

# Learning-Augmented Maximum Flow

Adam Polak<sup>a,b,\*</sup>, Maksym Zub<sup>b</sup>

<sup>a</sup>Max-Planck Institute for Informatics, Saarland Informatics Campus, 66123, Saarbrücken, Germany

<sup>b</sup>Jagiellonian University, ul. Łojasiewicza 6, 30-348, Kraków, Poland

---

## Abstract

We propose a framework for speeding up maximum flow computation by using predictions. A prediction is a flow, i.e., an assignment of non-negative flow values to edges, which satisfies the flow conservation property, but does not necessarily respect the edge capacities of the actual instance (since these were unknown at the time of learning). We present an algorithm that, given an  $m$ -edge flow network and a predicted flow, computes a maximum flow in  $O(m\eta)$  time, where  $\eta$  is the  $\ell_1$  error of the prediction, i.e., the sum over the edges of the absolute difference between the predicted and optimal flow values. Moreover, we prove that, given an oracle access to a distribution over flow networks, it is possible to efficiently PAC-learn a prediction minimizing the expected  $\ell_1$  error over that distribution. Our results fit into the recent line of research on learning-augmented algorithms, which aims to improve over worst-case bounds of classical algorithms by using predictions, e.g., machine-learned from previous similar instances. So far, the main focus in this area was on improving competitive ratios for online problems. Following Dinitz et al. (NeurIPS 2021), our results are among the firsts to improve the running time of an offline problem.

*Keywords:* combinatorial optimization, maximum flow, algorithms with predictions

---

## 1. Introduction

Computing a maximum  $s$ - $t$  flow in a flow network (i.e., in a directed graph with nonnegative edge capacities and designated source and sink nodes) is a basic problem in combinatorial optimization. It is a building block of a number of more advanced algorithms, with relevance both in theory (e.g., in graph algorithms and scheduling) and in practice (e.g., in computer vision and transport).

Imagine we are to solve multiple similar instances of the maximum flow problem, e.g., the instances are drawn at random from a distribution, or they are snapshots of a single underlying instance changing over time. Can we learn what a typical optimal solution looks like, and then use it to speed up further computations? Or, to put it differently, assume we have a solution – e.g., obtained from past data or computed by a very fast heuristic – that is not necessarily optimal, maybe not even feasible, but close to an optimal solution. How can we use such an imperfect solution to warm-start a maximum flow algorithm and get a better running time?

Warm-starting maximum flow algorithms has been studied in the past experimentally (e.g., in computer vision, where maximum flow is often used to compute minimum cuts in subsequent frames of a video [19, 18]). In contrast, we propose an approach with theoretical guarantees.<sup>1</sup>

Learning-augmented algorithms (also called *algorithms with predictions*) are the subject of a recent line of research that aims to improve over worst-case bounds of classical algorithms by using possibly imperfect predictions. So far, the main focus in this area was on improving competitive ratios for online problems. Dinitz et al. [6] took a first step to explore improving the running times of offline problems. They gave an algorithm for the weighted bipartite matching problem that uses a learned dual solution to improve over the running time of the classic Hungarian algorithm. Our approach draws inspiration from their work, but it differs significantly in two aspects. First, we learn a primal, not a dual solution. Second, we impose an additional restriction on the learned solution, namely the flow conservation prop-

---

\*Corresponding author

Email addresses: apolak@mpi-inf.mpg.de (Adam Polak), max.zub@student.uj.edu.pl (Maksym Zub)

<sup>1</sup>Interestingly, while the theoretical guarantees of our algorithms are linear in the  $\ell_1$  error of flow predictions, the empirical performance of the heuristic warm-starting approaches [19, 18] happens to be roughly linear in the  $\ell_0$  change in edge capacities.

erty. This restriction makes our learning problem harder and the subsequent algorithmic problem easier. In Section 1.2 we discuss these differences in greater depth.

### 1.1. Our results

We propose a framework for speeding up maximum flow computation by using predicted flow values. Here, by *prediction* we mean a flow, which satisfies the flow conservation property, but does not necessarily respect the edge capacities of the actual instance (since these were unknown at the time of learning). We present an algorithm that, given a flow network  $G = (V, E)$  with edge capacities  $c \in \mathbb{Z}_{\geq 0}^E$ , and a predicted flow  $f \in \mathbb{Z}_{\geq 0}^E$ , computes a maximum flow  $f_c^* \in \mathbb{Z}_{\geq 0}^E$  in  $O(|E| \cdot \eta)$  time, where  $\eta = \|f - f_c^*\|_1 = \sum_{e \in E} |f(e) - f_c^*(e)|$  is the  $\ell_1$  error of the prediction. Moreover, we prove that, given an oracle access to a (joint) distribution over edge capacities, it is possible to efficiently PAC-learn a prediction minimizing the expected  $\ell_1$  error over that distribution.

To formally state our results, let us first define the maximum flow problem and related concepts.

**Definition 1.** Given a directed graph  $G = (V, E)$ , source and sink nodes  $s, t \in V$ , and nonnegative integral edge capacities  $c : E \rightarrow \mathbb{Z}_{\geq 0}$ , the *maximum flow problem* asks to find a function  $f : E \rightarrow \mathbb{Z}_{\geq 0}$ , assigning nonnegative integral flow to the edges, that satisfies

- *capacity constraints*:  $\forall e \in E f(e) \leq c(e)$ , and
- *flow conservation*:

$$\forall v \in V \setminus \{s, t\} \sum_{u: (u,v) \in E} f(u, v) = \sum_{u: (v,u) \in E} f(v, u)$$

and maximizes the *flow value* defined as  $\text{val}(f) = \sum_{(s,u) \in E} f(s, u)$ . We denote by  $f_c^*$  a maximum flow for given capacities  $c$ .

Let us note that we have made the decision to focus on the integral version of the problem for two reasons. First, in many applications edge capacities are integral anyway, and hence there always exists an integral solution as well, see, e.g., [1]. Second, the error measure we work with, namely the  $\ell_1$  distance, is meaningless if one allows arbitrary scaling without changing the problem, as it would be the case for rational edge capacities. Finally, we remark that all the fastest maximum flow algorithms, including the almost-linear time one [4], are only *weakly polynomial*, i.e., their running times depend polylogarithmically on the maximum edge capacity value and, in particular, they work only with integral edge capacities.

In Section 2, we prove the following theorem giving an algorithm that can be seen as the Ford-Fulkerson method with a warm start.

**Theorem 2.** Given a directed graph  $G = (V, E)$ , source and sink  $s, t \in V$ , edge capacities  $c : E \rightarrow \mathbb{Z}_{\geq 0}$ , and a predicted flow function  $f : E \rightarrow \mathbb{Z}_{\geq 0}$  satisfying flow conservation, one can compute a maximum  $(s, t)$ -flow in  $G$ , in time

$$O(|E| \cdot \|f - f_c^*\|_1).$$

Note that the above bound holds simultaneously for every maximum flow  $f_c^*$ , which might not be unique. In other words, the prediction is good if it is close to at least one optimal solution.

One of the sought-after properties of learning-augmented algorithms is *robustness*, i.e., retaining worst-case guarantees of classic algorithms even for arbitrarily bad predictions. However, in the case of running time bounds, robustness comes essentially for free (up to a multiplicative factor of 2, vanishing in the asymptotic notation). Indeed, one can always run step-by-step an algorithm with predictions alongside the fastest known classic algorithm, stopping when either of them stops. Therefore, Theorem 2 paired with the recent  $O(|E|^{1+o(1)})$  time algorithm for the maximum flow problem [4] actually leads to a robust learning-augmented algorithm with running time

$$O(|E| \cdot \min\{\|f - f_c^*\|_1, |E|^{o(1)}\}).$$

Naturally, a similar bound but with the  $\ell_\infty$  error (instead of the  $\ell_1$  error) would be more desirable, but let us argue that such a bound is unlikely. Indeed, the maximum flow problem with unit edge capacities is not known to be significantly easier than the general problem, yet for that special case the  $\ell_\infty$  prediction error can trivially be bounded by a constant, and hence a meaningful bound of the desired form would mean that it is an easier problem. Let us also remark that our choice of the  $\ell_1$  distance as the prediction error metric is consistent with related works [6, 3].

The proof of Theorem 2 that we give at the beginning of Section 2 uses a simple algorithm that calls the Ford-Fulkerson method (see Section 1.2) as a black-box but also spends  $O(|E| \cdot \|f - f_c^*\|_1)$  time on top of that. Later, in Section 2.1, we prove the following stronger result that implies Theorem 2.

**Theorem 3.** Given a directed graph  $G = (V, E)$ , source and sink  $s, t \in V$ , edge capacities  $c : E \rightarrow \mathbb{Z}_{\geq 0}$ , and a predicted flow function  $f : E \rightarrow \mathbb{Z}_{\geq 0}$  satisfying flow conservation, one can reduce computing a maximum  $(s, t)$ -flow in  $G$  to computing maximum flows in two

116 graphs, each with  $O(|E|)$  edges and the maximum flow  
 117 value bounded by  $O(\|f - f_c^*\|_1)$ . The reduction works in  
 118  $O(|E|)$  time.

119 Composing Theorem 3 with the Ford-Fulkerson  
 120 method gives the same worst-case running time bound  
 121 as Theorem 2. However, the alternative approach has  
 122 the advantage that all the computations that take more  
 123 than linear time can be delegated to one of many avail-  
 124 able highly optimized implementations of maximum  
 125 flow algorithms, and therefore this approach might be  
 126 more efficient in practice

127 Next, we want to argue that predictions required by  
 128 the above algorithm can be efficiently learned, in a PAC-  
 129 learning sense. We assume that the underlying graph,  
 130 as well as the choice of the source and sink nodes, are  
 131 fixed. (This assumption is almost without loss of gener-  
 132 ality, because one can take the underlying graph to be a  
 133 clique, with capacities zero for nonexistent edges; that,  
 134 however, may cause a running time overhead, because  
 135 of the increased number of edges.) Our goal is to prove  
 136 that, given a joint distribution over edge capacities, we  
 137 can efficiently learn a flow approximately minimizing  
 138 the expected  $\ell_1$  error over that distribution. We do it in  
 139 two steps. First, in Section 3, we prove Theorem 4, giv-  
 140 ing an algorithm that finds a prediction that minimizes  
 141 the error on a given set of samples. It works under the  
 142 assumption<sup>2</sup> that  $f_{c_i}^*$ 's are unique (or, arbitrarily fixed)  
 143 solutions to the maximum flow problem instances given  
 144 by capacities  $c_i$ 's. Then, in Section 4, we prove The-  
 145 orem 5, which states that if the number of samples is  
 146 large enough, then such a prediction for these samples  
 147 also approximately minimizes the error for the whole  
 148 distribution.

**Theorem 4.** *Given a directed graph  $G = (V, E)$ , with  
 source and sink  $s, t \in V$ , and a collection of  $k$  lists of  
 edge capacities  $c_1, c_2, \dots, c_k \in \mathbb{Z}_{\geq 0}^E$ , one can find an in-  
 tegral flow prediction that minimizes the prediction er-  
 ror on this collection, i.e.,*

$$\hat{f} = \arg \min \left\{ \frac{1}{k} \sum_{i \in [k]} \|f - f_{c_i}^*\|_1 \mid \right. \\ \left. f : E \rightarrow \mathbb{Z}_{\geq 0} \text{ satisfying flow conservation} \right\},$$

149 in time  $O(T(k \cdot |E|))$ , where  $T(m) \leq m^{1+o(1)}$  denotes the  
 150 min-cost flow complexity in graphs with  $m$  edges.

<sup>2</sup>This is a standard assumption for PAC-learnability results in the literature on (static) algorithms with predictions, see, e.g., [6, 25, 5].

**Theorem 5.** *Let  $G = (V, E)$  be a directed graph, with source and sink  $s, t \in V$ , and let  $c_1, c_2, \dots, c_k \in \mathbb{Z}_{\geq 0}^E$ , for  $k = \Theta(c_{\max}^2 |E|^3 \log(c_{\max} |E|))$ , be independent samples from a distribution  $\mathcal{D}$ , where  $c_{\max} = \max_{c \in \text{supp}(\mathcal{D}), e \in E} c(e)$ . Let  $\hat{f} \in \mathbb{Z}_{\geq 0}^E$  be a flow that minimizes the prediction error on this collection of samples, as in Theorem 4. Then, with high probability over the choice of the samples, the expected  $\ell_1$  error of  $\hat{f}$  over  $\mathcal{D}$  is approximately minimum possible, i.e.,*

$$\mathbb{E}_{c \sim \mathcal{D}} \|\hat{f} - f_c^*\|_1 \leq \min_f \mathbb{E}_{c \sim \mathcal{D}} \|f - f_c^*\|_1 + O(1),$$

151 where the minimum is taken over functions  $f \in \mathbb{Z}_{\geq 0}^E$   
 152 satisfying the flow conservation property.

## 153 1.2. Related work

154 *Maximum flow algorithms.* There are numerous algo-  
 155 rithms for the maximum flow problem. The Ford-  
 156 Fulkerson method [12] is a starting point for many  
 157 of them, and its vanilla version runs in weakly poly-  
 158 nomial  $O(|E| \cdot \text{val}(f_c^*))$  time for integral edge capaci-  
 159 ties. The strongly polynomial time algorithms, which  
 160 also work for rational edge capacities, can be roughly  
 161 split into three groups: augmenting paths algorithms  
 162 (e.g., [9, 7, 13]), push-relabel algorithms (e.g., [14]),  
 163 and pseudoflow algorithms (e.g., [15]). Each of these  
 164 groups contains algorithms with running time  $\tilde{O}(|V| \cdot |E|)$   
 165 that are widely used in practice, see, e.g., [2, 11]. A long  
 166 line of research on Laplacian solvers and interior-point  
 167 methods, initiated by [26], culminated recently with a  
 168 (weakly polynomial) almost-linear  $O(|E|^{1+o(1)})$  time al-  
 169 gorithm [4].

170 In the light of this new development, it may seem  
 171 that our learning-augmented algorithm is only relevant  
 172 for very small prediction errors, namely  $\|f - f_c^*\|_1 \leq$   
 173  $|E|^{1+o(1)}$ . However, at this point it is not yet clear if the  
 174 new almost-linear time algorithm will lead to practical  
 175 developments.<sup>3</sup>

176 *Learning-augmented algorithms.* The idea of using  
 177 predictions to improve the performance of algorithms  
 178 is not a new one, see, e.g., [22]. However, the recent  
 179 systematic study of such methods – under the umbrella  
 180 term of *learning-augmented algorithms*, or simply *algo-*  
 181 *rithms with predictions* – seems to have started with the  
 182 works of Lykouris and Vassilvitskii [21], and Purohit,  
 183 Svitkina, and Kumar [24]. Since then, the field devel-  
 184 oped rapidly, see [23] for a survey. So far, the majority

<sup>3</sup>See <https://codeforces.com/blog/entry/100510> for a relevant discussion with an author of [4].

185 of the works focus on online algorithms, where predic- 232  
 186 tions help reduce uncertainty about the yet unseen part 233  
 187 of the instance. There are, however, also works on, e.g., 234  
 188 data structures [20], streaming algorithms [17], and sub- 235  
 189 linear algorithms [8]. Apart from a simple example of  
 190 binary search [21], until recently there were no works on  
 191 improving algorithms running times using predictions. 236  
 192 This has changed with the work of Diniz et al. [6], the 237  
 193 recent followup works of Chen et al. [3], and Sakaue 238  
 194 and Oki [25], as well as the work of Ergun et al. [10] 239  
 195 (which is however concerned with approximation algo- 240  
 196 rithms and uses a significantly different approach). 241

197 *Learning-augmented weighted bipartite matching.* A 242  
 198 direct inspiration for our approach is the work of Diniz 243  
 199 et al. [6]. They study the maximum weighted bipar- 244  
 200 tite matching problem and predict the dual<sup>4</sup> solution. 245  
 201 They give a learning-augmented algorithm that solves 246  
 202 the matching problem in  $O(|E| \sqrt{|V|} \cdot \min\{\eta, \sqrt{|V|}\})$  time, 247  
 203 where  $\eta$  is the  $\ell_1$  error of the predicted dual solution – 248  
 204 our Theorem 2 is an analogue of that result. They also 249  
 205 show that, given an oracle access to a joint distribution 250  
 206 over edge weights, one can efficiently learn a prediction 251  
 207 minimizing the expected  $\ell_1$  error over the distribution – 252  
 208 our Theorems 4 and 5 are together an analogue of that 253  
 209 result. 254

210 The most apparent difference between their approach 255  
 211 and ours is that they use a predicted dual solution and we 256  
 212 use a predicted primal solution. The reason they state 257  
 213 for focusing on the dual solution is that the primal so- 258  
 214 lution is very volatile to small changes in the input. Let 259  
 215 us note that this argument clearly applies to weighted 260  
 216 problems (in particular, e.g., to the minimum cost flow 261  
 217 problem) but it is not clear if it also applies to the maxi- 262  
 218 mum flow problem. Moreover, it is also not clear if the 263  
 219 dual solution is indeed less volatile, even for weighted 264  
 220 problems. 265

221 The second important difference is that they do not 266  
 222 impose any constraints on predictions, while we require 267  
 223 that the predicted solution satisfies the flow conserva- 268  
 224 tion property. This difference has the following conse- 269  
 225 quences. First, their learning algorithm can be very 270  
 226 simple – the best possible prediction is just a coordinate- 271  
 227 wise median over the solutions for the samples – while 272  
 228 we need to solve the minimum cost flow problem in- 273  
 229 stead. Second, turning a prediction into a feasible so- 274  
 230 lution is also harder for us, as we want to maintain the 275  
 231 flow conservation property. On the other hand, once we 276

232 have a feasible solution, the remaining part of our maxi-  
 233 mum flow algorithm is simple and easy to analyse, in  
 234 contrast with their tailored primal-dual analysis for the  
 235 analogous part of their algorithm.

236 *Other algorithmic speedups using predictions.* In their  
 237 recent work Chen et al. [3] improve the running time  
 238 of Diniz et al. for the matching problem, and extend  
 239 their framework to a couple of other graph problems:  
 240 the negative weights single-source shortest paths prob-  
 241 lem, the degree-constrained subgraph problem, and the  
 242 minimum cost 0-1 flow problem. For all these problems  
 243 they use predicted dual solutions. They also propose  
 244 general learnability theorems, which imply a result sim-  
 245 ilar to Lemma 8 (see also the discussion below the proof  
 246 of Lemma 8 for a comparison of such results with The-  
 247 orem 5).

248 Sakaue and Oki [25] propose a general framework for  
 249 augmenting discrete optimization problems with pre-  
 250 dictions. Their approach leads to faster algorithms for  
 251 three problems: weighted bipartite matching, weighted  
 252 matroid intersection, and discrete energy minimization  
 253 for computer vision. It is notable that their guarantees  
 254 hold with respect to the  $\ell_\infty$  prediction error, compared  
 255 to the  $\ell_1$  error in previous works [6, 3] and ours. In-  
 256 terestingly, their framework allows working with both  
 257 primal and dual predictions. They analyse what proper-  
 258 ties of a problem make it more suitable for one type of  
 259 predictions or the other, and they end up using primal  
 260 predictions for the discrete energy minimization prob-  
 261 lem.

262 Ergun et al. [10] study  $k$ -means clustering, and use  
 263 predictions to achieve in near-linear time approxima-  
 264 tion factors that are impossible to achieve without pre-  
 265 dictions even in polynomial time. They also propose a  
 266 very general learnability framework, based on relating  
 267 the VC-dimension of a class of functions to their com-  
 268 putational complexity, which implies bounds similar to  
 269 ours, but with a worse sample complexity.

270 *Dropping flow conservation constraints.* In the concu-  
 271 rrent and independent work [5], Davies, Moseley, Vas-  
 272 silvitskii, and Wang propose a similar framework for  
 273 speeding up maximum flow computations by using pre-  
 274 dictions. The key difference is that in their setup the  
 275 prediction does not necessarily need to satisfy the flow  
 276 conservation constraints. They prove analogs of our  
 277 Theorems 2–5 for their notion of prediction. In particu-  
 278 lar, they also use the  $\ell_1$  distance between the predicted  
 279 and the maximum flow as the prediction error, and they  
 280 achieve the same running time as in our Theorem 2.

<sup>4</sup>Recall that the matching problem can be formulated as a linear program, and every linear program has a corresponding dual program.

281 Their analog of our Theorem 2 also follows the two 319  
 282 step approach consisting of a feasibility step and an 320  
 283 optimization step. Their optimization step is identical to 321  
 284 ours, the key technical difference lies in the feasibility 322  
 285 step. In this step, our algorithm only needs to fix vi- 323  
 286 olated capacity constraints, while their algorithm fixes 324  
 287 both capacity and flow conservation constraints. It first 325  
 288 fixes the capacity constraints by just decreasing the flow 326  
 289 on oversaturated edges – possibly violating additional 327  
 290 flow conservation constraints, which however are going 328  
 291 to be fixed next. Then, it fixes the flow conservation 329  
 292 constraints using a reduction to a maximum flow com- 330  
 293 putation in an auxiliary graph – similar to our proof of 331  
 294 Theorem 3.

295 On the other hand, learning (Theorem 4) in their setup 332  
 296 is much easier. Since the prediction for each edge can 333  
 297 be optimized independently, it boils down to selecting – 334  
 298 for each edge – the median of optimal flow values along 335  
 299 that edge among samples. 336

### 300 1.3. Limitation and open problem

301 We do prove that a prediction with a small  $\ell_1$  error 340  
 302 can be used to speed up maximum flow computation, 341  
 303 and that given a distribution over flow networks one can 342  
 304 learn a prediction minimizing the  $\ell_1$  error. However, 343  
 305 we do not answer the question of what makes a distri- 344  
 306 bution have such a minimum that is actually small. 345  
 307 There seems to be no standard approach to address this 346  
 308 type of representation error question, and the related 347  
 309 works [6, 10, 3] do not address it either.

## 310 2. Warm-starting Ford-Fulkerson

**Theorem 2.** *Given a directed graph  $G = (V, E)$ , source 348  
 and sink  $s, t \in V$ , edge capacities  $c : E \rightarrow \mathbb{Z}_{\geq 0}$ , and 349  
 a predicted flow function  $f : E \rightarrow \mathbb{Z}_{\geq 0}$  satisfying flow 350  
 conservation, one can compute a maximum  $(s, t)$ -flow in 351  
 $G$ , in time*

$$O(|E| \cdot \|f - f_c^*\|_1).$$

311 *Proof.* At first, the predicted flow  $f$  does not necessarily 352  
 312 satisfy the capacity constraints imposed by  $c$ , i.e., for 353  
 313 some edges  $e \in E$  it might happen that  $f(e) > c(e)$ . 354  
 314 The algorithm consists of two steps. In the first step, 355  
 315 it turns  $f$  into  $\bar{f}$  that satisfies the capacity constraints, 356  
 316 while maintaining the flow conservation property. In 357  
 317 other words,  $\bar{f}$  is a feasible flow. Then, in the second 358  
 318 step, the algorithm augments  $\bar{f}$  to an optimal flow.

*First step: feasibility.* Recall that every integral flow 319  
 decomposes into cycles and  $s$ - $t$  paths<sup>5</sup> (see, e.g., [1, 320  
 Theorem 3.5]). The algorithm initializes  $\bar{f} = f$ . While 321  
 there is an edge  $e \in E$  with  $\bar{f}(e) > c(e)$ , the algorithm 322  
 uses, e.g., depth-first search to find a cycle or an  $s$ - $t$  323  
 path containing  $e$  (at least one of them is guaranteed to 324  
 exist because of the integral flow decomposition), and 325  
 decreases the flow  $\bar{f}$  along this cycle/path by one unit. 326  
 This keeps the invariant that  $\bar{f}$  satisfies the flow con- 327  
 servation property. When the process is done,  $\bar{f}$  also 328  
 satisfies all the capacity constraints. 329

*Second step: optimization.* Now, the algorithm con- 330  
 structs the *residual network* with respect to  $\bar{f}$ , i.e., 331  
 the flow network  $G_{\bar{f}} = (V, E_{\bar{f}})$  with edge set  $E_{\bar{f}} = 332  
 \{(u, v) \mid (u, v) \in E \text{ or } (v, u) \in E\}$  and residual capacities<sup>6</sup> 333  
 $c_{\bar{f}}(u, v) = (c(u, v) - \bar{f}(u, v)) + \bar{f}(v, u)$ . Here, for nota- 334  
 tional simplicity, we assume that  $c(u, v) = \bar{f}(u, v) = 0$  if 335  
 $(u, v) \notin E$ . Then, the algorithm runs the Ford-Fulkerson 336  
 method [12] on  $G_{\bar{f}}$  to find a maximum flow  $f_{c_{\bar{f}}}^*$  in time 337  
 $O(|E| \cdot \text{val}(f_{c_{\bar{f}}}^*))$ . Finally,  $\bar{f} + f_{c_{\bar{f}}}^*$  is a maximum flow for 338  
 the original edge capacities  $c$ , see, e.g., [1, Property 2.6]. 339

*Running time analysis.* Let  $\delta = \sum_{e \in E} \max\{f(e) - c(e), 0\}$  340  
 be the total amount by which the flow prediction vio- 341  
 lates the capacity constraints. The algorithm makes at 342  
 most  $\delta$  iterations in the first step, and each iteration de- 343  
 creases the flow value  $\text{val}(\bar{f})$  by at most one. We con- 344  
 clude that the first step runs in  $O(|E| \cdot \delta)$  time, and that 345  
 $\text{val}(f) - \text{val}(\bar{f}) \leq \delta$ . 346

The second step of the algorithm runs in time

$$\begin{aligned} O(|E| \cdot \text{val}(f_{c_{\bar{f}}}^*)) &= O(|E| \cdot (\text{val}(f_c^*) - \text{val}(\bar{f}))) = \\ &= O(|E| \cdot ((\text{val}(f_c^*) - \text{val}(f)) + (\text{val}(f) - \text{val}(\bar{f}))))). \end{aligned}$$

347 Let  $\eta = \|f - f_c^*\|_1$  denote the prediction error. It is easy 348  
 to see that  $|\text{val}(f_c^*) - \text{val}(f)| \leq \eta$ , and that  $\delta \leq \eta$ , so, in 349  
 particular,  $\text{val}(f) - \text{val}(\bar{f}) \leq \delta \leq \eta$ . Therefore, the run- 350  
 ning time of both steps of the algorithm can be bounded 351  
 by  $O(|E| \cdot \eta)$ .  $\square$

### 352 2.1. Alternative variant of the first step

353 In this section we give an alternative variant of the 354  
 first step of the above algorithm. The asymptotic run- 355  
 ning time remains the same, but, as we explained when

<sup>5</sup>I.e., there exists a collection  $p_1, \dots, p_k$  such that each  $p_i$  is either a (simple) cycle or a (simple)  $s$ - $t$  path in  $G$ , and  $f(e) = \#\{i \in [k] \mid e \in p_i\}$  for every edge  $e \in E$ .

<sup>6</sup>The amount of extra flow that can be sent from  $u$  to  $v$  equals the sum of remaining capacity for edge  $(u, v)$ , i.e.,  $c(u, v) - \bar{f}(u, v)$ , and the flow sent from  $v$  to  $u$ , which can be reversed, i.e.,  $\bar{f}(v, u)$ . Usually we expect only one of these two summands to be non-zero.

356 introducing Theorem 3 in Section 1.1, this alternative  
 357 algorithm might be more efficient in practice, because it  
 358 allows delegating most of the work to any existing well-  
 359 optimized maximum flow solver. See Algorithm 1.

---

**Algorithm 1:** Computing maximum flow

---

**Input:** flow network  $G = (V, E)$ , source  $s$ , target  
 $t$ , edge capacities  $c$ , predicted flow  $f$   
**Output:** maximum flow for edge capacities  $c$   
 /\* Step 1: Feasibility \*/  
 $\tilde{E} \leftarrow \emptyset$ ;  $\tilde{G} \leftarrow (V \cup \{\tilde{s}, \tilde{t}\}, \tilde{E})$ ;  $\delta \leftarrow 0$   
**foreach**  $(u, v) \in E$  **do**  
    $\tilde{E} \leftarrow \tilde{E} \cup \{(v, u)\}$   
   **if**  $f(u, v) \leq c(u, v)$  **then**  
      $\tilde{c}(v, u) \leftarrow f(u, v)$   
   **else**  
      $\tilde{c}(v, u) \leftarrow c(u, v)$   
      $\delta \leftarrow \delta + f(u, v) - c(u, v)$   
      $\tilde{E} \leftarrow \tilde{E} \cup \{(\tilde{s}, u), (v, \tilde{t})\}$   
      $\tilde{c}(\tilde{s}, u), \tilde{c}(v, \tilde{t}) \leftarrow f(u, v) - c(u, v)$   
 $\tilde{E} \leftarrow \tilde{E} \cup \{(s, t)\}$ ;  $\tilde{c}(s, t) \leftarrow \delta$   
 $\tilde{f} \leftarrow$  max flow from  $\tilde{s}$  to  $\tilde{t}$  in  $\tilde{G}$  with capacities  $\tilde{c}$   
**foreach**  $(u, v) \in E$  **do**  
    $\tilde{f}(u, v) \leftarrow \min(f(u, v), c(u, v)) - \tilde{f}(v, u)$   
 /\* Step 2: Optimization \*/  
 $G_{\tilde{f}} \leftarrow$  residual network with respect to flow  $\tilde{f}$   
 $f_{c_{\tilde{f}}}^* \leftarrow$  max flow from  $s$  to  $t$  in  $G_{\tilde{f}}$   
**return**  $\tilde{f} + f_{c_{\tilde{f}}}^*$

---

360 Consider graph  $\tilde{G} = (V, \tilde{E})$  with  $\tilde{E} = \{(v, u) \in V \times V \mid$   
 361  $(u, v) \in E\}$ , i.e., a copy of  $G$  with reversed edges. Set  
 362 capacities to  $\tilde{c}(v, u) = f(u, v)$ . Note that the first step  
 363 of the original algorithm essentially finds an integral  $t$ - $s$   
 364 flow  $\tilde{f}$  in  $\tilde{G}$  such that

- 365 (i) if  $f(u, v) > c(u, v)$ , then  $\tilde{f}(v, u) \geq f(u, v) - c(u, v)$ ,  
 366 for every  $(u, v) \in E$ ;  
 367 (ii)  $\text{val}(\tilde{f}) \leq \delta$ .

368 At the end of the first step  $\tilde{f}(u, v) = f(u, v) - \tilde{f}(v, u)$ . In  
 369 this section we give a different way to compute such  $\tilde{f}$ .

370 Add to  $\tilde{G}$  edge  $(s, t)$ , and set  $\tilde{c}(s, t) = \delta$ . Now,  
 371 the problem of finding a  $t$ - $s$  flow satisfying (i) and (ii) be-  
 372 comes the problem of finding a *circulation*<sup>7</sup> satisfying  
 373 (i). This problem – of finding a circulation with lower  
 374 bounds – can be reduced to a problem of finding a

<sup>7</sup>A circulation is defined similarly to a flow. The only exception is  
 that there are no designated source and sink nodes, and hence the flow  
 conservation property has to be satisfied for all the nodes of the graph  
 (see, e.g., [1, Section 1.2]).

375 maximum flow (without lower bounds) in a graph with  
 376 the maximum flow value equal to the sum of all lower  
 377 bounds, see, e.g., [1, Section 6.7]. The reduction works  
 378 as follows.

379 First, add to  $\tilde{G}$  two new nodes  $\tilde{s}$  and  $\tilde{t}$ . Next, for each  
 380 edge  $e = (u, v) \in E$  that violates the capacity constraint  
 381 let  $\delta_e = f(u, v) - c(u, v) > 0$  be the excess flow for  
 382 this edge; add to  $\tilde{G}$  two edges,  $(\tilde{s}, u)$  and  $(v, \tilde{t})$ , set their  
 383 capacities to  $\tilde{c}(\tilde{s}, u) = \tilde{c}(v, \tilde{t}) = \delta_e$ , and decrease the ca-  
 384 pacity of edge  $(v, u)$  by  $\delta_e$ , so that  $\tilde{c}(v, u) = c(u, v)$ . This  
 385 ends the description of the graph constructed in the re-  
 386 duction.

387 Note that the total capacity of edges leaving  $\tilde{s}$  equals  
 388 the total capacity of edges entering  $\tilde{t}$ , which also equals  
 389  $\delta = \sum_{e \in E} \max\{f(e) - c(e), 0\}$ . As we will see in a mo-  
 390 ment, the existence of a flow saturating these edges is  
 391 equivalent to the existence of a circulation satisfying the  
 392 lower bounds – which is guaranteed to exist because the  
 393 original first step of the algorithm finds such a circula-  
 394 tion.

395 After constructing  $\tilde{G}$  as above, the alternative first  
 396 step proceeds as follows. The algorithm computes a  
 397 maximum  $\tilde{s}$ - $\tilde{t}$  flow  $\tilde{f}$  in  $\tilde{G}$ , using a Ford-Fulkerson  
 398 method. Then, for each edge  $e = (u, v) \in E$  that vi-  
 399 olates the capacity constraint (in the original graph  $G$ ),  
 400 the algorithm first removes the saturated edges  $(\tilde{s}, u)$  and  
 401  $(v, \tilde{t})$  from  $\tilde{G}$ . Note that now nodes  $u$  and  $v$  do not sat-  
 402 isfy the flow conservation property, namely node  $v$  has  
 403 an excess of  $\delta_e$  units of incoming flow and node  $u$  has  
 404 a deficit of  $\delta_e$  units of incoming flow. The algorithm re-  
 405 stores the flow conservation property by increasing flow  
 406  $\tilde{f}(v, u)$  by  $\delta_e$  units, and therefore it ensures that the lower  
 bound for this edge is satisfied. This procedure essen-  
 tially proves the equivalence of the existence of a flow  
 saturating the sink and source edges and the existence  
 of a suitable circulation.

This ends the description of the alternative algorithm.  
 Let us analyse its running time. Graph  $\tilde{G}$  has  $O(|E|)$   
 edges and can be constructed in  $O(|E|)$  time. Since  
 $\text{val}(\tilde{f}) = \delta$ , the Ford-Fulkerson method runs in  $O(|E| \cdot \delta)$   
 time. Finally, transforming  $\tilde{f}$  to  $\tilde{f}$  takes  $O(|E|)$  time.  
 Therefore, the total running time of  $O(|E| \cdot \delta)$  remains  
 unchanged compared to the original first step of the al-  
 gorithm. This proves Theorem 3.

**Theorem 3.** *Given a directed graph  $G = (V, E)$ , source  
 and sink  $s, t \in V$ , edge capacities  $c : E \rightarrow \mathbb{Z}_{\geq 0}$ ,  
 and a predicted flow function  $f : E \rightarrow \mathbb{Z}_{\geq 0}$  satisfying  
 flow conservation, one can reduce computing a maxi-  
 mum  $(s, t)$ -flow in  $G$  to computing maximum flows in two  
 graphs, each with  $O(|E|)$  edges and the maximum flow  
 value bounded by  $O(\|f - f_c^*\|_1)$ . The reduction works in*

426  $O(|E|)$  time.

427 Finally, we remark that a similar trick – for handling  
428 edges initialized with a flow exceeding their capacities  
429 – was already proposed, albeit without provable running  
430 time guarantees, in the context of repeatedly solving  
431 similar minimum cut instances in a computer vision ap-  
432 plication [19]. That trick however only allows comput-  
433 ing the maximum flow value and a corresponding mini-  
434 mum cut, but not the flow itself.

### 435 3. Learning a prediction that minimizes error

**Theorem 4.** *Given a directed graph  $G = (V, E)$ , with source and sink  $s, t \in V$ , and a collection of  $k$  lists of edge capacities  $c_1, c_2, \dots, c_k \in \mathbb{Z}_{\geq 0}^E$ , one can find an integral flow prediction that minimizes the prediction error on this collection, i.e.,*

$$\hat{f} = \arg \min \left\{ \frac{1}{k} \sum_{i \in [k]} \|f - f_{c_i}^*\|_1 \mid f : E \rightarrow \mathbb{Z}_{\geq 0} \text{ satisfying flow conservation} \right\},$$

436 in time  $O(T(k \cdot |E|))$ , where  $T(m) \leq m^{1+o(1)}$  denotes the  
437 min-cost flow complexity in graphs with  $m$  edges.

438 *Proof.* The first step of the learning algorithm (see Al-  
439 gorithm 2) is to compute a maximum flow  $f_{c_i}^*$  for each  
440  $i \in [k]$ . This step can be completed in  $k \cdot T(|E|) \leq$   
441  $O(T(k \cdot |E|))$  time in total.

442 Now, the goal is to find an integral flow  $f$  (satis-  
443 fying the flow conservation property) that minimizes  
444  $\sum_{i \in [k]} \sum_{e \in E} |f(e) - f_{c_i}^*(e)| = \sum_{e \in E} \sum_{i \in [k]} |f(e) - f_{c_i}^*(e)|$ . For  
445 an edge  $e \in E$ , let  $\text{cost}_e(x) = \sum_{i \in [k]} |x - f_{c_i}^*(e)|$  denote  
446 the contribution of  $f(e)$  to the minimization objective,  
447 which now can be written simply as  $\sum_{e \in E} \text{cost}_e(f(e))$ .

448 Let us analyse how the function  $\text{cost}_e(x)$  behaves.  
449 Let  $x_1 \leq x_2 \leq \dots \leq x_k$  denote the sorted elements  
450 of the (multi-)set  $\{f_{c_1}^*(e), f_{c_2}^*(e), \dots, f_{c_k}^*(e)\}$ . Clearly,  
451  $\text{cost}_e(0) = \sum_{i \in [k]} x_i$ . For  $x \in [0, x_1]$ , the contribution  
452  $\text{cost}_e(x)$  is a decreasing linear function with slope  $-k$ .  
453 For  $x \in [x_1, x_2]$ , the slope is  $-k + 2$ . More generally,  
454 for  $x \in [x_i, x_{i+1}]$  the slope is  $2i - k$ , because increasing  
455 the flow by  $\delta$  increases also by  $\delta$  each of the first  $i$  sum-  
456 mands, and decreases by the same amount each of the  
457 remaining  $(k - i)$  summands in the sum  $\sum_{i \in [k]} |x - x_i|$ .  
458 Hence,  $\text{cost}_e(x)$  is piecewise linear and convex, and the  
459 overall goal is to find a flow minimizing a separable  
460 piecewise linear convex cost function.

461 The above problem can be reduced to the standard  
462 minimum cost flow problem [1, Chapter 14]. The re-  
463 duction works as follows. For notational simplicity, let

464  $x_0 = 0$  and  $x_{k+1} = +\infty$ . Replace each edge  $e \in E$  with  
465  $k + 1$  parallel edges  $e_0, \dots, e_k$ , and let edge  $e_i$  have ca-  
466 pacity  $x_{i+1} - x_i$  and cost (of sending one unit of flow)  
467 equal to  $2i - k$ . It is easy to observe that any opti-  
468 mal solution to the minimum cost flow problem in the  
469 constructed multigraph uses some prefix of the cheapest  
470 parallel edges for each  $e \in E$ , and the total cost of such  
471 prefix behaves exactly like  $\text{cost}_e$ . Since all the intro-  
472 duced capacities are integral, it is guaranteed that there  
473 exists an optimal integral solution. The multigraph has  
474  $(k + 1) \cdot |E|$  edges, hence the minimum cost flow can be  
475 found in  $O(T(k \cdot |E|))$ .  $\square$

---

#### Algorithm 2: Learning prediction from samples

---

**Input:** flow network  $G = (V, E)$ , and sampled edge capacity functions  $c_1, c_2, \dots, c_k$

**Output:** flow  $f$  minimizing  $\frac{1}{k} \sum_{i \in [k]} \|f - f_{c_i}^*\|_1$

**foreach**  $i \in \{1, 2, \dots, k\}$  **do**

    compute maximum flow  $f_{c_i}^*$  for capacities  $c_i$

**foreach**  $e \in E$  **do**

$x_1 \leq \dots \leq x_k \leftarrow$  sorted set  $\{f_{c_1}^*(e), \dots, f_{c_k}^*(e)\}$

$x_0 \leftarrow 0, x_{k+1} \leftarrow +\infty$

    replace  $e$  with  $k + 1$  parallel edges  $e_0, \dots, e_k$

**foreach**  $i \in \{0, \dots, k\}$  **do**

        capacity of  $e_i \leftarrow x_{i+1} - x_i$

        cost of  $e_i \leftarrow 2i - k$

**return** minimum cost flow in the modified graph

---

### 476 4. Sample complexity

**Theorem 5.** *Let  $G = (V, E)$  be a directed graph, with source and sink  $s, t \in V$ , and let  $c_1, c_2, \dots, c_k \in \mathbb{Z}_{\geq 0}^E$ , for  $k = \Theta(c_{\max}^2 |E|^3 \log(c_{\max} |E|))$ , be independent samples from a distribution  $\mathcal{D}$ , where  $c_{\max} = \max_{c \in \text{supp}(\mathcal{D}), e \in E} c(e)$ . Let  $\hat{f} \in \mathbb{Z}_{\geq 0}^E$  be a flow that minimizes the prediction error on this collection of samples, as in Theorem 4. Then, with high probability over the choice of the samples, the expected  $\ell_1$  error of  $\hat{f}$  over  $\mathcal{D}$  is approximately minimum possible, i.e.,*

$$\mathbb{E}_{c \sim \mathcal{D}} \|\hat{f} - f_c^*\|_1 \leq \min_f \mathbb{E}_{c \sim \mathcal{D}} \|f - f_c^*\|_1 + O(1),$$

477 where the minimum is taken over functions  $f \in \mathbb{Z}_{\geq 0}^E$   
478 satisfying the flow conservation property.

For a flow prediction  $f \in \mathbb{Z}_{\geq 0}^E$ , let us use

$$\begin{aligned} \text{cost}_{c_1, \dots, c_k}(f) &= \frac{1}{k} \sum_{i \in [k]} \|f - f_{c_i}^*\|_1, \\ \text{cost}_{\mathcal{D}}(f) &= \mathbb{E}_{c \sim \mathcal{D}} \|f - f_c^*\|_1 \end{aligned}$$

479 to denote the  $\ell_1$  error of  $f$  on the samples and on the  
480 distribution, respectively. We will use Hoeffding's in-  
481 equality to prove that the number of samples in The-  
482 orem 5 is large enough for  $\text{cost}_{c_1, \dots, c_k}(f)$  to be a good  
483 approximation of  $\text{cost}_{\mathcal{D}}(f)$ , with high probability for all  
484  $f$ 's simultaneously.

**Theorem 6** (Hoeffding's inequality [16]). *Let  $X_1, \dots, X_k$  be independent random variables with values from 0 to  $U$ , and let  $S = X_1 + \dots + X_k$  denote their sum. Then, for all  $t > 0$ ,*

$$P(|S - \mathbb{E}S| \geq t) \leq 2 \cdot \exp(-2t^2/kU^2).$$

485 To use the inequality, first we need a bound on the  
486 values of the considered functions.

487 **Lemma 7.** *Any flow prediction  $f$  that minimizes the er-  
488 ror must satisfy  $\|f\|_1 \leq 2c_{\max}|E|$ .*

489 Let us note that Lemma 7 is actually nontrivial. Even  
490 though  $\|f_{c_i}^*\|_{\infty} \leq c_{\max}$  for every  $i \in [k]$ , it may happen  
491 that  $\|f\|_{\infty} > c_{\max}$  because of the flow conservation con-  
492 straint, e.g., when multiple disjoint paths end at a single  
493 node and force a single edge going out of that node to  
494 have a flow larger than  $c_{\max}$ .

495 *Proof of Lemma 7.* For every  $c \in \text{supp}(\mathcal{D})$ , we have  
496  $\|c\|_1 \leq c_{\max}|E|$ , and, since  $0 \leq f_c^* \leq c$ , then also  
497  $\|f_c^*\|_1 \leq c_{\max}|E|$ . Moreover, by the triangle inequal-  
498 ity,  $\|f - f_c^*\|_1 + \|f_c^*\|_1 \geq \|f\|_1$ , and thus  $\|f - f_c^*\|_1 \geq$   
499  $\|f\|_1 - c_{\max}|E|$ . If  $\|f\|_1 > 2c_{\max}|E|$ , then  $\|f - f_c^*\|_1 >$   
500  $c_{\max}|E|$  for every  $c \in \text{supp}(\mathcal{D})$ , and thus also  $\text{cost}_{\mathcal{D}}(f) =$   
501  $\mathbb{E}_{c \sim \mathcal{D}} \|f - f_c^*\|_1 > c_{\max}|E|$ .

502 At the same time, if we consider the all-zero vector  
503 as a flow prediction, we have  $\|0 - f_c^*\|_1 = \|f_c^*\|_1 \leq$   
504  $c_{\max}|E|$ , for every  $c \in \text{supp}(\mathcal{D})$ , and thus also  $\text{cost}_{\mathcal{D}}(0) =$   
505  $\mathbb{E}_{c \sim \mathcal{D}} \|0 - f_c^*\|_1 \leq c_{\max}|E|$ . It follows that  $f$  could not  
506 minimize the error if  $\|f\|_1 > 2c_{\max}|E|$ .  $\square$

507 Now we are ready to apply Hoeffding's inequality in  
508 order to prove the following lemma.

**Lemma 8.** *With high probability over the choice of the  
samples, for all  $f \in \mathbb{Z}_{\geq 0}^E$  satisfying the flow conserva-  
tion property and such that  $\|f\|_1 \leq 2c_{\max}|E|$  it holds that*

$$|\text{cost}_{c_1, \dots, c_k}(f) - \text{cost}_{\mathcal{D}}(f)| \leq 1.$$

*Proof.* For a fixed  $f$ , satisfying the conditions of the  
lemma, let  $X_i = \frac{1}{k} \|f - f_{c_i}^*\|_1$ . We have that  $\|f - f_{c_i}^*\|_1 \leq$   
 $\|f\|_1 + \|f_{c_i}^*\|_1 \leq (2+1) \cdot c_{\max}|E|$ , so the random variable  $X_i$   
has values from 0 to  $3c_{\max}|E|/k$ . Clearly,  $\text{cost}_{c_1, \dots, c_k}(f) =$   
 $X_1 + \dots + X_k$ , and  $\mathbb{E} \text{cost}_{c_1, \dots, c_k}(f) = \text{cost}_{\mathcal{D}}(f)$ . Applying  
Hoeffding's inequality, with  $t = 1$ , we get that

$$\begin{aligned} P(|\text{cost}_{c_1, \dots, c_k}(f) - \text{cost}_{\mathcal{D}}(f)| \geq 1) \\ \leq 2 \cdot \exp\left(-\frac{2}{k \cdot (3c_{\max}|E|/k)^2}\right) \\ = \exp(-\Theta(k/(c_{\max}|E|)^2)) \\ = \exp(-\Theta(|E| \log(c_{\max}|E|))). \end{aligned}$$

Let  $\mathcal{F}$  denote the set of all  $f$ 's satisfying the  
conditions of the lemma. We can upper-bound the  
number of such  $f$ 's by  $|\mathcal{F}| \leq (2c_{\max}|E| + 1)^{|E|} =$   
 $\exp(\Theta(|E| \log(c_{\max}|E|)))$ . To finish the proof, note that

$$\begin{aligned} |\mathcal{F}| \cdot \text{poly}(c_{\max}|E|) \\ \leq \exp(\Theta(|E| \log(c_{\max}|E|))) \cdot \exp(\Theta(\log(c_{\max}|E|))) \\ = \exp(\Theta(|E| \log(c_{\max}|E|))), \end{aligned}$$

509 and hence we can take the union bound to conclude  
510 that with high probability it holds that  $|\text{cost}_{c_1, \dots, c_k}(f) -$   
511  $\text{cost}_{\mathcal{D}}(f)| \leq 1$  for all  $f \in \mathcal{F}$  simultaneously.  $\square$

512 Let us remark that the above proof of Lemma 8 cru-  
513 cially relies on the fact that it suffices to consider a finite  
514 set of possible flow predictions – because they are inte-  
515 gral and bounded – and therefore we can use the union  
516 bound. Dinitz et al. [6, Section 3.3 in their supplemental  
517 material] give a proof of an analogous result regarding  
518 learning optimal dual solution for the weighted bipar-  
519 tite matching problem. Their proof is more complex  
520 than ours, it uses the notion of pseudo-dimension, but  
521 thanks to that it works also for fractional predictions.  
522 We note that it is possible to prove alternative version of  
523 Lemma 8, in the spirit of Dinitz et al., that would gen-  
524 eralize to fractional flows but would lose a small factor  
525  $\log|E|$  in the sample complexity.

526 With Lemma 8 at hand, it takes a standard argument  
527 to prove Theorem 5.

*Proof of Theorem 5.* Let  $\hat{f}$  and  $\tilde{f}$  be flow predictions  
that minimize the error on the samples and on the whole  
distribution, respectively. By Lemma 7,  $\|\hat{f}\|_1, \|\tilde{f}\|_1 \leq$   
 $2c_{\max}|E|$ , and hence Lemma 8 applies. Note that it is  
crucial that Lemma 8 holds with high probability for all  
 $f$ 's, because  $\hat{f}$  is chosen after the samples are drawn  
from  $\mathcal{D}$ . We finish the proof with the following chain of



inequalities.

$$\begin{aligned} \text{cost}_{\mathcal{D}}(\hat{f}) &\stackrel{\text{Lemma 8}}{\leq} \text{cost}_{c_1, \dots, c_k}(\hat{f}) + 1 \\ &\leq \text{cost}_{c_1, \dots, c_k}(\bar{f}) + 1 \leq \text{cost}_{\mathcal{D}}(\bar{f}) + 2. \end{aligned}$$

because  $\hat{f}$  minimizes the error on  $c_1, \dots, c_k$       Lemma 8

□

## Acknowledgments

We would like to thank Alexandra Lassota, Sai Ganesh Nagarajan, and Moritz Venzin for helpful discussion.

## Funding

Part of this work was done while Adam Polak was at École Polytechnique Fédérale de Lausanne supported by the Swiss National Science Foundation projects *Latice Algorithms and Integer Programming* (185030) and *Complexity of Integer Programming* (CRFS-2\_207365). During the preparation of this paper Maksym Zub was a participant of the tutoring programme under the Excellence Initiative at the Jagiellonian University.

None of the funding sources had any involvement in study design, in the writing of the report, or in the decision to submit the article for publication.

## Declaration of interests

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

## References

- [1] Ravindra K. Ahuja, Thomas L. Magnanti, and James B. Orlin. *Network flows – theory, algorithms and applications*. Prentice Hall, 1993.
- [2] Yuri Boykov and Vladimir Kolmogorov. An experimental comparison of min-cut/max-flow algorithms for energy minimization in vision. *IEEE Trans. Pattern Anal. Mach. Intell.*, 26(9):1124–1137, 2004. doi:10.1109/TPAMI.2004.60.
- [3] Justin Y. Chen, Sandeep Silwal, Ali Vakilian, and Fred Zhang. Faster fundamental graph algorithms via learned predictions. In *International Conference on Machine Learning, ICML 2022*, volume 162 of *Proceedings of Machine Learning Research*, pages 3583–3602. PMLR, 2022. URL: <https://proceedings.mlr.press/v162/chen22v.html>.

- [4] Li Chen, Rasmus Kyng, Yang P. Liu, Richard Peng, Maximilian Probst Gutenberg, and Sushant Sachdeva. Maximum flow and minimum-cost flow in almost-linear time. In *63rd IEEE Annual Symposium on Foundations of Computer Science, FOCS 2022*, pages 612–623. IEEE, 2022. doi:10.1109/FOCS54457.2022.00064.
- [5] Sami Davies, Benjamin Moseley, Sergei Vassilvitskii, and Yuyan Wang. Predictive flows for faster ford-fulkerson. In *International Conference on Machine Learning, ICML 2023*, volume 202 of *Proceedings of Machine Learning Research*, pages 7231–7248. PMLR, 2023. URL: <https://proceedings.mlr.press/v202/davies23b.html>.
- [6] Michael Dinitz, Sungjin Im, Thomas Lavastida, Benjamin Moseley, and Sergei Vassilvitskii. Faster matchings via learned duals. In *Advances in Neural Information Processing Systems*, volume 34, pages 10393–10406. Curran Associates, Inc., 2021. URL: <https://papers.nips.cc/paper/2021/hash/5616060fb8ae85d93f334e7267307664-Abstract.html>.
- [7] Yefim Dinitz. Algorithm for solution of a problem of maximum flow in networks with power estimation. *Soviet Math. Dokl.*, 11:1277–1280, 1970.
- [8] Talya Eden, Piotr Indyk, Shyam Narayanan, Ronitt Rubinfeld, Sandeep Silwal, and Tal Wagner. Learning-based support estimation in sublinear time. In *9th International Conference on Learning Representations, ICLR 2021*. OpenReview.net, 2021. URL: <https://openreview.net/forum?id=tilovEHA3YS>.
- [9] Jack R. Edmonds and Richard M. Karp. Theoretical improvements in algorithmic efficiency for network flow problems. *J. ACM*, 19(2):248–264, 1972. doi:10.1145/321694.321699.
- [10] Jon C. Ergun, Zhili Feng, Sandeep Silwal, David P. Woodruff, and Samson Zhou. Learning-augmented k-means clustering. In *The Tenth International Conference on Learning Representations, ICLR 2022*. OpenReview.net, 2022. URL: <https://openreview.net/forum?id=X8cLTHexYy>.
- [11] Barak Fishbain, Dorit S. Hochbaum, and Stefan Müller. A competitive study of the pseudoflow algorithm for the minimum s-t cut problem in vision applications. *J. Real Time Image Process.*, 11(3):589–609, 2016. doi:10.1007/s11554-013-0344-3.
- [12] L. R. Ford and D. R. Fulkerson. Maximal flow through a network. *Canadian Journal of Mathematics*, 8:399–404, 1956. doi:10.4153/CJM-1956-045-5.
- [13] Zvi Galil and Amnon Naamad. An  $O(EV \log^2 V)$  algorithm for the maximal flow problem. *J. Comput. Syst. Sci.*, 21(2):203–217, 1980. doi:10.1016/0022-0000(80)90035-5.
- [14] Andrew V. Goldberg and Robert Endre Tarjan. A new approach to the maximum flow problem. In *Proceedings of the 18th Annual ACM Symposium on Theory of Computing, 1986*, pages 136–146. ACM, 1986. doi:10.1145/12130.12144.
- [15] Dorit S. Hochbaum. The pseudoflow algorithm: A new algorithm for the maximum-flow problem. *Oper. Res.*, 56(4):992–1009, 2008. doi:10.1287/opre.1080.0524.
- [16] Wassily Hoeffding. Probability inequalities for sums of bounded random variables. *Journal of the American Statistical Association*, 58(301):13–30, 1963. doi:10.1080/01621459.1963.10500830.
- [17] Chen-Yu Hsu, Piotr Indyk, Dina Katabi, and Ali Vakilian. Learning-based frequency estimation algorithms. In *7th International Conference on Learning Representations, ICLR 2019*. OpenReview.net, 2019. URL: <https://openreview.net/forum?id=r1lohoCqY7>.
- [18] Olivier Juan and Yuri Boykov. Active graph cuts. In *2006 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR 2006)*, pages 1023–1029. IEEE Computer Society, 2006. doi:10.1109/CVPR.2006.47.

- 629 [19] Pushmeet Kohli and Philip H. S. Torr. Efficiently solving dynamic markov random fields using graph cuts. In *10th IEEE International Conference on Computer Vision (ICCV 2005)*, pages 922–929. IEEE Computer Society, 2005. doi:10.1109/ICCV.2005.81.
- 630  
631  
632  
633
- 634 [20] Tim Kraska, Alex Beutel, Ed H. Chi, Jeffrey Dean, and Neoklis Polyzotis. The case for learned index structures. In *Proceedings of the 2018 International Conference on Management of Data, SIGMOD Conference 2018*, pages 489–504. ACM, 2018. doi:10.1145/3183713.3196909.
- 635  
636  
637  
638
- 639 [21] Thodoris Lykouris and Sergei Vassilvitskii. Competitive caching with machine learned advice. *J. ACM*, 68(4):24:1–24:25, 2021. doi:10.1145/3447579.
- 640  
641
- 642 [22] Mohammad Mahdian, Hamid Nazerzadeh, and Amin Saberi. Allocating online advertisement space with unreliable estimates. In *Proceedings 8th ACM Conference on Electronic Commerce (EC-2007)*, pages 288–294. ACM, 2007. doi:10.1145/1250910.1250952.
- 643  
644  
645  
646
- 647 [23] Michael Mitzenmacher and Sergei Vassilvitskii. Algorithms with predictions. In Tim Roughgarden, editor, *Beyond the Worst-Case Analysis of Algorithms*, pages 646–662. Cambridge University Press, 2020. doi:10.1017/9781108637435.037.
- 648  
649  
650
- 651 [24] Manish Purohit, Zoya Svitkina, and Ravi Kumar. Improving online algorithms via ML predictions. In *Advances in Neural Information Processing Systems 31: Annual Conference on Neural Information Processing Systems 2018, NeurIPS 2018*, pages 9684–9693, 2018. URL: <https://proceedings.neurips.cc/paper/2018/hash/73a427badebe0e32caa2e1fc7530b7f3-Abstract.html>.
- 652  
653  
654  
655  
656  
657
- 658 [25] Shinsaku Sakaue and Taihei Oki. Discrete-convex-analysis-based framework for warm-starting algorithms with predictions. In *NeurIPS*, 2022. URL: [http://papers.nips.cc/paper\\_files/paper/2022/hash/844e61124d9e1f58632bf0c8968ad728-Abstract-Conference.html](http://papers.nips.cc/paper_files/paper/2022/hash/844e61124d9e1f58632bf0c8968ad728-Abstract-Conference.html).
- 659  
660  
661  
662  
663
- 664 [26] Daniel A. Spielman and Shang-Hua Teng. Nearly-linear time algorithms for graph partitioning, graph sparsification, and solving linear systems. In *Proceedings of the 36th Annual ACM Symposium on Theory of Computing, 2004*, pages 81–90. ACM, 2004. doi:10.1145/1007352.1007372.
- 665  
666  
667  
668