

1 We thank all reviewers for their careful reading, their insightful and constructive comments, and for recognizing the  
 2 importance of this research area (R1,R2,R3,R4), the effort we put in bringing together tools from different areas  
 3 (R1,R2,R3,R4), and the soundness of theoretical results (R1,R4). We answer the main concerns below, but will  
 4 incorporate all feedback (typos, presentation, additional references) in the final version.

5 **(R2) Unclear connections between the finite planning and the infinite horizon performance  $V$ .** We decided to  
 6 study separately the errors stemming from the approximation of  $V^r$  by  $\hat{V}^r$  for any controls  $\mathbf{u}$  (in Thm. 3); and those  
 7 caused by the optimization of  $\hat{V}^r$  with a finite budget  $K$  (in Thm. 2). Following your suggestion, we propose to  
 8 reformulate Thm. 3 to also account for the effect of finite planning, through the planned action  $a_K$  (see next point).

9 **(R2) Reliance of the main result on the PE assumption.** We resorted to this assumption to get asymptotic near-  
 10 optimal performance when  $N \rightarrow \infty$ . It is quite common in the control literature, as recognized by R4. However, as you  
 11 pointed out, our non-asymptotic analysis actually yields a stronger input-dependent result. Accordingly, we propose to  
 12 remove the PE assumption from Thm. 3, which reverts the  $\log N/N$  term in the bound back to the input-dependent  
 13 term  $\beta_N(\delta)^2/\lambda_{\min}(G_{N,\lambda})$ . Thus, our main result becomes:

$$\hat{V}^r(a_K) \leq V^r(a^*) \leq V(a^*) \leq \hat{V}^r(a_K) + \underbrace{\Delta_\omega}_{\text{robustness to perturbations}} + \underbrace{\mathcal{O}\left(\frac{\beta_N(\delta)^2}{\lambda_{\min}(G_{N,\lambda})}\right)}_{\text{estimation error}} + \underbrace{\mathcal{O}\left(K^{-\frac{\log 1/\gamma}{\log \kappa}}\right)}_{\text{planning error}}$$

14 Then, we will mention in a corollary that under the (restrictive) PE assumption, this estimation error term reduces to the  
 15 more explicit form  $\mathcal{O}\left(\frac{\log(N^{d/2}/\delta)}{N}\right)$  which ensures asymptotic near-optimality when  $N \rightarrow \infty$  (and  $K \rightarrow \infty$ ).

16 **(R2) Unclear connection between the discrete measurement noise and the continuous time noise process.** In-  
 17 deed, the reasoning of the remark on line 135 is insufficient, since it only bounds  $\omega(t)$  at discrete time steps  $(t_n)_{n \in \mathbb{N}}$   
 18 while the interval prediction requires bounds  $\underline{\omega}(t) \leq \omega(t) \leq \bar{\omega}(t)$  that hold at all times  $t \geq 0$ . We overlooked it, and  
 19 are very grateful to R2 for pointing it out. Thankfully, this gap can easily be fixed. The most straightforward way is to  
 20 directly assume bounded perturbations, such a hypothesis is both realistic and ubiquitous in the robust MPC literature  
 21 [17, 4, 6, 28, 22, 29]. Another solution, in the spirit of our remark, is to further specify the continuous-time noise  
 22 process  $\omega(t)$ : in place of Assumpt. 2, the disturbance  $\omega(t)$  can be modeled as a Wiener process  $W_t$ . This has two effects:  
 23 first, the increments  $W_{t+dt} - W_t$  are independently and normally distributed, which allows keeping the estimation  
 24 procedure of Thm. 1 by simply replacing  $y_n$  and  $\Phi_n$  by the differences  $\tilde{y}_n = y_n - y_{n-1}$  and  $\tilde{\Phi}_n = \Phi_n - \Phi_{n-1}$  in the  
 25 expressions of  $\theta_{N,\lambda}$  and  $G_{N,\lambda}$ . Second, the running maximum  $\max_{0 \leq s \leq t} W_s$  of a Wiener process is Half-normally  
 26 distributed, which allows to define bounds  $\bar{\omega}_n$  on  $W_t$  that hold with probability  $\delta_n$  for all  $0 \leq t \leq t_n$ . Then, a union  
 27 bound over all  $[0, t_n]$  intervals can be applied to effectively bound  $\omega(t)$  for all  $t \geq 0$ . That being said, we would rather  
 28 include the first solution (bounded perturbations assumption) in the final version for simplicity.

29 **(R2) Suboptimality result defined as regret.** The term *simple* (as opposed to cumulative) *regret* has been used in  
 30 the planning literature [e.g. 19] as a synonym for *suboptimality*. We agree to use the latter, less ambiguous, instead.

31 **(R2,R3) Relevance of the multi-model extension.** We added this section to mitigate the assumption of linear  
 32 dynamics and uncertainty, which R3 finds quite restrictive. Allowing multiple models enables to make multi-modal  
 33 predictions, which we consider a substantial improvement: it was *e.g.* necessary for our simulated driving experiment.

34 **(R3) The discretization of controls might result in some loss of performance.** Indeed. However, as suggested in  
 35 line 67, this issue could be circumvented by resorting to more complex planning algorithms for continuous actions,  
 36 such as SOPC [10] which also comes with a regret bound. We settled for discrete actions for ease of presentation.

37 **(R4) Comparison against other methods.** We have wondered about what methods to compare against. In our related  
 38 work, we made the case that three major families of robust control (robust stabilization, robust constraint satisfaction, and  
 39 minimax LQ) cannot be applied to our environments, which are not stabilization tasks, involve non-convex constraints  
 40 and cannot be solved with quadratic costs. Turning to non-robust RL methods, we argued in lines 249-252 that a  
 41 model-based approach with the same dynamics priors and planning algorithm (our *nominal* baseline) would make the  
 42 fairest comparison. We have since run an additional comparison to DQN, which converged in 30k samples –causing 2k  
 43 collisions– to a policy that still suffers a 6% collision rate. This discussion will be included in the final version.

44 **(R1) More varied environments.** We share your enthusiasm about trying more applications. Note however that  
 45 the driving experiment with a state space of dimension 44 is already quite ambitious compared to usual experimental  
 46 settings in similar works, often limited to systems of dimension 2, 3 or 4 [3,4,11,12,17,28,29,34,38].