

# Distributionally Robust Optimization over Time

Kevin-Martin Aigner<sup>\*1</sup>, Andreas Bärmann<sup>1</sup>, Kristin Braun<sup>1</sup>, Frauke Liers<sup>1</sup>, Sebastian Pokutta<sup>2</sup>, Oskar Schneider<sup>3</sup>, Kartikey Sharma<sup>2</sup> and Sebastian Tschuppik<sup>1</sup>

<sup>1</sup>Friedrich-Alexander University Erlangen-Nürnberg, Erlangen, Germany

<sup>2</sup>Zuse Institute Berlin, Berlin, Germany

<sup>3</sup>Fraunhofer Institute for Integrated Circuits IIS, Nürnberg, Germany

12th May 2023

Stochastic Optimization (SO) is a classical approach for optimization under uncertainty that typically requires knowledge about the probability distribution of uncertain parameters. As the latter is often unknown, Distributionally Robust Optimization (DRO) provides a strong alternative that determines the best guaranteed solution over a set of distributions (ambiguity set). In this work, we present an approach for DRO over time that uses online learning and scenario observations arriving as a data stream to learn more about the uncertainty. Our robust solutions adapt over time and reduce the cost of protection with shrinking ambiguity. For various kinds of ambiguity sets, the robust solutions converge to the SO solution. Our algorithm achieves the optimization and learning goals without solving the DRO problem exactly at any step. We also provide a regret bound for the quality of the online strategy which converges at a rate of  $\mathcal{O}(\log T/\sqrt{T})$ , where  $T$  is the number of iterations. Furthermore, we illustrate the effectiveness of our procedure by numerical experiments on mixed-integer optimization instances from popular benchmark libraries and give practical examples stemming from telecommunications and routing. Our algorithm is able to solve the DRO over time problem significantly faster than standard reformulations.

**Keywords:** distributionally robust optimization; learning over time; online gradient descent

## 1 Introduction

Many practical optimization problems deal with uncertainties in the input parameters, and it is important to compute optima that are protected against them. Two prime methodologies to deal with uncertainty in optimization problems are *Stochastic Optimization (SO)* and *Robust Optimization (RO)*. SO considers all uncertain parameters to be random variables, and its solution approaches usually rely on the knowledge of the probability distribution. Classically, SO aims to find solutions that are optimal in expectation (or more generally with respect to chance constraints or different risk-measures, see Birge and Louveaux (1997)). RO is typically used when knowledge about the probability distribution is not at hand or a better guarantee of feasibility is desired (Ben-Tal et al. 2009). It strives to find solutions which perform best against adversarial realizations of the uncertain parameters from a predefined uncertainty set.

Even if the underlying probability distributions are not at hand, they can often be estimated from historical data. These estimates are naturally also affected by uncertainty. As such, recent research has focused on compromising between SO and RO in order to obtain better protection against uncertainty while controlling the conservatism of robust solutions. In particular,

---

<sup>\*</sup>Corresponding author. E-mail address: kevin-martin.aigner@fau.de

*Distributionally Robust Optimization (DRO)* aims to solve a “robust” version of a stochastic optimization problem under incomplete knowledge about the probability distribution. The benefit of DRO is that the solutions are fully protected against the uncertainty and thus outperform non-robust solutions with respect to worst-case performance.

Current research in DRO is primarily aimed at developing efficient solution techniques for static DRO problems, where the optimization problem is solved for a given and fixed ambiguity set. However, in many practical applications, additional information about the uncertainty becomes available over time. For example, in situations where planning processes need to be repeated over time, each plan may want to incorporate the outcomes of the previous decisions. Applications abound for such processes, for example in taxi or in ambulance planning. Applications also occur in iterative assigning landing time windows to aircraft such that security distances are kept at an airport at all times even in case of disturbances which may lead to frequent reassignments. For more details on the air traffic application, see (Heidt et al. 2016, Kapolke et al. 2016). Naturally, it would be beneficial to include any new data into a DRO approach as soon as it becomes available.

In this article, we present a DRO approach that iteratively incorporates such information over time. Specifically, we provide an online learning algorithm that solves DRO problems with limited initial knowledge about the uncertainty, but which can leverage additional incoming data. This allows the optimal solutions to adapt to the uncertainty and gradually reduce the cost of protection. To this end, we use scenario observations arriving as a data stream to construct and update the ambiguity sets. These sets contain the true data generating distribution with high confidence and converge to it over time. We also show that the solution to the DRO problem converges to the true SO solution, as the ambiguity sets shrink to the true distribution, and hence the online algorithm also converges to the SO solution. However, the primary goal is to use the online algorithm to solve the DRO problem.

The main feature of our work is an integrated procedure that can iteratively solve the DRO problem while learning reliable and time-dependent ambiguity sets. We show that our online algorithm outperforms prior methods. We also compare different approaches to construct data-driven ambiguity sets. In computational experiments, we evaluate our algorithm on state-of-the-art benchmark libraries and realistic case studies. We demonstrate that our online method leads to significantly reduced computation times with only marginal sacrifices in solution quality.

## **Problem Statement.**

We consider the problem of minimizing a function  $f: \mathcal{X} \times \mathcal{S} \rightarrow \mathbb{R}$  over some (possibly non-linear and/or mixed-integer) feasible set  $\mathcal{X}$ . We focus on the case of finitely many scenarios which are contained in the set  $\mathcal{S} := \{s_1, \dots, s_{|\mathcal{S}|}\}$ . Note that this is a natural modeling assumption that appears in several applications as many real-world random events are best represented via discrete scenarios. However, our approach is also able to treat the case of continuous random variables by sampling a sufficient large discrete scenario set from the probability distribution. Finite approximations such as the sample average approximation (Kleywegt et al. 2002) and other similar scenario reduction techniques are standard in stochastic optimization. They lead to algorithmic tractability in more general settings, which is necessary for any realistically sized problem. Furthermore, the use of a finite number of scenarios allows us to treat the probability distribution as a vector and hence leverage methods from first-order optimization to solve the problem of finding the optimal distribution. The setting of infinite-dimensional ambiguity sets for continuous distributions goes beyond the focus of this paper.

With this in mind, we define  $\mathcal{P}_0 := \{p \in [0, 1]^{|\mathcal{S}|} \mid \sum_{k=1}^{|\mathcal{S}|} p_k = 1\}$  as the  $|\mathcal{S}|$ -dimensional probability simplex. We start with the probability simplex as our initial ambiguity set as it imposes no restrictions on the distributions. However, we are not restricted to this choice and one can initialize with sets constructed using already available historical data. This would lead to less conservative solutions in practice without changing our theoretical results fundamentally.

Each point  $p \in \mathcal{P}_0$  represents a probability distribution over the scenarios  $s \in \mathcal{S}$ . Given a probability distribution of scenarios  $p^* \in \mathcal{P}_0$ , one can solve the following SO problem:

$$J^* := \min_{x \in \mathcal{X}} \mathbb{E}_{s \sim p^*} [f(x, s)] = \sum_{k=1}^S f(x, s_k) p_k^*. \quad (\text{SO})$$

However, if there is limited information about the probability vector  $p^*$ , we can limit the impact of uncertainty by solving the distributionally robust counterpart of the SO problem, namely

$$\min_{x \in \mathcal{X}} \max_{p \in \mathcal{P}_0} \mathbb{E}_{s \sim p} [f(x, s)] = \min_{x \in \mathcal{X}} \max_{s \in \mathcal{S}} f(x, s). \quad (\text{DRO})$$

This equality holds because the worst-case probability is realized by some unit-vector  $(0, \dots, 0, 1, 0, \dots, 0) \in \mathcal{P}_0$ . For this maximally large ambiguity set, the solutions to the DRO problem may be overly conservative. Furthermore, there is typically data available in the form of observed realizations of the uncertain parameters. The primary goal of this work is to develop an online algorithm to solve the above DRO problem over time while progressively integrating additional information. This is achieved by refining the ambiguity sets over time as we learn more about the uncertainty, e.g., with new realizations. The use of DRO along with online optimization limits the impact of adverse realizations. Simultaneously, learning ensures that we can adapt and increase our confidence as we gather more information.

We assume that information about the probability distribution  $p^*$  is revealed in the form of i.i.d. realizations over time. As such, we solve a sequence of DRO problems with progressively shrinking ambiguity sets  $\mathcal{P}_t$ . At each time step  $t = 1, \dots, T$  over a given horizon, we construct the ambiguity set  $\mathcal{P}_t$  according to the confidence regions estimated from scenario observations up to time  $t$ . Using these sets, we solve

$$\hat{J}_t := \min_{x \in \mathcal{X}} \max_{p \in \mathcal{P}_t} \mathbb{E}_{s \sim p} [f(x, s)]. \quad (\text{DRO}_t)$$

In order to solve problem  $\text{DRO}_t$  efficiently, our online optimization approach alternates between a gradient step (for  $p$ ) and solving the minimization problem (for  $x$ ). The task is to find a solution  $x_t$  for each round  $t$ . This approach can also be interpreted as a game over  $T$  rounds, where a player tries to make optimal decisions against an adversary who chooses the probability distribution from which to draw the uncertain realization.

In our analysis, we calculate the gap between the worst-case performance of  $x_t$  (generated by the online algorithm) over the set  $\mathcal{P}_t$  and the performance of the optimal DRO solution at time  $t$ . We prove that the average gap (reminiscent of the notion of dynamic regret) is bounded as

$$\frac{1}{T} \sum_{t=1}^T \left( \underbrace{\max_{p \in \mathcal{P}_t} \mathbb{E}_{s \sim p} [f(x_t, s)]}_{\text{our online cost at time } t} - \underbrace{\min_{x \in \mathcal{X}} \max_{p \in \mathcal{P}_t} \mathbb{E}_{s \sim p} [f(x, s)]}_{\text{optimal cost in hindsight at } t} \right) \leq \mathcal{O} \left( \sqrt{\frac{h(T)}{T}} \right),$$

with high probability. Here,  $h(T)$  is a bound on the path length of the distributions, i.e.,  $\sum_{t=1}^T \frac{1}{2} \|p_t - q_t\|_2^2 \leq h(T)$  for  $p_t \in \mathcal{P}_{t-1}$  and  $q_t \in \mathcal{P}_t$ . We show that  $h(T) \in \mathcal{O}(\log^2 T)$  for the different categories of ambiguity sets that we consider. This bound controls the difference between the performance of our online method and an exact DRO solver. It certifies that our approach successfully approximates the DRO solution with an average gap that decreases over time. We also show that the DRO solution converges to the true SO solution as the ambiguity sets converge to the true distribution in the limit.

## Related Work.

Current research in distributional robustness is mostly concerned with the appropriate choice of ambiguity sets to obtain guarantees on solution quality (Delage and Ye 2010, Van Parys

et al. 2021). It also focuses on the derivation of algorithmically tractable reformulations for the resulting robust counterparts, see e.g. Wiesemann et al. (2014), as well as Calafiore and Ghaoui (2006). Ambiguity sets can be constructed by imposing constraints on expectation (Chen et al. 2020), covariance (Delage and Ye 2010), mode (Hanasusanto et al. 2015), etc. Another option is to use distance metrics such as the Wasserstein metric (Esfahani and Kuhn 2018),  $\phi$ -divergence (Bayraksan and Love 2015),  $f$ -divergence (Duchi et al. 2021), kernel-based distances (Kirschner et al. 2020), hypothesis tests (Bertsimas et al. 2018), etc.

Ambiguity sets can also be defined by confidence bounds (Rahimian and Mehrotra 2019), and many popular ambiguity sets (Bertsimas et al. 2018, Kirschner et al. 2020) have associated probabilities of containing the true distribution. In this work, we construct the sets as the combination of a simplex and a bounded set defined by either confidence intervals,  $\ell_2$ -norm or kernel based metrics. These bounded sets function as confidence regions.

Online learning is an established field which provides algorithms for solving problems over time. For a broad introduction see Hazan et al. (2016). Recently, this approach has been leveraged to solve robust optimization problems (Ben-Tal et al. 2015, Ho-Nguyen and Kılınç-Karzan 2018, Pokutta and Xu 2021, Chen et al. 2017). Online learning has also been applied to DRO problems where the ambiguity set is constructed from data. Namkoong and Duchi (2016) leverage online optimization for a DRO problem with ambiguity sets defined by  $f$ -divergences. They use an alternating mirror descent algorithm and provide regret bounds for the same. The authors of Qi et al. (2021) propose a duality-free online stochastic method for a class of DRO problems with KL-divergence regularization on the dual variables. Different algorithms are proposed and analyzed in Levy et al. (2020) for DRO problems with convex objectives with conditional value at risk and  $\chi^2$ -divergence ambiguity sets. These existing works combine DRO and online learning in order to solve a single DRO problem. This paper looks at a different problem. The key difference between these works and ours is that we consider a planning problem that has to be solved repeatedly in time with growing knowledge about the uncertain parameters. Therefore, we establish an online framework to solve a series of DRO problem with changing ambiguity sets as new information arrives. This allows us to obtain robust online solutions while learning about the uncertainty.

In Kirschner et al. (2020), the authors also focus on a DRO problem in an online learning context. They allow for uncertainty in the parameters and only have noisy blackbox access to the objective function. To solve the DRO problem while learning the objective and the ambiguity set, the authors solve a large program with convex constraints at each stage. This was extended to more general ambiguity sets constructed with  $\phi$ -divergences in Husain et al. (2022). In contrast, our work primarily focuses on obtaining distributionally robust optimal solutions in an online fashion while learning the true distribution. As a key advantage we point out that our algorithm does not require the solution of the entire DRO problem at any step but only computes gradient steps in the ambiguity space. Furthermore, we consider multiple ambiguity sets and leverage the online gradient descent algorithm to allow for faster computation and better applicability in real-world settings.

Another similar work is by Sessa et al. (2020). Therein, the authors present an online learning approach with a multiplicative weight algorithm in order to compute strict robust mixed-strategies over a decision set. In contrast to our work however, they do not consider a DRO setting. Finally, Bärmann et al. (2018) and Bärmann et al. (2017) consider the problem of learning unknown objectives in an online fashion but without any uncertainty in the models.

### Contribution.

The two key differences in our work which distinguish it from regular online optimization are (i) use of DRO while learning from data and (ii) solving the DRO problem approximately. The first ensures that our solutions are robust to uncertainty in the knowledge of the true probability distribution. The second allows us to obtain robust solutions without solving the DRO problem

exactly at each step. Specifically, the key contributions of our work are:

*Online Learning Algorithm for DRO.* We provide an online algorithm to solve the DRO problem. It also learns the uncertainty from scenario observations over time, shrinking the ambiguity sets. This allows for rapid computation of the DRO solutions along with their adaptation. Thus, reducing the cost of protection.

*Stochastic Consistency.* We also prove that the solution of the DRO problem converges to the SO problem. Since our online algorithm solves the DRO problem, thus it too converges to the solution of the SO problem.

*High Probability Regret Bounds.* We prove that the cumulative regret between the solutions generated by our online method and the exact DRO solution at each time step shrinks at a rate of  $\mathcal{O}(\log T/\sqrt{T})$  with high probability.

*Flexibility of Uncertainty Models.* We consider 3 different ambiguity models: (i) confidence intervals, (ii)  $\ell_2$ -norm sets and (iii) kernel based ambiguity sets. These allow our approach to adapt to the application.

*Computational Results.* We provide a computational study on mixed-integer benchmark instances and on real world problem examples. Specifically, we compare on the MIPLIB and QPLIB libraries and further illustrate our results with two realistic applications from telecommunications and routing. In both cases our approach is considerably faster than solving the full distributionally robust counterparts.

## Outline

Section 2 presents the foundations of *Data-Driven DRO*. We present our algorithms and theoretical results on *DRO over Time* in Section 3. Finally, in Section 4, we evaluate our methods on benchmark instances.

## 2 Data-Driven DRO

In this section, we introduce ambiguity sets which form a key part of DRO problems. We also introduce the dual reformulation which is a standard way for solving robust or DRO problems.

### 2.1 Ambiguity Sets

DRO can provide robust protection against scenarios generated by any distribution inside an ambiguity set  $\mathcal{P}$ . However, depending on the size of the set e.g., if  $\mathcal{P} = \mathcal{P}_0$ , the protection may be too conservative. We can reduce this conservativeness by integrating the information gained from the data generated by the true distribution and thus shrinking the ambiguity set. We construct the ambiguity sets as the intersection of the probability simplex and a data-driven set which contains the true distribution (with high confidence). These data-driven sets can be prescribed by multi-dimensional confidence intervals or metrics such as the  $\ell_2$ -norm, kernel based deviation etc. These metrics are selected as they can provide guarantees about containing the true distribution. Their dependence on data also ensures that with more data, the sets converge to the true distribution.

Our data-driven ambiguity sets  $\mathcal{P}_t$ , where  $t \in \mathbb{N}$  equals the number of data points, can be of the following two forms

$$\mathcal{P}_t := \{p \in \mathcal{P} \mid l_t \leq p \leq u_t\} \text{ or } \mathcal{P}_t := \{p \in \mathcal{P} \mid d(p, \hat{p}_t) \leq \epsilon_t\}.$$

Here  $l_t, u_t \in [0, 1]^{|S|}$  are lower and upper bounds of confidence intervals,  $\hat{p}_t$  is the empirical distribution estimator for  $p^*$  and  $d(\cdot, \cdot)$  and  $\epsilon_t$  denote a distance metric and its respective bound. In both cases, the values of the parameters  $l_t, u_t$  and  $\epsilon_t$  are selected such that the true distribution lies inside the sets  $\mathcal{P}_t$  with high probability i.e.,

$$\mathbb{P}(p^* \in \mathcal{P}_t) \geq 1 - \delta_t,$$

where  $\delta_t \in (0, 1)$ . One key requirement for the ambiguity sets  $\mathcal{P}_t$  over time  $t = 0, \dots, T$  is that they contain the true distribution inside *all of them* with high probability, that means

$$p^* \in \bigcap_{t=0, \dots, T} \mathcal{P}_t,$$

with a probability of at least  $1 - \delta$ . This is achieved by ensuring that each set  $\mathcal{P}_t$  in round  $t$  contains the true distribution  $p^*$  with probability at least  $1 - \delta_t$  such that  $\sum_{t=1}^{\infty} \delta_t < \infty$ . In this paper, we choose  $\delta_t = \frac{6\delta}{\pi^2 t^2}$  for some predefined  $\delta \in (0, 1)$ . With this in mind, we have the following lemma.

**Lemma 2.1** *For ambiguity sets constructed with confidence  $\delta_t := \frac{6\delta}{\pi^2 t^2}$  and  $\delta \in (0, 1)$ , it follows that the true data generating distribution  $p^* \in \bigcap_{t=0, \dots, T} \mathcal{P}_t$  with a probability of at least  $1 - \delta$ .*

*Proof:* Given a sequence of events  $A_t$ , we can estimate the probability of their intersection with Boole's inequality as follows,

$$\mathbb{P}\left(\bigcap_{t=1, \dots, T} A_t\right) = 1 - \mathbb{P}\left(\bigcup_{t=1, \dots, T} A_t^c\right) \geq 1 - \sum_{t=1, \dots, T} \mathbb{P}(A_t^c).$$

Let  $A_t$  be the event that the true distribution  $p^*$  lies inside the uncertainty set  $\mathcal{P}_t$ . By the definition of an ambiguity set with confidence  $\delta_t$ , we know that  $\mathbb{P}(A_t) \geq 1 - \delta_t$ . This means that  $\mathbb{P}(A_t^c) \leq \delta_t$ . This inequality and the limit of the over-harmonic series (2-series:  $\sum_{t=1}^{\infty} \frac{1}{t^2} = \frac{\pi^2}{6}$ ) allows us to show that the probability of the event  $p^* \in \bigcap_{t=1 \dots T} \mathcal{P}_t$  (here without  $t = 0$ ) is at least

$$1 - \sum_{t=1}^T \delta_t \geq 1 - \sum_{t=1}^{\infty} \delta_t = 1 - \frac{6\delta}{\pi^2} \sum_{t=1}^{\infty} \frac{1}{t^2} = 1 - \delta.$$

Since

$$\bigcap_{t=1, \dots, T} \mathcal{P}_t \subset \mathcal{P} = \mathcal{P}_0 \Rightarrow \bigcap_{t=1, \dots, T} \mathcal{P}_t = \bigcap_{t=0, \dots, T} \mathcal{P}_t,$$

we conclude the proof (here with  $t = 0$ ).  $\square$

Note that the above result continues to hold as  $T \rightarrow \infty$ . The above lemma shows that if the confidence probability  $1 - \delta_t$  increases fast enough, the inclusion of the true distribution  $p^*$  in individual ambiguity sets  $\mathcal{P}_t$  is sufficient to guarantee uniform inclusion over all sets.

## 2.2 Choice of Ambiguity Sets

The choice of ambiguity sets depends on the application, the historical information available and level of protection desired. We discuss different ambiguity sets that are easily applicable and provide probabilistic performance guarantees. These different ways of constructing data-driven ambiguity sets demonstrate the flexibility of our approach.

**Confidence Intervals.** Confidence intervals for multinomial distributions can be calculated using various methods (see Wang (2008) for a survey). We use the analytic formula from Fitzpatrick and Scott (1987) for the construction of multi-dimensional intervals  $\mathcal{I}_t \subseteq [0, 1]^{|S|}$  via

$$\mathcal{I}_{kt} = [l_{kt}, u_{kt}] := \left[ \hat{p}_{kt} - \frac{z_{\delta_t}}{2\sqrt{t}}, \hat{p}_{kt} + \frac{z_{\delta_t}}{2\sqrt{t}} \right], \quad (1)$$

in each round  $t = 1, \dots, T$ . Here,  $\hat{p}_t$  is the maximum likelihood estimator for  $p^*$  and  $z_{\frac{\delta_t}{2}}$  denotes the upper  $(1 - \frac{\delta_t}{2})$ -percentile of the standard normal distribution. The corresponding ambiguity sets can then be the intersection of the confidence intervals and the probability simplex, i.e.,

$$\mathcal{P}_t := \mathcal{P}_0 \cap \mathcal{I}_t = \{p \in \mathcal{P}_0 \mid l_t \leq p \leq u_t\}, \quad t = 1, \dots, T.$$

The parameters  $l_t$  and  $u_t$  can be updated as mentioned in the definition of  $\mathcal{I}_{kt}$ . The parameter  $\hat{p}_t$  is the empirical probability distribution and can be updated by counting the observations of each scenario up to  $t$ . Confidence intervals work well in practice and are algorithmically preferable as they only impose linear constraints on the ambiguity sets. This provides significant scalability to problems which use confidence intervals. However, one disadvantage of them is that the probability guarantees that they provide only hold asymptotically. As a consequence, we extend our method to the following two sets which provide finite sample guarantees while being easy to reformulate.

**$\ell_2$ -norm sets.** For the finite sample setting, we have the following guarantee as proven in Weisman et al. (2003),

$$\mathbb{P}\left(\|\hat{p}_t - p^*\|_1 \leq \sqrt{\frac{2|\mathcal{S}|\log 2/\delta_t}{t}}\right) \geq 1 - \delta_t.$$

Given this inequality, along with the observation that  $\|p\|_2 \leq \|p\|_1$  we can construct the following set

$$\mathcal{P}_t = \{p \in \mathcal{P} \mid \|p - \hat{p}_t\|_2 \leq \epsilon_t\},$$

with  $\epsilon_t := \sqrt{\frac{2|\mathcal{S}|\log 2/\delta_t}{t}}$ , which provides the containment guarantee  $1 - \delta_t$ . For the above ambiguity set,  $\hat{p}_t$  and  $\epsilon_t$  are the two parameters which define the set. These can be updated by counting the scenario observations and as per the definition of  $\epsilon_t$  respectively.

**Kernel based ambiguity sets.** Kernel based ambiguity sets are another alternative that provide similar finite sample probability guarantees while allowing for flexibility in how the different scenarios are weighted. Given a kernel function  $k_M(s_i, s_j) : \mathcal{S} \times \mathcal{S} \rightarrow \mathbb{R}$  defined over the scenarios with kernel matrix  $M$ , we have the following probability guarantee (Kirschner et al. 2020),

$$\mathbb{P}\left(\|\hat{p}_t - p^*\|_M \leq \frac{\sqrt{C}}{\sqrt{t}}(2 + \sqrt{2\log(1/\delta_t)})\right) \geq 1 - \delta_t,$$

if  $k_M(s_i, s_j) \leq C$  for all scenarios. Here,  $\|q\|_M := \sqrt{q^\top M q}$ . With this inequality, we can construct an ambiguity set similar to the  $\ell_2$ -norm case with  $\epsilon_t := \frac{\sqrt{C}}{\sqrt{t}}(2 + \sqrt{2\log(1/\delta_t)})$ . Like  $\ell_2$  norm ambiguity sets, kernel based sets require two key parameters  $\hat{p}_t$  and  $\epsilon_t$ . Both can be updated either by counting scenarios or the definition of  $\epsilon_t$  for kernel sets.

Using these ambiguity sets, the resulting  $(\text{DRO}_t)$  forms a min-max problem, where the inner maximization problem optimizes a linear objective over a finite-dimensional convex feasible set. Therefore,  $(\text{DRO}_t)$  can be equivalently reformulated using duality theory as is commonly applied in convex robust optimization. These reformulations are discussed in the following section.

### 2.3 Solving DRO via Reformulation

When the ambiguity set is constructed with confidence intervals, the inner maximization problem forms a linear program. Then, by strong duality, for confidence intervals,  $(\text{DRO}_t)$  is equi-

valent to

$$\begin{aligned}
& \min_{x,z,\alpha,\beta} z - \langle l_t, \alpha \rangle + \langle u_t, \beta \rangle \\
& \text{s.t. } z - \alpha_k + \beta_k \geq f(x, s_k) \quad \forall k = 1, \dots, |\mathcal{S}|, \\
& \alpha, \beta \geq 0, \\
& x \in \mathcal{X}, z \in \mathbb{R}, \alpha, \beta \in \mathbb{R}^{|\mathcal{S}|}.
\end{aligned} \tag{2}$$

Here, the dual variables  $\alpha_k$  and  $\beta_k$  price the uncertainty for scenario  $s_k \in \mathcal{S}$ . This problem is of the same problem class as (SO), however larger in size. The reformulated DRO problem grows linearly with the number of scenarios which may become prohibitive if the cardinality of  $\mathcal{S}$  is large. Thus, the difficulty of solving (2) depends on the complexity of  $f$  and the cardinality of  $\mathcal{S}$ .

For the  $\ell_2$ -norm and kernel based ambiguity sets, the dual reformulation is given by

$$\begin{aligned}
& \min_{x,z,r} \sum_{k=1}^{|\mathcal{S}|} \hat{p}_{kt} (f(x, s_k) + r_k) + \epsilon_t \|\mathbf{f}(x) + z\mathbf{e} + \mathbf{r}\|_{A^{-1}} \\
& \text{s.t. } x \in \mathcal{X}, r_k \geq 0,
\end{aligned} \tag{3}$$

where the matrix  $A = I$  for the  $\ell_2$ -norm case and  $M$  for the kernel based sets. The vectors  $\mathbf{f}$ ,  $\mathbf{e}$ , and  $\mathbf{r}$  denote the functions  $f(x, s_k)$  for all  $k$ , the vector of all ones and the set of all  $r_k$  respectively. Although norm constraints typically remain algorithmically tractable, solving this reformulation is practically more difficult than solving SO.

## 2.4 DRO over Time

We now discuss our baseline framework of DRO over Time. The iterative procedure outlined in Algorithm 1 solves a sequence of reformulated DRO problems while learning from data. Prior to the first round, we assume that we have no data about the scenarios and therefore initialize the ambiguity set with the full probability simplex. As more information comes in, the ambiguity set is updated.

---

### Algorithm 1 DRO over Time

---

- 1: **Input:** functions  $f(\cdot, s)$  for  $s \in \mathcal{S}$ , feasible set  $\mathcal{X}$ , initial ambiguity set  $\mathcal{P}_0$
  - 2: **Output:** sequence of DRO solutions  $x_1, \dots, x_T$
  - 3: **for**  $t = 1$  **to**  $T$  **do**
  - 4:    $x_t \leftarrow$  solve Problem (2) or (3) for  $\mathcal{P}_{t-1}$
  - 5:    $\mathcal{P}_t \leftarrow$  observe data and update set parameters such as  $\hat{p}_t, l_t, u_t$  and  $\epsilon_t$  as per the type of ambiguity set.
  - 6: **end for**
- 

The following theorems show the convergence of the DRO solution  $\hat{J}_t$  to the true solution of the SO problem  $J^*$ . The proofs are provided in the electronic companion as they are adaptations from Esfahani and Kuhn (2018) to our setting. In the latter, analogous results have been proven for ambiguity sets constructed with the Wasserstein metric.

**Theorem 2.2** *If the feasible set  $\mathcal{X}$  is compact, then the optimal value of  $DRO_t$  converges over time to the optimal value of SO with probability 1, i.e.,*

$$\lim_{t \rightarrow \infty} \hat{J}_t = J^* \text{ with probability 1.}$$

**Theorem 2.3** *Let  $\{x_t\}_{t=1}^{\infty}$  be a sequence of optimal solutions to the problem  $DRO_t$ . If the feasible set  $\mathcal{X}$  is compact and the function  $f(x, s)$  is continuous in  $x$ , then any accumulation point of  $\{x_t\}_{t=1}^{\infty}$  is almost surely an optimal solution to the problem SO.*



These two results show the validity of the DRO over Time paradigm and guarantee that with a sufficient amount of data, Algorithm 1 converges to the solution of the SO problem with the true distribution. This is important as this indicates the importance of gathering more data to obtain better solutions.

In Algorithm 1, the DRO problem has to be solved repeatedly in each round. This is not viable for large problems as it requires long computation times. We remedy this by introducing a new algorithm which approximates the DRO problem over time and updates the ambiguity sets via online learning. For this online algorithm, we can show a dynamic regret bound that bounds the error of the approximation with sublinear expression in the number of rounds.

### 3 Online Robust Optimization

In this section, we introduce the online learning and optimization algorithm which is main contribution of our work. As in Ben-Tal et al. (2015) and Pokutta and Xu (2021), we consider robust optimization as a game between two players. The online algorithm can then be roughly described as alternating between solving the optimization problem for each player given the solution of the other. Thus, for each round  $t = 1, \dots, T$ , we decompose the min-max problem into two subproblems (one for each player) and perform the following steps:

1. First, the  $p$ -player determines  $p_t$  via an appropriate algorithm applied to problem

$$\max_{p \in \mathcal{P}_{t-1}} \mathbb{E}_{s \sim p} [f(x_{t-1}, s)],$$

based on the solution  $x_{t-1}$  from the previous round.

2. Then, the  $x$ -player computes  $x_t$  as a solution using the previously calculated  $p_t$

$$\min_{x \in \mathcal{X}} \mathbb{E}_{s \sim p_t} [f(x, s)].$$

Due to the fact that we consider probability distributions over a finite scenario set, the optimization problem of the  $p$ -player is finite-dimensional and therefore the Online Gradient Descent (Zinkevich 2003) is a canonical choice here as a learning algorithm. Given  $p_{t-1}$  and  $x_{t-1}$ , the update rule consists of a descent step

$$\tilde{p}_t = p_{t-1} + \eta \nabla_p \mathbb{E}_{s \sim p_{t-1}} [f(x_{t-1}, s)],$$

with step size  $\eta > 0$  and a subsequent projection step to ensure feasibility

$$p_t = \arg \min_{p \in \mathcal{P}_{t-1}} \frac{1}{2} \|p - \tilde{p}_t\|^2.$$

The probability distribution in the next iteration is therefore given as the unique solution

$$p_t = \arg \min_{p \in \mathcal{P}_{t-1}} \langle -\eta \nabla_p \mathbb{E}_{s \sim p_{t-1}} [f(x_{t-1}, s)], p \rangle + \frac{1}{2} \|p - p_{t-1}\|^2.$$

Note that the  $x$ -player uses  $p_t$  in round  $t$  to estimate  $x_t$  and solves a standard SO problem which is easier to solve than the reformulated DRO program. This  $p_t$  is then used to estimate  $p_{t+1}$  in round  $t + 1$ . As such, the value of  $x_t$  depends on  $p_t$  in round  $t$  (thus  $x_t$  and  $p_t$  are not conditionally independent in round  $t$ ). Therefore, as per Pokutta and Xu (2021), the learner for  $p_t$  has to be a strong learner in order to ensure sublinear regret. This paper also shows that the online gradient descent algorithm satisfies the necessary conditions for it to be a strong learner.

### 3.1 Algorithm

We provide a pseudo code of our method for DRO over Time via online robust optimization in Algorithm 2. It combines alternating solutions of the min and max problems with the update of the ambiguity sets. For the sake of simplicity, we assume that our algorithm starts without any knowledge of the probability distribution over the scenarios and therefore we initialize the ambiguity set as the full probability simplex in step 1. The algorithm can be easily modified to incorporate any historical information. The initialization of  $p_0 \in \mathcal{P}_0$  and  $x_0 \in \mathcal{X}$  in step 4 and step 5 can be chosen arbitrarily and does not effect our theoretical results. Each round  $t = 1, \dots, T$  starts with the update of  $p_t$  via projected gradient descent and of  $x_t$  as the solution of an SO problem. At the end of the round, we observe new data in form of (i.i.d.) scenario observations and update the ambiguity set as explained in Section 2.

---

#### Algorithm 2 DRO over Time with Online Projected Gradient Descent

---

- 1: **Input:** functions  $f(\cdot, s)$  for  $s \in \mathcal{S}$ , feasible set  $\mathcal{X}$ , initial ambiguity set  $\mathcal{P}_0$
  - 2: **Output:**  $x_1, \dots, x_T$
  - 3: Set  $\mathcal{P}_0 = \{p \in [0, 1]^{|\mathcal{S}|} \mid \sum_{k=1}^{|\mathcal{S}|} p_k = 1\}$
  - 4: Set  $p_0 = \left(\frac{1}{|\mathcal{S}|}, \dots, \frac{1}{|\mathcal{S}|}\right) \in [0, 1]^{|\mathcal{S}|}$
  - 5: Set  $x_0 = \min_{x \in \mathcal{X}} f(x, s_1)$
  - 6: **for**  $t = 1$  **to**  $T$  **do**
  - 7:    $\tilde{p}_t \leftarrow p_{t-1} + \eta \nabla_p \mathbb{E}_{s \sim p_{t-1}} [f(x_{t-1}, s)]$
  - 8:    $p_t \leftarrow \arg \min_{p \in \mathcal{P}_{t-1}} \frac{1}{2} \|p - \tilde{p}_t\|^2$
  - 9:    $x_t \leftarrow \arg \min_{x \in \mathcal{X}} \mathbb{E}_{s \sim p_t} [f(x, s)]$
  - 10:    $\mathcal{P}_t \leftarrow$  observe data and update set parameters such as  $\hat{p}_t, l_t, u_t$  and  $\epsilon_t$  as per the type of ambiguity set.
  - 11: **end for**
- 

Algorithm 2 provides a sequence of solutions  $x_t$  for each time step  $t = 0, \dots, T$ . In order to prove that the quality of the solutions  $x_t$  improves over time, we bound the average gap over time between the worst case performance of  $x_t$  and the optimal worst case (exact DRO) solution. Since the feasible set of the  $p$ -player changes over time it is not possible to apply existing techniques for regret bounds on min-max problems to our setting as these techniques have primarily focused on stationary ambiguity sets. In this paper, we extend these existing techniques to the case of shrinking ambiguity sets by leveraging the fact that all the ambiguity sets contain a common distribution (the true distribution).

The two main ingredients for the theoretical analysis are a constrained gradient ( $\nabla_p \mathbb{E}_{s \sim p} [f(x, s)]$ ) and a constrained path length ( $\sum_{t=1}^T \frac{1}{2} \|p_t - q_t\|^2 \leq h(T)$  for all  $p_t \in \mathcal{P}_{t-1}, q_t \in \mathcal{P}_t$ ) for the online gradient descent to work on non-stationary feasible sets. The former is a classical assumption for steepest descent algorithms while the latter is commonly used in dynamic regret bounds for online algorithms, see e.g. (Zhao et al. 2020, Mokhtari et al. 2016).

For constant ambiguity sets, it is known that for the static regret (which compares the performance difference of online solutions to a single best action in hindsight) bounds of the form  $\mathcal{O}(1/\sqrt{T})$  can be derived (Pokutta and Xu 2021, Besbes et al. 2015). For our case, a careful analysis of the algorithm leads to a dynamic regret bound of  $\mathcal{O}(\sqrt{h(T)}/\sqrt{T})$  that is presented in Theorem 3.1, with the corresponding bounding terms  $h(T)$  for the path lengths being proven afterwards. This is consistent with other findings about dynamic regret bounds in the literature on online learning, see (Mokhtari et al. 2016). We are able to achieve a bound for this new setting with shrinking ambiguity sets.

**Theorem 3.1 (Dynamic regret bound)** *Let  $f : \mathcal{X} \times \mathcal{S} \rightarrow \mathbb{R}$  be uniformly bounded, i.e., for all  $(x, s) \in \mathcal{X} \times \mathcal{S}$ , there exists a constant  $G > 0$  such that  $|f(x, s)| \leq G$ . Let  $\eta := \sqrt{\frac{2h(T)}{G^2 T |\mathcal{S}|}}$  where*

$\sum_{t=1}^T \frac{1}{2} \|p_t - q_t\|^2 \leq h(T)$  for  $p_t \in \mathcal{P}_{t-1}$  and  $q_t \in \mathcal{P}_t$ . The output  $(x_1, \dots, x_T)$  from Algorithm 2 with confidence update  $\delta_t := \frac{6\delta}{\pi^2 t^2}$  and  $\delta \in (0, 1)$  fulfills

$$\frac{1}{T} \sum_{t=1}^T \left( \max_{p \in \mathcal{P}_t} \mathbb{E}_{s \sim p} [f(x_t, s)] - \min_{x \in \mathcal{X}} \max_{p \in \mathcal{P}_t} \mathbb{E}_{s \sim p} [f(x, s)] \right) \leq G \sqrt{\frac{2|\mathcal{S}|h(T)}{T}} + \frac{2G}{T},$$

with probability at least  $1 - \delta$ .

*Proof of Theorem 3.1* Define  $g_t(p) := -\mathbb{E}_{s \sim p} [f(x_t, s)]$ . An online gradient descent iteration is given by

$$p_{t+1} = \arg \min_{p \in \mathcal{P}_t} \langle \eta \nabla g_t(p_t), p \rangle + \frac{1}{2} \|p - p_t\|^2,$$

with the variational inequality

$$\langle \eta \nabla g_t(p_t), u_t - p_{t+1} \rangle + \langle p_{t+1} - p_t, u_t - p_{t+1} \rangle \geq 0, \text{ for all } u_t \in \mathcal{P}_t$$

as the optimality criteria. Classical theory for gradient descent (by rearranging the previous inequality and using Cauchy-Schwarz) yields

$$\langle \eta \nabla g_t(p_t), p_t - u_t \rangle \leq \frac{1}{2} \|p_t - u_t\|^2 - \frac{1}{2} \|p_{t+1} - u_t\|^2 + \frac{\eta^2}{2} \|\nabla g_t(p_t)\|^2.$$

Summation over rounds  $t = 1, \dots, T$  results in the following inequality for all  $u_t \in \mathcal{P}_t$ :

$$\sum_{t=1}^T \langle \eta \nabla g_t(p_t), p_t - u_t \rangle \leq \sum_{t=1}^T \frac{\eta^2}{2} \|\nabla g_t(p_t)\|^2 + \sum_{t=1}^T \left( \frac{1}{2} \|p_t - u_t\|^2 - \frac{1}{2} \|p_{t+1} - u_t\|^2 \right).$$

Next, we bound the terms on the right side starting with

$$\sum_{t=1}^T \left( \frac{1}{2} \|p_t - u_t\|^2 - \frac{1}{2} \|p_{t+1} - u_t\|^2 \right) \leq \sum_{t=1}^T \frac{1}{2} \|p_t - u_t\|^2.$$

Note that this is, in contrast to classical steepest descent theory, not a telescoping sum. We make use of the fact that  $\sum_{t=1}^T \frac{1}{2} \|p_t - u_t\|^2 \leq h(T)$  for all  $u_t \in \mathcal{P}_t$  and all rounds  $t = 1, \dots, T$  with probability at least  $1 - \delta$  as per Lemma 3.3. We provide the proofs for these path length bounds in the next extra subsection. Furthermore, since  $g_t(p) = -\mathbb{E}_{s \sim p} [f(x_t, s)]$  is linear in  $p$ , we use the gradient bound

$$\|\nabla g_t(p_t)\|^2 = \sum_{k=1}^{|\mathcal{S}|} |f(x_t, s_k)|^2 \leq |\mathcal{S}| G^2,$$

for all  $t = 1, \dots, T$ , because  $f$  is bounded. Thus we get

$$\sum_{t=1}^T \langle \nabla g_t(p_t), p_t - u_t \rangle \leq \frac{h(T)}{\eta} + \frac{\eta}{2} T |\mathcal{S}| G^2,$$

for all  $u_t \in \mathcal{P}_t$ , for all  $t = 1, \dots, T$  with probability at least  $1 - \delta$ . Choosing the bound minimizing step size (minimize right-hand side with respect to  $\eta$ )  $\eta := \sqrt{\frac{2h(T)}{G^2 |\mathcal{S}| T}}$  yields

$$\sum_{t=1}^T \langle \nabla g_t(p_t), p_t - u_t \rangle \leq G \sqrt{2|\mathcal{S}|h(T)T}.$$

Since  $g_t(p) = -\mathbb{E}_{s \sim p} [f(x_t, s)]$  is linear in  $p$  for all  $t = 1, \dots, T$ , it follows

$$\sum_{t=1}^T (\mathbb{E}_{s \sim u_t} [f(x_t, s)] - \mathbb{E}_{s \sim p_t} [f(x_t, s)]) = \sum_{t=1}^T \langle \nabla g_t(p_t), p_t - u_t \rangle \leq G \sqrt{2|\mathcal{S}|h(T)T}.$$

Now we choose in each round  $t = 1, \dots, T$  the worst-case  $u_t := \arg \max_{p \in \mathcal{P}_t} \mathbb{E}_{s \sim p} [f(x_t, s)] \in \mathcal{P}_t$  and recall the definition of  $x_t = \arg \min_{x \in \mathcal{X}} \mathbb{E}_{s \sim p_t} [f(x, s)]$  to obtain

$$\sum_{t=1}^T (\mathbb{E}_{s \sim u_t} [f(x_t, s)] - \mathbb{E}_{s \sim p_t} [f(x_t, s)]) = \sum_{t=1}^T \left( \max_{p \in \mathcal{P}_t} \mathbb{E}_{s \sim p} [f(x_t, s)] - \min_{x \in \mathcal{X}} \mathbb{E}_{s \sim p_t} [f(x, s)] \right).$$

Since  $p_t \in \mathcal{P}_{t-1}$ , we know that  $\min_{x \in \mathcal{X}} \mathbb{E}_{s \sim p_t} [f(x, s)] \leq \min_{x \in \mathcal{X}} \max_{p \in \mathcal{P}_{t-1}} \mathbb{E}_{s \sim p} [f(x, s)]$  for all  $t = 1, \dots, T$  and thus we can conclude

$$\sum_{t=1}^T \left( \max_{p \in \mathcal{P}_t} \mathbb{E}_{s \sim p} [f(x_t, s)] - \min_{x \in \mathcal{X}} \max_{p \in \mathcal{P}_{t-1}} \mathbb{E}_{s \sim p} [f(x, s)] \right) \leq G \sqrt{2|\mathcal{S}|h(T)T},$$

with a probability of at least  $1 - \delta$ . We add and subtract  $\min_{x \in \mathcal{X}} \max_{p \in \mathcal{P}_t} \mathbb{E}_{s \sim p} [f(x, s)]$  on the LHS. Rearranging the terms like this allows us to write the LHS as

$$\begin{aligned} & \sum_{t=1}^T \left( \max_{p \in \mathcal{P}_t} \mathbb{E}_{s \sim p} [f(x_t, s)] - \min_{x \in \mathcal{X}} \max_{p \in \mathcal{P}_t} \mathbb{E}_{s \sim p} [f(x, s)] \right) \\ & + \sum_{t=1}^T \min_{x \in \mathcal{X}} \max_{p \in \mathcal{P}_t} \mathbb{E}_{s \sim p} [f(x, s)] - \min_{x \in \mathcal{X}} \max_{p \in \mathcal{P}_{t-1}} \mathbb{E}_{s \sim p} [f(x, s)]. \end{aligned}$$

The last two terms telescope. Bringing them to the RHS and using the upper bound  $G$  on  $|f(x, s)|$  we can conclude

$$\sum_{t=1}^T \left( \max_{p \in \mathcal{P}_t} \mathbb{E}_{s \sim p} [f(x_t, s)] - \min_{x \in \mathcal{X}} \max_{p \in \mathcal{P}_t} \mathbb{E}_{s \sim p} [f(x, s)] \right) \leq G \sqrt{2|\mathcal{S}|h(T)T} + 2G \quad \text{w.p. } 1 - \delta.$$

Dividing by  $T$  on both sides completes the proof.  $\square$

A second regret bound is also proven in the electronic companion which allows for better dependence on the number of scenarios  $|\mathcal{S}|$  but at a cost of worse dependence on the total iteration count  $T$ . This is achieved by replacing the path length  $\sum_{t=1}^T \frac{1}{2} \|p_t - u_t\|^2 \leq h(T)$  by  $\sum_{t=2}^T \|p_t - q_t\| \leq h'(T)$ , which leads to different bounds. Numerically, it is better in the beginning i.e., for small  $t$  but has worse asymptotic behavior.

As stated in the following Corollary, we also observe that for all considered ambiguity sets the dynamic regret bound converges to zero.

**Corollary 3.2 (Convergence of Regret)** *If  $\lim_{T \rightarrow \infty} h(T)/T = 0$ , the dynamic regret converges to 0 with probability  $1 - \delta$  i.e.,*

$$\lim_{T \rightarrow \infty} \frac{1}{T} \sum_{t=1}^T \left( \max_{p \in \mathcal{P}_t} \mathbb{E}_{s \sim p} [f(x_t, s)] - \min_{x \in \mathcal{X}} \max_{p \in \mathcal{P}_t} \mathbb{E}_{s \sim p} [f(x, s)] \right) = 0.$$

*Proof of Corollary 3.2* We know that

$$\begin{aligned} & \limsup_{T \rightarrow \infty} \frac{1}{T} \sum_{t=1}^T \left( \max_{p \in \mathcal{P}_t} \mathbb{E}_{s \sim p} [f(x_t, s)] - \min_{x \in \mathcal{X}} \max_{p \in \mathcal{P}_t} \mathbb{E}_{s \sim p} [f(x, s)] \right) \\ & \leq \limsup_{T \rightarrow \infty} \left( G \sqrt{\frac{2h(T)|\mathcal{S}|}{T}} + \frac{2G}{T} \right) = 0, \end{aligned}$$

with probability at least  $1 - \delta$ . Thus, we can write,

$$\limsup_{T \rightarrow \infty} \frac{1}{T} \sum_{t=1}^T \left( \max_{p \in \mathcal{P}_t} \mathbb{E}_{s \sim p} [f(x_t, s)] - \min_{x \in \mathcal{X}} \max_{p \in \mathcal{P}_t} \mathbb{E}_{s \sim p} [f(x, s)] \right) \leq 0. \quad (4)$$

To prove the lower bound, let  $\bar{x}_t$  be the optimal solution to the problem  $\min_{x \in \mathcal{X}} \max_{p \in \mathcal{P}_t} \mathbb{E}_{s \sim p} [f(x, s)]$ . Then we can write the inner term in the left hand side (LHS) in equation (4) as

$$\max_{p \in \mathcal{P}_t} \mathbb{E}_{s \sim p} [f(x_t, s)] - \max_{p \in \mathcal{P}_t} \mathbb{E}_{s \sim p} [f(\bar{x}_t, s)].$$

We know that  $\bar{x}_t$  is the optimal solution to the problem  $\min_{x \in \mathcal{X}} \max_{p \in \mathcal{P}_t} \mathbb{E}_{s \sim p} [f(x, s)]$ , this means that

$$\begin{aligned} \max_{p \in \mathcal{P}_t} \mathbb{E}_{s \sim p} [f(x_t, s)] &\geq \min_{x \in \mathcal{X}} \max_{p \in \mathcal{P}_t} \mathbb{E}_{s \sim p} [f(x, s)] \\ &= \max_{p \in \mathcal{P}_t} \mathbb{E}_{s \sim p} [f(\bar{x}_t, s)]. \end{aligned}$$

Thus we get

$$\max_{p \in \mathcal{P}_t} \mathbb{E}_{s \sim p} [f(x_t, s)] - \max_{p \in \mathcal{P}_t} \mathbb{E}_{s \sim p} [f(\bar{x}_t, s)] \geq 0. \quad (5)$$

The above lower bound and the limsup-bound (4) together prove the result.  $\square$

From the above result and Theorem 3.1, we can observe that the dynamic regret, i.e., the average gap between the best solution in hindsight of each round and the solution evaluated in the algorithm, decreases at a rate of  $\mathcal{O}(\sqrt{h(T)}/\sqrt{T})$  and tends to zero. Therefore, we have a performance guarantee (with sublinear regret) when solving the DRO problems approximately over time. At the same time, Algorithm 2 is applicable to large-sized problems in contrast to the reformulated DRO problems (2) or (3).

### 3.2 Bounded Path Lengths

The following lemma illustrates the high-probability path lengths for the confidence intervals, the kernel-based and  $l_2$ -norm ambiguity sets.

**Lemma 3.3** *Given ambiguity sets of the form specified in Section 2, we have*

$$\frac{1}{2} \sum_{t=1}^T \|p_t - q_t\|^2 \leq h(T),$$

for all  $p_t \in \mathcal{P}_{t-1}, q_t \in \mathcal{P}_t$  with probability at least  $1 - \delta$ . The functions  $h(T)$  for different categories of ambiguity sets are as given:

1. **Confidence Intervals:**

$$h(T) = 8|\mathcal{S}| \log(\pi T)(2 + \log T).$$

2. **Kernel based ambiguity sets:**

$$h(T) = \frac{1}{2} \left( 2 + \frac{4\sqrt{C}}{\lambda} \right)^2 + \frac{32C}{\lambda^2} \log \frac{\pi T}{\sqrt{6\delta}} (1 + \log T),$$

where  $\lambda$  denotes the smallest eigenvalue of the kernel matrix  $M$ .

3.  **$l_2$ -norm ambiguity sets:**

$$h(T) = 8|\mathcal{S}| \log \frac{\pi T}{\sqrt{3\delta}} (2 + \log T).$$

*Proof of Lemma 3.3.*

**Confidence Intervals.** We show in the electronic companion (Lemma 6.6) that the ambiguity sets  $\mathcal{P}_t$  derived from (1) with confidence update  $\delta_t := \frac{6\delta}{\pi^2 t^2}$  and  $\delta \in (0, 1)$  for all rounds  $t = 1, \dots, T$  fulfill

$$\sup_{x \in \mathcal{P}_0, y \in \mathcal{P}_1} \|x - y\| \leq \sqrt{16|\mathcal{S}| \log \pi} \quad \text{and} \quad \sup_{x \in \mathcal{P}_{t-1}, y \in \mathcal{P}_t} \|x - y\| \leq \frac{\sqrt{16|\mathcal{S}| \log(\pi(t-1))}}{\sqrt{t-1}},$$

with a probability of at least  $1 - \delta$ . This allows for calculating the function  $h(T)$ :

$$\begin{aligned} \frac{1}{2} \sum_{t=1}^T \|p_t - q_t\|^2 &\leq \frac{1}{2} 16|\mathcal{S}| \log \pi + \frac{1}{2} \sum_{t=2}^T 16|\mathcal{S}| \frac{\log(\pi(t-1))}{t-1} \\ &\leq 8|\mathcal{S}| \log \pi + 8|\mathcal{S}| \log(\pi(T-1)) \sum_{t=1}^{T-1} \frac{1}{t} \\ &\leq 8|\mathcal{S}| \log \pi + 8|\mathcal{S}| \log(\pi(T-1))(1 + \log(T-1)) \\ &\leq 8|\mathcal{S}| \log(\pi T)(2 + \log T). \end{aligned}$$

The second and third inequalities are from bounding  $t$  and from observing that  $\sum_{t=1}^{T-1} (1/t) \leq 1 + \log(T-1)$ .

**Kernel based ambiguity sets.** We show in the electronic companion (Lemma 6.9) that given an ambiguity set of the form  $\mathcal{P}_t = \{p \in \mathcal{P} \mid \|p - \hat{p}\|_M \leq \epsilon_t\}$  with  $\epsilon_t := \frac{\sqrt{C}}{\sqrt{t}}(2 + \sqrt{2 \log(1/\delta_t)})$  with  $\delta_t = \frac{6\delta}{\pi^2 t^2}$  we have for  $t \geq 2$ ,

$$\begin{aligned} \sup_{x \in \mathcal{P}_{t-1}, y \in \mathcal{P}_t} \|x - y\|_2 &\leq \frac{8\sqrt{C}}{\lambda\sqrt{t-1}} \sqrt{\log(\pi t / \sqrt{6\delta})}, \\ \text{and } \sup_{x \in \mathcal{P}_0, y \in \mathcal{P}_1} \|x - y\|_2 &\leq 2 + \frac{4\sqrt{C}}{\lambda} \quad \text{for } t = 1, \end{aligned}$$

with probability at least  $1 - \delta$ . Calculating the function  $h(T)$ , we have

$$\begin{aligned} \frac{1}{2} \sum_{t=1}^T \|p_t - q_t\|_2^2 &\leq \frac{1}{2} \left(2 + \frac{4\sqrt{C}}{\lambda}\right)^2 + \sum_{t=2}^T \frac{32C}{\lambda^2(t-1)} \log(\pi t / \sqrt{6\delta}) \\ &\leq \frac{1}{2} \left(2 + \frac{4\sqrt{C}}{\lambda}\right)^2 + \frac{32C}{\lambda^2} \log(\pi T / \sqrt{6\delta}) \sum_{t=2}^T \frac{1}{t-1} \\ &\leq \frac{1}{2} \left(2 + \frac{4\sqrt{C}}{\lambda}\right)^2 + \frac{32C}{\lambda^2} \log(\pi T / \sqrt{6\delta})(1 + \log T). \end{aligned}$$

Here, the first inequality arises from Lemma 6.9. The second and third inequalities are from bounding  $t$  and from observing that  $\sum_{t=1}^{T-1} (1/t) \leq 1 + \log(T-1)$ .

**$\ell_2$ -norm ambiguity sets.** We show in the electronic companion (Lemma 6.8) that given an ambiguity set of the form  $\mathcal{P}_t = \{p \in \mathcal{P} \mid \|p - \hat{p}\|_2 \leq \epsilon_t\}$  with  $\epsilon_t := \sqrt{\frac{2|\mathcal{S}| \log(2/\delta_t)}{t}}$  and  $\delta_t = \frac{6\delta}{\pi^2 t^2}$ , we have

$$\sup_{x \in \mathcal{P}_0, y \in \mathcal{P}_1} \|x - y\|_2 \leq 4\sqrt{|\mathcal{S}| \log(\pi / \sqrt{3\delta})} \quad \text{and} \quad \sup_{x \in \mathcal{P}_{t-1}, y \in \mathcal{P}_t} \|x - y\|_2 \leq 4\sqrt{\frac{|\mathcal{S}| \log(\pi(t-1) / \sqrt{3\delta})}{t-1}},$$

with probability at least  $1 - \delta$ . For bound  $h(T)$ , it follows

$$\begin{aligned} \frac{1}{2} \sum_{t=1}^T \|p_t - q_t\|^2 &\leq \frac{1}{2} (4\sqrt{|\mathcal{S}| \log(\pi/\sqrt{3\delta})})^2 + \frac{1}{2} \sum_{t=2}^T 16|\mathcal{S}| \frac{\log(\pi(t-1)/\sqrt{3\delta})}{t-1} \\ &\leq 8|\mathcal{S}| \log(\pi/\sqrt{3\delta}) + 8|\mathcal{S}| \log(\pi T/\sqrt{3\delta}) \sum_{t=2}^T \frac{1}{t-1} \\ &\leq 8|\mathcal{S}| \log(\pi T/\sqrt{3\delta}) (2 + \log T). \end{aligned}$$

□

From Lemma 3.3, we know that the shrinking ambiguity sets yield path length bounds  $h(T)$  that increase at a rate of  $\mathcal{O}(\log^2 T)$ . Therefore, the dynamic regret bound converges to zero with  $\mathcal{O}(\log T/\sqrt{T})$ . It is only slightly worse than the typical  $\mathcal{O}(1/\sqrt{T})$  bounds. In summary, this novel algorithm yields the remarkable benefit that it integrates changing ambiguity sets (which represent growing knowledge of the uncertainty) into regret bounds for online robust optimization.

## 4 Numerical Results

We illustrate the performance of Algorithm 2 through numerical experiments on mixed-integer linear and quadratic programs (MIPs & MIQPs) as well as on two real-world applications, namely distributionally robust network design (Ghayoori and Leon-Garcia 2013) and optimal route choice. These problems allow to demonstrate the wide applicability of our novel approach. The latter does not assume anything about the structure of the original problem apart from the fact that it should be algorithmically tractable, with an available solution approach. We are thus able to apply DRO over time to both discrete and continuous optimization problems to account for various application types.

All computations are carried out using a Python 3.8.5 implementation on machines with Intel Core i7 CPU 2.80 GHz processor and 16 GB RAM. Each of these has four cores of 3.5 GHz each and 32 GB RAM. We utilized SCIP 7.0 as the MIP and MIQP solver (Gamrath et al. 2020) and IPOPT 3.13.3 (Wächter and Biegler 2006) to compute the projections in Algorithm 2 as the solution of a convex optimization problem. For the ambiguity sets, we choose  $1 - \delta = 0.9$ . Different values for this parameter impact the size of the resulting sets, but have no significant impact on the ability of Algorithm 2 to learn DRO solutions. For the kernel-based ambiguity sets, we use the Gaussian kernel function  $k_M(s_i, s_j) = \exp\left(-\frac{\|s_i - s_j\|_2^2}{2}\right)$ .

### 4.1 Benchmark Instances

To validate the performance of the method, we use publicly available instances from well known benchmark libraries MIPLIB (Gleixner et al. 2021) and QPLIB (Furini et al. 2019), which contain a collection of MIPs and MIQPs, respectively. In all test cases, the uncertain objective function  $f$  has the form  $f(x, s) = x^\top Qx + (c + s)^\top x + d$ , with  $Q \in \mathbb{R}^{n \times n}$ ,  $c \in \mathbb{R}^n$ ,  $d \in \mathbb{R}$ ,  $x \in \mathcal{X}$  and  $s \in \mathcal{S}$ . For linear problems, we have  $Q = 0$ . We generate different cost scenarios  $s \in \mathcal{S} \subset \mathbb{R}^n$ ,  $n > 0$  by perturbing the coefficients of variables in the objective function randomly by up to 50%.

We sort the problems in the libraries by increasing number of variables and choose the first 15 MIPs and the first 10 convex MIQPs with bounded variables that can be solved within an hour of CPU time. All instances are listed in the electronic companion. The number of scenarios in the linear test cases is set to  $|\mathcal{S}| \in \{10, 50\}$ . As the MIQP instances are more difficult to solve, we set  $|\mathcal{S}| = 2$ .

Our theoretical result in Theorem 3.1 bounds the average gap between the best solution in hindsight of each round and our online solution. Therefore, we numerically evaluate the dynamic

regret from Theorem 3.1 and the worst-case expected value of solutions from Algorithm 2 in comparison with the DRO solution obtained via reformulation (2). In each round, one realization of the uncertain parameter is revealed. These realizations are drawn randomly according to the initially generated probability distribution. Thus, the number of data points equals the number of rounds.

Figure 1(a) shows the worst-case (or guaranteed) expected objective value in each round for the online algorithm for  $T = 10000$  rounds for instance *blend2* from MIPLIB with  $|\mathcal{S}| = 10$ . Therein, the results for all three types of ambiguity sets (confidence intervals,  $\ell_2$ -norm, kernel-based) are shown under the assumed ambiguity set in each period. The lines represent the corresponding average values over time. One can observe that the confidence intervals (red) yield the fastest convergence and lead to solutions with least costs. In Figure 1(b) and Figure 1(c), we plot the error between the worst-case protection of the online solutions and the DRO solutions or to the SO solutions, respectively. It can be observed that the distance between the lines decreases over rounds. This means that the average error in solving the DRO problem with Algorithm 2 shrinks and we learn the robust solution rapidly. Both solutions also converge to the true stochastic solution. Since, the confidence interval ambiguity sets yield a better performance and shrink faster, we focus on interval sets for the remaining experiments.

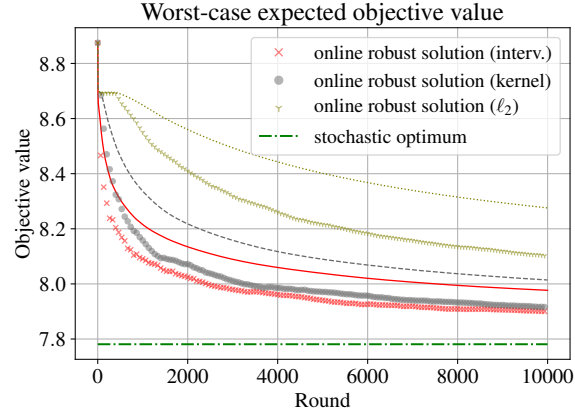
Next, we evaluate the performance of Algorithm 2. Table 1 shows the average running times per iteration over all instances for different scenarios and ambiguity sets. It is obvious from Table 1 that for large and difficult problems (like e.g. nonlinear mixed-integer optimization problems), using Algorithm 2 allows for significant time savings. This becomes more and more pronounced for increasing number of iterations, as the time savings multiply by the number of rounds.

	$ \mathcal{S} $	ONLINE ROBUST	EXACT DRO
MIP (I)	10	52.4s	115.8s
MIP ( $\ell_2$ )	10	49.4s	127.5s
MIP (K)	10	56.3s	129.5s
MIP (I)	50	57.7s	176.7s*
MIP ( $\ell_2$ )	50	60.4s	206.1s*
MIP (K)	50	67.0s	244.4s*
MIQP (I)	2	170.2s	271.4s*
MIQP ( $\ell_2$ )	2	186.3s	329.5s*
MIQP (K)	2	188.6s	359.6s*

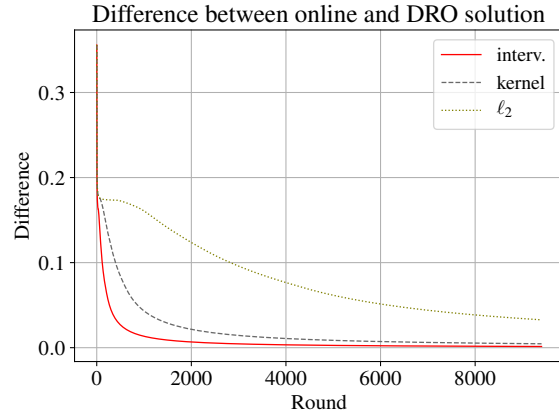
Table 1: Average running times per iteration. I:Interval, K:Kernel,  $\ell_2$ : $\ell_2$ -norm. (\*) The DRO problems *k16x240b* and *10004* could not be solved within a runtime limit of one hour.

In more detail, in 14 out of 15 MIP instances and in 9 out of 10 MIQP ones, Algorithm 2 was able to run an iteration on average significantly faster than solving reformulation (2). The impact of different ambiguity sets on the solution times is negligible for our approach. However, increasing the number of scenarios amplifies the size of the reformulated DRO problem and thus results in challenging problems for which the online algorithm is considerably more efficient. The main advantage of Algorithm 2 is that we avoid solving the full  $\text{DRO}_t$  problem in each round, which can be algorithmically challenging. Instead, only the solution of an SO and a convex projection problem is calculated in the learning decomposition approach. The computational results show that our approach is able to generate high-quality solutions within a short running time.

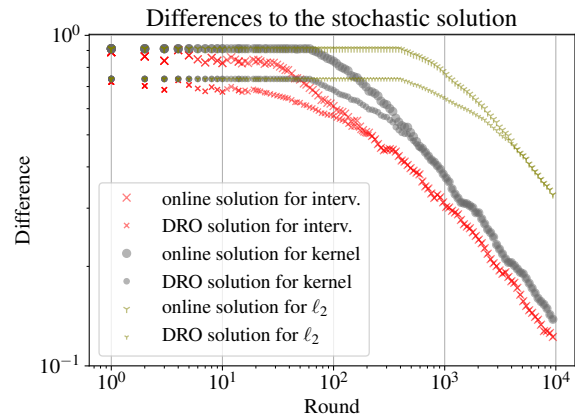




(a) Evolution of worst-case expected objective value for all three types of ambiguity sets in each round. In addition to the actual values, the worst-case average is given as a line.



(b) Convergence of online average to DRO average (dynamic regret).



(c) Convergence of online robust and DRO solutions to the stochastic solution.

Figure 1: Results for *blend2* with  $|\mathcal{S}| = 10$  and  $T = 10000$ .

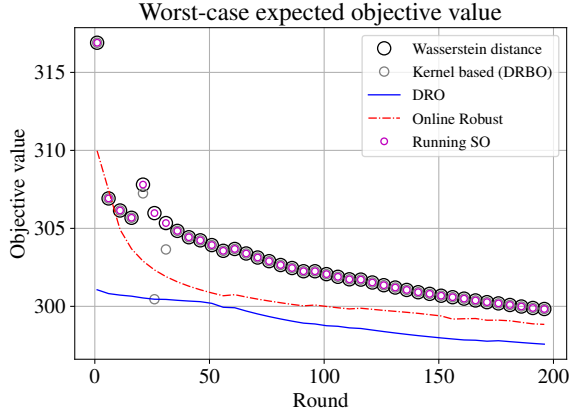


Figure 2: Results for *supportcase16* ( $|\mathcal{S}| = 20, T = 100$ ).

#### 4.1.1 Comparison with Other Methods

In this section, we compare the performance of our novel algorithms against other approaches that can solve DRO problems. Specifically, we focus on the ambiguity sets outlined by Esfahani and Kuhn (2018) (Wasserstein) and Kirschner et al. (2020) (Distributionally Robust Bayesian Optimization or DRBO), which leverage Wasserstein and Kernel based ambiguity sets respectively. We implement the Wasserstein and DRBO approaches as outlined in the respective papers with an exact computation of the corresponding distributionally-robust optimization problem. In the numerical experiments, we observe that the interval ambiguity sets yield a solution with a comparable worst-case expected objective as the solutions generated by the Wasserstein and DRBO methods. We also compare our results with the worst-case performance of the stochastic optimum using the Maximum Likelihood Estimator (MLE) in each round (*running SO*). This solution convergences to the true stochastic optimum in the limit ( $T \rightarrow \infty$ ) but is not protected against ambiguity and has no worst-case guarantees in contrast to distributionally robust solutions.

	$ \mathcal{S}  = 10$	$ \mathcal{S}  = 50$
DRO	2.2s	3.5s
Wasserstein	0.4s	1.1s
DRBO	7.9s	30.9s
Online robust	0.3s	0.4s
Running SO	0.2s	0.2s

Table 2: Avg. running times for *supportcase16*,  $T = 100$ .

As a representative example, Figure 2 illustrates these results for the *supportcase16* instance from MIPLIB with  $|\mathcal{S}| = 20$  and  $T = 100$ . For all methods, the worst-case expected objective value shrinks over time with more data. The online and exact solutions yield similar worst-case protection as the other methods. However, the online robust solution can be calculated more efficiently than the other DRO methods, cf. Table 2. Only the stochastic solutions with the MLE can be computed faster. However they do not have solution quality guarantees under ambiguity and may lead to a bad worst-case objective. Thus in total, the online robust method is preferable. The average running times of all instances are given in Table 3 and support the observations. Indeed, using the online robust approach, the instances all can be solved quickly within the time limit, whereas it takes considerably longer for the other approaches that may

already reach the time limit for some instances.

	MIP $ \mathcal{S}  = 10$	MIP $ \mathcal{S}  = 50$	MIQP $ \mathcal{S}  = 2$
DRO	45.6s	55.9s*	271.4s*
Wasserstein	52.3s	59.1s	299.9s
DRBO	42.7s**	66.1s**	738.3s*
Online robust	26.8s	27.1s	170.2s
Running SO	26.6s	26.9s	172.6s

Table 3: Avg. running times for benchmark instances. (\*) One testcase could not be solved within one hour. (\*\*) Four testcases could not be solved within one hour.

## 4.2 Network Design under Uncertainty

In addition to solving classical benchmark instances as performed in the last section, we next study the solution of a practically very relevant and highly challenging combinatorial optimization problem under uncertainty. Given an undirected graph  $\mathcal{G} = (V, E)$ , the goal of robust network design (Cacchiani et al. 2016) is to compute a minimal cost network topology together with the corresponding edge capacities  $f \in \mathbb{Z}_+^{|E|}$  in order to fulfill a given demand  $b \in \mathbb{Z}^{|V|}$ . The demand is assumed to be uncertain and an element of the scenario set  $\mathcal{S} := \{b_1, \dots, b_{|\mathcal{S}|}\} \subset \mathbb{Z}^{|V|}$  with (unknown) probability vector  $p^* \in [0, 1]^{|\mathcal{S}|}$ . For every edge  $\{i, j\} \in E$  and every scenario  $s \in \mathcal{S}$ , we are given costs  $c_{ijs} > 0$  and flow capacities  $d_{ijs}, \bar{d}_{ij} > 0$ . In this practical application, it can be assumed that additional information on the demand distributions become available over time, so that a DRO over time approach is a very natural modelling choice.

The optimization problem for distributionally robust network design is then given by

$$\begin{aligned}
& \min_f \max_{p \in \mathcal{P}} \sum_{s \in \mathcal{S}} \sum_{\{i,j\} \in E} c_{ijs} (f_{ijs} + f_{jis}) p_s \\
& \text{s.t.} \quad \sum_{j:\{j,i\} \in E} f_{jis} - \sum_{j:\{i,j\} \in E} f_{ijs} = b_{is} \quad \forall i \in V, s \in \mathcal{S}, \\
& \quad \sum_{s \in \mathcal{S}} (f_{ijs} + f_{jis}) \leq \bar{d}_{ij} \quad \forall \{i, j\} \in E, \\
& \quad f_{ijs} + f_{jis} \leq d_{ijs} \quad \forall \{i, j\} \in E, s \in \mathcal{S}, \\
& \quad f_{ijs} \in \mathbb{Z}_+ \quad \forall \{i, j\} \in E, s \in \mathcal{S}.
\end{aligned}$$

Here, the objective minimizes the flow costs. The first constraint ensures that demands are satisfied. The second and third constraints limit the flows to be within arc capacities and the final constraint ensures integral flows.

We evaluate the novel approach on the instances *res8* ( $|V| = 50, |E| = 77$ ), *w1\_100* ( $|V| = 100, |E| = 207$ ) and *w1\_200* ( $|V| = 200, |E| = 775$ ) from Altın et al. (2007) and construct the scenarios as follows: On half the nodes, we place balanced random demands from  $\{-10, \dots, 10\}$ . We also restrict all flow capacities to  $d_{ijs} = 10$ , where edge costs and coupling capacity bounds are uniformly chosen from  $c_{ijs} \in \{1, \dots, 10\}$  and  $\bar{d}_{ij} \in \{|\mathcal{S}|, \dots, 10|\mathcal{S}|\}$ .

In Figure 3, the worst-case expectation of solutions for instance *res8* with  $|\mathcal{S}| = 10$  using confidence intervals is illustrated over 10000 rounds on logarithmic axes. The online robust solution rapidly converges to the DRO solution. Though starting with conservative outcomes due to limited data, the online solutions improve very quickly. The running time benefit of the online robust approach is clearly visible in Table 4. For larger instances, it is about 272

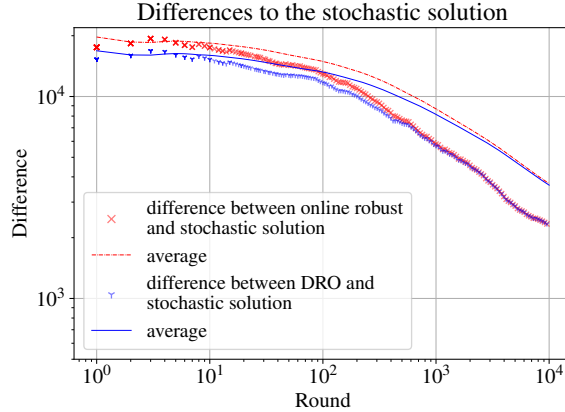


Figure 3: Results for *res8* with  $|\mathcal{S}| = 10$  and  $T = 10000$ .

times faster than solving reformulation (2). In summary, our method is able to generate robust solutions with shorter running time.

	$ \mathcal{S} $	ONLINE ROBUST	EXACT DRO
RES8	10	0.2 s	0.5 s
RES8	50	0.6 s	11.6 s
w1_100	10	0.3 s	32.0 s
w1_100	50	1.5 s	95.6 s
w1_200	10	1.2 s	38.7 s
w1_200	50	4.7 s	1282.2 s

Table 4: Average running times per iteration using confidence intervals.

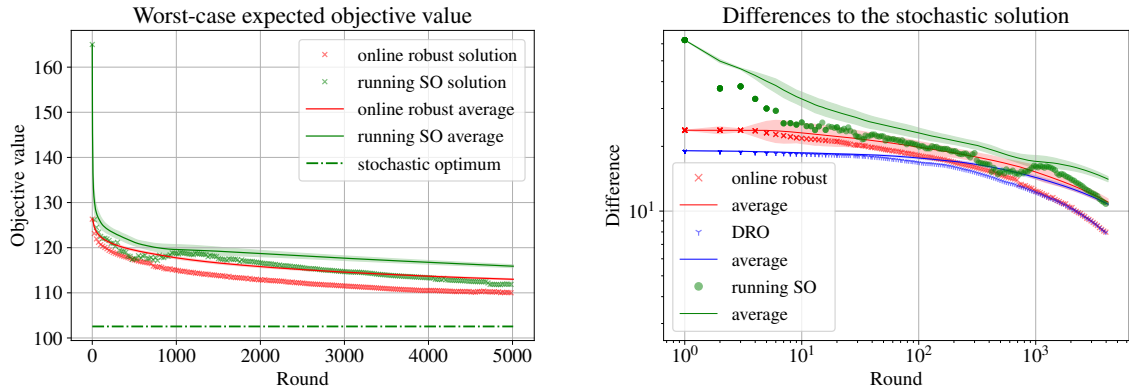
#### 4.2.1 Learning Optimal Route Choice

In this section, we consider the problem of choosing the shortest paths in a street network where the travel times on the arcs are affected by random deviations. It is natural to assume that the driver gradually adapts the route according to the observed travel times in order to reach the destination in the shortest possible (expected) time, making a DRO over time approach a good modeling choice.

This means that the driver solves the shortest-path problem on a directed graph  $G = (V, A)$  with uncertain travel times  $c: A \rightarrow \mathbb{R}_+$ . Let  $v_1, v_2 \in V$  be the origin and the destination, respectively. We assume that there is a finite set of traffic scenarios  $\mathcal{S} = \{c_1, c_2, \dots, c_{|\mathcal{S}|}\} \subset \mathbb{R}_+^{|A|}$  which correspond to different realizations of the travel times on the arcs, each materializing with an unknown probability  $p_k^* \in [0, 1]$ ,  $k = 1, \dots, |\mathcal{S}|$ .

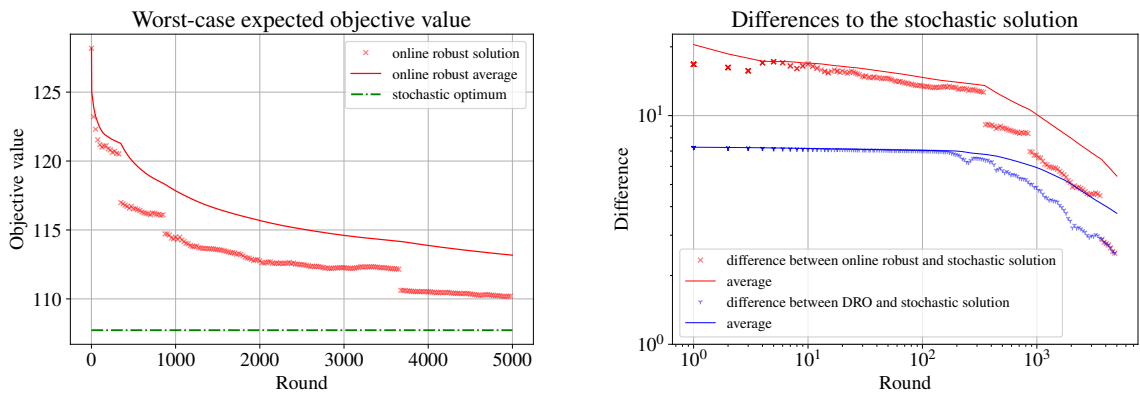
In each round  $t = 1, \dots, T$ , (e.g. every morning when driving to work), the driver chooses a  $v_1$ - $v_2$ -route given by the vector  $x_t \in \{0, 1\}^{|A|}$ , which models the edges traveled, along the path chosen, in that round. The expected travel time in a round  $t \in T$  is then given by  $\sum_{s \in \mathcal{S}} p_s^*(c_s, x_t)$ . As the true scenario distribution is unknown, the driver is assumed to solve the distributionally robust shortest-path problem in an online fashion, i.e. using Algorithm 2.

In the following, we analyse the outcome of this experiment on an aggregated version of the real-world city network of Chicago. It is available as instance *ChicacoSketch* in Ben Stabler’s library of transportation networks (Stabler et al. 2018) and has 933 nodes and 2950 arcs (of which we ignore the 387 nodes representing “zones” as well as their incident arcs). In this data set, each arc  $a$  has a certain free-flow time  $c_{\text{free},a}$ , which we assume to be the uncongested



(a) Evolution of worst-case expected objective for on- (b) Convergence of online robust, DRO solutions and line robust and stochastic solutions with MLE (running running SO solutions averaged over ten examples. SO) averaged over ten examples.

Figure 4: The outcome of Algorithm 2 for learning an optimal route choice in terms of solution quality over time for  $T = 5000$ .



(a) Evolution of worst-case expected objective for a (b) Convergence of online robust and DRO solutions for single example.

Figure 5: The outcome of Algorithm 2 for single runs that correspond to the routes given in Figure 6.

travel time. In addition, we generate nine congestion scenarios by perturbing  $c_{\text{free},a}$ . We first choose  $v_1$  and  $v_2$  such that the driven path would span the entire extract of the city map. Now, for all arcs  $a$  we uniformly draw  $c_{s,a} \sim [0, 2c_{\text{free},a}]$ . Finally, we uniformly draw a random “true” probability distribution  $p^* \sim \mathcal{P}$ .

For the above setup, we use Algorithm 2 in order to let the driver iteratively adapt to the dynamically changing travel times. In Figure 4, we illustrate solution quality over time. In order to show the stability of our algorithm, we additionally repeated the experiments ten times and plot their mean solution quality as well as their standard deviation. We observe that the online average and DRO average jointly converge towards the expected value of the minimum expected travel time. The regret tends to zero over the long run. The non-monotonic behavior of the solution determined using the maximum likelihood estimator (running SO) is explained by the fact that the solutions have no protection against uncertainty. Therefore, the worst-case performance of these solutions fluctuates.

As an illustrative example, in Figures 5(a) and 5(b), one specific run is plotted to evaluate it in more detail. In these plots, one can see that in each of the rounds 1, 349, 860 and 3661, long stretches of the chosen path are abruptly improved due to the discrete nature of the decisions. There are also visible jumps in the online robust and exact DRO solution quality.

In Figure 6, we depict how route choice evolve over the rounds  $t \in T$ ; introducing a new picture whenever a structurally new solution is found. At the beginning, the driver takes the nominally optimal path, i.e. , the optimal path w.r.t. the unperturbed cost vector. We see two subtours in which the path differs from the true stochastic solution. In the following iterations, there remain two deviating subtours, but they change slightly in round 1 and round 349. In round 1879, one of these subtours disappears, and after round 3661 the solution coincides with the true stochastic optimum.

Altogether, this example shows how DRO can be used to improve performance in the face of uncertainty by leveraging information arriving over time. At the beginning, with no or little information available, hedging against uncertainty necessarily means to implement conservative solutions. However, as more information on the uncertainty is gathered over time, the solution quality improves as protection against uncertainty is less costly.

## 5 Conclusion

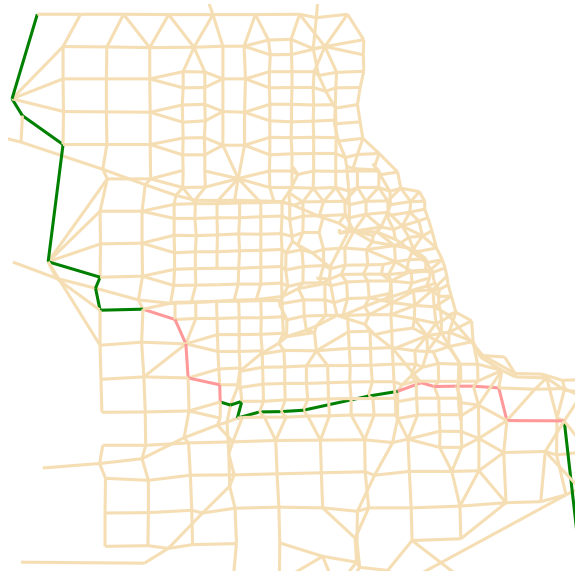
We introduce a novel method for decision-making under uncertainty over time, employing a combination of distributionally robust optimization and online learning. In each iteration, our algorithm solves a stochastic optimization problem in combination with the online gradient descent algorithm. We show that our online algorithm converges to the exact solution of the DRO problem with an increasing amount of iterations. We also show that DRO solution converges to the true SO optimum in the limit. The theoretical and numerical results demonstrate the effectiveness of this method. Indeed, it obtains high-quality robust solutions with short computational times. Though, our work is tailored to discrete scenarios of finite dimension and requires solving the outer minimization problem exactly, it can be applied to a wide range of practical problems with varying uncertainty models and objective function structures. Furthermore, our flexible framework can be extended, either by incorporating more stochastically expressive ambiguity sets (e.g. for continuous distributions) or using online methods also for the decision problem of the  $x$ -player.

## 6 Acknowledgments

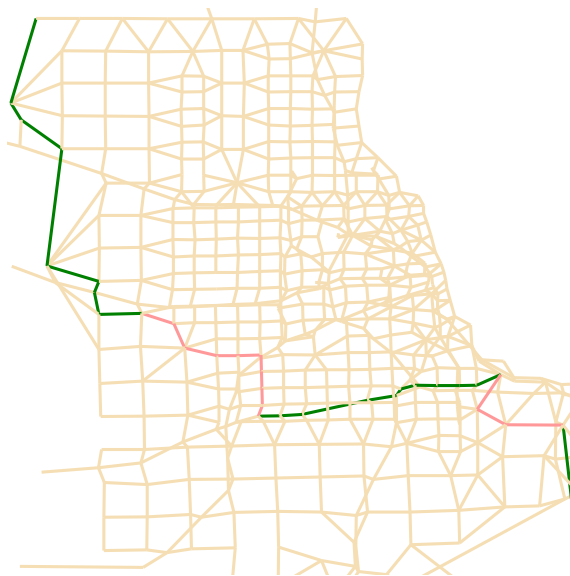
The authors thank the DFG for their support within Projects B06 and B10 in CRC TRR 154, as well as within Project-ID 416229255 - SFB 1411. This work has been supported by grant 03El1036A from the Federal Ministry for Economic Affairs and Energy, Germany. Furthermore,



(a) True stochastic solution



(b) Round 0



(c) Round 1





we acknowledge financial support by the Bavarian Ministry of Economic Affairs, Regional Development and Energy through the Center for Analytics – Data – Applications (ADA-Center) within the framework of “BAYERN DIGITAL II”.

## References

- Altın, A., Amaldi, E., Belotti, P., and Pınar, M. Ç. (2007). Provisioning virtual private networks under traffic uncertainty. *Networks*, 49(1):100–115.
- Bärmann, A., Martin, A., Pokutta, S., and Schneider, O. (2018). An online-learning approach to inverse optimization. *arXiv preprint arXiv:1810.12997*.
- Bärmann, A., Pokutta, S., and Schneider, O. (2017). Emulating the expert: Inverse optimization through online learning. In *ICML*, pages 400–410. PMLR.
- Bayraksan, G. and Love, D. K. (2015). Data-driven stochastic programming using phi-divergences. In *The Operations Research Revolution*, pages 1–19. INFORMS.
- Ben-Tal, A., Ghaoui, L. E., and Nemirovski, A. (2009). *Robust Optimization*. Princeton.
- Ben-Tal, A., Hazan, E., Koren, T., and Mannor, S. (2015). Oracle-based robust optimization via online learning. *Oper. Res.*, 63(3):628–638.
- Bertsimas, D., Gupta, V., and Kallus, N. (2018). Data-driven robust optimization. *Math. Prog.*, 167(2):235–292.
- Besbes, O., Gur, Y., and Zeevi, A. (2015). Non-stationary stochastic optimization. *Oper. Res.*, 63(5):1227–1244.
- Birge, J. R. and Louveaux, F. (1997). *Introduction to Stochastic Programming*. Springer Series in Operations Research. Springer.
- Cacchiani, V., Jünger, M., Liers, F., Lodi, A., and Schmidt, D. (2016). Single-commodity robust network design with finite and hose demand sets. *Math. Prog.*, 157(1):297 – 342.
- Calafiore, G. C. and Ghaoui, L. E. (2006). On distributionally robust chance-constrained linear programs. *J. Optim. Th. App.*, 130(1):1–22.
- Chen, R. S., Lucier, B., Singer, Y., and Syrgkanis, V. (2017). Robust optimization for non-convex objectives. *Adv. NeurIPS*, 30:4708–4717.
- Chen, Z., Sim, M., and Xiong, P. (2020). Robust stochastic optimization made easy with rsome. *Management Sci.*, 66(8):3329–3339.
- Delage, E. and Ye, Y. (2010). Distributionally robust optimization under moment uncertainty with application to data-driven problems. *Oper. Res.*, 58(3):595–612.
- Duchi, J. C., Glynn, P. W., and Namkoong, H. (2021). Statistics of robust optimization: A generalized empirical likelihood approach. *Math. Oper. Res.*, 46(3):946–969.
- Esfahani, P. M. and Kuhn, D. (2018). Data-driven distributionally robust optimization using the wasserstein metric: Performance guarantees and tractable reformulations. *Math. Prog.*, 171(1):115–166.
- Fitzpatrick, S. and Scott, A. (1987). Quick simultaneous confidence intervals for multinomial proportions. *J. Am. Stat. Ass.*, 82(399):875–878.
- Furini, F., Traversi, E., Belotti, P., Frangioni, A., Gleixner, A., Gould, N., Liberti, L., Lodi, A., Misener, R., Mittelmann, H., Sahinidis, N. V., Vigerske, S., and Wiegele, A. (2019). Qplib: A library of quadratic programming instances. *Math. Prog. Comp.*, 11(2):237 – 265.
- Gamrath, G., Anderson, D., Bestuzheva, K., Chen, W.-K., Eifler, L., Gasse, M., Gemander, P., Gleixner, A., Gottwald, L., Halbig, K., Hendel, G., Hojny, C., Koch, T., Le Bodic, P., Maher, S. J., Matter, F., Miltenberger, M., Mühmer, E., Müller, B., Pfetsch, M. E., Schlösser, F., Serrano, F., Shinano, Y., Tawfik, C., Vigerske, S., Wegscheider, F., Weninger, D., and Witzig, J. (2020). The SCIP Optimization Suite 7.0. ZIB-Report 20-10, Zuse Institute Berlin.
- Ghayoori, A. and Leon-Garcia, A. (2013). Robust network design. In *2013 IEEE Inter. Conf. Comm. (ICC)*, pages 2409–2414. IEEE.
- Gleixner, A., Hendel, G., Gamrath, G., Achterberg, T., Bastubbe, M., Berthold, T., Christophel, P., Jarck, K., Koch, T., Linderoth, J., et al. (2021). Miplib 2017: data-driven compilation of the 6th mixed-integer programming library. *Math. Prog. Comp.*, 13(3):443–490.

- Hanasusanto, G. A., Roitch, V., Kuhn, D., and Wiesemann, W. (2015). A distributionally robust perspective on uncertainty quantification and chance constrained programming. *Math. Prog.*, 151(1):35–62.
- Hazan, E. et al. (2016). Introduction to online convex optimization. *Found. and Trends Optim.*, 2(3-4):157–325.
- Heidt, A., Helmke, H., Kapolke, M., Liers, F., and Martin, A. (2016). Robust runway scheduling under uncertain conditions. *Journal of Air Transport Management*, 56(PA):28–37.
- Ho-Nguyen, N. and Kilinç-Karzan, F. (2018). Online first-order framework for robust convex optimization. *Oper. Res.*, 66(6):1670–1692.
- Husain, H., Nguyen, V., and Hengel, A. v. d. (2022). Distributionally robust bayesian optimization with  $\phi$ -divergences. *arXiv preprint arXiv:2203.02128*.
- Kapolke, M., Fürstenau, N., Heidt, A., Liers, F., Mittendorf, M., and Weiß, C. (2016). Pre-tactical optimization of runway utilization under uncertainty. *Journal of Air Transport Management*, 56(PA):48–56.
- Kirschner, J., Bogunovic, I., Jegelka, S., and Krause, A. (2020). Distributionally robust bayesian optimization. In *Inter. Conf. on Artif. Intel. and Stat.*, pages 2174–2184. PMLR.
- Kleywegt, A. J., Shapiro, A., and Homem-de Mello, T. (2002). The sample average approximation method for stochastic discrete optimization. *SIAM J. Optim.*, 12(2):479–502.
- Levy, D., Carmon, Y., Duchi, J. C., and Sidford, A. (2020). Large-scale methods for distributionally robust optimization. *Adv. NeurIPS*, 33:8847–8860.
- Mohajerin Esfahani, P. and Kuhn, D. (2018). Data-driven distributionally robust optimization using the wasserstein metric: Performance guarantees and tractable reformulations. *Math. Prog.*, 171(1):115–166.
- Mokhtari, A., Shahrapour, S., Jadbabaie, A., and Ribeiro, A. (2016). Online optimization in dynamic environments: Improved regret rates for strongly convex problems. In *2016 IEEE 55th Conference on Decision and Control (CDC)*, pages 7195–7201.
- Namkoong, H. and Duchi, J. C. (2016). Stochastic gradient methods for distributionally robust optimization with f-divergences. *Adv. NeurIPS*, 29:2208–2216.
- Pokutta, S. and Xu, H. (2021). Adversaries in online learning revisited: with applications in robust optimization and adversarial training. *arXiv preprint arXiv:2101.11443*.
- Qi, Q., Guo, Z., Xu, Y., Jin, R., and Yang, T. (2021). An online method for a class of distributionally robust optimization with non-convex objectives. *Adv. NeurIPS*, 34:10067–10080.
- Rahimian, H. and Mehrotra, S. (2019). Distributionally robust optimization: A review. *arXiv preprint arXiv:1908.05659*.
- Sessa, P. G., Bogunovic, I., Kamgarpour, M., and Krause, A. (2020). Mixed strategies for robust optimization of unknown objectives. In *Inter. Conf. on Artif. Intel. and Stat.*, pages 2970–2980. PMLR.
- Stabler, B., Bar-Gera, H., and Sall, E. (2018). Transportation networks for research. Available at: <https://github.com/bstabler/TransportationNetworks>.
- Van Parys, B. P., Esfahani, P. M., and Kuhn, D. (2021). From data to decisions: Distributionally robust optimization is optimal. *Management Sci.*, 67(6):3387–3402.
- Wächter, A. and Biegler, L. T. (2006). On the implementation of an interior-point filter line-search algorithm for large-scale nonlinear programming. *Math. Prog.*, 106(1):25–57.
- Wang, H. (2008). Exact confidence coefficients of simultaneous confidence intervals for multinomial proportions. *J. of Multi. Anal.*, 99(5):896–911.
- Weissman, T., Ordentlich, E., Seroussi, G., Verdu, S., and Weinberger, M. J. (2003). Inequalities for the  $l_1$  deviation of the empirical distribution. *Hewlett-Packard Labs, Tech. Rep.*
- Wiesemann, W., Kuhn, D., and Sim, M. (2014). Distributionally robust convex optimization. *Oper. Res.*, 62(6):1358–1376.
- Zhao, P., Zhang, Y.-J., Zhang, L., and Zhou, Z.-H. (2020). Dynamic regret of convex and smooth functions. *Adv. NeurIPS*, 33:12510–12520.
- Zinkevich, M. (2003). Online convex programming and generalized infinitesimal gradient ascent. In *ICML*, pages 928–936.

## Appendix

Here we present the proofs omitted from the main body of the paper as well as additional numerical experiments.

### 6.1 Proofs of Theorem 2.2 and Theorem 2.3

In the following, we prove that the sequence of solutions to the problem  $(DRO_t)$  converge to the true stochastic optimization (SO) problem. These results are an adaptation of the proofs presented in Mohajerin Esfahani and Kuhn (2018) for our setting.

**Lemma 6.1** *Given the true distribution  $p^*$  and the ambiguity set  $\mathcal{P}_t$  at any time  $t$ , we have*

$$\mathbb{P}[p^* \in \mathcal{P}_t] \geq 1 - \delta_t \quad \text{for all } t = 1, \dots, T.$$

*Proof:* This is true by the construction of the set  $\mathcal{P}_t$ , which is such that it contains the true distribution with probability at least  $1 - \delta_t$ .  $\square$

**Lemma 6.2 (Finite sample guarantee)** *Given a solution  $x_t$  to the problem  $DRO_t$ , we prove that*

$$\mathbb{P}[\mathbb{E}_{s \sim p^*}[f(x_t, s)] \leq \widehat{J}_t] \geq 1 - \delta_t \quad \text{for all } t = 1, \dots, T.$$

*Proof:* From Lemma 6.1, we know that  $p^* \in \mathcal{P}_t$  with probability at least  $1 - \delta_t$ . Thus, we have

$$\mathbb{E}_{s \sim p^*}[f(x_t, s)] \leq \max_{p \in \mathcal{P}_t} \mathbb{E}_{s \sim p}[f(x_t, s)],$$

with probability at least  $1 - \delta_t$ .

The right hand side (RHS) term in the above equation is the definition of  $\widehat{J}_t$ . Thus,

$$\mathbb{E}_{s \sim p^*}[f(x_t, s)] \leq \widehat{J}_t,$$

with probability of at least  $1 - \delta_t$ .  $\square$

**Lemma 6.3 (Borel-Cantelli Lemma)** *Let  $E_1, E_2, \dots$  be a sequence of events. If  $\sum_{i=1}^{\infty} P(E_i) < \infty$  then*

$$P[\text{an infinite number of } E_i \text{ occur}] = 0.$$

**Lemma 6.4 (Convergence of Distributions)** *Given the ambiguity set  $\mathcal{P}_t$ , we prove that*

$$\lim_{t \rightarrow \infty} \sup_{p \in \mathcal{P}_t} \|p - p^*\|_2 = 0 \quad \text{with probability 1.}$$

*Proof:* From Lemma 6.7, 6.8 and 6.9 we know that for any of the three given types of ambiguity sets there exists a function  $r(t)$  which satisfies

$$\mathbb{P}\left[\sup_{p \in \mathcal{P}_t} \|p - p^*\|_2 \leq r(t)\right] \geq 1 - \delta_t,$$

and  $\lim_{t \rightarrow \infty} r(t) = 0$ .

This means that

$$\mathbb{P}\left[\sup_{p \in \mathcal{P}_t} \|p - p^*\|_2 - r(t) > 0\right] \leq \delta_t.$$

By construction, it follows that  $\sum_{t=1}^{\infty} \delta_t < \infty$  (as  $\delta_t = \frac{6\delta}{\pi^2 t^2}$ ). Then the Borel-Cantelli Lemma 6.3 implies that

$$\mathbb{P}\left[\lim_{t \rightarrow \infty} \sup_{p \in \mathcal{P}_t} \|p - p^*\|_2 - r(t) \leq 0\right] = 1.$$

Since  $\lim_{t \rightarrow \infty} r(t) = 0$  and  $\|p - p^*\|_2 \geq 0$ , this means that

$$\lim_{t \rightarrow \infty} \sup_{p \in \mathcal{P}_t} \|p - p^*\|_2 = 0 \text{ with probability 1.}$$

□

*Proof of Theorem 2.2* We know that  $x_t \in \mathcal{X}$  and  $J^* \leq \mathbb{E}_{s \sim p^*}[f(x_t, s)]$  as  $x_t$  is a suboptimal solution. Applying Lemma 6.2, we obtain

$$\mathbb{P} \left[ J^* \leq \mathbb{E}_{s \sim p^*}[f(x_t, s)] \leq \widehat{J}_t \right] \geq \mathbb{P} [p^* \in \mathcal{P}_t] \geq 1 - \delta_t.$$

Since  $\sum_{t=1}^{\infty} \delta_t < \infty$ , by the Borel-Cantelli lemma,

$$\mathbb{P} \left[ J^* \leq \lim_{t \rightarrow \infty} \mathbb{E}_{s \sim p^*}[f(x_t, s)] \leq \lim_{t \rightarrow \infty} \widehat{J}_t \right] = 1.$$

Let  $\gamma \geq 0$ . Since  $\mathcal{X}$  is compact, there exists a  $\gamma$ -optimal solution  $x^\gamma$  to the stochastic problem, i.e.,

$$\mathbb{E}_{s \sim p^*}[f(x^\gamma, s)] \leq J^* + \gamma.$$

Let  $p_t^\gamma \in \mathcal{P}_t$  be a  $\gamma$ -optimal distribution to  $x^\gamma$ , i.e.

$$\sup_{p \in \mathcal{P}_t} \mathbb{E}_{s \sim p}[f(x^\gamma, s)] \leq \mathbb{E}_{s \sim p_t^\gamma}[f(x^\gamma, s)] + \gamma.$$

Then we can write

$$\begin{aligned} & \limsup_{t \rightarrow \infty} \widehat{J}_t \\ & \leq \limsup_{t \rightarrow \infty} \sup_{p \in \mathcal{P}_t} \mathbb{E}_{s \sim p}[f(x^\gamma, s)] \\ & \leq \limsup_{t \rightarrow \infty} \mathbb{E}_{s \sim p_t^\gamma}[f(x^\gamma, s)] + \gamma \\ & = \limsup_{t \rightarrow \infty} \mathbb{E}_{s \sim p^*}[f(x^\gamma, s)] + \sum_{s \in \mathcal{S}} f(x^\gamma, s)(p_{st}^\gamma - p_s^*) + \gamma \\ & \leq \limsup_{t \rightarrow \infty} \mathbb{E}_{s \sim p^*}[f(x^\gamma, s)] + G\|p_t^\gamma - p^*\|_2 + \gamma \\ & = \mathbb{E}_{s \sim p^*}[f(x^\gamma, s)] + \gamma \quad \text{w.p. 1} \\ & = J^* + 2\gamma \quad \text{w.p. 1,} \end{aligned}$$

where the first inequality holds because of the definition of  $\widehat{J}_t$ , the second inequality holds because of the definition of  $p_t^\gamma$ , the first equality holds as we add and subtract  $p^*$  and the third inequality holds as  $|f(x, s)| \leq G$  for all  $(x, s) \in \mathcal{X} \times \mathcal{S}$  by assumption. The final two equalities hold because of Lemma 6.4 and the definition of  $x^\gamma$  respectively.

With this we conclude that  $\limsup_{t \rightarrow \infty} \widehat{J}_t \leq J^*$ . Along with the earlier assertion of  $J^* \leq \lim_{t \rightarrow \infty} \widehat{J}_t$ , we can now complete the proof that  $\widehat{J}_t \rightarrow J^*$  via the sandwich argument. □

*Proof of Theorem 2.3* Let  $\{s_t\}_{t=1}^{\infty}$  be any sequence of scenario realizations such that  $\lim_{t \rightarrow \infty} \widehat{J}_t = J^*$ . By Theorem 2.2, we have  $J^* \leq \mathbb{E}_{s \sim p^*}[f(x_t, s)] \leq \widehat{J}_t$  with probability 1. By the same theorem, we also know that  $\lim_{t \rightarrow \infty} \widehat{J}_t = J^*$  w.p. 1. Then, we can write

$$\liminf_{t \rightarrow \infty} \mathbb{E}_{s \sim p^*}[f(x_t, s)] \leq \liminf_{t \rightarrow \infty} \widehat{J}_t = J^*. \quad (6)$$

Consider any limit point of the sequence  $\{x_t\}_{t=1}^{\infty}$ . Since the set  $\mathcal{X}$  is compact, then there exists a limit point of  $\{x_t\}_{t=1}^{\infty}$  which lies in  $\mathcal{X}$ . WLOG let  $x^*$  be that point and

$\liminf_{t \rightarrow \infty} x_t = x^*$ .

Then we have

$$\begin{aligned}
J^* &\leq \mathbb{E}_{s \sim p^*}[f(x^*, s)] \\
&= \mathbb{E}_{s \sim p^*}[\liminf_{t \rightarrow \infty} f(x_t, s)] \\
&= \sum_{s \in \mathcal{S}} \liminf_{t \rightarrow \infty} f(x_t, s) p_s^* \\
&= \liminf_{t \rightarrow \infty} \sum_{s \in \mathcal{S}} f(x_t, s) p_s^* \leq J^*,
\end{aligned}$$

where the first inequality holds because  $x^* \in \mathcal{X}$ , the second inequality holds as  $\liminf_{t \rightarrow \infty} x_t = x^*$  and because  $f(x, s)$  is continuous in  $x$ . The second equality exploits that  $\mathcal{S}$  is finite and the final inequality holds because of (6). Thus, we have  $\mathbb{E}_{s \sim p^*}[f(x^*, s)] = J^*$  which completes the proof.  $\square$

## 6.2 Proofs of Dynamic regret bounds

In order to prove the dynamic regret bound of Theorem 3.1, we first show that the ambiguity sets are shrinking at a rate of  $\mathcal{O}(\sqrt{\log T/T})$  and contain the true data generating distribution with a high confidence. The latter is stated in the following Lemma.

A crucial point for a shrinking dynamic regret bound is that the ambiguity sets are shrinking over time. Our ambiguity sets are constructed with increasing confidence probabilities for the multinomial distribution. We show that the ambiguity sets are shrinking even though the confidence  $1 - \delta_t$  is increasing ( $\delta_t = \frac{6\delta}{\pi^2 t^2}$ ).

**Lemma 6.5** *For the upper  $(1 - \frac{\delta_t}{2})$ -percentile  $z_{\frac{\delta_t}{2}}$  of the standard normal distribution with confidence update*

$\delta_t := \frac{6\delta}{\pi^2 t^2}$  and  $\delta \in (0, 1)$ , it follows that

$$z_{\frac{\delta_t}{2}}^2 \leq 4 \log(\pi t),$$

for all rounds  $t = 1, \dots, T$ .

*Proof:* By the definition of the standard normal distribution, for the upper percentile  $1 - \frac{\delta_t}{2}$ , we have the Gaussian tail bound

$$\begin{aligned}
1 - \frac{\delta_t}{2} &\leq e^{-\frac{1}{2} z_{\frac{\delta_t}{2}}^2} \\
\implies \log(1 - \frac{\delta_t}{2}) &\leq -\frac{z_{\frac{\delta_t}{2}}^2}{2} \\
\implies z_{\frac{\delta_t}{2}}^2 &\leq -2 \log(1 - \frac{\delta_t}{2}) \\
&= -2 \log\left(\frac{2\pi^2 t^2 - 6\delta}{2\pi^2 t^2}\right) \\
&= 2 \log\left(\frac{\pi^2 t^2}{\pi^2 t^2 - 3\delta}\right),
\end{aligned}$$

for all rounds  $t = 1, \dots, T$ . Since  $\pi^2 t^2 - 3\delta \geq 1$  for all rounds  $t = 1, \dots, T$  and  $\delta \in (0, 1)$  (try  $t = 1$  and  $\delta = 1$ ), we are able write

$$z_{\frac{\delta_t}{2}}^2 \leq 2 \log(\pi^2 t^2) \leq 4 \log(\pi t),$$

for all  $t = 1, \dots, T$ .  $\square$

We now prove that the ambiguity sets  $\mathcal{P}_t$  shrink at a rate of  $\mathcal{O}(\sqrt{\log t/t})$  even as the confidence requirement  $1 - \delta_t$  increases.

Here, we provide a proof for ambiguity sets as defined by confidence intervals. The proofs for the other sets are provided in the electronic companion.

**Lemma 6.6 (Confidence Interval Sets)** *The ambiguity sets  $\mathcal{P}_t$  derived from (1) with confidence update  $\delta_t := \frac{6\delta}{\pi^2 t^2}$  and  $\delta \in (0, 1)$  for all rounds  $t = 1, \dots, T$  fulfill*

$$\sup_{x \in \mathcal{P}_0, y \in \mathcal{P}_1} \|x - y\| \leq \sqrt{16|\mathcal{S}| \log \pi} \text{ and } \sup_{x \in \mathcal{P}_{t-1}, y \in \mathcal{P}_t} \|x - y\| \leq \frac{\sqrt{16|\mathcal{S}| \log(\pi(t-1))}}{\sqrt{t-1}},$$

for all  $t = 2, \dots, T$  with a probability of at least  $1 - \delta$ .

*Proof:* As we know from Lemma 2.1 that  $p^* \in \bigcap_{t=0, \dots, T} \mathcal{P}_t$  with a probability of at least  $1 - \delta$ , we can compute for  $t = 1$ :

$$\begin{aligned} \sup_{x \in \mathcal{P}_0, y \in \mathcal{P}_1} \|x - y\| &= \sup_{x \in \mathcal{P}_0, y \in \mathcal{P}_1} \|x - p^* + p^* - y\| \\ &\leq \sup_{x \in \mathcal{P}_0} \|x - p^*\| + \sup_{y \in \mathcal{P}_1} \|p^* - y\| \\ &\leq \sup_{x, p \in \mathcal{P}_0} \|x - p\| + \sup_{p, y \in \mathcal{P}_1} \|p - y\| \\ &\leq \sqrt{2} + \sqrt{4|\mathcal{S}| \log \pi} \leq \sqrt{16|\mathcal{S}| \log \pi}, \end{aligned}$$

with a probability of at least  $1 - \delta$  because of  $\sup_{x, y \in \mathcal{P}_0} \|x - y\| = \sqrt{(1-0)^2 + (0-1)^2} = \sqrt{2}$  and Lemma 6.7.

Similarly for  $t = 2, \dots, T$ :

$$\begin{aligned} \sup_{x \in \mathcal{P}_{t-1}, y \in \mathcal{P}_t} \|x - y\| &= \sup_{x \in \mathcal{P}_{t-1}, y \in \mathcal{P}_t} \|x - p^* + p^* - y\| \\ &\leq \sup_{x \in \mathcal{P}_{t-1}} \|x - p^*\| + \sup_{y \in \mathcal{P}_t} \|p^* - y\| \\ &\leq \sup_{x, p \in \mathcal{P}_{t-1}} \|x - p\| + \sup_{p, y \in \mathcal{P}_t} \|p - y\| \\ &\leq 2 \frac{\sqrt{4|\mathcal{S}| \log(\pi(t-1))}}{\sqrt{t-1}}, \end{aligned}$$

with a probability of at least  $1 - \delta$ .  $\square$

**Lemma 6.7** *The ambiguity sets  $\mathcal{P}_t$  derived from (1) with confidence update  $\delta_t := \frac{6\delta}{\pi^2 t^2}$  and  $\delta \in (0, 1)$  for all rounds  $t = 1, \dots, T$  fulfill*

$$\sup_{x, y \in \mathcal{P}_t} \|x - y\| \leq \frac{\sqrt{4|\mathcal{S}| \log(\pi t)}}{\sqrt{t}}.$$

*Proof:* We can compute using Lemma 6.5:

$$\sup_{x, y \in \mathcal{P}_t} \|x - y\|^2 \leq \sum_{k=1}^{|\mathcal{S}|} (u_{kt} - l_{kt})^2 = \sum_{k=1}^{|\mathcal{S}|} \frac{z_{\delta_t}^2}{t} \leq \frac{4|\mathcal{S}| \log(\pi t)}{t}.$$

$\square$

Finally, In the following result, we prove that the distance between elements from consecutive  $\ell_2$ -norm and kernel based ambiguity sets  $\mathcal{P}_{t-1}$  and  $\mathcal{P}_t$  also shrinks as measured in the  $\ell_2$ -norm while providing tighter guarantees of inclusion on the true distribution..

**Lemma 6.8 ( $\ell_2$ -norm Sets)** Given an ambiguity set of the form  $\mathcal{P}_t = \{p \in \mathcal{P} \mid \|p - \hat{p}\|_2 \leq \epsilon_t\}$  with  $\epsilon_t := \sqrt{\frac{2|\mathcal{S}|\log 2/\delta_t}{t}}$  and  $\delta_t = \frac{6\delta}{\pi^2 t^2}$ , then for  $|\mathcal{S}| \geq 2$  we have

$$\sup_{x \in \mathcal{P}_0, y \in \mathcal{P}_1} \|x - y\|_2 \leq 4\sqrt{|\mathcal{S}|\log(\pi/\sqrt{3\delta})} \text{ and } \sup_{x \in \mathcal{P}_{t-1}, y \in \mathcal{P}_t} \|x - y\|_2 \leq 4\sqrt{\frac{|\mathcal{S}|\log(\pi(t-1)/\sqrt{3\delta})}{t-1}},$$

with probability at least  $1 - \delta$ .

*Proof:* For the case  $t = 1$  we have

$$\begin{aligned} \sup_{x \in \mathcal{P}_0, y \in \mathcal{P}_1} \|x - y\|_2 &\leq \sup_{x \in \mathcal{P}_0} \|x - p^*\|_2 + \sup_{y \in \mathcal{P}_1} \|p^* - y\|_2 \\ &\leq 2 + 2\sqrt{|\mathcal{S}|\log(\pi/\sqrt{3\delta})} \\ &\leq 4\sqrt{|\mathcal{S}|\log(\pi/\sqrt{3\delta})}. \end{aligned}$$

The last inequality occurs because  $\sqrt{|\mathcal{S}|\log(\pi/\sqrt{3\delta})} > 1$  for  $|\mathcal{S}| \geq 2$ . Now for the case  $t > 1$ , given the true distribution  $p^*$ , we can write,

$$\begin{aligned} \|x - y\|_2 &\leq \|x - p^*\|_2 + \|p^* - y\|_2 \\ &\leq 2\epsilon_{t-1}. \end{aligned}$$

Thus,

$$\sup_{x \in \mathcal{P}_{t-1}, y \in \mathcal{P}_t} \|x - y\|_2 \leq 4\sqrt{\frac{|\mathcal{S}|\log(\pi(t-1)/\sqrt{3\delta})}{t-1}}.$$

Here, the first inequality arises from triangle inequality. The second inequality is due to the fact that the true distribution is contained inside all sets  $\mathcal{P}_t$  for  $t = 1, \dots, T$  with probability at least  $1 - \delta$ .  $\square$

**Lemma 6.9 (Kernel Based Sets)** Given an ambiguity set of the form  $\mathcal{P}_t = \{p \in \mathcal{P} \mid \|p - \hat{p}\|_M \leq \epsilon_t\}$  with  $\epsilon_t := \frac{\sqrt{C}}{\sqrt{t}}(2 + \sqrt{2\log(1/\delta_t)})$  with  $\delta_t = \frac{6\delta}{\pi^2 t^2}$  we have for  $t \geq 2$ ,

$$\begin{aligned} \sup_{x \in \mathcal{P}_{t-1}, y \in \mathcal{P}_t} \|x - y\|_2 &\leq \frac{8\sqrt{C}}{\lambda\sqrt{t-1}}(\sqrt{\log(\pi t/\sqrt{6\delta})}), \\ \text{and } \sup_{x \in \mathcal{P}_0, y \in \mathcal{P}_1} \|x - y\|_2 &\leq 2 + \frac{4\sqrt{C}}{\lambda} \text{ for } t = 1, \end{aligned}$$

with probability at least  $1 - \delta$ .

*Proof:* For  $t = 1$ , we have

$$\begin{aligned} \|x - y\|_2 &\leq \|x - p^*\|_2 + \|p^* - y\|_2 \\ &\leq 2 + \frac{1}{\lambda}\|p^* - y\|_M \\ &\leq 2 + \frac{4\sqrt{C}}{\lambda}. \end{aligned}$$

The second inequality comes from the definition of the set  $\mathcal{P}_0$  and the definition of the norm  $\|\cdot\|_M$ .

For the case  $t \geq 2$ , we get.

$$\begin{aligned}
\|x - y\|_2 &\leq \|x - p^*\|_2 + \|p^* - y\|_2 \\
&\leq \frac{1}{\lambda} (\|x - p^*\|_M + \|p^* - y\|_M) \\
&\leq \frac{4}{\lambda} \frac{\sqrt{C}}{\sqrt{t-1}} (1 + \sqrt{\log(\pi(t-1)/\sqrt{6\delta})}) \\
&\leq \frac{4}{\lambda} \frac{\sqrt{C}}{\sqrt{t-1}} (1 + \sqrt{\log(\pi t/\sqrt{6\delta})}).
\end{aligned}$$

Here, the first inequality comes from the triangle inequality, the second from the fact that  $\sqrt{x^\top M x} \geq \lambda \|x\|_2$  for any positive definite matrix  $M$  with the minimum eigen value  $\lambda$ . The final inequality arises from the construction of the ambiguity set which contains the true distribution with probability at least  $1 - \delta$ . Now note that  $\pi^2 t^2 / 6\delta > 3$  for  $t \geq 2$ . Thus, we can write

$$\|x - y\|_2 \leq \frac{8}{\lambda} \frac{\sqrt{C}}{\sqrt{t-1}} \sqrt{\log(\pi t/\sqrt{6\delta})}.$$

□

The properties of the confidence interval ambiguity sets mentioned previously and those of the  $\ell_2$ -norm and kernel based sets prove in Lemmas 6.8 and 6.9 along with the cumulative path length bounds from Theorem 6.11 enable us to prove the following shrinking dynamic regret.

**Theorem 6.10 (Dynamic regret bound)** *Let  $f : \mathcal{X} \times \mathcal{S} \rightarrow \mathbb{R}$  be uniformly bounded, i.e., for all  $(x, s) \in \mathcal{X} \times \mathcal{S}$ , a constant  $G > 0$  exists such that  $|f(x, s)| \leq G$ . Let  $\eta := \sqrt{\frac{3+2h'(T)}{TG^2|\mathcal{S}|}}$  with  $\sum_{t=2}^T \|p - q\| \leq h'(T)$  for  $p \in \mathcal{P}_{t-1}$  and  $q \in \mathcal{P}_t$ . The output  $(x_1, \dots, x_T)$  from Algorithm with confidence update  $\delta_t := \frac{6\delta}{\pi^2 t^2}$  and  $\delta \in (0, 1)$  fulfills*

$$\frac{1}{T} \sum_{t=1}^T \left( \max_{p \in \mathcal{P}_t} \mathbb{E}_{s \sim p} [f(x_t, s)] - \min_{x \in \mathcal{X}} \max_{p \in \mathcal{P}_t} \mathbb{E}_{s \sim p} [f(x, s)] \right) \leq G \sqrt{\frac{3|\mathcal{S}| + 2|\mathcal{S}|h'(T)}{T}} + \frac{2G}{T},$$

with probability at least  $1 - \delta$ .

*Proof of Theorem 6.10* Define  $g_t(p) := -\mathbb{E}_{s \sim p} [f(x_t, s)]$ . A gradient descent iteration is given by

$$p_{t+1} = \arg \min_{p \in \mathcal{P}_t} \langle \eta \nabla g_t(p_t), p \rangle + \frac{1}{2} \|p - p_t\|^2,$$

with optimality criteria  $\langle \eta \nabla g_t(p_t), u_t - p_{t+1} \rangle + \langle p_{t+1} - p_t, u_t - p_{t+1} \rangle \geq 0$  for all  $u_t \in \mathcal{P}_t$ . Classical theory for gradient descent yields

$$\langle \eta \nabla g_t(p_t), p_t - u_t \rangle \leq \frac{1}{2} \|p_t - u_t\|^2 - \frac{1}{2} \|p_{t+1} - u_t\|^2 + \frac{\eta^2}{2} \|\nabla g_t(p_t)\|^2.$$

Summation over rounds  $t = 1, \dots, T$  results in the following inequality for all  $u_t \in \mathcal{P}_t$ :

$$\sum_{t=1}^T \langle \eta \nabla g_t(p_t), p_t - u_t \rangle \leq \sum_{t=1}^T \frac{\eta^2}{2} \|\nabla g_t(p_t)\|^2 + \sum_{t=1}^T \left( \frac{1}{2} \|p_t - u_t\|^2 - \frac{1}{2} \|p_{t+1} - u_t\|^2 \right). \quad (7)$$



Next, we rearrange the terms on the RHS as

$$\begin{aligned}
& \sum_{t=1}^T \left( \frac{1}{2} \|p_t - u_t\|^2 - \frac{1}{2} \|p_{t+1} - u_t\|^2 \right) \\
&= \frac{1}{2} \sum_{t=1}^T \left( \|p_t\|^2 - \|p_{t+1}\|^2 \right) + \frac{1}{2} \sum_{t=1}^T 2(p_{t+1} - p_t) \cdot u_t \\
&= \frac{1}{2} \sum_{t=1}^T \left( \|p_t\|^2 - \|p_{t+1}\|^2 \right) + \sum_{t=1}^T (p_{t+1} - p_t) \cdot u_t.
\end{aligned}$$

Now consider the expression  $\sum_{t=1}^T (p_{t+1} - p_t) \cdot u_t$ . We can write it as

$$\begin{aligned}
& \sum_{t=1}^T (p_{t+1} - p_t) \cdot u_t \\
&= (p_2 - p_1) \cdot u_1 + (p_3 - p_2) \cdot u_2 + \cdots + (p_{T+1} - p_T) \cdot u_T \\
&= p_{T+1} \cdot u_T - p_1 \cdot u_1 + \sum_{t=2}^T (u_{t-1} - u_t) \cdot p_t.
\end{aligned}$$

For the other expression we have

$$\frac{1}{2} \sum_{t=1}^T \left( \|p_t\|^2 - \|p_{t+1}\|^2 \right) = \frac{1}{2} \|p_1\|^2 - \frac{1}{2} \|p_{T+1}\|^2.$$

This then allows us to write

$$\begin{aligned}
& \sum_{t=1}^T \langle \eta \nabla g_t(p_t), p_t - u_t \rangle \leq \sum_{t=1}^T \frac{\eta^2}{2} \|\nabla g_t(p_t)\|^2 + \frac{1}{2} \|p_1\|^2 \\
& \quad - \frac{1}{2} \|p_{T+1}\|^2 + p_{T+1} \cdot u_T - p_1 \cdot u_1 + \sum_{t=2}^T (u_{t-1} - u_t) \cdot p_t.
\end{aligned}$$

We know that  $\|p\| \leq 1$  and for any 2 probability vectors  $p$  and  $q$  we have that  $0 \leq p \cdot q \leq 1$ . Thus we can write

$$\sum_{t=1}^T \langle \eta \nabla g_t(p_t), p_t - u_t \rangle \leq \sum_{t=1}^T \frac{\eta^2}{2} \|\nabla g_t(p_t)\|^2 + \frac{1}{2} + 1 + \sum_{t=2}^T (u_{t-1} - u_t) \cdot p_t.$$

Note that

$$\|\nabla g_t(p_t)\|^2 = \sum_{k=1}^{|\mathcal{S}|} |f(x_t, s_k)|^2 \leq |\mathcal{S}| G^2,$$

This, along with the fact that  $\sum_{t=2}^T (u_{t-1} - u_t) \cdot p_t \leq \sum_{t=2}^T \|u_{t-1} - u_t\| \|p_t\| \leq \sum_{t=2}^T \|u_{t-1} - u_t\|$ , allows us to write

$$\sum_{t=1}^T \langle \nabla g_t(p_t), p_t - u_t \rangle \leq \frac{\eta T}{2} |\mathcal{S}| G^2 + \frac{3}{2\eta} + \frac{1}{\eta} \sum_{t=2}^T \|u_{t-1} - u_t\|.$$

By assumption,  $\sum_{t=2}^T \|u_{t-1} - u_t\| \leq h'(T)$  with probability at least  $1 - \delta$  for some function  $h'(\cdot)$ . Then we have

$$\sum_{t=1}^T \langle \nabla g_t(p_t), p_t - u_t \rangle \leq \frac{\eta T}{2} |\mathcal{S}| G^2 + \frac{3}{2\eta} + \frac{h'(T)}{\eta},$$

with probability at least  $1 - \delta$ . Choosing the optimal  $\eta = \sqrt{\frac{3+2h'(T)}{TG^2|\mathcal{S}|}}$ , we can write our result as

$$\sum_{t=1}^T \langle \nabla g_t(p_t), p_t - u_t \rangle \leq 2\sqrt{\frac{T}{2}|\mathcal{S}|G^2 \cdot \left(\frac{3}{2} + h'(T)\right)}.$$

Then we can write

$$\sum_{t=1}^T \langle \nabla g_t(p_t), p_t - u_t \rangle \leq G|\mathcal{S}|^{\frac{1}{2}}\sqrt{3T + 2Th'(T)}.$$

Since  $g_t(p) = -\mathbb{E}_{s \sim p} [f(x_t, s)]$  is linear in  $p$  for all  $t = 1, \dots, T$ , it follows

$$\begin{aligned} & \sum_{t=1}^T (\mathbb{E}_{s \sim u_t} [f(x_t, s)] - \mathbb{E}_{s \sim p_t} [f(x_t, s)]) \\ &= \sum_{t=1}^T (g_t(p_t) - g_t(u_t)) = \sum_{t=1}^T \langle \nabla g_t(p_t), p_t - u_t \rangle \\ &\leq G|\mathcal{S}|^{\frac{1}{2}}\sqrt{3T + 2Th'(T)}. \end{aligned}$$

Now we choose in each round  $t = 1, \dots, T$  the worst-case  $u_t := \arg \max_{p \in \mathcal{P}_t} \mathbb{E}_{s \sim p} [f(x_t, s)] \in \mathcal{P}_t$  and recall  $x_t = \arg \min_{x \in \mathcal{X}} \mathbb{E}_{s \sim p_t} [f(x, s)]$  to obtain

$$\sum_{t=1}^T (\mathbb{E}_{s \sim u_t} [f(x_t, s)] - \mathbb{E}_{s \sim p_t} [f(x_t, s)]) = \sum_{t=1}^T \left( \max_{p \in \mathcal{P}_t} \mathbb{E}_{s \sim p} [f(x_t, s)] - \min_{x \in \mathcal{X}} \mathbb{E}_{s \sim p_t} [f(x, s)] \right).$$

Since  $p_t \in \mathcal{P}_{t-1}$ , we know that  $\min_{x \in \mathcal{X}} \mathbb{E}_{s \sim p_t} [f(x, s)] \leq \min_{x \in \mathcal{X}} \max_{p \in \mathcal{P}_{t-1}} \mathbb{E}_{s \sim p} [f(x, s)]$  for all  $t = 1, \dots, T$  and thus we can conclude

$$\sum_{t=1}^T \left( \max_{p \in \mathcal{P}_t} \mathbb{E}_{s \sim p} [f(x_t, s)] - \min_{x \in \mathcal{X}} \max_{p \in \mathcal{P}_{t-1}} \mathbb{E}_{s \sim p} [f(x, s)] \right) \leq G|\mathcal{S}|^{\frac{1}{2}}\sqrt{3T + 2Th'(T)},$$

with a probability of at least  $1 - \delta$ . We add and subtract  $\min_{x \in \mathcal{X}} \max_{p \in \mathcal{P}_t} \mathbb{E}_{s \sim p} [f(x, s)]$  on the LHS. Rearranging the terms this allows us to write the LHS as

$$\begin{aligned} & \sum_{t=1}^T \left( \max_{p \in \mathcal{P}_t} \mathbb{E}_{s \sim p} [f(x_t, s)] - \min_{x \in \mathcal{X}} \max_{p \in \mathcal{P}_t} \mathbb{E}_{s \sim p} [f(x, s)] \right) \\ &+ \sum_{t=1}^T \min_{x \in \mathcal{X}} \max_{p \in \mathcal{P}_t} \mathbb{E}_{s \sim p} [f(x, s)] - \min_{x \in \mathcal{X}} \max_{p \in \mathcal{P}_{t-1}} \mathbb{E}_{s \sim p} [f(x, s)]. \end{aligned}$$

Observing that the last two terms telescope, bringing them to the RHS and using the upper bound  $G$  on  $|f(x, s)|$  we can conclude

$$\sum_{t=1}^T \left( \max_{p \in \mathcal{P}_t} \mathbb{E}_{s \sim p} [f(x_t, s)] - \min_{x \in \mathcal{X}} \max_{p \in \mathcal{P}_t} \mathbb{E}_{s \sim p} [f(x, s)] \right) \leq G|\mathcal{S}|^{\frac{1}{2}}\sqrt{3T + 2Th'(T)} + 2G \quad \text{w.p. } 1 - \delta.$$

Dividing by  $T$  on both sides completes the proof.  $\square$

**Theorem 6.11** *Given ambiguity sets of the form specified in Section 2, we have*

$$\frac{1}{2} \sum_{t=1}^T \|p_t - q_t\|^2 \leq h(T) \quad \text{and} \quad \sum_{t=2}^T \|p_t - q_t\| \leq h'(T) \quad \text{for all } p_t \in \mathcal{P}_{t-1}, q_t \in \mathcal{P}_t,$$

with probability at least  $1 - \delta$ . The functions  $h(T)$  and  $h'(T)$  for different categories of ambiguity sets are as given below:

1. *Confidence Intervals:*

$$h(T) = 8|\mathcal{S}| \log(\pi T)(2 + \log T)$$

$$h'(T) = 8\sqrt{|\mathcal{S}|T \log(\pi T)}$$

2. *Kernel based ambiguity sets, where  $\lambda$  denotes the smallest eigenvalue of the kernel matrix  $M$ :*

$$h(T) = \frac{1}{2} \left( 2 + \frac{4\sqrt{C}}{\lambda} \right)^2 + \frac{32C}{\lambda^2} \log(\pi T / \sqrt{6\delta})(1 + \log T)$$

$$h'(T) = \frac{16\sqrt{C}}{\lambda} \sqrt{T \log(\pi T / \sqrt{6\delta})}$$

3.  *$\ell_2$ -norm ambiguity sets:*

$$h(T) = 8|\mathcal{S}| \log \frac{\pi T}{\sqrt{3\delta}} (2 + \log T)$$

$$h'(T) = 8\sqrt{|\mathcal{S}|T \log \frac{\pi T}{\sqrt{3\delta}}}$$

**Confidence Intervals.** Calculating the function  $h(T)$ , we have

$$\begin{aligned} \frac{1}{2} \sum_{t=1}^T \|p_t - q_t\|^2 &\leq \frac{1}{2} 16|\mathcal{S}| \log \pi + \frac{1}{2} \sum_{t=2}^T 16|\mathcal{S}| \frac{\log(\pi(t-1))}{t-1} \\ &\leq 8|\mathcal{S}| \log \pi + 8|\mathcal{S}| \log(\pi(T-1)) \sum_{t=1}^{T-1} \frac{1}{t} \\ &\leq 8|\mathcal{S}| \log \pi + 8|\mathcal{S}| \log(\pi(T-1))(1 + \log(T-1)) \\ &\leq 8|\mathcal{S}| \log(\pi T)(2 + \log T). \end{aligned}$$

Here, the first inequality arises from Lemma 6.6. The second and third inequalities are from bounding  $t$  and from observing that  $\sum_{t=1}^{T-1} (1/t) \leq 1 + \log(T-1)$ .

Now for the function  $h'(T)$ , we can calculate

$$\begin{aligned} \sum_{t=2}^T \|p_t - q_t\| &\leq \sum_{t=2}^T 4 \frac{\sqrt{|\mathcal{S}| \log(\pi(t-1))}}{\sqrt{t-1}} \\ &\leq 4|\mathcal{S}|^{\frac{1}{2}} \sqrt{\log(\pi T)} \sum_{t=2}^T \frac{1}{\sqrt{t-1}} \\ &\leq 8|\mathcal{S}|^{\frac{1}{2}} \sqrt{\log(\pi T)} \sqrt{T}. \end{aligned}$$

Here, the first inequality arises from Lemma 6.6. The second and third inequalities are from bounding  $t$  and from observing that  $\sum_{t=2}^T (1/\sqrt{t-1}) \leq 2\sqrt{T-1} \leq 2\sqrt{T}$ .

**Kernel based ambiguity sets.** Calculating the function  $h(T)$ , we have

$$\begin{aligned} \frac{1}{2} \sum_{t=1}^T \|p_t - q_t\|_2^2 &\leq \frac{1}{2} \left( 2 + \frac{4\sqrt{C}}{\lambda} \right)^2 + \sum_{t=2}^T \frac{32C}{\lambda^2(t-1)} \log(\pi(t-1)/\sqrt{6\delta}) \\ &\leq \frac{1}{2} \left( 2 + \frac{4\sqrt{C}}{\lambda} \right)^2 + \frac{32C}{\lambda^2} \log(\pi T / \sqrt{6\delta}) \sum_{t=2}^T \frac{1}{t-1} \\ &\leq \frac{1}{2} \left( 2 + \frac{4\sqrt{C}}{\lambda} \right)^2 + \frac{32C}{\lambda^2} \log(\pi T / \sqrt{6\delta})(1 + \log T). \end{aligned}$$

Here, the first inequality arises from Lemma 6.9. The second and third inequalities are from bounding  $t$  and from observing that  $\sum_{t=1}^{T-1} (1/t) \leq 1 + \log(T-1) \leq 1 + \log T$ .

Now, for the function  $h'(T)$ , we have

$$\begin{aligned} \sum_{t=2}^T \|p_t - q_t\|_2 &\leq \sum_{t=2}^T \frac{8}{\lambda} \frac{\sqrt{C}}{\sqrt{t-1}} \sqrt{\log(\pi(t-1)/\sqrt{6\delta})} \\ &\leq \frac{8\sqrt{C}}{\lambda} \sqrt{\log(\pi T/\sqrt{6\delta})} \sum_{t=2}^T \frac{1}{\sqrt{t-1}} \\ &\leq \frac{16\sqrt{C}}{\lambda} \sqrt{T \log(\pi T/\sqrt{6\delta})}. \end{aligned}$$

Here, the first inequality arises from Lemma 6.9. The second and third inequalities are from bounding  $t$  and from observing that  $\sum_{t=1}^{T-1} (1/\sqrt{t-1}) \leq 2\sqrt{T-1} \leq 2\sqrt{T}$ .

**$\ell_2$ -norm ambiguity sets.** Calculating the function  $h(T)$ , we have

$$\begin{aligned} \frac{1}{2} \sum_{t=1}^T \|p_t - q_t\|^2 &\leq \frac{1}{2} \left( 4\sqrt{|\mathcal{S}| \log(\pi/\sqrt{3\delta})} \right)^2 + \frac{1}{2} \sum_{t=2}^T 16|\mathcal{S}| \frac{\log(\pi(t-1)/\sqrt{3\delta})}{t-1} \\ &\leq 8|\mathcal{S}| \log(\pi/\sqrt{3\delta}) + 8|\mathcal{S}| \log(\pi T/\sqrt{3\delta}) \sum_{t=2}^T \frac{1}{t-1} \\ &\leq 8|\mathcal{S}| \log(\pi T/\sqrt{3\delta}) (2 + \log T). \end{aligned}$$

Here, the first inequality arises from Lemma 6.8. The second and third inequalities are proven similar to the case of  $h(T)$  for the interval sets.

Now for the function  $h'(T)$ , we can calculate

$$\begin{aligned} \sum_{t=2}^T \|p_t - q_t\| &\leq \sum_{t=2}^T 4\sqrt{\frac{|\mathcal{S}| \log(\pi(t-1)/\sqrt{3\delta})}{t-1}} \\ &\leq 4\sqrt{|\mathcal{S}| \log(\pi T/\sqrt{3\delta})} \sum_{t=2}^T \sqrt{\frac{1}{t-1}} \\ &\leq 8\sqrt{|\mathcal{S}| T \log(\pi T/\sqrt{3\delta})}. \end{aligned}$$

Here, the first inequality arises from Lemma 6.8. The second and third inequalities are proven similar to the case of  $h'(T)$  for the interval sets.  $\square$

## 6.3 Numerical Experiments

In this section, we provide the details for the numerical experiments conducted in Section 4. We also provide an additional set of experiments on an optimal routing problem to further illustrate our algorithms.

### 6.3.1 Benchmark Instances

All mixed-integer linear optimization problems used in our numerical experiments can be found in Table 5. The entries show the number of variables and constraints for each problem. The same holds for the quadratic problems listed in Table 6.

Name	Variables				Constraints
	All	Bin.	Int.	Cont.	
blend2	353	239	25	89	274
flugpl	18	0	11	7	19
gr4x6	48	24	0	24	34
neos-1430701	312	156	0	156	668
noswot	128	75	25	28	182
prod1	250	149	0	101	208
prod2	301	200	0	101	211
ran13x13	338	169	0	169	195
supportcase14	304	304	0	0	234
supportcase16	319	319	0	0	130
beavma	390	195	0	195	372
k16x240b	480	240	0	240	256
neos-3610040	430	85	0	345	335
neos-3611689	421	88	0	333	323
timtab1CUTS	397	77	94	226	371

Table 5: Overview of MIP instances

Name	Variables			Constraints	
	All	Bin.	Cont.	All	Quadr.
7579	300	100	200	203	1
10001	485	426	59	296	1
10002	485	426	59	296	1
10004	1058	999	59	867	1
10010	269	262	7	147	1
10003	1058	999	59	867	1
10008	845	713	132	416	1
10009	605	473	132	246	1
10011	1390	1258	132	873	1
10012	967	835	132	538	1

Table 6: Overview of MIQP instances