \leq d_\sigma - \ell(v'_i, v_\sigma) - 1$, and
- 549 2. $2\ell(v'_i, v_\sigma) + 1 \leq S'$.

550 Our algorithm considers all of the reachable requests that have not been covered by the
551 walk as yet, chooses the one with the highest reward, and takes a detour to visit and cover
552 the request, before returning to σ'_i and resuming the walk. In order to deal with time window
553 errors, our algorithm starts the walk a little early, or on time, or a little late, as in the proof
554 of Lemma 5. The algorithm is described below formally.

■ **Algorithm 1** Online algorithm for TSP-TW with predictions.

Offline input: Graph G , predicted requests I' , walk $W' \in \mathcal{W}(G)$, service times S' .

Online input: True requests I .

Output: Walk $W \in \mathcal{W}(G)$.

- 1: Let $K = L_{\min}/2$. Select ϵ uniformly at random from $\{-1, 0, 1\}$.
- 2: Define the set of covered requests $C = \emptyset$.
- 3: **for** $i \leftarrow 1$ to $|\text{Cov}(W', I', S')|$ **do**
- 4: Let t'_i denote the time at which W' visits σ'_i .
- 5: Set $t_i \leftarrow t'_i + \epsilon K$.
- 6: Visit v'_i at time t_i .
- 7: Construct the set $R_i(W, t_i)$ of requests in I reachable at time t_i .
- 8: **if** $R_i(W, t_i) \setminus C = \emptyset$ **then**
- 9: Do nothing.
- 10: **else**
- 11: Let $\hat{\sigma}$ be the highest reward request in $R_i(W, t_i) \setminus C$.
- 12: Visit $v_{\hat{\sigma}}$; spend one unit of idle time at $v_{\hat{\sigma}}$; return to v'_i .
- 13: Set $C \leftarrow C \cup \{\hat{\sigma}\}$.

555 We begin our analysis by noting that the walk W constructed by the algorithm is always
 556 able to visit the vertices v'_i corresponding to requests $\sigma'_i \in \text{Cov}(W', I', S')$ feasibly at the
 557 desired times t_i . This is because, by construction, the length of the detours that the walk
 558 W takes in Step 12 is always at most S' – the amount of idle time W' spends at v'_i – by
 559 virtue of the fact that $\hat{\sigma} \in R_i(W, t_i)$ and therefore, $2\ell(v'_i, v_{\hat{\sigma}}) + 1 \leq S'$. Therefore, all of the
 560 requests $\hat{\sigma}$ visited in Step 12 are indeed visited by the walk W .

561 We now relate the total reward covered by W to the reward contained in the true
 562 requests σ_i corresponding to $\sigma'_i \in \text{Cov}(W', I', S')$. To do so, we first note that with constant
 563 probability each such request is reachable by W .

564 \triangleright **Claim 20.** For each i , $\sigma_i \in R_i(W, t_i)$ with probability at least $1/3$.

565 **Proof.** Recall that by definition we have $\sigma'_i = M(\sigma_i)$ and so, $2\ell(v'_i, v_i) + 1 \leq 2\Lambda_M + 1 \leq S'$.
 566 So the request σ always satisfies the second requirement in the definition of the reachable
 567 set $R_i(W, t_i)$. Let us now consider the first requirement and recall that $t_i = t'_i + \epsilon K$ where
 568 $\epsilon \in \{-1, 0, 1\}$. We will now argue that $t_i \in [r_i, d_i - 1 - \ell(v'_i, v_i)]$ for at least one of the three
 569 choices of ϵ . The claim then follows from the uniformly random choice of ϵ .

- 570 1. If $t'_i \in [r_i, d_i - 1 - \ell(v'_i, v_i)]$, then the claim holds for $\epsilon = 0$ and $t_i = t'_i$.
- 571 2. Suppose that $t'_i < r_i$. Then, for $\epsilon = 1$ we have that $t_i = t'_i + K \geq r'_i + \tau_M \geq r_i$, and also
 572 $t_i = t'_i + K < r_i + K < d_i - L_{\min} + L_{\min}/2$ and thus $t_i = t'_i + K \leq d_i - 1 - \ell(v'_i, v_i)$ since
 573 $\ell(v'_i, v_i) \leq \Lambda_M \leq L_{\min}/2$. Thus, in this case we have $t_i = t'_i + K \in [r_i, d_i - 1 - \ell(v'_i, v_i)]$
 574 with the choice of $\epsilon = 1$.
- 575 3. Finally, suppose that $t'_i > d_i - 1 - \ell(v'_i, v_i)$. Then, for $\epsilon = -1$ we have that $t_i = t'_i - K \leq$
 576 $d'_i - S' - \tau_M \leq d_i - S'$ and thus $t_i - K \leq d_i - 1 - \ell(v'_i, v_i)$ since $S' \geq 2\Lambda_M + 1 \geq \ell(v'_i, v_i) + 1$.
 577 Also, $t_i = t'_i - K > d_i - 1 - \ell(v_i, v'_i) - L_{\min}/2 > r_i + L_{\min}/2 - \ell(v_i, v'_i) - 1$ and thus
 578 $t'_i - K \geq r_i$, since $\ell(v_i, v'_i) \leq \Lambda_M \leq L_{\min}/2$. Thus, in this case we have obtained that
 579 $t_i = t'_i - K \in [r_i, d_i - 1 - \ell(v'_i, v_i)]$ with the choice of $\epsilon = -1$.

580 ◀

581 We are now ready to prove Lemma 7 via a matching-type argument. To account for the
 582 reward covered by the walk W constructed by the algorithm, we will employ a standard
 583 charging scheme. Every time the algorithm takes a detour to cover some true request $\hat{\sigma}$ from
 584 a predicted request σ'_i in Step 12, we will credit half of the earned reward $\pi_{\hat{\sigma}}$ to $\hat{\sigma}$ itself, and
 585 half of the reward to the request σ_i . Formally, let $\text{Cr}(\sigma)$ denote the total credit received by
 586 $\sigma \in I$. Then during Step 12 we will increment both $\text{Cr}(\hat{\sigma})$ and $\text{Cr}(\sigma_i)$ by $\pi_{\hat{\sigma}}/2$.

587 Now consider some $\sigma_i \in I$ corresponding to $\sigma'_i \in \text{Cov}(W', I', S')$. By Claim 20, this
 588 request is in $R_i(W, t_i)$ with probability at least $1/3$. If at time t_i , the request has already
 589 been covered by W , then we get $\text{Cr}(\sigma_i) \geq \pi_i/2$. Otherwise, we pick a $\hat{\sigma} \in R_i(W, t_i)$ with
 590 $\pi_{\hat{\sigma}} \geq \pi_i$, and therefore, once again we get $\text{Cr}(\sigma_i) \geq \pi_i/2$.

591 Putting everything together, we get

$$\begin{aligned}
 592 \quad \mathbb{E}[\text{Rew}(W, I, S)] &= \mathbb{E} \left[\sum_{\sigma \in I} \text{Cr}(\sigma) \right] \geq \sum_{i: \sigma'_i \in \text{Cov}(W', I', S')} \mathbb{E} \left[\frac{\pi_i}{2} \mathbb{1}[\sigma_i \in R_i(W, t_i)] \right] \\
 593 &\geq \frac{1}{3} \cdot \sum_{i: \sigma'_i \in \text{Cov}(W', I', S')} \frac{\pi_i}{2} \\
 594 &\geq \frac{1}{6} \cdot \frac{1}{\rho_M} \cdot \sum_{i: \sigma'_i \in \text{Cov}(W', I', S')} \pi'_i \\
 595 &= \frac{1}{6\rho_M} \cdot \text{Rew}(W', I', S')
 \end{aligned}$$

596 This completes the proof of the lemma.

597 **7 Lower bounds**

598 In this section, we present lower bounds that complement our results. First, we will motivate
 599 the need for predictions in Section 7.1. Then, in Section 7.2 we will show that the competitive
 600 ratio of TW-TSP with predictions must scale linearly with the error in locations. Finally, in
 601 Section 7.3 we argue the need for non-zero service times in the definition of TW-TSP with
 602 predictions.

603 **7.1 Lower bounds for online TW-TSP without predictions**

604 We argue that Online TW-TSP does not admit any reasonable competitive ratio in the
 605 absence of predictions. In the case of deterministic algorithms where their entire behavior
 606 is predictable, simple instances with only 2 vertices and appropriately small time-windows
 607 suffice to argue that no bounded guarantee for the approximation ratio is achievable.

608 **► Lemma 21.** *The competitive ratio of any deterministic online algorithm for Online*
 609 *TW-TSP on instances with $L_{\min} \leq D$ is unbounded.*

610 **Proof.** Let DET be any deterministic algorithm and let G be the line graph with just two
 611 vertices v_1, v_2 connected via an edge of length D . Since DET is deterministic, we can assume
 612 knowledge of its position at any step t as soon as we have specified all requests with release
 613 time $\leq t$. We will now construct a request sequence I that uses the information.

614 For the first D time-steps, we don't release any request. Then, at $t = D$, let $v_D \in \{v_1, v_2\}$
 615 be the position that DET's walk is currently at and likewise let v'_D be the other vertex of
 616 G . We construct the first request to be $\sigma = (v'_D, D, D + L, 1)$ for any $L \leq D$. Clearly, DET
 617 cannot service this request as even for $L = D$ it arrives on v'_D at deadline and cannot service

618 it for one step. Then, we don't release any new request for the next $2D$ steps, and at $t = 3D$
 619 we repeat the same process, by requesting the vertex that DET doesn't currently occupy.
 620 Likewise, we repeat the same process at $t = 5D$, $t = 7D$ etc. Independently of the size of our
 621 request sequence, the total reward collected by DET is 0.

622 On the other hand, it is not hard to see that if our request sequence I has N requests in
 623 total, then $\text{OPT}(G, I, 1) = N$. This is due to the fact that requests are spaced $2D$ -steps from
 624 each other, and thus an optimal offline algorithm that had knowledge of the entire sequence
 625 in advance would always be able to arrive at each request on time, servicing it within its
 626 respective time-window. ◀

627 For randomized algorithms, a slight improvement can be achieved. In particular, the
 628 randomized algorithm that picks a vertex uniformly at random and then remains idle on it
 629 for the entire sequence achieves a competitive ratio of $1/n$. It turns out that this is actually
 630 the best possible ratio that a randomized algorithm can achieve on Online TW-TSP:

631 ▶ **Theorem 8.** *The competitive ratio of any randomized online algorithm for Online TW-TSP*
 632 *on instances with $L_{\min} \leq D$ is at most $1/n$.*

633 **Proof.** Let $G(V, E, \ell)$ be the uniform complete graph of $n = |V|$ vertices, where all edges have
 634 a length of D . Fix any integer N and let vertices $v_1, v_2, \dots, v_N \in V$ be drawn independently
 635 and uniformly at random. Next, we fix any window length $L \leq D$ and consider the (random)
 636 request sequence on these vertices $I = \{\sigma_i\}_{i=1}^N$ for $\sigma_i = (v_i, (2i-1)D, (2i-1)D + L, 1)$. From
 637 Yao's mininmax principle, a lower bound on the (expected) competitive ratio of deterministic
 638 algorithms on this randomized instance will imply the same lower bound for randomized
 639 algorithms.

640 Since the time-windows are spaced $2D$ -away from each other, it is not hard to see that
 641 for any realization of I , $\text{OPT}(G, I, 1) = N$. On the other hand, since $D \geq L$, we get that the
 642 only way to service a request is to be on its corresponding vertex on release time. Since the
 643 vertices are random, for any deterministic algorithm this happens with probability precisely
 644 $1/n$, and thus the expected reward of any deterministic algorithm on this instance is N/n ,
 645 proving the claim. ◀


646 7.2 Tight dependence on location error

647 In this section we show that a linear dependency on the location error is unavoidable for any
 648 randomized online algorithm for the TW-TSP with predictions, even assuming exponential
 649 computational power. In other words, we formally prove Theorem 10, which we re-state for
 650 the reader's convenience:

651 ▶ **Theorem 10.** *For any $S > 0$, there exists an instance (G, I, I', M) of the TW-TSP with*
 652 *predictions satisfying $\tau_M = 0$, $\rho_M = 1$, and $\Lambda_M = S$ such that the competitive ratio of any*
 653 *randomized online algorithm taking the tuple (G, I', Λ_M) as offline input and I as online*
 654 *input asymptotically approaches $1/(S + 1)$.*

655 **Proof.** Fix any $S > 0$ and let K, C and N be integer parameters that will be specified later.
 656 We construct a graph $G = \cup_{i=0}^{N-1} G_i$ that consists of N copies G_0, \dots, G_{N-1} of the complete
 657 graph on C vertices with all edge lengths equal to S , arranged in a way so that each vertex
 658 in G_i connects to each vertex in G_{i+1} with an edge of length KS . A pictorial example for
 659 small values of C and N is shown in Figure 1.

660 Next, we select independently and uniformly at random one vertex v_i from each sub-
 661 graph G_i and construct the (randomized) request sequence $I = \{\sigma_i\}_{i=0}^{N-1}$ where we denote



■ **Figure 1** An example of G for $C = 3$ and $N = 4$. Black edges have a length of S and red edges have a length of KS .

662 $\sigma_i = (v_i, r_i, r_i + KS, 1)$ for $r_i = i \cdot (KS + 1)$. As for the predictions, we simply construct a
 663 second instance I' in the same manner and offer it as an offline prediction for I . Observe that
 664 for all possible constructions of I and I' it holds that there exists a matching M between
 665 them with $\tau_M = 0$, $\rho_M = 1$ and $\Lambda_M = S$, namely the matching that pairs together vertices
 666 from the same sub-graphs. This prediction provides no information to the algorithm other than the fact that the (true) instance I was constructed via the above randomized approach.
 667 Also, notice that the location error Λ_M can be arbitrarily small compared to the window
 668 length $L = KS$ by choosing appropriately large K .

670 It is easy to see that for all possible realizations of I it holds that $\text{OPT}(G, I, 1) = N$.
 671 Indeed, consider the walk that starts from the (random) vertex v_0 , remains idle for 1 step and
 672 then visits vertex v_1 , remains idle for one step, visits v_2 , etc. Such a walk would visit each
 673 vertex v_i at step $t = i \cdot KS + i = r_i$ and thus would service all the requests in I , achieving
 674 a total reward of N . Next, we will show that the expected reward of any deterministic
 675 algorithm on the random sequence I approaches N/S . Using Yao's minimax principle, this
 676 will immediately translate to a lower bound that approaches $1/S$ for randomized algorithms,
 677 completing the proof of the theorem.

678 Fix any deterministic algorithm for TW-TSP with predictions on instance (G, I, I', M) .
 679 Note that since the predictions I' supply zero information, it suffices to analyze the algorithm
 680 as a deterministic online algorithm for Online TW-TSP on instance I . The key observation
 681 is that due to the fact that both the time windows of the requests and the edges that connect
 682 different sub-graphs have a length of KS , it is impossible for the algorithm to know on what
 683 vertex v_i of subgraph G_i the request is going to arrive before visiting some possibly different
 684 vertex of the same subgraph. In particular, if the algorithm is in subgraph G_j for $j < i$ at
 685 time r_i , then it cannot reach vertex v_i before the time window of request i ends. Thus, in
 686 order for any deterministic algorithm to serve the request on sub-graph G_i , it first has to
 687 visit some vertex v'_i of G_i . If it so happens that $v'_i = v_i$ then it can immediately service the
 688 request, otherwise it has to travel a distance of S in order to reach v_i .

689 We partition the set of requests into two sets N_+ and N_- based on whether the determ-
 690 inistic algorithm happens to arrive on the correct vertex of the sub-graph or not. All the
 691 requests in N_+ can be serviced without any extra delay, exactly as done by the optimal walk.
 692 On the other hand, servicing a request in N_- requires the algorithm to spend an extra time
 693 of S in order to transition to the right vertex. Since the time-windows have a length of KS ,
 694 this can be done at most K times before the algorithm runs out of slack to spare. When this
 695 happens, the algorithm would have to skip the next S requests in order to recover enough

696 slack to fix its next mistake.

697 Formally, consider the last request in N_- that the algorithm feasibly serves, call it l .
 698 Let A be the number of requests in N_- the algorithm serves through extra delay prior to l ,
 699 and let B be the number of requests the algorithm skips in N_+ or N_- prior to l . Then in
 700 order for the algorithm to have reached l before its time window ends, it must be the case
 701 that the total extra delay incurred by the algorithm, namely $AS - B$, is no more than KS .
 702 Rearranging we get:

$$703 \quad A(S + 1) - (A + B) \leq KS, \quad \text{or,} \quad A \leq \frac{A + B}{S + 1} + \frac{KS}{S + 1} < \frac{N}{S + 1} + K$$

704 Thus, we get that the total expected reward gathered by the algorithm is at most

$$705 \quad \mathbb{E}[|N_+|] + A + 1 \leq \mathbb{E}[|N_+|] + \frac{N}{S + 1} + K + 1$$

706 Finally, using the fact that $\mathbb{E}[|N_+|] = N/C$, we get that the competitive ratio of any
 707 deterministic algorithm on the random instance I is at most

$$708 \quad \frac{1}{C} + \frac{K + 1}{N} + \frac{1}{S + 1}$$

709 Choosing $N \gg K$ and C sufficiently large, we can make this competitive ratio asymptotically
 710 approach $1/(S + 1)$ as desired. \blacktriangleleft

711 7.3 Comparing the optimal with and without service times

712 As we saw in Lemma 4, the gap between $\text{OPT}(G, I, 1)$ and $\text{OPT}(G, I, S)$ depends linearly on
 713 the service time S . In this section, we study the gap between $\text{OPT}(G, I, 1)$ and $\text{OPT}(G, I, 0)$,
 714 showing that there exists a much sharper separation between them.

715 \blacktriangleright **Theorem 22.** *For any integers L and D , $L \leq D$, there exists an offline instance (G, I) of*
 716 *the TW-TSP where D is the diameter of the network and every request has a time window*
 717 *length equal to L , such that $\text{OPT}(G, I, 1) \leq \frac{L}{D+1} \cdot \text{OPT}(G, I, 0)$.*

718 **Proof.** Consider the line graph G with vertices v_0 through v_D connected sequentially via
 719 edges of length 1. Clearly, the diameter of G is D . Next, consider the sequence of $(D + 1)$ -
 720 requests $I = \{\sigma_i\}_{i=0}^D$ where $\sigma_i = (v_i, i, L + i, 1)$. Observe that $\text{OPT}(G, I, 0) = D + 1$ as
 721 simply following the walk from v_0 to v_D will cover all the requests if the service costs are 0.

722 On the other hand, it is not very hard to see that $\text{OPT}(G, I, 1) = L$. First, observe that
 723 without loss we can assume that the optimal walk starts on v_0 . If not, let σ_i be the *first*
 724 request served by the optimal. Since σ_i cannot be served prior to step $r_i = i$, we could
 725 instead start at v_0 and take i steps in order to reach v_i and then follow the original walk,
 726 achieving precisely the same reward.

727 Our argument is completed by the simple observation that any walk that starts from v_0
 728 and has served x requests with service cost 1 *cannot* visit vertex v_j prior to step $j + x$. Thus,
 729 as soon as x becomes L all the future time-windows will be missed. \blacktriangleleft

730 Theorem 22 states that $\text{OPT}(G, I, 0)$ is a much stronger benchmark than $\text{OPT}(G, I, 1)$.
 731 However, it doesn't exclude the possibility of designing an algorithm that is competitive
 732 against this stronger benchmark $\text{OPT}(G, I, 0)$. We will next show that no randomized online
 733 algorithm with predictions can obtain a better than linear competitive ratio against this
 734 benchmark. Intuitively, requiring the algorithm to spend non-zero time at each request

735 also allows the algorithm to recover from possible mistakes due to the prediction errors by
 736 skipping requests that the optimum services with a delay of 1. A similar approach is not
 737 possible in the 0 service time setting.

738 ► **Theorem 9.** *For any $S > 0$, there exists an instance (G, I, I', M) of the TW-TSP with
 739 predictions and service times 0, satisfying $\tau_M = 0$, $\rho_M = 1$, and $\Lambda_M = S$, such that any
 740 randomized online algorithm taking the tuple (G, I', Λ_M) as offline input and I as online
 741 input achieves a reward no larger than $O(1/n) \cdot \text{OPT}(G, I, 0)$. Here n is the number of
 742 vertices in G .*

743 **Proof.** We construct the same graph $G = \cup_{i=0}^{N-1} G_i$ as in Theorem 10, that consists of N
 744 copies of the complete graph with edge lengths S , connected sequentially with edges of length
 745 KS (see Figure 1). As for the request sequence, we once again select independently and
 746 uniformly at random one vertex v_i from each sub-graph G_i and construct the (randomized)
 747 request sequence $I = \{\sigma_i\}_{i=0}^{N-1}$ where $\sigma_i = (v_i, iKS, (i+1)KS - 1, 1)$; notice that release
 748 times are slightly different from Theorem 10 to account for the absence of service costs.

749 For the predictions, we simply construct a second instance I' in the same manner and
 750 offer it as an offline prediction for I . By matching the requests on the same sub-graph
 751 together, we get $\tau_M = 0$, $\rho_M = 1$ and $\Lambda_M = S$. As it clearly holds that $\text{OPT}(G, I, 0) = N$
 752 for any realization of I , to prove our theorem it suffices to argue that any deterministic
 753 algorithm for TW-TSP with predictions on instance (G, I, I', M) gets an expected reward of
 754 $\mathcal{O}(N/n)$ and apply Yao's *minimax principle*. Furthermore, since the prediction I' does not
 755 provide any information on I , it suffices to bound the reward of deterministic algorithms for
 756 Online TW-TSP on (online) instance I .

757 Fix any deterministic algorithm for Online TW-TSP on instance I . Observe that since
 758 edges between different sub-graphs have a length of KS and time-windows have a length of
 759 $KS - 1$, it is impossible for the algorithm to service some request σ_i unless at $t = r_i$ it is
 760 already in some vertex of sub-graph G_i . For the same reason, it is not possible to service any
 761 request σ_j after servicing some other request σ_i with $i > j$. Finally, since all requests can
 762 be reached by their release time if the algorithm starts from a vertex in G_0 , we can assume
 763 without loss that this is indeed the case.

764 We partition our set of N requests into two sets N_+ and N_- based on whether the
 765 deterministic algorithm *happens* to arrive on the correct vertex of the sub-graph before the
 766 request's release time or not. From the random construction of our instance, we have that
 767 $\mathbb{E}[|N_+|] \leq N/C$. For requests in N_- , if the algorithm wishes to service them it has to take
 768 a detour of length S in order to reach the correct vertex. Our proof relies on the fact that
 769 after servicing K requests in N_- , the algorithm can no longer service any other request. To
 770 see this, let v_j be the K -th request in N_- that was serviced by the deterministic algorithm.
 771 As we can already established, any request v_i with $i < j$ can no longer be serviced. On
 772 the other hand, since without loss the algorithm starts at a vertex of G_0 , just reaching a
 773 vertex in G_i while taking K detours of length S requires at least $iKS + K > d_i$ time-steps.
 774 Putting everything together, we get that the expected reward of the algorithm is at most
 775 $N/C + K = \mathcal{O}(N/n)$ by setting $K = \mathcal{O}(1)$ and $N = 2K$, since $n = NC$. ◀

776 ——— References ———

- 777 1 Antonios Antoniadis, Themis Gouleakis, Pieter Kleer, and Pavel Kolev. Secretary and online
 778 matching problems with machine learned advice. In *Annual Conference on Neural Information
 779 Processing Systems (NeurIPS)*, 2020.
- 780 2 Esther M. Arkin, Joseph B. M. Mitchell, and Giri Narasimhan. Resource-constrained geometric
 781 network optimization. In *Symposium on Computational Geometry (SoCG)*, 1998.

- 782 3 Yossi Azar, Stefano Leonardi, and Noam Touitou. Flow time scheduling with uncertain
783 processing time. In *Symposium on the Theory of Computing (STOC)*, 2021.
- 784 4 Yossi Azar, Debmalya Panigrahi, and Noam Touitou. Online graph algorithms with predictions.
785 In *ACM-SIAM Symposium on Discrete Algorithms (SODA)*, 2021.
- 786 5 Yossi Azar and Adi Vardi. TSP with time windows and service time. *CoRR*, abs/1501.06158,
787 2015.
- 788 6 Étienne Bamas, Andreas Maggiori, and Ola Svensson. The primal-dual method for learning
789 augmented algorithms. In *Annual Conference on Neural Information Processing Systems*
790 (*NeurIPS*), 2020.
- 791 7 Nikhil Bansal, Avrim Blum, Shuchi Chawla, and Adam Meyerson. Approximation algorithms
792 for deadline-TSP and vehicle routing with time-windows. In *Symposium on the Theory of*
793 *Computing (STOC)*, 2004.
- 794 8 Reuven Bar-Yehuda, Guy Even, and Shimon Shahar. On approximating a geometric prize-
795 collecting traveling salesman problem with time windows. *J. Algorithms*, 55:76–92, 2005.
- 796 9 Magnus Berg, Joan Boyar, Lene M. Favrholdt, and Kim S. Larsen. Online minimum spanning
797 trees with weight predictions. In *Workshop on Algorithms and Data Structures*, 2023.
- 798 10 Giulia Bernardini, Alexander Lindermayr, Alberto Marchetti-Spaccamela, Nicole Megow, Leen
799 Stougie, and Michelle Sweering. A universal error measure for input predictions applied to
800 online graph problems. In *Annual Conference on Neural Information Processing Systems*
801 (*NeurIPS*), 2022.
- 802 11 Avrim Blum, Shuchi Chawla, David Karger, Terran Lane, and Maria Minkoff. Approximation
803 algorithms for orienteering and discounted-reward TSP. In *Foundations of Computer Science*
804 (*FOCS*), 2003.
- 805 12 Niv Buchbinder, Yaron Fairstein, Konstantina Mellou, Ishai Menache, and Joseph (Seffi)
806 Naor. Online virtual machine allocation with lifetime and load predictions. In *International*
807 *Conference on Measurement and Modeling of Computer Systems (SIGMETRICS)*, 2021.
- 808 13 Shuchi Chawla and Dimitris Christou. Online TSP with predictions. *CoRR*, abs/2304.01958,
809 2024.
- 810 14 Chandra Chekuri, Nitish Korula, and Martin Pál. Improved algorithms for orienteering and
811 related problems. In *ACM-SIAM Symposium on Discrete Algorithms (SODA)*, 2008.
- 812 15 Chandra Chekuri and Amit Kumar. Maximum coverage problem with group budget con-
813 straints and applications. In *International Workshop on Approximation, Randomization, and*
814 *Combinatorial Optimization (APPROX)*, 2004.
- 815 16 Ke Chen and Sarel Har-Peled. The orienteering problem in the plane revisited. In *Symposium*
816 *on Computational Geometry (SoCG)*, 2006.
- 817 17 Michael Dinitz, Sungjin Im, Thomas Lavastida, Benjamin Moseley, and Sergei Vassilvitskii.
818 Faster matchings via learned duals. In *Annual Conference on Neural Information Processing*
819 *Systems (NeurIPS)*, 2021.
- 820 18 Paul Dütting, Silvio Lattanzi, Renato Paes Leme, and Sergei Vassilvitskii. Secretaries with
821 advice. In *ACM Conference on Economics and Computation (EC)*, 2021.
- 822 19 Jon C. Ergun, Zhili Feng, Sandeep Silwal, David Woodruff, and Samson Zhou. Learning-
823 augmented k -means clustering. In *International Conference on Learning Representations*
824 (*ICLR*), 2022.
- 825 20 Thomas Wilhelm Erlebach, Murilo Santos de Lima, Nicole Megow, and Jens Schloter. Learning-
826 augmented query policies for minimum spanning tree with uncertainty. In *Embedded Systems*
827 *and Applications (ESA)*, 2022.
- 828 21 Jie Gao, Su Jia, Joseph S. B. Mitchell, and Lu Zhao. Approximation algorithms for time-
829 window TSP and prize collecting TSP problems. In *Workshop on the Algorithmic Foundations*
830 *of Robotics (WAFR)*, 2016.
- 831 22 Themis Gouleakis, Konstantinos Lakis, and Golnoosh Shahkarami. Learning-augmented
832 algorithms for online TSP on the line. *CoRR*, abs/2206.00655, 2022.

- 833 23 Hsiao-Yu Hu, Hao-Ting Wei, Meng-Hsi Li, Kai-Min Chung, and Chung-Shou Liao. Online
834 TSP with predictions. *CoRR*, abs/2206.15364, 2022.
- 835 24 Sungjin Im, Ravi Kumar, Mahshid Montazer Qaem, and Manish Purohit. Non-clairvoyant
836 scheduling with predictions. In *ACM Symposium on Parallelism in Algorithms and Architectures*
837 (*SPAA*), 2021.
- 838 25 Zhihao Jiang, Debmalya Panigrahi, and Kevin Sun. Online algorithms for weighted paging
839 with predictions. In *International Colloquium on Automata, Languages, and Programming*
840 (*ICALP*), 2020.
- 841 26 Yoshiyuki Karuno and Hiroshi Nagamochi. 2-approximation algorithms for the multi-vehicle
842 scheduling problem on a path with release and handling times. *Discret. Appl. Math.*, 129:433–
843 447, 2003.
- 844 27 Silvio Lattanzi, Thomas Lavastida, Benjamin Moseley, and Sergei Vassilvitskii. Online
845 scheduling via learned weights. In *ACM-SIAM Symposium on Discrete Algorithms (SODA)*,
846 2020.
- 847 28 Thomas Lavastida, Benjamin Moseley, R. Ravi, and Chenyang Xu. Learnable and instance-
848 robust predictions for online matching, flows and load balancing. In *European Symposium on*
849 *Algorithms (ESA)*, 2021.
- 850 29 Alexander Lindermayr and Nicole Megow. Algorithms with predictions. [https://algorithms-
851 with-predictions.github.io/](https://algorithms-with-predictions.github.io/).
- 852 30 Thodoris Lykouris and Sergei Vassilvitskii. Competitive caching with machine learned advice.
853 *J. ACM*, 68(4):24:1–24:25, 2021.
- 854 31 Dhruv Rohatgi. Near-optimal bounds for online caching with machine learned advice. In
855 *ACM-SIAM Symposium on Discrete Algorithms (SODA)*, 2020.
- 856 32 John N. Tsitsiklis. Special cases of traveling salesman and repairman problems with time
857 windows. *Networks*, 22(3):263–282, 1992.
- 858 33 Chenyang Xu and Benjamin Moseley. Learning-augmented algorithms for online steiner tree.
859 *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI)*, 36:8744–8752, 2022.