

1 We thank the reviewers for their helpful comments which we will address in our final submission, as outlined below.

2 **Reviewer 1:** We are going to improve the expansion and heuristic sections in the final submission.

3 A naive implementation of the expansion step can be computationally expensive, since we have to compute the potential
4 expanders for multiple values of p until we find one such that $|g_t(W_t^\epsilon, A_t(p))| > 0$. In our implementation, we keep
5 track of the nodes such that $|g_t(W_t^\epsilon, A_t(p))| = 0$ to make this step efficient. Further improvements are possible by
6 exploiting properties of the kernel (or metric $d(\cdot, \cdot)$), which often encodes that expandability is a local property. Similar
7 locality considerations can greatly improve the efficiency in many parts of the algorithm.

8 We keep track of \bar{S}_t^o and \bar{S}_t^p using the graph library `networkx` and by adding/removing edges as we acquire new data
9 points. Since at each iteration only few edges are removed/added, this step is not computationally expensive. We are
10 going to add details of an efficient implementation to the appendix and will release our code.

11 **Reviewer 2:** We agree that there are similarities between GOOSE and the method by Berkenkamp 2017 (B17), in
12 that we use the same statistical analysis tools to build accurate confidence intervals and guarantee reachability-based
13 safety constraints during exploration. However, there are fundamental differences between the two works.

14 The most important difference is that of exploration. While B17 provides safety guarantees in the continuous domain,
15 for the exploration analysis they discretize the domain. Thus, in our context, their method can be thought of as an
16 extension of SAFEOPT by Sui 2015 and SAFEMDP by Turchetta 2016 to the continuous domain. While the three
17 methods focus on safety in different settings and thus use different tools to construct pessimistic estimates of the safe
18 set, they all collect data within their respective safe set estimates by following the same strategy that provably explores
19 the entire safe set. In particular, they all expand the safe set by reducing the maximal uncertainty within the current
20 safe set¹, which is easy to analyze but can be very data inefficient in practice. Thus, B17 addresses the problem of
21 safe identification of dynamical systems with continuous state-action spaces by learning about the transition dynamics
22 uniformly over the domain. This is evident from their exploration guarantees in Theorem 4 iii), which hold exclusively
23 for the exploration strategy introduced in equation (6) and which is the same as in SAFEOPT and SAFEMDP.

24 In contrast, GOOSE is a safety add-on layer with strong safety and completeness guarantees that focuses on improving
25 the sample efficiency in exploration. In particular, our analysis allows *any* goal-directed exploration strategy to be used
26 in order to efficiently learn about the part of the domain that is relevant toward the achievement of a given goal. Since
27 the method B17 fundamentally builds on a discrete exploration analysis, it might be possible to use the more efficient
28 exploration scheme of GOOSE in their setting too. While this would require more analysis due to inherent challenges
29 in continuous domains, GOOSE provides a first step in this direction. Thus, we think that our work is complementary
30 to B17 rather than overlapping with it.

31 There are several other differences between the two approaches: (i) Different sources of uncertainty and, therefore, of
32 risk. GOOSE considers safety-critical external environments with known transition dynamics, while B17 addresses
33 uncertainty in the dynamics but does not account for external factors. (ii) Different constraints: B17 focuses on stability
34 constraints, a specific type of constraint that is relevant in dynamical systems. GOOSE looks at level sets of generic
35 unknown functions, which can model safety constraint in a variety of IML scenarios that may not involve dynamical
36 systems, including BO and active learning. (iii) Different assumptions: GOOSE assumes a known model and, under this
37 assumption, stability constraints have been extensively studied. B17 assumes that a Lyapunov function, whose choice
38 implicitly affects the quality of the estimate of the asymptotically stable region, is given. Moreover, B17 implicitly
39 assumes that state action pairs are safely reachable, which is not required for GOOSE.

40 **Reviewer 3:** γ_t denotes the information capacity of the safety constraint, which is an information theoretic complexity
41 measure of the encoded function class. We use it to quantify how many data points we require to learn the function up to a
42 certain accuracy. We will clarify this in the final submission.

43 Without additional assumptions on the environment, it is not possible to provide a bound that does not require complete
44 exploration. As a counter example, consider a graph \mathcal{G} that is a chain of nodes, where each node only provides
45 information about the safety of the next one in the chain. Thus, if the unsafe IML oracle suggests a node at the end of
46 the chain, we must individually learn about each node in the chain to expand the safe set, which is fundamentally what
47 the bound in Theorem 1 encodes. In the general case, this is intuitive if one thinks of SMDP by Turchetta 2016 as the
48 safe equivalent of breadth first search and GOOSE as the equivalent of A^* . Depending on the graph, the location of the
49 target and the heuristic, both breadth first search and A^* may need to visit all the nodes in the graph to find a path to the
50 target. Similarly, in our worst case analysis, it may be necessary for both GOOSE and SMDP to learn about the safety
51 of all nodes in the graph before evaluating the oracle suggestion.

¹SAFEOPT additionally considers maximizers and expanders for efficiency, but their analysis focuses on complete exploration.