

537 **A Illustration of RCL**

538 We illustrate the online optimization process of RCL in Fig. 1.

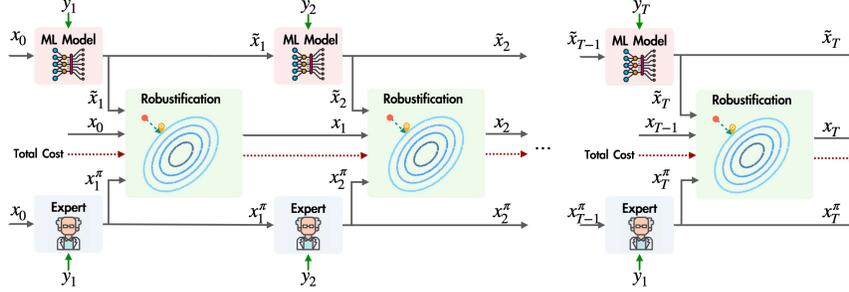


Figure 1: Robustness-constrained online optimization using RCL. The expert algorithm and ML model run independently. At each time $t = 1, \dots, T$, RCL projects the ML prediction \tilde{x}_t into a robustified action set.

539 **B Case Study: Battery Management for EV Charging Stations**

540 We now explore the performance of RCL using a case study focused on battery management in electric
541 vehicle (EV) charging stations [48]. We first formulate the problem as an instance of SOCO, and then
542 present the baseline algorithms. Finally, we discuss the performance of RCL. Our results highlight
543 the advantage of RCL in terms of robustness guarantees compared to pure ML models, as well as the
544 benefit of training a robustification-aware ML model in terms of the average cost.

545 **B.1 Problem Formulation**

546 Batteries are routinely used in EV charging stations to handle the rapidly fluctuating charging demands
547 and protect the connected grid. Thus, properly managing battery charging/discharging decisions is
548 crucial for reliability, lifespan, and safety of batteries and grids.

549 We consider the management of N batteries. At each time step t , suppose that $x_t \in \mathbb{R}_+^N$ represents the
550 State of Charge (SoC) and $u_t \in \mathbb{R}^N$ represents the battery charging/discharging schedule, depending
551 on the sign of u_t (i.e., positive means charging, and vice versa). The canonical form of the battery
552 dynamics can be written as $x_{t+1} = Ax_t + Bu_t - w_t$, where A is a $N \times N$ matrix which models
553 the self-degradation of the N -battery system, B is a $N \times N$ matrix which represents the charging
554 efficiency of each battery unit, w_t is a $N \times 1$ vector which denotes the current demand in terms of
555 the charging rate (kW) of all the EVs connected to the charging stations. Assuming that the initial
556 SoC as x_0 , the goal is to control the batteries to minimize the difference between the current SoC
557 of all batteries and a nominal value \bar{x} , plus a charging/discharging cost to account for battery usage
558 [49, 50], which can be expressed mathematically as $\min_{u_1, u_2, \dots, u_{T+1}} \sum_{t=1}^{T+1} \|x_t - \bar{x}\|^2 + b\|u_t\|^2$.

559 This problem falls into SOCO based on the reduction framework described in [49]. Specifi-
560 cally, at time step $t + 1$, we can expand x_{t+1} based on the battery dynamics as $x_{t+1} = A^t x_1 +$
561 $\sum_{j=1}^t A^{t-j} B u_j - \sum_{j=1}^t A^{t-j} w_j$. We define the context parameter as $y_t = \bar{x} - A^t x_1 + \sum_{i=1}^t A^{t-i} w_i$
562 and the action as $a_t = \sum_{i=1}^t A^{t-i} B u_i$. Then, assuming an identity matrix B (ignoring charging
563 loss), the optimization problem becomes $\min_{a_1, \dots, a_T} \|x_1 - \bar{x}\|^2 + b\|u_T\|^2 + \sum_{t=1}^T \|a_t - y_t\|^2 +$
564 $b\|a_t - A a_{t-1}\|^2$. Given an initial value of x_1 , this problem can be further simplified and reformulated
565 as

$$\min_{a_1, a_2, \dots, a_T} \sum_{t=1}^T \frac{1}{b} \|a_t - y_t\|^2 + \|a_t - A a_{t-1}\|^2, \quad (7)$$

566 which is in a standard SOCO form by considering y_t as the context and a_t as the action at time t .

567 To validate the effectiveness of RCL, we use a public dataset [51] provided by ElaadNL, a Dutch EV
568 charging infrastructure company. We collect a dataset containing transaction records from ElaadNL
569 charging stations in the Netherlands from January to June of 2019. Each transaction record contains

570 the energy demand, transaction start time and charging time. As the data does not specify the details
571 of battery units, we consider the battery units as a single combine battery by summing up the energy
572 demand within each hour to obtain the hourly energy demand.

573 We use the January to February data as the training dataset, March to April data as the validation
574 dataset for tuning the hyperparameters such as learning rate, and May to June as the testing dataset.
575 We consider each problem instance as one day ($T = 24$ hours, plus an initial action). Thus, a sliding
576 window of 25 is applied, moving one hour ahead each time, on the raw data to generate 1416 problem
577 instances, where the first demand of each instance is used as the initial action of all the algorithms.
578 We set $b = 10$ and $A = I$ for the cost function in Eqn. (7).

579 All the algorithms use the same ML architecture, when applicable, with the same initialized weights
580 in our experiments for fair comparison. To be consistent with the literature [52, 53], all the ML
581 models are trained offline. Specifically, we use a recurrent neural network (RNN) model that contains
582 2 hidden layers, each with 8 neurons, and implement the model using PyTorch. We train the RNN for
583 140 epochs with a batch size of 50. When the RNN model is trained as a standalone optimizer in a
584 robustification-oblivious manner, the training process takes around 1 minute on a 2020 MacBook
585 Air with 8GB memory and a M1 chipset. When RNN is trained in a robustification-aware manner, it
586 takes around 2 minutes. The testing process is almost instant and takes less than 1 second.

587 B.2 Baseline Algorithms

588 By default, RCL uses a robustification-aware ML model due to the advantage of average cost perfor-
589 mance compared to a robustification-oblivious model. We compare RCL with several representative
590 baseline algorithms as summarized below.

591 • **Offline Optimal Oracle (OPT)**: This is the optimal offline algorithm that has all the contextual
592 information and optimally solves the problem.

593 • **Regularized Online Balanced Descent (ROBD)**: ROBD is the state-of-the-art order-optimal online
594 algorithm with the best-known competitive ratio for our SOCO setting [54, 49]. The parameters of
595 ROBD are all optimally set according to [49]. By default, RCL uses ROBD as its expert for robustness.

596 • **Hitting Cost Minimizer (HitMin)**: HitMin is a special instance of ROBD by setting the parameters
597 such that it greedily minimizing the hitting cost at each time. This minimizer can be empirically
598 effective and hence also used in ROBD as a regularizer.

599 • **Machine Learning Only (ML)**: ML is trained as a standalone optimizer in a robustification-oblivious
600 manner. It does not use robustification during online optimization.

601 • **Expert-Calibrated Learning (EC-L2O)**: It is an ML-augmented algorithm that applies to our SOCO
602 setting by using an ML model to regularize online actions without robustness guarantees [55]. We
603 set its parameters based on the validation dataset to have the optimal average performance with an
604 empirical competitive ratio less than $(1 + \lambda)CR^\pi$.

605 • **RCL with a robustification-oblivious ML model (RCL-0)**: To differentiate the two forms of RCL, we
606 use RCL to refer to RCL with a robustification-aware ML model and RCL-0 for the robustification-
607 oblivious ML model, where “-0” represents robustification-obliviousness.

608 To highlight our key contribution to the SOCO literature, the baseline algorithms we choose are
609 representative of the state-of-the-art expert algorithms, effective heuristics, and ML-augmented
610 algorithms for the SOCO setting we consider. While there are a few other ML-augmented algorithms
611 for SOCO [56, 57, 58], they do not apply to our problem as they consider unsquared switching costs
612 in a metric space and exploit the natural triangular inequality. Adapting them to the squared switching
613 costs is non-trivial.

614 B.3 Results

615 We now present the results of our case study and begin with the case in which the hitting cost function
616 (parameterized by y_t) is immediately known without feedback delay. The results for the case with
617 feedback delay are presented in Section B.4. Throughout the discussion, the reported values are
618 normalized with respect to those of the respective OPT. The average cost (**AVG**) and competitive
619 ratio (**CR**) are all empirical results reported on the testing dataset.

	RCL				RCL-0				ML	EC-L2O	ROBD	HitMin
	$\lambda=0.6$	$\lambda=1$	$\lambda=3$	$\lambda=5$	$\lambda=0.6$	$\lambda=1$	$\lambda=3$	$\lambda=5$				
AVG	1.4704	1.1144	1.0531	1.0441	1.4780	1.2432	1.0855	1.0738	1.0668	1.1727	1.6048	1.2003
CR	1.7672	1.2905	1.4405	1.3014	2.2103	2.4209	2.4200	3.0322	3.2566	2.0614	1.7291	2.0865

Table 1: Competitive ratio and average cost comparison of different algorithms.

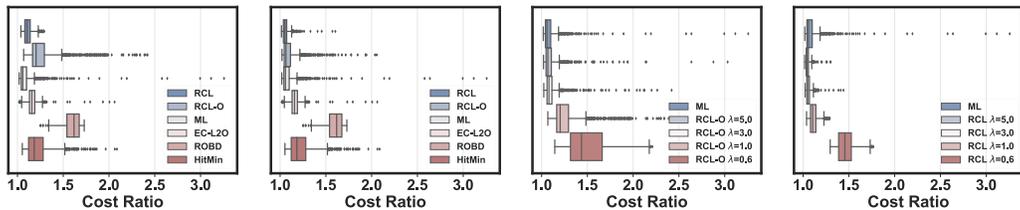
620 By Theorem 4.1, there is a trade-off (governed by $\lambda > 0$) between exploiting ML predictions for good
621 average performance and following the expert for robustness. Here, we begin with the default setting
622 of $\lambda = 1$ and investigate the impact of different choices of λ on both RCL and RCL-0 in Section B.3.3.

623 B.3.1 The performance of RCL

624 As shown in Table I, with $\lambda = 1$, both RCL and RCL-0 have a good average cost, but RCL has a
625 lower average cost than RCL-0 and is outperformed only by ML in terms of the average cost. RCL
626 and RCL-0 have the same competitive ratio (i.e., $(1 + \lambda)$ times the competitive ratio of ROBD).
627 Empirically, RCL has the lowest competitive ratio than all the other algorithms, demonstrating the
628 practical power of RCL for robustifying, potentially untrusted, ML predictions. In this experiment,
629 RCL outperforms ROBD in terms of the empirical competitive ratio because it exploits the good
630 ML predictions for those problem instances that are adversarial to ROBD. This result complements
631 Theorem 4.1, where we show theoretically that RCL can outperform ROBD in terms of the cost by
632 properly setting λ .

633 By comparison, ML performs well on average by exploiting the historical data, but has the highest
634 competitive ratio due to its expected lack of robustness. The two alternative baselines, EC-L2O and
635 HitMin, are empirically good on average and also in the worst case, but they do not have guaranteed
636 robustness. On the other hand, ROBD is very robust, but its average cost is also the worst among all
637 the algorithms under consideration.

638 We further show in Fig. 2(a) the box plots for cost ratios with $\lambda = 1$, providing a detailed view of
639 the algorithms' performance. The key message is that RCL obtains the best of both worlds — a good
640 average cost and a good competitive ratio (empirically even better than the expert ROBD).



(a) ROBD as the expert (b) HitMin as the expert (c) RCL-0 w/ different λ (d) RCL w/ different λ

Figure 2: Cost ratio distributions ($\lambda = 1$ by default).

641 B.3.2 Utilizing HitMin as the expert

642 RCL is flexible and can work with any expert online algorithm, even an expert that does not have good
643 or bounded competitive ratios. Thus, it is interesting to see how RCL performs given an alternative
644 expert. For example, in Table I, HitMin empirically outperforms ROBD in terms of the average,
645 although it is not as robust as ROBD. Thus, using $\lambda = 1$, we leverage HitMin as the expert for RCL
646 and RCL-0, and show the cost ratio distributions in Fig. 2(b). Comparing Fig. 2(b) with Fig. 2(a),
647 we see that RCL and RCL-0 both have many low cost ratios by using HitMin as the expert, but the
648 worst case for RCL is not as good as when using ROBD as the expert. For example, the average cost
649 and competitive ratio are 1.0515 and 1.6035, respectively, for RCL. This result is not surprising, as
650 the new expert HitMin has a better average performance but worse competitive ratio than the default
651 expert ROBD.

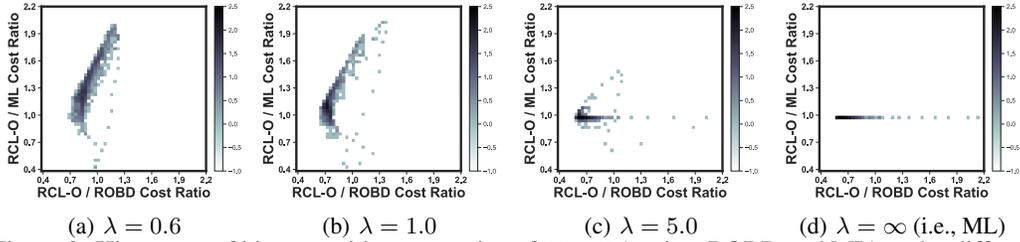


Figure 3: Histogram of bi-competitive cost ratios of RCL-0 (against ROBD and ML) under different λ . For better visualization, the color map represents logarithmic values of the cost ratio histogram with a base of 10.

652 B.3.3 Impact of λ

653 Theorem 4.1 shows the point that we need to set a large enough λ in order to provide enough flexibility
 654 for RCL to exploit good ML predictions. With a small $\lambda > 0$, despite the stronger competitiveness
 655 against the expert, it is possible that RCL may even empirically perform worse than both the ML
 656 model and the expert. Thus, we now investigate the impact of λ .

657 We see from Table 1 that the empirical average cost and competitive ratio of RCL are both worse with
 658 $\lambda = 0.6$ than with the default $\lambda = 1$. More interestingly, by setting $\lambda = 5$, the average cost of RCL is
 659 even lower than that of ML. This is because ML in our experiment performs fairly well on average.
 660 Thus, by setting a large $\lambda = 5$, RCL is able to exploit the benefits of good ML predictions for many
 661 typical cases, while using the expert ROBD as a safeguard to handle a few bad problem instances for
 662 which ML cannot perform well. Also, the empirical competitive ratio of RCL is better with $\lambda = 5$
 663 than with $\lambda = 3$, supporting Theorem 4.1 that a larger λ may not necessarily increase the competitive
 664 ratio as RCL can exploit good ML predictions. In addition, given each λ , RCL outperforms RCL-0,
 665 which highlights the importance of training the ML model in a robustification-aware manner to avoid
 666 the mismatch between training and testing objectives.

667 We also show in Fig. 2(c) and Fig. 2(d) the cost ratio distributions for RCL-0 and RCL, respectively,
 668 under different λ . The results reaffirm our main Theorem 4.1 as well as the importance of training
 669 the ML model in a robustification-aware manner.

670 Next, we show the bi-competitive cost ratios of RCL-0 against both the expert ROBD and the ML pre-
 671 dictions. We focus on RCL-0 as its ML model is trained as a standalone optimizer, whereas RCL uses
 672 a robustification-aware ML model that is not specifically trained to produce good pre-robustification
 673 predictions. According to Theorem 4.1, RCL-0 obtains a potentially better competitiveness against
 674 ML but a worse competitive against the expert ROBD when λ increases, and vice versa. To further
 675 validate the theoretical analysis, we test RCL-0 with different λ and obtain the 2D histogram of its
 676 bi-competitive cost ratios against ROBD and ML, respectively. The results are shown in Fig. 3. In
 677 agreement with our analysis, the cost ratio of RCL-0 against ROBD never exceeds $(1 + \lambda)$ for any
 678 $\lambda > 0$. Also, with a small $\lambda = 0.6$, the cost ratio of RCL-0 against ROBD concentrates around 1,
 679 while it does not exploit the benefits of ML predictions very well. On the other hand, with a large
 680 $\lambda = 5$, the cost ratio of RCL-0 against ROBD can be quite high, although it follows (good) ML
 681 predictions more closely for better average performance. Most importantly, by increasing $\lambda > 0$, we
 682 can see the general trend that RCL-0 follows the ML predictions more closely while still being able
 683 to guarantee competitiveness against ROBD. Again, this confirms the key point of our main insights
 684 in Theorem 4.1.

685 B.3.4 Larger distributional shifts

686 In our dataset, ML performs very well on average as the testing distribution matches well with its
 687 training distribution. To consider more challenging cases as a stress test, we manually increase the
 688 testing distributional shifts by adding random noise following $\mathcal{N}(0, \sigma)$ to a certain fraction p_c of the
 689 testing samples. Note that, as we intentionally stress test RCL and RCL-0 under a larger distributional
 690 shift, their ML models remain unchanged as in the default setting and are not re-trained by adding
 691 noisy data to the training dataset.

		$p_c=0.05$			$p_c=0.1$			$p_c=0.2$		
		$\sigma=0.06$	$\sigma=0.08$	$\sigma=0.1$	$\sigma=0.06$	$\sigma=0.08$	$\sigma=0.1$	$\sigma=0.06$	$\sigma=0.08$	$\sigma=0.1$
AVG	RCL	1.1331	1.1444	1.1556	1.1487	1.1693	1.1904	1.1827	1.2254	1.2697
	RCL-0	1.2425	1.2436	1.2462	1.2416	1.2434	1.2478	1.2370	1.2394	1.2469
	ML	1.0722	1.0778	1.0855	1.0770	1.0874	1.1018	1.0858	1.1053	1.1325
	EC-L2O	1.1728	1.1737	1.1754	1.1731	1.1750	1.1784	1.1727	1.1757	1.1815
	ROBD	1.6048	1.6048	1.6048	1.6048	1.6049	1.6049	1.6048	1.6048	1.6049
	HitMin	1.2112	1.2195	1.2302	1.2202	1.2357	1.2557	1.2410	1.2724	1.3127
CR	RCL	2.5028	2.9697	3.2247	2.6553	3.0283	3.2711	2.5714	3.0123	3.1653
	RCL-0	2.4209	2.4209	2.4209	2.4209	2.4209	2.4209	2.4209	2.4209	2.4209
	ML	6.5159	8.9245	11.6627	4.4025	6.5090	9.4168	5.5798	7.3956	9.3903
	EC-L2O	3.4639	4.6034	5.9666	2.6545	3.6740	5.1129	2.9766	3.7713	4.6983
	ROBD	1.7291	1.7291	1.7291	1.7291	1.7291	1.7291	1.7291	1.7291	1.7298
	HitMin	4.8573	6.7746	8.8383	3.1492	4.8253	7.0405	5.0632	6.9699	8.9246

Table 2: Average cost and competitive ratio comparison of different algorithms. We study the effect of introducing out-of-distribution (OOD) samples. Within the testing dataset, we randomly select a fraction of p_c of samples and add some random noise following $\mathcal{N}(0, \sigma)$ to contaminate these data samples (whose input values are all normalized within $[0, 1]$).

692 With the default $\lambda = 1$, we show the average cost and competitive ratio results in Table 2. We see
693 that ROBD is very robust and little affected by the distributional shifts. In terms of the competitive
694 ratio, ML, HitMin and EC-L2O are not robust, resulting in a large competitive ratio when we add
695 more noisy samples. The average cost performance of RCL is empirically better than that of RCL-0
696 in almost all cases, except for a slight increase in the practically very rare case where 20% samples
697 are contaminated with large noise. On the other hand, as expected, the competitive ratios of RCL and
698 RCL-0 both increase as we add more noise. While RCL has a higher competitive ratio than RCL-0
699 empirically in the experiment, they both have the same guaranteed $(1 + \lambda)$ competitiveness against
700 ROBD regardless of how their ML models are trained. Also, their competitive ratios are both better
701 than other algorithms, showing the effectiveness of our novel robustification process.

702 B.4 Results with Feedback Delay

703 We now turn to the case when there is a one-step feedback delay, i.e., the context parameter y_t is not
704 known to the agent until time $t + 1$. For this setting, we consider the best-known online algorithm
705 iROBD [49] as the expert that handles the feedback delay with a guaranteed competitive ratio with
706 respect to OPT. The other baseline online algorithms — ROBD, EC-L2O, and HitMin— presented in
707 Section B.2 require the immediate revelation of y_t without feedback delay and hence do not directly
708 apply to this case. Thus, for comparison, we use the predicted context, denoted by \hat{y}_t , with up to 15%
709 prediction errors in the baseline online algorithms, and reuse the algorithm names (e.g., EC-L2O
710 uses predicted \hat{y}_t as if it were the true context for decision making). We train ML using the same
711 architecture as in Section B.3 with the exception that only delayed context is provided as input for
712 both training and testing. The reported values are normalized with respect to those of the respective
713 offline optimal algorithm OPT. The average cost (AVG) and competitive ratio (CR) are all empirical
714 results reported on the testing dataset.

715 We show the results in Table 3 and Fig. 4. We see that with the default $\lambda = 1$, both RCL and RCL-0
716 have a good average cost, but RCL has a lower average cost than RCL-0 and is outperformed only
717 by ML in terms of the average cost. RCL and RCL-0 have the same competitive ratio guarantee (i.e.,
718 $(1 + \lambda)$ times the competitive ratio of iROBD). Nonetheless, RCL has the lowest competitive ratio
719 than all the other algorithms, demonstrating the power of RCL to leverage both ML prediction and the
720 robust expert. In this experiment, both RCL and RCL-0 outperform iROBD in terms of the empirical
721 competitive ratio because they are able to exploit the good ML predictions for those problem instances
722 that are difficult for iROBD.

723 By comparison, ML performs well on average by exploiting the historical data, but has a high
724 competitive ratio. The alternative baselines — ROBD, EC-L2O and HitMin— use predicted context
725 \hat{y}_t as the true context. Except for the good empirical competitive ratio of ROBD, they do not have
726 good average performance or guaranteed robustness due to their naively trusting the predicted context
727 (that can potentially have large prediction errors). Note that the empirical competitive ratio of ROBD
728 with predicted context is still much higher than that with the true context in Table 1. These results
729 reinforce the point that blindly using ML predictions (i.e., predicted context in this example) without

	RCL				RCL-0				ML	EC-L2O	iROBD	HitMin	ROBD
	$\lambda=0.6$	$\lambda=1$	$\lambda=3$	$\lambda=5$	$\lambda=0.6$	$\lambda=1$	$\lambda=3$	$\lambda=5$					
AVG	1.5011	1.3594	1.2874	1.2899	1.5134	1.3690	1.2949	1.3026	1.2792	1.4112	2.3076	2.6095	2.5974
CR	2.9797	2.4832	3.2049	3.9847	2.9797	2.4832	3.3367	4.3040	8.4200	15.1928	4.7632	26.0264	2.8478

Table 3: Competitive ratio and average cost comparison of different algorithms with feedback delay.

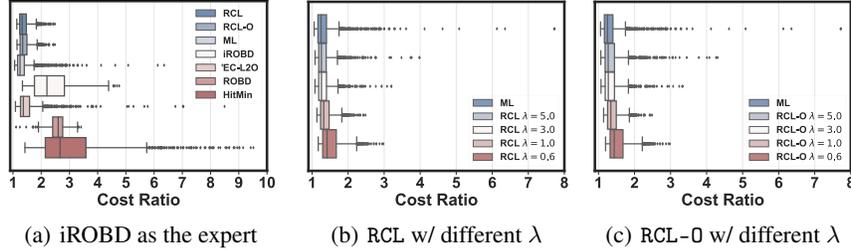


Figure 4: Cost ratio distributions with feedback delay ($\lambda = 1$ by default)

730 additional robustification can lead to poor performance in terms of both average cost and worst-case
731 cost ratio.

732 We further show in Fig. 4 the box plots for cost ratios of different algorithms, providing a detailed
733 view of the algorithms' performance. The key message is that RCL obtains the best of both worlds —
734 a good average cost and a good competitive ratio. Moreover, we see that by setting $\lambda = 1$, we provide
735 enough freedom to RCL to exploit the benefits of ML predictions while also ensuring worst-case
736 robustness. Thus, like in the no-delay case in Table 1 and Fig. 2, the empirical competitive ratio of
737 RCL with $\lambda = 1$ is even lower than that with $\lambda = 0.6$.

738 C Proof of Theorems and Corollaries in Section 4

739 C.1 Proof of Theorem 4.1 (Cost Ratio)

740 To prove Theorem 4.1, we first give some technical lemmas about the smoothness of cost functions
741 from Lemma C.1 to Lemma C.3.

Lemma C.1 (Lemma 4 in [59]). *Assume $f(x)$ is β smooth, for any $\lambda > 0$, we have*

$$f(x) \leq (1 + \lambda)f(y) + \left(1 + \frac{1}{\lambda}\right)\frac{\beta}{2}\|x - y\|^2 \quad \forall x, y \in \mathcal{X}$$

742 **Lemma C.2.** *Assume $f(x)$ is β_1 smooth and $d(x)$ is β_2 smooth, then $f(x) + d(x)$ is $\beta_1 + \beta_2$ smooth.*

743 **Lemma C.3.** *Suppose that the hitting cost $f(x, y_t)$ is β_h -smooth with respect to x , The switching
744 cost is $d(x_t, x_{t-1}) = \frac{1}{2}\|x_t - \delta(x_{t-p:t-1})\|^2$, where $\delta(\cdot)$ is L_i -Lipschitz with respect to x_{t-i} . Then
745 for any two action sequences $x_{1:T}$ and $x'_{1:T}$, we must have*

$$\text{cost}(x_{1:T}) - (1 + \lambda)\text{cost}(x'_{1:T}) \leq \frac{\beta + (1 + \sum_{k=1}^p L_k)^2}{2} \left(1 + \frac{1}{\lambda}\right) \|x_{1:T} - x'_{1:T}\|^2, \quad \forall \lambda > 0 \quad (8)$$

746 *Proof.* The objective to be bounded can be decomposed as

$$\begin{aligned} & \text{cost}(x_{1:T}) - (1 + \lambda)\text{cost}(x'_{1:T}) \\ &= \left(\sum_{t=1}^T f(x_t, y_t) - (1 + \lambda)f(x'_t, y_t) \right) + \frac{1}{2} \left(\sum_{t=1}^T \|x_t - \delta(x_{t-p:t-1})\|^2 - (1 + \lambda)\|x'_t - \delta(x'_{t-p:t-1})\|^2 \right) \end{aligned} \quad (9)$$

747 Since hitting cost is β_h -smooth, then

$$\sum_{t=1}^T f(x_t, y_t) - (1 + \lambda)f(x'_t, y_t) \leq \frac{\beta_h}{2} \left(1 + \frac{1}{\lambda}\right) \sum_{t=1}^T \|x_t - x'_t\|^2 \quad (10)$$

748 Besides, based on the Lipschitz assumption of function $\delta(\cdot)$, we have

$$\begin{aligned}
& \|x_t - \delta(x_{t-p:t-1})\|^2 - (1 + \lambda)\|x'_t - \delta(x'_{t-p:t-1})\|^2 \\
& \leq (1 + \frac{1}{\lambda})\|(x_t - x'_t) + (\delta(x_{t-p:t-1}) - \delta(x'_{t-p:t-1}))\|^2 \\
& \leq (1 + \frac{1}{\lambda})\left(\|x_t - x'_t\| + \|\delta(x_{t-p:t-1}) - \delta(x'_{t-p:t-1})\|\right)^2 \\
& \leq (1 + \frac{1}{\lambda})\left(\|x_t - x'_t\| + \sum_{k=1}^p L_k\|x_{t-k} - x'_{t-k}\|\right)^2 \\
& \leq (1 + \frac{1}{\lambda})\left(1 + \sum_{k=1}^p L_k\right)\left(\|x_t - x'_t\|^2 + \sum_{k=1}^p L_k\|x_{t-k} - x'_{t-k}\|^2\right)
\end{aligned} \tag{11}$$

749 Summing up the switching costs of all time steps together, we have

$$\begin{aligned}
& \sum_{t=1}^T \|x_t - \delta(x_{t-p:t-1})\|^2 - (1 + \lambda)\|x'_t - \delta(x'_{t-p:t-1})\|^2 \\
& \leq (1 + \frac{1}{\lambda})\left(1 + \sum_{k=1}^p L_k\right)\sum_{t=1}^T\left(\|x_t - x'_t\|^2 + \sum_{k=1}^p L_k\|x_{t-k} - x'_{t-k}\|^2\right) \\
& \leq (1 + \frac{1}{\lambda})\left(1 + \sum_{k=1}^p L_k\right)\sum_{t=1}^T\left(1 + \sum_{k=1}^p L_k\right)\|x_t - x'_t\|^2 \\
& = (1 + \frac{1}{\lambda})\left(1 + \sum_{k=1}^p L_k\right)^2\sum_{t=1}^T\|x_t - x'_t\|^2
\end{aligned} \tag{12}$$

750 Substituting Eqn. (12) and Eqn. (10) into Eqn. (9), we finish the proof. \square

751 Now we propose Lemma C.4 based on these above lemmas, which ensures the feasibility of robustness
752 constraint in Eqn. (1)

753 **Lemma C.4.** *Let π be any expert algorithm for the SOCO problem with multi-step feedback delays
754 and multi-step switching costs, for any $\lambda \geq 0$ and $\lambda \geq \lambda_0 \geq 0$, the total cost by the projected actions
755 x_t must satisfy $\text{cost}(x_{1:T}) \leq (1 + \lambda)\text{cost}(x_{1:T}^\pi)$*

756 *Proof.* We prove by induction that the constraints in Eqn. (1) are satisfied for each t . For $t = 1$, since
757 we assume the initial actions are the same ($x_{-p+1:0} = x_{-p+1:0}^\pi$), it is obvious that $x = x_1^\pi$ satisfies
758 the robustness constraints Eqn. (1).

759 Then for any time step $t \geq 2$, suppose it holds at $t - 1$ that

$$\begin{aligned}
& \sum_{\tau \in \mathcal{A}_{t-1}} f(x_\tau, y_\tau) + \sum_{\tau \in \mathcal{A}_{t-1} \cup \mathcal{B}_{t-1}} d(x_\tau, x_{\tau-p:\tau-1}) + \sum_{\tau \in \mathcal{B}_{t-1}} H(x_\tau, x_\tau^\pi) + G(x, x_{t-p:t-1}, x_{t-p:t}^\pi) \\
& \leq (1 + \lambda)\left(\sum_{\tau \in \mathcal{A}_{t-1}} f(x_\tau^\pi, y_\tau) + \sum_{\tau \in \mathcal{A}_{t-1} \cup \mathcal{B}_{t-1}} d(x_\tau^\pi, x_{\tau-p:\tau-1}^\pi)\right)
\end{aligned} \tag{13}$$

760 Now the robustness constraints Eqn. (1) is satisfied if we prove $x_t = x_t^\pi$ satisfies the constraints in
761 Eqn. (1) at time step t . Since for the sets \mathcal{A} and \mathcal{B} , we have

$$(\mathcal{A}_t \cup \mathcal{B}_t) \setminus (\mathcal{A}_{t-1} \cup \mathcal{B}_{t-1}) = \{t\}, \quad \mathcal{A}_{t-1} \subseteq \mathcal{A}_t, \tag{14}$$

762 so it holds that

$$\sum_{\tau \in \mathcal{A}_t \cup \mathcal{B}_t} d(x_\tau, x_{\tau-p:\tau-1}) - \sum_{\tau \in \mathcal{A}_{t-1} \cup \mathcal{B}_{t-1}} d(x_\tau, x_{\tau-p:\tau-1}) = d(x_t, x_{t-p:t-1}) \tag{15}$$

763 By Lemma [C.1](#), we have

$$\begin{aligned}
& d(x_t^\pi, x_{t-p:t-1}) - (1 + \lambda)d(x_t^\pi, x_{t-p:t-1}^\pi) \\
& \leq \frac{1}{2} \left(1 + \frac{1}{\lambda}\right) \|\delta(x_{t-p:t-1}) - \delta(x_{t-p:t-1}^\pi)\|^2 \\
& \leq \frac{1}{2} \left(1 + \frac{1}{\lambda}\right) \left(\sum_{i=1}^p L_i \|x_{t-i} - x_{t-i}^\pi\|\right)^2
\end{aligned} \tag{16}$$

764 Denote $\alpha = 1 + \sum_{k=1}^p L_k$. For the reservation cost, we have

$$\begin{aligned}
& G(x_{t-1}, x_{t-p-1:t-2}, x_{t-p-1:t-1}^\pi) - G(x_t^\pi, x_{t-p:t-1}, x_{t-p:t}^\pi) \\
& = \frac{\alpha(1 + \frac{1}{\lambda_0})}{2} \left(\sum_{k=1}^p \sum_{i=0}^{p-k} L_{k+i} \|x_{t-i-1} - x_{t-i-1}^\pi\|^2 - \sum_{k=1}^p \sum_{i=1}^{p-k} L_{k+i} \|x_{t-i} - x_{t-i}^\pi\|^2 \right) \\
& = \frac{\alpha(1 + \frac{1}{\lambda_0})}{2} \left(\sum_{k=0}^{p-1} \sum_{i=1}^{p-k} L_{k+i} \|x_{t-i} - x_{t-i}^\pi\|^2 - \sum_{k=1}^p \sum_{i=1}^{p-k} L_{k+i} \|x_{t-i} - x_{t-i}^\pi\|^2 \right) \\
& = \frac{\alpha(1 + \frac{1}{\lambda_0})}{2} \left(\sum_{k=0}^{p-1} \sum_{i=1}^{p-k} L_{k+i} \|x_{t-i} - x_{t-i}^\pi\|^2 - \sum_{k=1}^{p-1} \sum_{i=1}^{p-k} L_{k+i} \|x_{t-i} - x_{t-i}^\pi\|^2 \right) \\
& = \frac{\alpha(1 + \frac{1}{\lambda_0})}{2} \sum_{i=1}^p L_i \|x_{t-i} - x_{t-i}^\pi\|^2
\end{aligned} \tag{17}$$

765 Continuing with Eqn. [\(17\)](#), we have

$$\begin{aligned}
& G(x_{t-1}, x_{t-p-1:t-2}, x_{t-p-1:t-1}^\pi) - G(x_t^\pi, x_{t-p:t-1}, x_{t-p:t}^\pi) = \frac{\alpha(1 + \frac{1}{\lambda_0})}{2} \sum_{i=1}^p L_i \|x_{t-i} - x_{t-i}^\pi\|^2 \\
& \geq \frac{(1 + \frac{1}{\lambda_0})(\sum_{i=1}^p L_i)^2}{2} \sum_{i=1}^p \frac{L_i}{\sum_{i=1}^p L_i} \|x_{t-i} - x_{t-i}^\pi\|^2 \\
& \geq \frac{(1 + \frac{1}{\lambda_0})(\sum_{i=1}^p L_i)^2}{2} \left(\sum_{i=1}^p \frac{L_i}{\sum_{i=1}^p L_i} \|x_{t-i} - x_{t-i}^\pi\| \right)^2 \\
& = \frac{1}{2} \left(1 + \frac{1}{\lambda_0}\right) \left(\sum_{i=1}^p L_i \|x_{t-i} - x_{t-i}^\pi\| \right)^2 \geq \frac{1}{2} \left(1 + \frac{1}{\lambda}\right) \left(\sum_{i=1}^p L_i \|x_{t-i} - x_{t-i}^\pi\| \right)^2
\end{aligned} \tag{18}$$

766 where the second inequality holds by Jensen's inequality. Therefore, combining with [\(16\)](#), we have

$$d(x_t^\pi, x_{t-p:t-1}) + G(x_t^\pi, x_{t-p:t-1}, x_{t-p:t}^\pi) \leq G(x_{t-1}, x_{t-p-1:t-2}, x_{t-p-1:t-1}^\pi) + (1 + \lambda)d(x_t^\pi, x_{t-p:t-1}^\pi) \tag{19}$$

767 By Eqn. [\(19\)](#), we have

$$\begin{aligned}
& G(x_t^\pi, x_{t-p:t-1}, x_{t-p:t}^\pi) + \sum_{\tau \in \mathcal{A}_t \cup \mathcal{B}_t} d(x_\tau, x_{\tau-p:\tau-1}) - \sum_{\tau \in \mathcal{A}_{t-1} \cup \mathcal{B}_{t-1}} d(x_\tau, x_{\tau-p:\tau-1}) \\
& \leq G(x_{t-1}, x_{t-p-1:t-2}, x_{t-p-1:t-1}^\pi) + (1 + \lambda) \left(\sum_{\tau \in \mathcal{A}_t \cup \mathcal{B}_t} d(x_\tau^\pi, x_{\tau-p:\tau-1}^\pi) - \sum_{\tau \in \mathcal{A}_{t-1} \cup \mathcal{B}_{t-1}} d(x_\tau^\pi, x_{\tau-p:\tau-1}^\pi) \right)
\end{aligned} \tag{20}$$

768 Now we define a new set $\mathcal{D}_t = \mathcal{A}_t \setminus \mathcal{A}_{t-1}$, which denotes the timestep set for the newly received
769 context parameters at t .

770 **Case 1:** If $t \in \mathcal{D}_t$, then $B_{t-1} \setminus B_t = \mathcal{D}_t \setminus \{t\}$, then we have

$$\begin{aligned} & \left(\sum_{\tau \in \mathcal{A}_t} f(x_\tau, y_\tau) + \sum_{\tau \in \mathcal{B}_t} H(x_\tau, x_\tau^\pi) \right) - \left(\sum_{\tau \in \mathcal{A}_{t-1}} f(x_\tau, y_\tau) + \sum_{\tau \in \mathcal{B}_{t-1}} H(x_\tau, x_\tau^\pi) \right) \\ &= \sum_{\tau \in \mathcal{D}_t} f(x_\tau, y_\tau) - \sum_{\tau \in \mathcal{D}_t \setminus \{t\}} H(x_\tau, x_\tau^\pi) = f(x_t^\pi, y_t) + \sum_{\tau \in \mathcal{D}_t \setminus \{t\}} f(x_\tau, y_\tau) - \sum_{\tau \in \mathcal{D}_t \setminus \{t\}} H(x_\tau, x_\tau^\pi) \end{aligned} \quad (21)$$

771 Since hitting cost $f(\cdot, y_t)$ is β_h -smooth, we have

$$\begin{aligned} & \sum_{\tau \in \mathcal{D}_t \setminus \{t\}} f(x_\tau, y_\tau) - \sum_{\tau \in \mathcal{D}_t \setminus \{t\}} (1 + \lambda) f(x_\tau^\pi, y_\tau) \\ & \leq \frac{\beta_h(1 + \frac{1}{\lambda})}{2} \sum_{\tau \in \mathcal{D}_t \setminus \{t\}} \|x_\tau^\pi - x_\tau\|^2 \leq \sum_{\tau \in \mathcal{D}_t \setminus \{t\}} H(x_\tau, x_\tau^\pi) \end{aligned} \quad (22)$$

772 Substituting Eqn. (22) back to Eqn. (21), we have

$$\begin{aligned} & \left(\sum_{\tau \in \mathcal{A}_t} f(x_\tau, y_\tau) + \sum_{\tau \in \mathcal{B}_t} H(x_\tau, x_\tau^\pi) \right) - \left(\sum_{\tau \in \mathcal{A}_{t-1}} f(x_\tau, y_\tau) + \sum_{\tau \in \mathcal{B}_{t-1}} H(x_\tau, x_\tau^\pi) \right) \\ & \leq (1 + \lambda) \left(\sum_{\tau \in \mathcal{A}_t} f(x_\tau, y_\tau) - \sum_{\tau \in \mathcal{A}_{t-1}} f(x_\tau, y_\tau) \right) \end{aligned} \quad (23)$$

773 **Case 2:** If $t \notin \mathcal{D}_t$, then $(B_{t-1} \cup \{t\}) \setminus B_t = \mathcal{D}_t$ and we have

$$\begin{aligned} & \left(\sum_{\tau \in \mathcal{A}_t} f(x_\tau, y_\tau) + \sum_{\tau \in \mathcal{B}_t} H(x_\tau, x_\tau^\pi) \right) - \left(\sum_{\tau \in \mathcal{A}_{t-1}} f(x_\tau, y_\tau) + \sum_{\tau \in \mathcal{B}_{t-1}} H(x_\tau, x_\tau^\pi) \right) \\ &= \sum_{\tau \in \mathcal{D}_t} f(x_\tau, y_\tau) - \sum_{\tau \in \mathcal{D}_t} H(x_\tau, x_\tau^\pi) + H(x_t^\pi, x_t^\pi) \\ &= \sum_{\tau \in \mathcal{D}_t} f(x_\tau, y_\tau) - \sum_{\tau \in \mathcal{D}_t} H(x_\tau, x_\tau^\pi) \end{aligned} \quad (24)$$

774 Since hitting cost $f(\cdot, y_t)$ is β_h -smooth, we have

$$\sum_{\tau \in \mathcal{D}_t} f(x_\tau, y_\tau) - \sum_{\tau \in \mathcal{D}_t} (1 + \lambda) f(x_\tau^\pi, y_\tau) \leq \frac{\beta_h(1 + \frac{1}{\lambda})}{2} \sum_{\tau \in \mathcal{D}_t} \|x_\tau^\pi - x_\tau\|^2 \leq \sum_{\tau \in \mathcal{D}_t} H(x_\tau, x_\tau^\pi) \quad (25)$$

775 Since $\lambda \geq 0$, we substitute Eqn. (25) back to Eqn. (24), we have the same conclusion as Eqn. (23).

776 Adding Eqn. (13), Eqn. (20) and Eqn. (23) together, we can prove $x = x_t^\pi$ satisfies the constraints in

777 Eqn. (1). At time step T , we have

$$\begin{aligned} & \sum_{\tau \in \mathcal{A}_T} f(x_\tau, y_\tau) + \sum_{\tau \in \mathcal{A}_T \cup \mathcal{B}_T} d(x_\tau, x_{\tau-p:\tau-1}) + \sum_{\tau \in \mathcal{B}_T} (f(x_\tau, y_\tau) - (1 + \lambda) f(x_\tau^\pi, y_\tau)) \\ & \leq \sum_{\tau \in \mathcal{A}_T} f(x_\tau, y_\tau) + \sum_{\tau \in \mathcal{A}_T \cup \mathcal{B}_T} d(x_\tau, x_{\tau-p:\tau-1}) + \sum_{\tau \in \mathcal{B}_T} H(x_\tau, x_\tau^\pi) \\ & \leq (1 + \lambda) \left(\sum_{\tau \in \mathcal{A}_T} f(x_\tau^\pi, y_\tau) + \sum_{\tau \in \mathcal{A}_T \cup \mathcal{B}_T} d(x_\tau^\pi, x_{\tau-p:\tau-1}^\pi) \right) \end{aligned} \quad (26)$$

778 In other words

$$\sum_{\tau \in \mathcal{A}_T \cup \mathcal{B}_T} (f(x_\tau, y_\tau) + d(x_\tau, x_{\tau-p:\tau-1})) \leq (1 + \lambda) \sum_{\tau \in \mathcal{A}_T \cup \mathcal{B}_T} (f(x_\tau^\pi, y_\tau) + d(x_\tau, x_{\tau-p:\tau-1})) \quad (27)$$

779

□

780 In the next lemma, we bound the difference between the projected action and the ML predictions.

781 **Lemma C.5.** *Suppose hitting cost is β_h -smooth, given the expert policy π , ML predictions $\tilde{x}_{1:T}$, for*
 782 *any $\lambda > 0$ and $\lambda_1 > 0$, the total distance between actual actions $x_{1:T}$ and ML predictions $\tilde{x}_{1:T}$ are*
 783 *bounded,*

$$\sum_{i=1}^T \|x_t - \tilde{x}_t\|^2 \leq \sum_{i=1}^T \left(\left[\|\tilde{x}_t - x_t^\pi\| - \sqrt{K \left(d(x_t^\pi, x_{t-p:t-1}^\pi) + \sum_{\tau \in \mathcal{D}_t} f(x_\tau^\pi, y_\tau) \right)} \right]^+ \right)^2 \quad (28)$$

784 where $[\cdot]^+$ is the ReLU function and $K = \frac{2(\lambda - \lambda_0)}{\beta_h(1 + \frac{1}{\lambda_0}) + \alpha^2(1 + \frac{1}{\lambda_0})}$, $\alpha = 1 + \sum_{i=1}^p L_i$

785 *Proof.* Suppose we at $t - 1$ have the following inequality:

$$\begin{aligned} & \sum_{\tau \in \mathcal{A}_{t-1}} f(x_\tau, y_\tau) + \sum_{\tau \in \mathcal{A}_{t-1} \cup \mathcal{B}_{t-1}} d(x_\tau, x_{\tau-p:\tau-1}) + \sum_{\tau \in \mathcal{B}_{t-1}} H(x_\tau, x_\tau^\pi) + G(x, x_{t-p:t-1}, x_{t-p:t}^\pi) \\ & \leq (1 + \lambda) \left(\sum_{\tau \in \mathcal{A}_{t-1}} f(x_\tau^\pi, y_\tau) + \sum_{\tau \in \mathcal{A}_{t-1} \cup \mathcal{B}_{t-1}} d(x_\tau^\pi, x_{\tau-p:\tau-1}^\pi) \right) \end{aligned} \quad (29)$$

786 Remember that $\mathcal{D}_t = \mathcal{A}_t \setminus \mathcal{A}_{t-1}$ is the set of the time steps for the newly received context parameters
 787 at t . The robustness constraint in Eqn. (1) is satisfied if x_t satisfies the following inequality.

$$\begin{aligned} & \left(\sum_{\tau \in \mathcal{D}_t} f(x_\tau, y_\tau) + \sum_{\tau \in \mathcal{B}_t} H(x_\tau, x_\tau^\pi) - \sum_{\tau \in \mathcal{B}_{t-1}} H(x_\tau, x_\tau^\pi) \right) + d(x_t, x_{t-p:t-1}) + G(x_t, x_{t-p:t-1}, x_{t-p:t}^\pi) \\ & - G(x_{t-1}, x_{t-p-1:t-2}, x_{t-p-1:t-1}^\pi) \leq (1 + \lambda) \left(d(x_t^\pi, x_{t-p:t-1}^\pi) + \sum_{\tau \in \mathcal{D}_t} f(x_\tau^\pi, y_\tau) \right) \end{aligned} \quad (30)$$

788 For the switching cost, we have

$$\begin{aligned} & d(x, x_{t-p:t-1}) - (1 + \lambda_0)d(x_t^\pi, x_{t-p:t-1}^\pi) \\ & \leq \frac{1}{2} \left(1 + \frac{1}{\lambda_0}\right) \left(\|x - x_t^\pi\| + \|\delta(x_{t-p:t-1}) - \delta(x_{t-p:t-1}^\pi)\| \right)^2 \\ & \leq \frac{1}{2} \left(1 + \frac{1}{\lambda_0}\right) \left(\|x - x_t^\pi\| + \sum_{i=1}^p L_i \|x_{t-i} - x_{t-i}^\pi\| \right)^2 \\ & \leq \frac{\alpha(1 + \frac{1}{\lambda_0})}{2} \left(\|x - x_t^\pi\|^2 + \sum_{i=1}^p L_i \|x_{t-i} - x_{t-i}^\pi\|^2 \right) \end{aligned} \quad (31)$$

789 The first inequality comes from Lemma C.1, the second inequality comes from the L_i -Lipschitz
 790 assumption, and the third inequality is because $\alpha \geq 1$. Besides, from Eqn (17), we have

$$G(x_{t-1}, x_{t-p-1:t-2}, x_{t-p-1:t-1}^\pi) - G(x_t^\pi, x_{t-p:t-1}, x_{t-p:t}^\pi) = \frac{\alpha(1 + \frac{1}{\lambda_0})}{2} \sum_{i=1}^p L_i \|x_{t-i} - x_{t-i}^\pi\|^2 \quad (32)$$

791 Thus we have

$$\begin{aligned} & G(x, x_{t-p:t-1}, x_{t-p:t}^\pi) - G(x_{t-1}, x_{t-p-1:t-2}, x_{t-p-1:t-1}^\pi) \\ & = G(x, x_{t-p:t-1}, x_{t-p:t}^\pi) - G(x_t^\pi, x_{t-p:t-1}, x_{t-p:t}^\pi) + G(x_t^\pi, x_{t-p:t-1}, x_{t-p:t}^\pi) - G(x_{t-1}, x_{t-p-1:t-2}, x_{t-p-1:t-1}^\pi) \\ & = G(x, x_{t-p:t-1}, x_{t-p:t}^\pi) - G(x_t^\pi, x_{t-p:t-1}, x_{t-p:t}^\pi) - \frac{\alpha(1 + \frac{1}{\lambda_0})}{2} \sum_{i=1}^p L_i \|x_{t-i} - x_{t-i}^\pi\|^2. \end{aligned} \quad (33)$$

792 Combining with inequality (31), we have

$$\begin{aligned}
& G(x_t, x_{t-p:t-1}, x_{t-p:t}^\pi) - G(x_{t-1}, x_{t-p-1:t-2}, x_{t-p-1:t-1}^\pi) + d(x_t, x_{t-p:t-1}) - (1 + \lambda_0)d(x_t^\pi, x_{t-p:t-1}^\pi) \\
& \leq G(x_t, x_{t-p:t-1}, x_{t-p:t}^\pi) - G(x_t^\pi, x_{t-p:t-1}, x_{t-p:t}^\pi) + \frac{\alpha(1 + \frac{1}{\lambda_0})}{2} \|x_t - x_t^\pi\|^2 \\
& = \frac{\alpha(1 + \frac{1}{\lambda_0}) \sum_{k=1}^p L_k}{2} \|x_t - x_t^\pi\|^2 + \frac{\alpha(1 + \frac{1}{\lambda_0})}{2} \sum_{k=1}^p \|x_t - x_t^\pi\|^2 \\
& = \frac{\alpha^2(1 + \frac{1}{\lambda_0})}{2} \|x_t - x_t^\pi\|^2
\end{aligned} \tag{34}$$

793 Substituting Eqn. (34) back to Eqn. (30), we have

$$\begin{aligned}
& \sum_{\tau \in \mathcal{D}_t} (f(x_\tau, y_\tau) - (1 + \lambda_0)f(x_\tau^\pi, y_\tau)) + \sum_{\tau \in \mathcal{B}_t} H(x_\tau, x_\tau^\pi) - \sum_{\tau \in \mathcal{B}_{t-1}} H(x_\tau, x_\tau^\pi) \\
& + \frac{\alpha^2(1 + \frac{1}{\lambda_0})}{2} \|x - x_t^\pi\|^2 \leq (\lambda - \lambda_0) \left(d(x_t^\pi, x_{t-p:t-1}^\pi) + \sum_{\tau \in \mathcal{D}_t} f(x_\tau^\pi, y_\tau) \right)
\end{aligned} \tag{35}$$

794 **Case 1:** If $t \in \mathcal{D}_t$, then $B_{t-1} \setminus B_t = \mathcal{D}_t \setminus \{t\}$, then Eqn. (35) becomes

$$\begin{aligned}
& f(x_t, y_t) - (1 + \lambda_0)f(x_t^\pi, y_t) + \frac{\alpha^2(1 + \frac{1}{\lambda_0})}{2} \|x - x_t^\pi\|^2 \\
& + \sum_{\tau \in \mathcal{D}_t \setminus \{t\}} f(x_\tau, y_\tau) - (1 + \lambda_0)f(x_\tau^\pi, y_\tau) - H(x_\tau, x_\tau^\pi) \leq (\lambda - \lambda_0) \left(d(x_t^\pi, x_{t-p:t-1}^\pi) + \sum_{\tau \in \mathcal{D}_t} f(x_\tau^\pi, y_\tau) \right)
\end{aligned} \tag{36}$$

795 Since hitting cost is β_h -smooth, the sufficient condition for Eqn. (35) becomes

$$\frac{(\beta_h + \alpha^2)(1 + \frac{1}{\lambda_0})}{2} \|x - x_t^\pi\|^2 \leq (\lambda - \lambda_0) \left(d(x_t^\pi, x_{t-p:t-1}^\pi) + \sum_{\tau \in \mathcal{D}_t} f(x_\tau^\pi, y_\tau) \right) \tag{37}$$

796 Since the hitting cost is non-negative, the sufficient condition can be further simplified, which is

$$\frac{(\beta_h + \alpha^2)(1 + \frac{1}{\lambda_0})}{2} \|x - x_t^\pi\|^2 \leq (\lambda - \lambda_0) (f(x_t^\pi, y_t) + d(x_t^\pi, x_{t-p:t-1}^\pi)) \tag{38}$$

797 **Case 2:** If $t \notin \mathcal{D}_t$, then $(B_{t-1} \cup \{t\}) \setminus B_t = \mathcal{D}_t$, then Eqn. (35) becomes

$$\begin{aligned}
& \frac{\alpha^2(1 + \frac{1}{\lambda_0})}{2} \|x - x_t^\pi\|^2 + H(x, x_t^\pi) + \sum_{\tau \in \mathcal{D}_t} (f(x_\tau, y_\tau) - (1 + \lambda_0)f(x_\tau^\pi, y_\tau) - H(x_\tau, x_\tau^\pi)) \\
& \leq (\lambda - \lambda_0) \left(d(x_t^\pi, x_{t-p:t-1}^\pi) + \sum_{\tau \in \mathcal{D}_t} f(x_\tau^\pi, y_\tau) \right)
\end{aligned} \tag{39}$$

798 Since hitting cost is β_h -smooth, the sufficient condition for Eqn. (39) becomes

$$\frac{(\beta_h + \alpha^2)(1 + \frac{1}{\lambda_0})}{2} \|x - x_t^\pi\|^2 \leq (\lambda - \lambda_0) \left(d(x_t^\pi, x_{t-p:t-1}^\pi) + \sum_{\tau \in \mathcal{D}_t} f(x_\tau^\pi, y_\tau) \right) \tag{40}$$

Now we define

$$K = \frac{2(\lambda - \lambda_0)}{(\beta_h + \alpha^2)(1 + \frac{1}{\lambda_0})}$$

799 At time step t , if x'_t is the solution to this alternative optimization problem

$$\begin{aligned} x'_t &= \arg \min_x \frac{1}{2} \|x - \tilde{x}_t\|^2 \\ \text{s.t. } \|x - x_t^\pi\|^2 &\leq K \left(d(x_t^\pi, x_{t-p:t-1}^\pi) + \sum_{\tau \in \mathcal{D}_t} f(x_\tau^\pi, y_\tau) \right) \end{aligned} \quad (41)$$

800 The solution to this problem can be calculated asd

$$\begin{aligned} x'_t &= \theta x_t^\pi + (1 - \theta) \tilde{x}_t \\ \theta &= \left[1 - \frac{\sqrt{K \left(d(x_t^\pi, x_{t-p:t-1}^\pi) + \sum_{\tau \in \mathcal{D}_t} f(x_\tau^\pi, y_\tau) \right)}}{\|\tilde{x}_t - x_t^\pi\|} \right]^+ \end{aligned} \quad (42)$$

801 Then $\|x'_t - \tilde{x}_t\| = \left[\|\tilde{x}_t - x_t^\pi\| - \sqrt{K \left(d(x_t^\pi, x_{t-p:t-1}^\pi) + \sum_{\tau \in \mathcal{D}_t} f(x_\tau^\pi, y_\tau) \right)} \right]^+$. Since x'_t also
802 satisfies the original robustness constraint, we have $\|x_t - \tilde{x}_t\| \leq \|x'_t - \tilde{x}_t\|$ and we finish the proof.

803 \square

804 **Proof of Theorem 4.1**

805 Now summing up the distance through 1 to T, we have

$$\sum_{i=1}^T \|x_i - \tilde{x}_i\|^2 \leq \sum_{i=1}^T \left(\left[\|\tilde{x}_i - x_i^\pi\| - \sqrt{K \left(d(x_i^\pi, x_{i-p:i-1}^\pi) + \sum_{\tau \in \mathcal{D}_i} f(x_\tau^\pi, y_\tau) \right)} \right]^+ \right)^2 \quad (43)$$

806 Based on Lemma C.3 we have $\forall \lambda_2 > 0$,

$$\text{cost}(x_{1:T}) - (1 + \lambda_2) \text{cost}(\tilde{x}_{1:T}) \leq \frac{\beta + \alpha^2}{2} \left(1 + \frac{1}{\lambda_2}\right) \sum_{i=1}^T \|x_i - \tilde{x}_i\|^2. \quad (44)$$

807 Suppose the offline optimal action sequence is $x_{1:T}^*$, the optimal cost is $\text{cost}(x_{1:T}^*)$. Then we divide
808 both sides of Eqn. (44) by $\text{cost}(x_{1:T}^*)$, and get $\forall \lambda_2 > 0$,

$$\begin{aligned} \text{cost}(x_{1:T}) &\leq (1 + \lambda_2) \text{cost}(\tilde{x}_{1:T}) + \frac{\beta + \alpha^2}{2} \left(1 + \frac{1}{\lambda_2}\right) \sum_{i=1}^T \left(\left[\|\tilde{x}_i - x_i^\pi\| - \sqrt{K \left(d(x_i^\pi, x_{i-p:i-1}^\pi) + \sum_{\tau \in \mathcal{D}_i} f(x_\tau^\pi, y_\tau) \right)} \right]^+ \right)^2 \end{aligned} \quad (45)$$

809 By substituting $K = \frac{2(\lambda - \lambda_0)}{(\beta_h + \alpha^2)(1 + \frac{1}{\lambda_0})}$ back to Eqn (46), we have

$$\begin{aligned} \text{cost}(x_{1:T}) &\leq (1 + \lambda_2) \text{cost}(\tilde{x}_{1:T}) + \left(1 + \frac{1}{\lambda_2}\right) \sum_{i=1}^T \left[\frac{\beta + \alpha^2}{2} \|\tilde{x}_i - x_i^\pi\|^2 \right. \\ &\quad \left. - \frac{\lambda - \lambda_0}{1 + \frac{1}{\lambda_0}} \left(d(x_i^\pi, x_{i-p:i-1}^\pi) + \sum_{\tau \in \mathcal{D}_i} f(x_\tau^\pi, y_\tau) \right) \right]^+ \end{aligned} \quad (46)$$

810 By defining single step cost of the expert π as $\text{cost}_t^\pi = d(x_t^\pi, x_{t-p:t-1}^\pi) + \sum_{\tau \in \mathcal{D}_t} f(x_\tau^\pi, y_\tau)$ and the
811 auxiliary cost as $\Delta(\lambda) = \sum_{i=1}^T \left[\|\tilde{x}_i - x_i^\pi\|^2 - \frac{2(\lambda - \lambda_0)}{(\beta_h + \alpha^2)(1 + \frac{1}{\lambda_0})} \text{cost}_i^\pi \right]^+$

$$\text{cost}(x_{1:T}) \leq \left(\sqrt{\text{cost}(\tilde{x}_{1:T})} + \sqrt{\frac{\beta + \alpha^2}{2} \Delta(\lambda)} \right)^2 \quad (47)$$

812 Combined with Lemma C.4, we obtain the following bound, which finished this proof.

$$\text{cost}(x_{1:T}) \leq \min \left((1 + \lambda) \text{cost}(x_{1:T}^\pi), \left(\sqrt{\text{cost}(\tilde{x}_{1:T})} + \sqrt{\frac{\beta + \alpha^2}{2} \Delta(\lambda)} \right)^2 \right) \quad (48)$$

813 **C.2 Proof of Theorem 4.2**

814 *Proof.* We first give the formal definition of Rademacher complexity of the ML model space with
815 robustification.

816 **Definition 5** (Rademacher Complexity). *Let $\text{Rob}_\lambda(\mathcal{W}) = \{\text{Rob}_\lambda(h_W), W \in \mathcal{W}\}$ be the ML model
817 space with robustification constrained by (2). Given the dataset \mathcal{S} , the Rademacher complexity with
818 respect to $\text{Rob}_\lambda(\mathcal{W})$ is*

$$\text{Rad}_\mathcal{S}(\text{Rob}_\lambda(\mathcal{W})) = \frac{1}{|\mathcal{S}|} \mathbb{E}_\nu \left[\sup_{W \in \mathcal{W}} \left(\sum_{i \in \mathcal{S}} \nu_i \text{Rob}_\lambda(h_W(y^i)) \right) \right],$$

819 where y^i is the i -th sample in \mathcal{S} , and ν_1, \dots, ν_n are independently drawn from Rademacher distribu-
820 tion.

821 Since the cost functions are smooth, they are locally Lipschitz continuous for the bounded action
822 space, and we can apply the generalization bound based on Rademacher complexity [60] for the
823 space of robustified ML model $\text{Rob}_\lambda(h_W)$. Given any ML model h_W trained on dataset \mathcal{S} , with
824 probability at least $1 - \delta, \delta \in (0, 1)$,

$$\mathbb{E}_{\mathbb{P}_y}[\text{cost}_{1:T}] \leq \overline{\text{cost}}_\mathcal{S}(\text{Rob}_\lambda(h_W)) + 2\Gamma_x \text{Rad}_\mathcal{S}(\text{Rob}_\lambda(\mathcal{W})) + 3\bar{c} \sqrt{\frac{\log(2/\delta)}{|\mathcal{S}|}}, \quad (49)$$

825 where $\Gamma_x = \sqrt{T} |\mathcal{X}| [\beta_h + \frac{1}{2}(1 + \sum_{i=1}^p L_i)(1 + \sum_{i=1}^p L_i)]$ with $|\mathcal{X}|$ being the size of the action
826 space \mathcal{X} and β_h, L_i , and p as the smoothness constant, Lipschitz constant of the nonlinear term in the
827 switching cost, and the memory length as defined in Assumptions 1 and 2, and \bar{c} is the upper bound
828 of the total cost for an episode. We can get the average cost bound in Proposition 4.2.

829 Next, we prove that the Rademacher complexity of the ML model space with robustification is no
830 larger than the Rademacher complexity of the ML model space without robustification expressed as
831 $\{h_W, W \in \mathcal{W}\}$, i.e. we need to prove $\text{Rad}_\mathcal{S}(\text{Rob}_\lambda(\mathcal{W})) \leq \text{Rad}_\mathcal{S}(\mathcal{W})$. The Rademacher complexity
832 can be expressed by Dudley's entropy integral [61] as

$$\text{Rad}_\mathcal{S}(\text{Rob}_\lambda(\mathcal{W})) = \mathcal{O} \left(\frac{1}{\sqrt{|\mathcal{S}|}} \int_0^\infty \sqrt{\log \mathbb{N}(\epsilon, \text{Rob}_\lambda(\mathcal{W}), L_2(\mathcal{S}))} d\epsilon \right), \quad (50)$$

833 where $\mathbb{N}(\epsilon, \text{Rob}_\lambda(\mathcal{W}), L_2(\mathcal{S}))$ is the covering number [61] with respect to radius ϵ and the function
834 distance metric $\|h_1 - h_2\|_{L_2(\mathcal{S})} = \frac{1}{|\mathcal{S}|} \sum_{i \in \mathcal{S}} \|h_1(x_i) - h_2(x_i)\|^2$ where h_1 and h_2 are two functions
835 defined on the space including dataset \mathcal{S} . We can find that for any two different weights W_1 and W_2 ,
836 their corresponding post-robustification distance $\|\text{Rob}_\lambda(h_{W_1}) - \text{Rob}_\lambda(h_{W_2})\|_{L_2(\mathcal{S})}$ is no larger than
837 their pre-robustification distance $\|h_{W_1} - h_{W_2}\|_{L_2(\mathcal{S})}$. To see this, we discuss three cases given any
838 input sample y . If both $h_{W_1}(y)$ and $h_{W_2}(y)$ lie in the projection set, then $\text{Rob}_\lambda(h_{W_1})(y) = h_{W_1}(y)$
839 and $\text{Rob}_\lambda(h_{W_2})(y) = h_{W_2}(y)$. If $h_{W_1}(y)$ lies in the projection set while $h_{W_2}(y)$ is out of
840 the projection set, the projection operation based on the closed convex projection set will make
841 $\|\text{Rob}_\lambda(h_{W_1})(y) - \text{Rob}_\lambda(h_{W_2})(y)\|$ to be less than $\|h_{W_1}(y) - h_{W_2}(y)\|$. If both $h_{W_1}(y)$ and $h_{W_2}(y)$
842 lie out of the projection set, we still have $\|\text{Rob}_\lambda(h_{W_1})(y) - \text{Rob}_\lambda(h_{W_2})(y)\| \leq \|h_{W_1}(y) - h_{W_2}(y)\|$
843 since the projection set at each round is a closed convex set [62]. Therefore, after robustifi-
844 cation, the distance between two models with different weights will not become larger, i.e.
845 $\|\text{Rob}_\lambda(h_{W_1}) - \text{Rob}_\lambda(h_{W_2})\|_{L_2(\mathcal{S})} \leq \|h_{W_1} - h_{W_2}\|_{L_2(\mathcal{S})}$, which means RCL has a covering number
846 $\mathbb{N}(\epsilon, \text{Rob}_\lambda(\mathcal{W}), L_2(\mathcal{S}))$ no larger than that of the individual ML model $\mathbb{N}(\epsilon, \mathcal{W}, L_2(\mathcal{S}))$ for any ϵ .
847 Thus the Rademacher complexity with the robustification procedure does not increase.

848 By [63], the upper bound of Rademacher complexity with respect to the space of ML model
849 $\text{Rad}_\mathcal{S}(\text{Rob}_\lambda(\mathcal{W}))$ is in the order of $\mathcal{O}(\frac{1}{\sqrt{|\mathcal{S}|}})$. Since the Rademacher complexity with the robustifica-
850 tion procedure satisfies $\text{Rad}_\mathcal{S}(\text{Rob}_\lambda(\mathcal{W})) \leq \text{Rad}_\mathcal{S}(\mathcal{W})$, it also decreases with the dataset size in the
851 order of $\mathcal{O}(\frac{1}{\sqrt{|\mathcal{S}|}})$. \square

852 D Robustification-aware Training

853 Theorem 4.2 also shows the benefits of training the ML model in a robustification-aware manner.
 854 Specifically, by comparing the losses in (5) and (6), we see that using (6) as the robustification-aware
 855 loss for training W can reduce the term $\text{cost}_{\mathcal{S}}(\text{ROB}(h_W))$ in the average cost bound, which matches
 856 exactly with the training objective in (6). The robustification-aware approach is only beginning to be
 857 explored in the ML-augmented algorithm literature and non-trivial (e.g., unconstrained downstream
 858 optimization in [55]), especially considering that (1) is a constrained optimization problem with no
 859 explicit gradients.

860 Gradient-based optimizers such as Adam [64] are the de facto state-of-the-art algorithms for training
 861 ML models, offering better optimization results, convergence, and stability compared to those non-
 862 gradient-based alternatives [65]. Thus, it is crucial to derive the gradients of the loss function with
 863 respect to the ML model weight W given the added robustification step.

864 Next, we derive the gradients of x_t with respect to \tilde{x}_t . For the convenience of presentation, we use
 865 the basic SOCO setting with a single-step switching cost and no hitting cost delay as an example,
 866 while noting that the same technique can be extended to derive gradients in more general settings.
 867 Specifically, for this setting, the pre-robustification prediction is given by $\tilde{x}_t = h_W(\tilde{x}_{t-1}, y_t)$, where
 868 W denotes the ML model weight. Then, the actual post-robustification action x_t is obtained by
 869 projection in (1) by setting $q = 0$ and $p = 1$, given the ML prediction \tilde{x}_t , the expert's action x_t^π and
 870 cumulative cost $\text{cost}(x_{1:t})$ up to t , and the actual cumulative cost $\text{cost}(x_{1:t-1})$ up to $t - 1$.

871 The gradient of the loss function $\text{cost}(x_{1:T}) = \sum_{t=1}^T (f(x_t, y_t) + d(x_t, x_{t-1}))$ with respect to the
 872 ML model weight W is given by $\sum_{t=1}^T \nabla_W (f(x_t, y_t) + d(x_t, x_{t-1}))$. Next, we write the gradient
 873 of per-step cost with respect to W as follows:

$$\begin{aligned} & \nabla_W (f(x_t, y_t) + d(x_t, x_{t-1})) \\ &= \nabla_{x_t} (f(x_t, y_t) + d(x_t, x_{t-1})) \nabla_W x_t + \nabla_{x_{t-1}} (f(x_t, y_t) + d(x_t, x_{t-1})) \nabla_W x_{t-1} \quad (51) \\ &= \nabla_{x_t} (f(x_t, y_t) + d(x_t, x_{t-1})) \nabla_W x_t + \nabla_{x_{t-1}} d(x_t, x_{t-1}) \nabla_W x_{t-1}, \end{aligned}$$

874 where the gradients $\nabla_{x_t} (f(x_t, y_t) + d(x_t, x_{t-1}))$ and $\nabla_{x_{t-1}} d(x_t, x_{t-1})$ are trivial given the hitting
 875 and switching cost functions, and the gradient $\nabla_W x_{t-1}$ is obtained at time $t - 1$ in the same way as
 876 $\nabla_W x_t$. To derive $\nabla_W x_t$, by the chain rule, we have:

$$\nabla_W x_t = \nabla_{\tilde{x}_t} x_t \nabla_W \tilde{x}_t + \nabla_{\text{cost}(x_{1:t-1})} x_t \nabla_W \text{cost}(x_{1:t-1}), \quad (52)$$

877 where $\nabla_W \tilde{x}_t$ is the gradient of the ML output (following a recurrent architecture illustrated in Fig. 1
 878 in the appendix) with respect to the weight W and can be obtained recursively by using off-the-shelf
 879 BPTT optimizers [64], and $\nabla_W \text{cost}(x_{1:t-1}) = \sum_{\tau=1}^{t-1} \nabla_W (f(x_\tau, y_\tau) + d(x_\tau, x_{\tau-1}))$ can also be
 880 recursively calculated once we have the gradient in Eqn. (51). Nonetheless, it is non-trivial to
 881 calculate the two gradient terms in Eqn. (52), i.e., $\nabla_{\tilde{x}_t} x_t$ and $\nabla_{\text{cost}(x_{1:t-1})} x_t$, where x_t itself is the
 882 solution to the constrained optimization problem (1) unlike in the simpler unconstrained case [55].
 883 As we cannot explicitly write x_t in a closed form in terms of \tilde{x}_t and $\text{cost}(x_{1:t-1})$, we leverage the
 884 KKT conditions [66, 67, 68] to implicitly derive $\nabla_{\tilde{x}_t} x_t$ and $\nabla_{\text{cost}(x_{1:t-1})} x_t$ in the next proposition.

Proposition D.1 (Gradients by KKT conditions). *Let $x_t \in \mathcal{X}$ and $\mu \geq 0$ be the primal and dual solutions to the problem (1), respectively. The gradients of x_t with respect to \tilde{x}_t and $\text{cost}(x_{1:t-1})$ are*

$$\begin{aligned} \nabla_{\tilde{x}_t} x_t &= \Delta_{11}^{-1} [I + \Delta_{12} \text{Sc}(\Delta, \Delta_{11})^{-1} \Delta_{21} \Delta_{11}^{-1}], \\ \nabla_{\text{cost}(x_{1:t-1})} x_t &= \Delta_{11}^{-1} \Delta_{12} \text{Sc}(\Delta, \Delta_{11})^{-1} \mu, \end{aligned}$$

885 where $\Delta_{11} = I + \mu \left(\nabla_{x_t, x_t} f(x_t, y_t) + \left(1 + \left(1 + \frac{1}{\lambda_0} \right) (L_1^2 + L_1) \right) I \right)$, $\Delta_{12} = \nabla_{x_t} f(x_t, y_t) +$
 886 $(x_t - \delta(x_{t-1})) + \left(1 + \left(1 + \frac{1}{\lambda_0} \right) (L_1^2 + L_1) \right) (x_t - x_t^\pi)$, $\Delta_{21} = \mu \Delta_{12}^\top$, $\Delta_{22} = f(x_t, y_t) +$
 887 $d(x_t, x_{t-1}) + G(x_t, x_t^\pi) + \text{cost}(x_{1:t-1}) - (1 + \lambda) \text{cost}(x_{1:t}^\pi)$, and $\text{Sc}(\Delta, \Delta_{11}) = \Delta_{22} - \Delta_{21} \Delta_{11}^{-1} \Delta_{12}$
 888 is the Schur-complement of Δ_{11} in the blocked matrix $\Delta = \begin{bmatrix} \Delta_{11} & \Delta_{12} \\ \Delta_{21} & \Delta_{22} \end{bmatrix}$.

889 If the ML prediction \tilde{x}_t happens to lie on the boundary such that the inequality in (1) becomes an
 890 equality for $x = \tilde{x}_t$, then the gradient does not exist in this case and $\text{Sc}(\Delta, \Delta_{11})$ may not be full-
 891 rank. Nonetheless, we can still calculate the pseudo-inverse of $\text{Sc}(\Delta, \Delta_{11})$ and use Proposition D.1
 892 to calculate the subgradient. Such approximation is actually a common practice to address non-
 893 differentiable points for training ML models, e.g., using 0 as the subgradient of $\text{ReLU}(\cdot)$ at the zero
 894 point [64].

895 **Reference**

- 896 [48] W. Tang, S. Bi, and Y. Zhang. Online coordinated charging decision algorithm for electric
897 vehicles without future information. *IEEE Trans. Smart Grid*, 5:2810–2824, May 2014.
- 898 [49] Weici Pan, Guanya Shi, Yiheng Lin, and Adam Wierman. Online optimization with feedback
899 delay and nonlinear switching cost. *Proc. ACM Meas. Anal. Comput. Syst.*, 6(1), feb 2022.
- 900 [50] Tongxin Li, Ruixiao Yang, Guannan Qu, Yiheng Lin, Steven Low, and Adam Wierman.
901 Equipping black-box policies with model-based advice for stable nonlinear control. In
902 <https://arxiv.org/abs/2206.01341>, 2022.
- 903 [51] Chris Develder, Nasrin Sadeghianpourhamami, Matthias Strobbe, and Nazir Refa. Quantifying
904 flexibility in ev charging as dr potential: Analysis of two real-world data sets. In *2016 IEEE*
905 *International Conference on Smart Grid Communications (SmartGridComm)*, pages 600–605.
906 IEEE, 2016.
- 907 [52] Mohammad Ali Alomrani, Reza Moravej, and Elias B. Khalil. Deep policies for online bipartite
908 matching: A reinforcement learning approach. *CoRR*, abs/2109.10380, 2021.
- 909 [53] Bingqian Du, Zhiyi Huang, and Chuan Wu. Adversarial deep learning for online resource
910 allocation. *ACM Trans. Model. Perform. Eval. Comput. Syst.*, 6(4), feb 2022.
- 911 [54] Gautam Goel, Yiheng Lin, Haoyuan Sun, and Adam Wierman. Beyond online balanced descent:
912 An optimal algorithm for smoothed online optimization. In *NeurIPS*, volume 32, 2019.
- 913 [55] Pengfei Li, Jianyi Yang, and Shaolei Ren. Expert-calibrated learning for online optimization
914 with switching costs. *Proc. ACM Meas. Anal. Comput. Syst.*, 6(2), Jun 2022.
- 915 [56] Antonios Antoniadis, Christian Coester, Marek Elias, Adam Polak, and Bertrand Simon. Online
916 metric algorithms with untrusted predictions. In *ICML*, 2020.
- 917 [57] Nicolas Christianson, Tinashe Handina, and Adam Wierman. Chasing convex bodies and
918 functions with black-box advice. In *COLT*, 2022.
- 919 [58] Daan Rutten, Nico Christianson, Debankur Mukherjee, and Adam Wierman. Online optimiza-
920 tion with untrusted predictions. *CoRR*, abs/2202.03519, 2022.
- 921 [59] Guanya Shi, Yiheng Lin, Soon-Jo Chung, Yisong Yue, and Adam Wierman. Online optimization
922 with memory and competitive control. In *NeurIPS*, volume 33. Curran Associates, Inc., 2020.
- 923 [60] Peter L Bartlett and Shahar Mendelson. Rademacher and gaussian complexities: Risk bounds
924 and structural results. *Journal of Machine Learning Research*, 3(Nov):463–482, 2002.
- 925 [61] Martin Wainwright. Rademacher and empirical covering. [https://people.eecs.berkeley](https://people.eecs.berkeley.edu/~wainwrig/stat241b/lec20.pdf)
926 [edu/~wainwrig/stat241b/lec20.pdf](https://people.eecs.berkeley.edu/~wainwrig/stat241b/lec20.pdf), 2009.
- 927 [62] Constantine Caramanis and Sujay Sanghavi. Projection onto a convex set. [https://users](https://users.ece.utexas.edu/~cmcaram/EE381V_2012F/Lecture_3_Scribe_Notes.final.pdf)
928 [ece.utexas.edu/~cmcaram/EE381V_2012F/Lecture_3_Scribe_Notes.final.pdf](https://users.ece.utexas.edu/~cmcaram/EE381V_2012F/Lecture_3_Scribe_Notes.final.pdf),
929 2012.
- 930 [63] Peter L Bartlett, Dylan J Foster, and Matus J Telgarsky. Spectrally-normalized margin bounds
931 for neural networks. *Advances in neural information processing systems*, 30, 2017.
- 932 [64] Ian Goodfellow, Yoshua Bengio, and Aaron Courville. *Deep Learning*. MIT Press, 2016.
- 933 [65] Hyung Ju Suh, Max Simchowitz, Kaiqing Zhang, and Russ Tedrake. Do differentiable simulators
934 give better policy gradients? In *Proceedings of the 39th International Conference on Machine*
935 *Learning*, volume 162 of *Proceedings of Machine Learning Research*, pages 20668–20696.
936 PMLR, 17–23 Jul 2022.
- 937 [66] S. Boyd and L. Vandenberghe. *Convex Optimization*. Cambridge University Press, 2004.
- 938 [67] Brandon Amos. Tutorial on amortized optimization for learning to optimize over continuous
939 domains. *CoRR*, abs/2202.00665, 2022.

- 940 [68] Akshay Agrawal, Brandon Amos, Shane Barratt, Stephen Boyd, Steven Diamond, and J. Zico
941 Kolter. Differentiable convex optimization layers. In H. Wallach, H. Larochelle, A. Beygelzimer,
942 F. d'Alché-Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing*
943 *Systems*, volume 32. Curran Associates, Inc., 2019.