SAFER AUTONOMOUS DRIVING IN A STOCHASTIC, PARTIALLY-OBSERVABLE ENVIRONMENT BY HIERARCHICAL CONTINGENCY PLANNING

Anonymous authors
Paper under double-blind review

ABSTRACT

When learning to act in a stochastic, partially observable environment, an intelligent agent should be prepared to anticipate a change in its belief of the environment state, and be capable of adapting its actions on-the-fly to changing conditions. As humans, we are able to form contingency plans when learning a task with the explicit aim of being able to correct errors in the initial control, and hence prove useful if ever there is a sudden change in our perception of the environment which requires immediate corrective action.

This is especially the case for autonomous vehicles (AVs) navigating real-world situations where safety is paramount, and a strong ability to react to a changing belief about the environment is truly needed.

In this paper we explore an end-to-end approach, from training to execution, for learning robust contingency plans and combining them with a hierarchical planner to obtain a robust agent policy in an autonomous navigation task where other vehicles' behaviours are unknown, and the agent's belief about these behaviours is subject to sudden, last-second change. We show that our approach results in robust, safe behaviour in a partially observable, stochastic environment, generalizing well over environment dynamics not seen during training.

1 INTRODUCTION

Developing a controller for AVs to be used in a real-life navigation environment poses a number of challenges including perception, and modeling environment dynamics (Michelmore et al., 2020; Yurtsever et al., 2020). The limitations in vehicle sensors, as well as the stochastic nature of inter-vehicle interactions, introduce a level of uncertainty in autonomous navigation tasks which hinders the ability to control an agent from the standpoint of both safety, and adequate performance. Sensor information, as well as other drivers' intentions, are subject to sudden change and a robust navigation algorithm must be able to safely adapt on-the-fly to these unforeseen changes.

Stochastic environments can be modeled as Partially-Observable Markov Decision Processes (POMDPs) (Kochenderfer et al., 2015). Solving POMDPs is challenging, yet possible, for example using methods combining learning and planning such as Brechtel et al. (2014); Hoel et al. (2020a). Having access to a model of the environment dynamics allows us to use planning algorithms, such as tree-search methods (Browne et al., 2012), alongside learning to both increase sample efficiency, and have access to a better representation of the environment's state-space structure (Machado et al., 2017). In cases where planning is possible it is much easier for an agent to predict the outcome of its actions and hence better adapt to eventual changes in the state-space (McAllister and Rasmussen, 2017). However, a change in the values of the stochastic model parameters (e.g. a change in a vehicle's sensor accuracy, unplanned scene obstruction, or simply an unforeseen behaviour from another vehicle) may induce a sharp drop in the agent's performance due to its inability to generalize well to new environment parameters.

An approach for tackling this issue has been to design controllers to have a high capacity for generalization: instead of attempting to learn over the entire space of possible environments, we design an algorithm to act well enough over the whole space, having trained only on a tractable subset of environment configurations. One such example is known as contingency planning, whereby con-

tingency plans are put in place to specifically counter stochastic environments in which failing to prepare for possible problems in advance can be an expensive mistake (Pryor and Collins, 1996).

Our contribution. In this work we make the following contributions:

- We introduce a method for learning a contingency policy concurrently to the optimal policy during training, such that the former is well-adapted to serve as a contingency plan.
- We design a high-level controller using on a model-based planner which estimates the outcome probability for each low-level policy.
- We experimentally show with an autonomous navigation task that by combining these approaches we are able to achieve a much safer agent performance in the case of a stochastic environment, without being overly conservative in agent behaviour.

2 BACKGROUND

Deep reinforcement learning (RL) is a method for learning control algorithms in a unsupervised manner, and exploiting the strength of deep learning methods to optimize a control algorithm within an environment that may be too complex to handle with other methods. Recently, deep RL implementations have been popular for solving an array of Atari games Silver et al. (2018), where they have only recently achieved super-human performance with a single agent over the whole library of games (Badia et al., 2020), limiting the amount of ad-hoc hyper-parameter tuning needed for the RL algorithm to solve any specific environment.

Implementing an RL algorithm requires us to model the environment as a Markov Decision Process (MDP), defined by the following elements: A finite set of states $s \in \mathcal{S}$, indexed by the timestep at which they are encountered: s_t. A finite set of actions $a \in \mathcal{A}$, also indexed by their respective timesteps: a_t. A transition model $\mathcal{T}(s, a, s') : \mathcal{S} \times \mathcal{A} \times \mathcal{S} \to [0, 1]$, representing the probability of passing from s to s' after taking action a_t, $P(s_{t+1} = s'|s_t = s, a_t = a)$. An immediate reward function $R(s, a, s') : \mathcal{S} \times \mathcal{A} \times \mathcal{S} \to \mathbb{R}$ and discount factor $\gamma \in (0, 1]$, controlling the weight in value of states further along the Markov chain.

A POMDP is further augmented by an observation model $\mathcal{O}$, when we no longer have access to the true state s_t, but to an observation thereof $o_t = \mathcal{O}(s_t)$.

Actions in the POMDP are taken by a policy $\pi : \mathcal{O} \to \mathcal{A}$ mapping observations to actions. The value of an observation under a policy π, is given by the value function $V^\pi : \mathcal{O} \to \mathbb{R}$, which represents the expected future discounted sum of rewards, if policy π is followed from the true state s, of which we have an observation $o = O(s)$:

$$V^\pi(o) := \mathbb{E}\left[\sum_{t=0}^{\infty} \gamma^t R(s_t, a_t, s_{t+1})\right], \tag{1}$$
$$s_0 = s, o_t = O(s_t), a_t = \pi(o_t), s_{t+1} \sim \mathcal{T}(s_t, a_t).$$

Actions are chosen by the policy π, so as to maximize the action-value function $Q^\pi : \mathcal{O} \times \mathcal{A} \to \mathbb{R}$ which assigns values to actions according to the value of the states that are reached under π:

$$Q^\pi(o_t, a_t) := \mathbb{E}_{s_{t+1}}\left[R(s_t, a_t, s_{t+1}) + \gamma \cdot V^\pi(o_{t+1})\right], \tag{2}$$
$$s_{t+1} \sim \mathcal{T}(s_t, a_t), o_{t+1} = O(s_{t+1}).$$

We use Q^π in order to define the optimal policy, which we denote π^*, as the policy taking actions that maximizes Q^π: $\pi^*(o) := \arg\max_{a \in \mathcal{A}} Q^{\pi^*}(o, a)$. Equation (2) highlights that the action-value function is a sort of one-step look-ahead of the value of the next possible state s_{t+1}, in order to determine the value of actions in the current observation o_t. We denote the sequence of observations (trajectory through observation-space) visited by a policy π during an episode as:

$$\tau_\pi := \{o_t\}_{t \in [0,T]}, \tag{3}$$
$$s_0 \in \mathcal{S}, o_t = O(s_t), a_t = \pi(o_t), s_{t+1} \sim \mathcal{T}(s_t, a_t).$$

2.1 Training an Off-Policy Deep Q Network

The off-policy Q-learning algorithm (Mnih et al., 2015) is well-suited to learning a policy's Q-function in the discrete action space which we consider. Training off-policy allows us to train on past transitions stored in a replay buffer (Schaul et al., 2016) which both stabilizes training, and increases sample efficiency since a single transition sample may be used during multiple training steps. We perform a stochastic gradient descent on the MSE between the Q-estimate for the current step, and the discounted 1-step "lookahead" Q-estimate summed with the transition reward:

$$\mathcal{L}(o_t, a_t, r_t, o_{t+1}) = \left(Q_\theta(o_t, a_t) - \left[r_t + \gamma Q_\theta(o_{t+1}, \arg\max_{a'} Q_{\theta'}(o_{t+1}, a')\right]\right)^2,$$

where θ are the current parameters for the Q-network, and θ' are parameters that are 'frozen' for a certain amount of steps. This is known as double Q-learning using target networks (Van Hasselt et al., 2016) and reduces the variance of gradient updates.

3 Related Work

Hierarchical RL. Hierarchical reinforcement learning (Dayan and Hinton, 1993; Sutton et al., 1999) is a promising approach for helping RL algorithms generalize better in increasingly complex environments. Such an approach allows us to learn subtasks individually, and later combine them with a high-level policy to achieve a higher level of performance, and higher sample efficiency (Nachum et al., 2018; Zhang et al., 2021). This is powerful as it allows difficult tasks to be broken down into manageable, less difficult ones which are easier to learn (Konidaris and Barto, 2007; Andreas et al., 2017). Due to this, several works apply the hierarchical structure of control to navigation tasks, which lend themselves well to modular controllers (Fisac et al., 2018)

Contingency planning. Some approaches use an estimation of the confidence of the actions proposed by an agent's policy (Bouton et al., 2019; Clements et al., 2019; Hoel et al., 2020b) to determine whether or not an agent's policy is sufficiently good in the current environment state. When this is not the case, control is typically given to a separate, often open-loop controller looking to mitigate any possible negative behaviour if failure cannot be avoided otherwise (Dalal et al., 2018; Filos et al., 2020). One issue with open-loop contingency plans – or any open-loop policy in general – is that they do not take into account the closed-loop nature of most real-world environments, whose dynamics are dependent on the actions of the agent and may themselves fail if not implemented carefully.

Approaches such as Rhinehart et al. (2021) or Killing et al. (2021), tackle the problem of high environment uncertainty, by prioritizing an information gathering if the agent is too uncertain about its policy's outcome. These approaches work well, and choose to approach by default with caution, if ever there is missing or uncertain information in the agent's input space. Cunningham et al. (2015) uses a form of voting between candidate policies to determine which to act on, based on both performance and safety. However, these approaches are based on either already available policies, focusing on the task of selecting the best available one, or learning from human-trajectories. Our approach has a strong focus on learning policies in an unsupervised manner, with the explicit aim of using them as contingency plans.

4 Learning Contingency Policies

Problem statement. We wish to learn, on one hand, π^*, the optimal policy in our given environment, and on the other, π_1, a contingency policy able to navigate more safely through the environment, at the cost of performance, if ever there is high uncertainty linked to following the optimal policy.

4.1 Contingency Policies

Given the nature of a contingency plan, it must aim to exploit trajectories through different areas of state-space than trajectories from the optimal policy. Since we consider cases where typically local uncertainties in the navigation task disrupt the performance of the optimal policy, in order for

the contingency plan to stay viable it should have trajectories through sufficiently different areas of state-space to avoid areas of high uncertainty. For example a contingency plan to high-speed trajectories, will often be more conservative and favor low-speed policies.

Using reward augmentation. Our approach for learning π^* and π_1 concurrently, is to augment the reward function by a term which we will refer to as a *trajectory penalty*, R^{pen}, which is proportional to the similarity in expected observation-space trajectories between two policies. The aim of augmenting the reward in this way is for the policy of the contingency agent to be incentivized to solve the environment through different areas of state-space than the optimal policy. With this aim in mind we define a metric on agents' observation trajectories:

$$\mathcal{M}(\tau_{\pi_1}, \mathbb{E}\left[\tau_{\pi^*}\right]) := \int \left| \nu\left(\phi\left(\tau_{\pi_1}\right)\right) - \nu\left(\phi\left(\mathbb{E}\left[\tau_{\pi^*}\right]\right)\right) \right| d\phi(o), \tag{4}$$

where ϕ is an observation-feature function, and ν is the density function over observation features. It is possible to remain the most general possible with $\phi(o) = o$, however with some domain knowledge we can modify ϕ to retain only the features which best define the 'distinctness' of an agent's trajectory.

Designing the penalty term R^{pen} to be inversely proportional to the trajectory metric with respect to a reference agent, we can define it as:

$$R^{pen}_{\pi^*}(\tau_{\pi_1}) := -\frac{\alpha}{\mathcal{M}(\tau_{\pi_1}, \mathbb{E}\left[\tau_{\pi^*}\right]) + \delta}, \tag{5}$$

where α and δ determine the relative weight of the trajectory penalty, with respect to the regular reward function R. A lower value for α will hardly penalize the contingency policy for having a similar observation distribution to the reference agent, whereas higher weighting will make it seek a highly different trajectory, disregarding the original objective of the task given by the regular reward function. The value function for π_1 becomes:

$$V^{\pi_1}(o) := \mathop{\mathbb{E}}_{\tau \sim \pi_1}\left[R(\tau) + R^{pen}_{\pi^*}(\tau)\right], \quad \tau|_{t=0} = o,$$

where $\tau \sim \pi_1$ is the expected observation-space trajectory under π_1. We denote Π, the set of all policies (optimal and contingency) learned in the POMDP. Stochastic parameters for the POMDP are stored in a vector $\boldsymbol{\mu}$.

Contingency plans are needed. When building a hierarchical control framework, much of the generalization capabilities come from the correct identification of the environment parameters, and then executing the relevant policy (Pateria et al., 2021). When a new objective is introduced to the agent or environment parameters change causing new environment dynamics, making the optimal policy unable to perform well, we must ensure that the contingency policy can function well from the 'hand-off' point otherwise it will be ineffective. When the observation-space gets more complex, we may no longer assume that the learned contingency policy has sufficient knowledge over its domain such that it is able to correctly control the agent from any such 'hand-off' state.

This is linked to the exploration trade-off made during training, where we explicitly limit an agent's exploration of the available state-space in the interest of making learning tractable (Sutton and Barto, 2018). Due to limited exploration by design, this limits the domain on which the Q^π-function is well-learnt. This isn't an issue for navigating from the starting state to each objective, however if during execution there is a change in objective or environment dynamics, then there will be some states from which the contingency policy will not perform well, due to its untrained Q^π-function. For this not to be the case, we would have to train this policy to start from a wide range of possible starting points. This may be equivalently seen as a lack of exploration and insufficient training of the RL agent, however in more complex environments time and computing constraints prevent us from training the Q^π-function over the policy's entire domain.

Contingency policy's domain. In order for the contingency policy to be effective, we must ensure that it functions well in areas of observation-space where the agent would have been guided into by another policy. This can be done during training, by changing the initial starting state of the agent during episodes, forcing it to learn to find its objective from a specific zone in its domain.

It is possible to implement a random episode initialization, however this is not the aim of increasing the acting observation-space of the contingency policy because we wish specifically to learn areas of observation-space heavily explored by the optimal policy, being the most likely 'hand-off' point to the contingency policy. Having both optimal and contingency policies train concurrently allows us to sample random states in the replay buffer of the optimal, as initial states for the contingency policy. This is equivalent to letting the optimal agent take the first random number of steps in the environment before handing over control to the contingency policy, mirroring the intended use-case for this training paradigm.

Further, with domain knowledge we are able to guide the contingency policy to more specific areas of state-space. For example in an autonomous navigation task, in order to avoid collisions, we may want to a contingency policy to be good at re-directing specifically from high-speed trajectories, and focus on those high-speed states to be more effective.

4.2 Hierarchical Controller

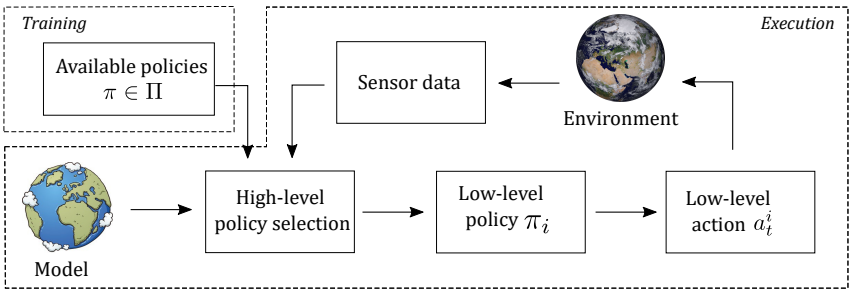

Figure 1: Structure of hierarchical controller composed of available policies and model-based planner (High-level policy selection).

Contingency plans alone are not sufficient: a higher-level controller is required, capable of selecting the best policy with respect to the current environment observation. In our approach, we combine the learned optimal and contingency policies with a model-based planner, in order to increase the robustness and safety of the acting agent with respect to environment uncertainty. Figure 1 shows the structure of the hierarchical controller, whose high-level policy selection is in fact the planner module responsible for estimating the safety of each of the available policies $\pi \in \Pi$, and selecting the one with the lowest estimated chance of failure.

4.3 Planner

The planner's role is to estimate the chance of failure for each $\pi \in \Pi$ available to the controller. The approach we retain for our purposes is straightforward: we perform roll-outs over the set of possible environment parameters $\boldsymbol{\mu}$ for each π, and return an estimated failure rate based on those roll-outs which will define which policy is selected by the planner. This process estimates the probability of policy failure given the current belief about environment stochastic parameters $\boldsymbol{\mu}$. Let $P(s_{fail}|\pi, \boldsymbol{\mu})$ denote the probability that policy π will fail for a given $\boldsymbol{\mu}$. The planner's goal is to select a policy according to:

$$\pi = \arg\min_{\pi \in \Pi} P(s_{fail}|\pi, \boldsymbol{\mu}^*),$$

where $\boldsymbol{\mu}^*$ are the true (unknown) parameter values. Since this value is unknown, we may approximate it by sampling $\boldsymbol{\mu}_i$'s which are closest to $\boldsymbol{\mu}^*$ according to the history of observations in the environment. Let $C(\pi, \boldsymbol{\mu}_i) \in [0, 1]$ represent whether or not there is a failure (collision) from policy π, after roll-out. Then:

$$P(s_{fail}|\pi, \boldsymbol{\mu}^*) \approx \frac{1}{M} \sum_{i=1}^{M} C(\pi, \boldsymbol{\mu}_i),$$

for M samples of $\boldsymbol{\mu}_i$. We note that the quality of the approximation relies on how well the $\boldsymbol{\mu}_i$'s are sampled. The closer they are to $\boldsymbol{\mu}^*$, the better the planner will be able to estimate the true probability of failure for each policy. We use a simple approach consisting in eliminating $\boldsymbol{\mu}_i$'s from

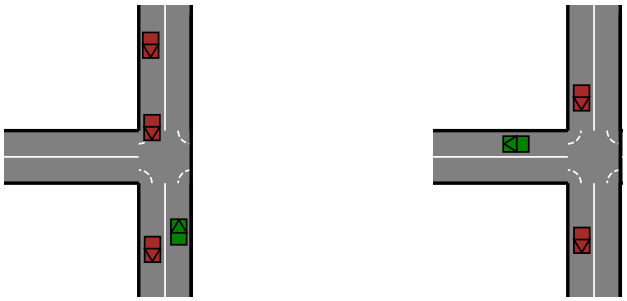

(a) $t = 3$s (b) $t = 5$s

Figure 2: Navigation task: ego makes a left turn across the intersection with oncoming traffic. Target vehicles may either be aggressive (i.e. disregarding presence of ego in intersection) or cooperative (i.e. slowing down if ego is close to intersection point).

the sampling pool, if targets are observed behaving in contradiction to the considered environment parameters (based on vehicle speed, in our navigation task).

5 EXPERIMENTS

5.1 SIMULATION ENVIRONMENT

We evaluate our approach in a common autonomous navigation task (Brechtel et al., 2014; Hubmann et al., 2018; Bouton et al., 2019; Bernhard et al., 2019; Rhinehart et al., 2021). Figure 2 shows two frames of the environment with oncoming vehicles in the intersection. In navigation tasks, the controllable agent is usually referred to as the ego whereas the other vehicles are referred to as targets. In this task the ego must adjust its speed in order to pass through the intersection without colliding with any of the oncoming targets. The optimal policy with respect to the reward function (6) is to pass through the intersection as fast as possible while avoiding collisions. We expect a contingency policy to be useful in this scenario, when there is a sudden change in estimated target behaviour by the planner and the ego will have to either slow down to let an aggressive target through, or speed up if it seems the target is slowing down too much which may also cause a collision.

Ego and target behaviours. The ego's action space is a set of longitudinal acceleration values: $a \in \mathcal{A} = \{-4, -2, -1, 0, 1, 2\}$ m/s^2. Target vehicles may have either a cooperative, or aggressive style of driving, this behaviour being unobservable by the ego. An aggressive target will speed through the intersection disregarding the presence of the ego vehicle, whereas a cooperative target will slow down if ever the ego vehicle approaches the intersection, although it will not come to a complete stop if ever ego halts in the middle of the road.

Reward function. To penalize collisions and encourage faster episode termination, the step-reward r_t is set-up as follows per time step t:

$$r_t = \begin{cases} -5 & \text{if collision} \\ -0.1 & \text{otherwise} \end{cases}. \tag{6}$$

Stochastic environment parameters. The environment dynamics depend on the targets' behaviours. The behaviour for each target vehicle is a random variable B_i representing degree of aggressiveness, which are collected in $\boldsymbol{\mu}$ where $\boldsymbol{\mu}|_i = B_i$ for $i \in [1, N]$ for N targets. In practice we use either $b_i = 0$ for a cooperative target, or $b_i = 1$ for an aggressive target.

5.2 IMPLEMENTATION DETAILS

Target behaviours. During training we set $b_i = 1, \ \forall i \in [1, N]$. Both optimal and contingency policies are trained on these environment parameters.

Trajectory penalty. In practice we replace the expectation operator $\mathbb{E}[\tau_\pi]$ in (4) by the mean value over the last 100 samples of π^*'s replay buffer. In this environment we use the speed of the ego vehicle as an observation feature $\phi(o) = \dot{x}_{ego}$. Trajectory penalty scaling factors used are: $\alpha = 3, \delta = 0.1$, which are fixed by a rough initial sweep.

Adjusting the return values for different initial states. We use a modified initialization for the contingency agent to increase its capabilities for compensating erroneous behaviour from a previous policy. We initialize 50% of the contingency agent's episodes in this way, during the entire training phase. Forcing the contingency policy to start the episode in states sampled from the optimal policy's replay buffer will incur greater trajectory penalties due to the higher trajectory proximity. However, we do not need to to add an explicit compensation term: the contingency agent is still learning its true action-value function Q^{π_1}. We take into consideration that the trajectory penalty term R^{pen} will reduce the mean computed performance score, so it is important to look at the training score and trajectory penalties separately when comparing performance from both policies. To compensate for varying episode length due to a different initial state, the contingency policy is attributed the cumulated rewards corresponding to the sampled initial state from the replay buffer.

6 RESULTS

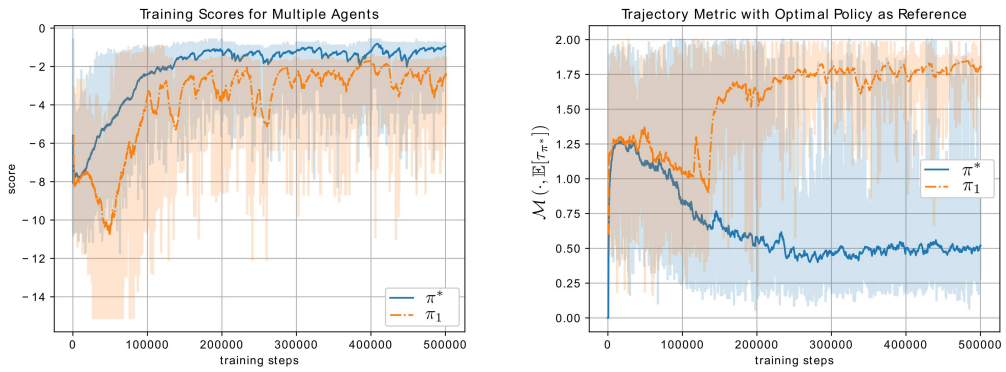

(a) (b)

Figure 3: (a) Training scores for both agents during training phase. This figure does not take into account the trajectory penalties R^{pen} for π_1, only the regular rewards R. (b) Evolution of computed trajectory metric term $\mathcal{M}(\cdot, \mathbb{E}[\tau_{\pi^*}])$ for optimal and contingency policies. Although computed for both policies, the resulting trajectory penalty is only attributed to π_1.

Training a contingency policy with replay buffer. Figure 3a shows the scores achieved by both the optimal and contingency policies during training. We can see π^* learning until reaching its optimal performance score after approximately 200k training steps. π_1's optimization objective depends on an additional R^{pen} term, which is not stable until the observation-density in the replay buffer of π^* becomes approximately static. This is why we see the initial training curve of π_1 lag behind that of π^*. Once π^*'s expected observation trajectory, $\mathbb{E}[\tau_{\pi^*}]$, has a low enough variance then π_1 can make some meaningful progress towards its contingency objective. Looking at π_1's training curve in our case this appears to happen around the 50k step mark. Even once π_1 has reached its top performance, it continues to be more jittery than π^*. This is due to the sampled replay buffer initializations.

Trajectory metric during training. Figure 3b shows the evolution of the trajectory metric (4). The values for π^* are shown for comparison, but not actually included in its reward during training. The shape for π^* is to be expected, as the latest trajectory will initially be different from the mean trajectory sampled from its own replay buffer. However as its parameters θ converge and exploration is weaned away, the episodic trajectory will get closer to the mean trajectory sampled from the replay buffer.

We notice that initially both policies have similar trajectories τ. This is due to the high degree of initial random exploration. There is a sharp rise in the metric for π_1 around the 150k mark, and we may deduce that at this point the most recent trajectories τ_{π_1} are sufficiently different to the observations present in π^*'s replay buffer such that the metric between the two increases. We may conclude that this is where the contingency policy is converging to a contingency plan that has a different trajectory from that of the optimal policy (notably lower-speed trajectories).

Hierarchical planner safety performance. To evaluate the overall performance of our hierarchical controller, we compute the number of successes vs. failures on the navigation task, over a range of environment parameter values $\boldsymbol{\mu}$ for different agents. Table 1 compiles the results for each controller, over an environment parameter sample size of $M = 200$.

Table 1: Controller performances for a range of agents. Collision rate gives ratio of number of successes vs. failures over all sampled environment configurations. Average score gives us the mean of scores obtained in cases where the agent does not fail.

Controllers	**Success rate**	**Average Score**
π^*	0.496	-1.200
$\boldsymbol{\pi_1}$ without replay buffer init.	0.885	-1.500
$\boldsymbol{\pi_1}$ with replay buffer init.	0.954	-1.466
H-control without replay buffer init.	0.890	-1.380
H-control with replay buffer init.	$\mathbf{1.000}$	$\mathbf{-1.238}$

From our results, we clearly deduce that a single agent has a hard time generalizing to new target behaviours, even though it may have achieved optimal performance within its training environment. Although π^* has the highest average score in cases when it does not fail, it does not generalize well and performs very poorly in unseen environment instances. In our navigation task, the contingency policy that is learnt corresponds to the ego having a more conservative driving attitude. This explains why the success rate is higher for π_1 with respect to π^*, although the average score decreases. When adding the replay buffer initialization to π_1, the success rate further increases due to the increased ability of the contingency agent to generalize to greater areas of observation-space.

Compared to the individual policies, we expect the hierarchical controller to perform better, due to its access to a model-based planner combined with both policies. Interestingly, we find that the contingency policy with replay buffer initialization has a higher success rate than the hierarchical controller using π_1 without the initialization. This highlights the importance of generalization when designing safe, robust algorithms. Finally, we see that the highest success rate is achieved by our proposed approach of combining optimal, and robust contingency policies. We are also able to keep a relatively high performance metric, due to π^* being available to the controller.

7 Conclusion

In conclusion, we have presented an approach for learning multiple policies in an autonomous navigation task and adapting the approach to specifically learn a robust contingency policy, which when combined with a model-based planner, is able to increase the robustness of the agent with respect to stochastic environment parameters. In our intersection use-case, we are able to reach a rate of no collisions for any (sampled) environment configuration, even when we only had access to a single one of these configurations during training. We acknowledge that the planner module has a simple behavioural prediction for target vehicles, and although it is sufficient in our simulation environment to obtain good performance, better detection of the environment's stochastic parameters $\boldsymbol{\mu}$ will increase the robustness of the overall agent, and ultimately the effectiveness of having an available contingency policy.

One main advantage of this hybrid approach is the ability to separate performance and safety in an RL framework. Whereas using a single reward function and relying on reward engineering to obtain the correct behaviour can be arbitrary, in this case we are able to optimize for performance in a complex environment, and ensure a high level safety in unseen instances of environment dynamics without having to tweak performance and safety terms in a single reward function.

REFERENCES

Jacob Andreas, Dan Klein, and Sergey Levine. Modular multitask reinforcement learning with policy sketches. In *Proceedings of the 34th International Conference on Machine Learning*, volume 70, pages 166–175, 06–11 Aug 2017.

Adrià Puigdomènech Badia, Bilal Piot, Steven Kapturowski, Pablo Sprechmann, Alex Vitvitskyi, Zhaohan Daniel Guo, and Charles Blundell. Agent57: Outperforming the Atari human benchmark. In *Proceedings of the 37th International Conference on Machine Learning*, volume 119, pages 507–517, Jul 2020.

Julian Bernhard, Stefan Pollok, and Alois Knoll. Addressing inherent uncertainty: Risk-sensitive behavior generation for automated driving using distributional reinforcement learning. In *IEEE Intelligent Vehicles Symposium (IV)*, pages 2148–2155, 2019.

M. Bouton, A. Nakhaei, K. Fujimura, and M. J. Kochenderfer. Safe reinforcement learning with scene decomposition for navigating complex urban environments. In *IEEE Intelligent Vehicles Symposium (IV)*, pages 1469–1476, 2019.

Sebastian Brechtel, Tobias Gindele, and Rüdiger Dillmann. Probabilistic decision-making under uncertainty for autonomous driving using continuous pomdps. In *17th International IEEE Conference on Intelligent Transportation Systems (ITSC)*, pages 392–399, 2014.

Cameron Browne, Edward Powley, Daniel Whitehouse, Simon Lucas, Peter Cowling, Philipp Rohlfshagen, Stephen Tavener, Diego Perez Liebana, Spyridon Samothrakis, and Simon Colton. A survey of monte carlo tree search methods. *IEEE Transactions on Computational Intelligence and AI in Games*, pages 1–43, 2012.

William R. Clements, Benoît-Marie Robaglia, Bastien Van Delft, Reda Bahi Slaoui, and Sébastien Toth. Estimating risk and uncertainty in deep reinforcement learning. *arXiv Preprint*, 2019.

Alexander G. Cunningham, Enric Galceran, Ryan M. Eustice, and Edwin Olson. Mpdm: Multipolicy decision-making in dynamic, uncertain environments for autonomous driving. In *IEEE International Conference on Robotics and Automation (ICRA)*, pages 1670–1677, 2015.

Gal Dalal, Krishnamurthy Dvijotham, Matej Vecerik, Todd Hester, Cosmin Paduraru, and Yuval Tassa. Safe exploration in continuous action spaces. *arXiv preprint*, 2018.

Peter Dayan and Geoffrey E Hinton. Feudal reinforcement learning. In *Advances in Neural Information Processing Systems*, volume 5, pages 271–278. Morgan-Kaufmann, 1993.

Angelos Filos, Panagiotis Tigkas, Rowan Mcallister, Nicholas Rhinehart, Sergey Levine, and Yarin Gal. Can autonomous vehicles identify, recover from, and adapt to distribution shifts? In *Proceedings of the 37th International Conference on Machine Learning*, volume 119, pages 3145–3153, 2020.

Jaime F. Fisac, Eli Bronstein, Elis Stefansson, Dorsa Sadigh, S. Shankar Sastry, and Anca D. Dragan. Hierarchical game-theoretic planning for autonomous vehicles. pages 9590–9596, 2018.

Carl-Johan Hoel, Katherine Driggs-Campbell, Krister Wolff, Leo Laine, and Mykel J. Kochenderfer. Combining planning and deep reinforcement learning in tactical decision making for autonomous driving. *IEEE Transactions on Intelligent Vehicles*, 5:294–305, 2020a.

Carl-Johan Hoel, Krister Wolff, and Leo Laine. Tactical decision-making in autonomous driving by reinforcement learning with uncertainty estimation. In *IEEE Intelligent Vehicles Symposium (IV)*, pages 1563–1569, 2020b.

Constantin Hubmann, Jens Schulz, Marvin Becker, Daniel Althoff, and Christoph Stiller. Automated driving in uncertain environments: Planning with interaction and uncertain maneuver prediction. *IEEE Transactions on Intelligent Vehicles*, 3:5–17, 2018.

Christoph Killing, Adam Villaflor, and John M. Dolan. Learning to robustly negotiate bi-directional lane usage in high-conflict driving scenarios. In *IEEE International Conference on Robotics and Automation (ICRA)*, pages 8090–8096, 2021.

Mykel J. Kochenderfer, Christopher Amato, Girish Chowdhary, Jonathan P. How, Hayley J. Davison Reynolds, Jason R. Thornton, Pedro A. Torres-Carrasquillo, N. Kemal Üre, and John Vian. *Decision Making Under Uncertainty: Theory and Application*. The MIT Press, 1st edition, 2015.

George Konidaris and Andrew Barto. Building portable options: Skill transfer in reinforcement learning. In *Proceedings of the Twentieth International Joint Conference on Artificial Intelligence*, 2007.

Marlos C. Machado, Marc G. Bellemare, and Michael Bowling. A Laplacian framework for option discovery in reinforcement learning. In *Proceedings of the 34th International Conference on Machine Learning*, volume 70, pages 2295–2304, 2017.

Rowan McAllister and Carl Edward Rasmussen. Data-efficient reinforcement learning in continuous state-action gaussian-pomdps. In *Advances in Neural Information Processing Systems*, volume 30, 2017.

Rhiannon Michelmore, Matthew Wicker, Luca Laurenti, Luca Cardelli, Yarin Gal, and Marta Kwiatkowska. Uncertainty quantification with statistical guarantees in end-to-end autonomous driving control. In *IEEE International Conference on Robotics and Automation (ICRA)*, pages 7344–7350, 2020.

Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Andrei A. Rusu, Joel Veness, Marc G. Bellemare, Alex Graves, Martin Riedmiller, Andreas K. Fidjeland, Georg Ostrovski, Stig Petersen, Charles Beattie, Amir Sadik, Ioannis Antonoglou, Helen King, Dharshan Kumaran, Daan Wierstra, Shane Legg, and Demis Hassabis. Human-level control through deep reinforcement learning. *Nature*, 2015.

Ofir Nachum, Shixiang (Shane) Gu, Honglak Lee, and Sergey Levine. Data-efficient hierarchical reinforcement learning. In *Advances in Neural Information Processing Systems*, volume 31, 2018.

Shubham Pateria, Budhitama Subagdja, Ah-hwee Tan, and Chai Quek. Hierarchical reinforcement learning: A comprehensive survey. *ACM Computing Surveys*, 54, 2021.

Louise Pryor and Gregg C. Collins. Planning for contingencies: A decision-based approach. *Journal for Artificial Intelligence Research*, 4:287–339, 1996.

Nicholas Rhinehart, Jeff He, Charles Packer, Matthew A. Wright, Rowan McAllister, Joseph E. Gonzalez, and Sergey Levine. Contingencies from observations: Tractable contingency planning with learned behavior models. In *2021 IEEE International Conference on Robotics and Automation*, pages 13663–13669, 2021.

Tom Schaul, John Quan, Ioannis Antonoglou, and David Silver. Prioritized experience replay. In *Proceedings of the International Conference on Learning Representations*, 2016.

David Silver, Thomas Hubert, Julian Schrittwieser, Ioannis Antonoglou, Matthew Lai, Arthur Guez, Marc Lanctot, Laurent Sifre, Dharshan Kumaran, Thore Graepel, Timothy Lillicrap, Karen Simonyan, and Demis Hassabis. A general reinforcement learning algorithm that masters chess, shogi, and go through self-play. *Science*, 362(6419):1140–1144, 2018.

Richard Sutton, Doina Precup, and Satinder Singh. Between mdps and semi-mdps: A framework for temporal abstaction in reinforcement learning. *Artificial Intelligence*, 1999.

Richard S. Sutton and Andrew G. Barto. *Reinforcement Learning: An Introduction*. The MIT Press, second edition, 2018.

Hado Van Hasselt, Arthur Guez, and David Silver. Deep reinforcement learning with double q-learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 30, 2016.

Ekim Yurtsever, Jacob Lambert, Alexander Carballo, and Kazuya Takeda. A survey of autonomous driving: Common practices and emerging technologies. *IEEE Access*, 8:58443–58469, 2020.

Jesse Zhang, Haonan Yu, and Wei Xu. Hierarchical reinforcement learning by discovering intrinsic options. In *International Conference on Learning Representations*, 2021.