

# Interpretable Goal Recognition in the Presence of Occluded Factors for Autonomous Vehicles

Josiah P. Hanna<sup>1,2,3</sup>, Arrasy Rahman<sup>1,2</sup>, Elliot Fosong<sup>1,2</sup>,  
Francisco Eiras<sup>1,4</sup>, Mihai Dobre<sup>1</sup>, John Redford<sup>1</sup>,  
Subramanian Ramamoorthy<sup>1,2</sup>, and Stefano V. Albrecht<sup>1,2</sup>

**Abstract**—Recognising the goals or intentions of observed vehicles is a key step towards predicting the long-term future behaviour of other agents in an autonomous driving scenario. When there are unseen obstacles or occluded vehicles in a scenario, goal recognition may be confounded by the effects of these unseen entities on the behaviour of observed vehicles. Existing prediction algorithms that assume rational behaviour with respect to inferred goals may fail to make accurate long-horizon predictions because they ignore the possibility that the behaviour is influenced by such unseen entities. We introduce the Goal and Occluded Factor Inference (GOFI) algorithm which bases inference on inverse-planning to jointly infer a probabilistic belief over goals and potential occluded factors. We then show how these beliefs can be integrated into Monte Carlo Tree Search (MCTS). We demonstrate that jointly inferring goals and occluded factors leads to more accurate beliefs with respect to the true world state and allows an agent to safely navigate several scenarios where other baselines take unsafe actions leading to collisions.

## I. INTRODUCTION

Predicting the behaviour of other vehicles is a fundamental challenge to developing autonomous vehicles that can drive safely and effectively in the real world. At a given moment in time, a controlled vehicle (hereafter called the *ego-vehicle*) must use a limited number of observations of the current and past states of other (non-ego) vehicles to infer their future locations and velocities, and how they will respond to future ego-vehicle actions. Accurate prediction allows the ego-vehicle to anticipate possible future collisions and plan actions that maximise the probability of safely and efficiently reaching its target location [2], [9], [19], [26].

Two major approaches to this problem are inverse-planning and deep-learning methods. Inverse-planning methods assume other vehicles behave near-rationally and predict optimal paths under inferred rewards or goals [2], [12]. Deep-learning methods use datasets of past vehicle trajectories and learn the conditional probability of future trajectory segments given past observations and road layouts [10], [14]. While black-box deep-learning-based methods have demonstrated an impressive ability to predict future vehicle behaviour, they lack interpretability and may perform poorly if test conditions differ from how they were trained [11], [22]. Since interpretability is a key aspect of trustworthy autonomous



(a) Ego-vehicle View



(b) Bird's Eye View

Fig. 1: Example driving scenario where the behaviour of an observed vehicle (#1) depends on a vehicle (#4) that the ego-vehicle (blue) cannot observe because it is occluded by a building.

systems [31], we focus on inverse-planning prediction methods. Predicted behaviours can then be explained as rational with respect to a certain goal or cost function.

Recently, Albrecht et al. [2] showed how goal recognition is important for long-horizon prediction. By inferring the goals (e.g., target locations) that a non-ego vehicle is attempting to achieve and assuming the vehicle behaves approximately rationally, the space of likely future paths for that vehicle is constrained. When goals are the main factor governing a vehicle's behaviour, goal recognition may be sufficient for explaining much of its future behaviour. However, in a partially observable setting, the future behaviour of non-ego vehicles may depend on occluded factors that the ego-vehicle never observes. Without considering these factors, an inverse-planning prediction method may incorrectly infer goals. For example, consider the scenario shown in Fig. 1. The ego-vehicle (blue) cannot see an oncoming vehicle (#4) but observes that vehicle #1 has stopped with a gap between it and the car in front of it. Without knowledge of the oncoming vehicle, the ego-vehicle may conclude that #1 has

<sup>1</sup>Five AI Ltd. U.K., [firstname.lastname@five.ai](mailto:firstname.lastname@five.ai)

<sup>2</sup>School of Informatics, University of Edinburgh, U.K.

<sup>3</sup>Computer Sciences Department, University of Wisconsin–Madison

<sup>4</sup>Department of Engineering Science, University of Oxford, U.K.

Correspondence to [stefano.albrecht@five.ai](mailto:stefano.albrecht@five.ai)

stopped to allow space for the ego-vehicle to merge in front of it. However, an alternative explanation for the behaviour of #1 is that it desires to turn right (and is failing to use turn signals) but is waiting for an unobserved vehicle to pass. Thus, goal recognition is confounded by the presence of the on-coming vehicle.

The key insight behind our work is that observations of a given vehicle are a function of both that vehicle’s goal and the presence or absence of potential occluded factors in the world. Thus, we can use the past behaviour of an observed vehicle to both infer its goal and the existence of these occluded factors. While previous work [1], [29] has proposed using other vehicles as “sensors” for occluded factor inference, this prior work neglects the effects of vehicle goals or intentions on inference.

In this paper, we first describe how occluded factors can confound prediction that relies on goal recognition. To address this issue, we then introduce the Goal and Occluded Factor Inference (GOFI) algorithm which resolves confounding by jointly inferring a probabilistic belief both over goals and occluded factors. We then show how this algorithm can be integrated into Monte Carlo Tree Search (MCTS) [15] to select actions for the ego-vehicle. Finally, we empirically evaluate this algorithm on autonomous driving scenarios where reasoning about both goals and occluded factors is necessary for safe and efficient driving. Our empirical results show that GOFI improves inference accuracy and completes a higher percentage of trials in the evaluated scenarios.

## II. PRELIMINARIES

We consider a set of vehicles,  $\mathcal{N}$ , in which each vehicle  $i \in \mathcal{N}$  has a *local state*  $s_t^i \in \mathcal{S}$  describing its pose and velocity at timestep  $t$ . We represent the ego-vehicle by index  $\varepsilon$ . The joint state at time  $t$  is denoted by  $s_t := (s_t^1, \dots, s_t^{|\mathcal{N}|}) \in \mathcal{S}^{|\mathcal{N}|}$ . The ego-vehicle  $\varepsilon$  has a potentially limited view of other vehicles in  $\mathcal{N}$ . At step  $t$ , we assume that  $\varepsilon$  processes a sensory input sequence  $o_{1:t}$  to obtain a current state-estimate  $\hat{s}_t \in \mathcal{S}^{|\mathcal{I}|}$ , where  $\mathcal{I}$  is the set of all vehicles of which the ego-vehicle is aware, i.e., vehicles observed in some time-step up to and including  $t$ . The processing required to obtain  $\hat{s}_t$  from  $o_{1:t}$  (i.e. state filtering [3]) is outside the scope of this paper.

We assume each vehicle  $i$  has a goal,  $g^i \in \mathcal{G}^i$ , from a set of possible goals  $\mathcal{G}^i$  and that  $g^i$  remains constant through a fixed driving scenario. In this work, goals correspond to target locations and velocities. At each timestep  $t$ , each vehicle  $i$  chooses an action  $a_t^i \in \mathcal{A}^i$  as a function of its goal and state-estimate. The true state of the world then evolves as an (unknown) function of the joint action.

The ego-vehicle may be unaware of vehicles, pedestrians, or other obstacles in the environment. We refer to such vehicles or obstacles as *occluded factors*. The possible presence or absence of occluded factors means that there are multiple world states consistent with  $\hat{s}_t$ . We define  $\mathcal{Z} \subset \mathcal{S}^{|\mathcal{N}|}$  as the set of states consistent with  $\hat{s}_t$  where each element  $z \in \mathcal{Z}$  represents a possible configuration of occluded factors. For example, in Fig. 1, the ego-vehicle can observe only cars 1–3, thus  $\hat{s}_t$  includes only their poses and velocities. However,

this state is also consistent with the state that extends  $\hat{s}_t$  by including a vehicle with pose at the location of car 4.

In general, the size of  $\mathcal{Z}$  may be infinite, thus for tractability, we restrict ourselves to considering a finite set. While the algorithm we introduce later is capable of reasoning over any finite subset of  $\mathcal{Z}$ , in this paper we assume there are  $k$  locations at which the presence of an unseen vehicle may be possible. In practice, suitable locations could be inferred using blindspot detection and/or a map of the static road layout. Then, we define the set  $\mathcal{Z} := \{z_{\mathbf{v}} \mid \mathbf{v} \in \{0, 1\}^k\}$ , where  $z_{\mathbf{v}}$  corresponds to a state consistent with  $\hat{s}$  but with the presence of an occluded factor at location  $j$  if  $v_j = 1$ . We refer to elements of  $\mathcal{Z}$  as *occluded factor instantiations*.

Our objective is to use the observed behaviour of each  $i \in \mathcal{I}$  to infer a probabilistic belief  $p(g, z | \hat{s}_{1:t}) \forall z \in \mathcal{Z}, g \in \mathcal{G}^i$ , i.e., the likelihood that vehicle  $i$  has goal  $g$  and there is an occluded factor instantiation  $z$  that influences its behaviour. We can then use these probabilities to predict the future behaviour of vehicle  $i$ , infer the existence of unseen vehicles or obstacles, and plan safer actions.

## III. RELATED WORK

*a) Goal Recognition:* Our work addresses the problem of inferring goals of autonomous vehicles as an intermediate step to predicting their future trajectory. Goal recognition algorithms powered by deep neural networks [8], [9], [19], [26] have demonstrated accurate prediction by supervised-learning on large datasets. However, these methods lack interpretability which is a challenge for their deployment in trustworthy autonomous systems [16]. The IGP2 algorithm of Albrecht et al. [2] is closest to the algorithm we introduce in Section IV; we base inference on inverse-planning for interpretability and then integrate the resulting probabilistic belief into MCTS. However, our work differs in that we infer the presence of potential occluded factors along with goals. As we show in our empirical analysis, a goal recognition only approach (representative of IGP2) produces incorrect beliefs in the presence of occluded factors. Outside of autonomous driving, goal (or plan) recognition has a long history in AI research [4], [28]. Ramirez and Geffner [23] introduced one of the first goal recognition via inverse-planning algorithms. This work inspired a number of later works which generalised the settings in which inverse-planning goal recognition could be performed [6], [24], [27], [32]. The most relevant to our work is the work of Ramirez and Geffner [25], however, this earlier work assumes the observer (i.e., ego-vehicle) can compute the belief-state of the observed agent (i.e., non-ego vehicles). In our setting, this assumption would amount to knowing if occluded factors are present.

*b) Occluded Factors in Autonomous Driving:* Recent works have considered how occlusion affects vehicle tracking [12], [20] and motion planning [7], [13], [30]. These problems are orthogonal to the problem of occluded factor inference from the behaviour of an observed vehicle. The works most closely related to ours are the works of Sun et al. [29] and Afolabi et al. [1] who infer occluded obstacles from the actions of observed vehicles. Sun et al. [29] infer

so-called *social information* (i.e., occluded factors) with an inverse-planning approach based on a reward function learned with inverse reinforcement learning. Afolabi et al. [1] use human behaviour to identify an occupancy-grid representation of an environment. Neither of these works consider goal recognition and occluded factor inference jointly. Instead of attempting to infer the existence of occluded factors, other works attempt cautious planning when occlusion is possible. For example, Morales et al. [21] use inverse reinforcement learning in the presence of potential occluded obstacles to learn a reward function describing how humans handle such scenarios. Zhang et al. develop a game-theoretic framework for generating behaviours that avoid collision with unobserved other vehicles [35]. While a cautious approach is reasonable when the full state of the world cannot be observed, our approach effectively observes possible occluded factors using the behaviour of other vehicles as a sensor.

#### IV. GOAL AND OCCLUDED FACTOR INFERENCE

In this section, we introduce an algorithm for inferring a probabilistic belief about a non-ego’s goals and occluded factors from the behaviour of that non-ego vehicle. We first motivate the need for joint inference of goals and occluded factors. We then introduce an inverse-planning algorithm for conducting this inference.

##### A. Occluded Factors Confound Goal Recognition

Our aim is to estimate  $\Pr(g, z | \hat{s}_{1:t}^i)$  for each observed vehicle  $i \in \mathcal{I}$ , each possible goal  $g \in \mathcal{G}^i$ , and each possible occluded factor instantiation  $z \in \mathcal{Z}$ . Prior work has shown how  $\Pr(z | \hat{s}_{1:t}^i)$  [1], [29] and  $\Pr(g | \hat{s}_{1:t}^i)$  [2] can be inferred in isolation for inverse-planning prediction of non-ego behaviour. Since  $\hat{s}_{1:t}$  contains information about both  $g^i$  and  $z$ , we might hope that we can simply put these prior approaches together to infer both  $g$  and  $z$  from the same set of past observations. Unfortunately, since past observations,  $\hat{s}_{1:t}^i$ , depend on both  $g^i$  and  $z$ , goals and occluded factors are dependent conditioned on past observations. While it is reasonable to assume that goals and occluded factors are independent a priori, conditioning on  $\hat{s}_{1:t}$  breaks independence and necessitates joint inference, i.e.,  $\Pr(z | \hat{s}_{1:t}^i) \cdot \Pr(g | \hat{s}_{1:t}^i) \neq \Pr(g, z | \hat{s}_{1:t}^i)$ .

Intuitively, an optimal trajectory for one goal may appear sub-optimal when influenced by an occluded factor that the ego-vehicle cannot observe. The result is that the observed trajectory may provide evidence *against* the non-ego’s true goal instead of *for* it. This case is illustrated by the example in Section I in which the non-ego vehicle pauses at the junction. The pause is sub-optimal if the goal is to turn, however, the pause is *not* sub-optimal for turning if we know there is an occluded vehicle.

##### B. An Algorithm for Goal and Occluded Factor Inference

To address this challenge, we introduce the **Goal and Occluded Factor Inference (GOFI)** algorithm. GOFI outputs a probabilistic belief over possible occluded factor states,  $\Pr(z | \hat{s}_{1:t}^i)$ , and goals of the non-ego,  $\Pr(g | z, \hat{s}_{1:t}^i)$ . These

beliefs can then be used downstream for ego-vehicle planning, as we show in Section V.

GOFI is based on the rational inverse-planning IGP2 algorithm introduced by Albrecht et al. [2]. However, since IGP2 does not model occluded factors, it may determine that a vehicle is acting irrationally with respect to one goal even if it is perfectly rational once occluded factors are taken into account. Our main algorithmic contribution is to extend this algorithm to handle such occluded factors in goal recognition and to infer occluded factors along with goals.

Rational inverse-planning computes the optimal plan for a non-ego vehicle to achieve a potential goal given a potential instantiation of occluded factors. The joint probability of that goal and occluded factor instantiation is computed based on the difference between the non-ego’s observed behaviour and this optimal plan. Specifically, let  $s_{1:T}^{*i}(g, z)$  denote the optimal trajectory for vehicle  $i$  from  $s_1^i$  to  $g$  and let  $c^*(g, z) \in \mathbb{R}_0^+$  be its cost under a pre-defined cost function. Let  $s_{1:T}^{+i}$  be the concatenation of the observed trajectory  $\hat{s}_{1:t}^i$  and the optimal trajectory from  $s_t^i$  to  $g$ ; let  $c^+(g, z) \in \mathbb{R}_0^+$  be its cost. GOFI begins by computing the unnormalised likelihood of the observed trajectory for an observed non-ego vehicle  $i$  conditioned on a particular  $g$  and  $z$  instantiation. Following Baker et al. [6] and Ramirez and Geffner [24], we parameterise the likelihood as a Boltzmann distribution with temperature  $\beta^{-1}$ :

$$L(\hat{s}_{1:t}^i | g, z) := \exp(\beta(c^*(g, z) - c^+(g, z))). \quad (1)$$

This computation is repeated for each potential goal and occluded factor instantiation. From the computed likelihoods, the joint probability of any occluded factor state and goal can be computed with Bayes’ rule as:

$$\Pr(g, z | \hat{s}_{1:t}^i) \propto L(\hat{s}_{1:t}^i | g, z) p(g) p(z). \quad (2)$$

We then obtain  $\Pr(z | \hat{s}_{1:t}^i)$  and  $\Pr(g | \hat{s}_{1:t}^i, z)$  via marginalisation and a further application of Bayes rule. Pseudocode for GOFI is given in Algorithm 1.

---

#### Algorithm 1 Goal and Occluded Factor Inference (GOFI)

---

**Input:** vehicle  $i$ , state estimates  $\hat{s}_{1:t}$ , possible goals  $\mathcal{G}^i$ , set of occluded factor  $\mathcal{Z}$

**Returns:** occluded factor probabilities  $\Pr(z | \hat{s}_{1:t}^i)$  and goal probabilities  $p(g^i | \hat{s}_{1:t}^i, z)$

- 1: Set prior probabilities  $p(g^i)$ ,  $p(z)$  (e.g. uniform)
  - 2: **for all**  $z \in \mathcal{Z}$  **do**
  - 3:     **for all**  $g^i \in \mathcal{G}^i$  **do**
  - 4:          $s_{1:T}^{*i} \leftarrow \text{PLANOPTIMAL}(\hat{s}_1^i, g^i, z)$
  - 5:          $c^* \leftarrow \text{cost}(s_{1:T}^{*i}, z)$
  - 6:          $s_{t+1:T}^{+i} \leftarrow \text{PLANOPTIMAL}(\hat{s}_t^i, g^i, z)$
  - 7:          $s_{1:t}^{+i} \leftarrow \hat{s}_{1:t}^i$
  - 8:          $c^+ \leftarrow \text{cost}(s_{1:T}^{+i}, z^i)$
  - 9:          $L(\hat{s}_{1:t}^i | g^i, z) \leftarrow \exp(\beta(c^* - c^+))$
  - 10:  $\Pr(z | \hat{s}_{1:t}^i) \propto \sum_g L(\hat{s}_{1:t}^i | g, z) p(g) p(z)$
  - 11:  $\Pr(g^i | \hat{s}_{1:t}^i, z) \leftarrow L(\hat{s}_{1:t}^i | g^i, z) p(g^i) p(z) / p(z | \hat{s}_{1:t}^i)$
  - 12: **Return**  $\Pr(z | \hat{s}_{1:t}^i)$ ,  $\Pr(g^i | \hat{s}_{1:t}^i, z)$
-

Implementing GOFI requires knowledge of the set of possible goals and occluded factors. The set of possible goals can be computed heuristically, e.g., using the static road layout and positing possible goals at landmarks such as the end of visible lanes. The set of possible occluded factors could be obtained by determining where blindspots are from a perception module and a map of the static road layout. The main limitation of GOFI is that it involves computing  $2|\mathcal{G}||\mathcal{Z}|$  plans which may be computationally demanding in heavily occluded scenarios such as in dense traffic. Fortunately, the structure of GOFI lends itself to a parallel implementation and thus the run time can (in principle) be reduced to close to the time for the longest running call to PLANOPTIMAL and subsequent cost function computation.

GOFI also makes use of a prior distribution over elements of each of these sets (Algorithm 1, Line 1). These prior distributions can be used to encode prior knowledge about goals and occluded factors (e.g., historical traffic data suggesting that in the evening most vehicles are turning into a nearby residential area) or incorporate information from perception. Signals from other vehicles can also be incorporated into priors: a vehicle which is signalling to turn is likely to have a goal that requires the turn (though it is *not* certain to turn, as the signal may be accidental). In the absence of such knowledge these priors can default to the uniform distribution.

A final consideration for GOFI is how to merge the ego’s beliefs when observing multiple vehicles. Goals for each vehicle can be inferred independently from each vehicle [2], however, the set of occluded factors is shared across vehicles. Thus, runs of GOFI for each vehicle may produce different beliefs  $\Pr(z|\hat{s}_{1,t})$ . Belief merging across multiple non-ego vehicles can be handled by using the posterior distribution over occluded factors given the observed trajectory of one vehicle as the prior distribution when processing the observed trajectory of a second vehicle.

## V. INTEGRATING INFERENCE INTO PLANNING

Ultimately, we seek to infer vehicle goals and occluded factor instantiations to improve ego-vehicle action selection. In this section we show how an inferred belief about goals and occluded factors can be integrated into MCTS [15] action selection. MCTS builds a search tree of possible ego-vehicle action sequences by running a fixed number of iterations, exploring possible action sequences. At the start of each iteration, we sample the occluded factor instantiation and goal instantiations for each observed non-ego vehicle from the belief output by GOFI.<sup>1</sup> The sampled instantiations are then treated as the ground-truth occluded factor instantiation and goal assignments for that iteration of MCTS. Using sampled beliefs as ground-truth is known as *determinisation* in the planning under partial observability literature [34]. Full pseudo-code for this procedure is given in Appendix I.

<sup>1</sup>When considering multiple non-ego vehicles, the occluded factor instantiation should be sampled from the merged belief across vehicles as described in the previous section.

MCTS can be used with either finite or real-valued action-spaces [17], [18]. To facilitate long-horizon planning, we follow Albrecht et al. [2] and plan using *macro-actions* where a macro-action is a sequence of pre-defined maneuvers such as turn, change-lane, or follow-lane. An example of a macro-action is “follow lane to exit” which entails a follow-lane maneuver followed by a turn. Maneuvers define trajectories consisting of a reference path and target speeds along the path, which are mapped to steering and acceleration controls via PID control or adaptive cruise control.

The key requirement for this planning process is to provide adequate behaviour models for both observed and non-ego vehicles, and inferred vehicles or obstacles corresponding to occluded factors. We assume observed vehicles follow the optimal trajectory to the sampled goal. We assume that inferred vehicles corresponding to occluded factors follow simple behaviours such as constant velocity lane-following or a stationary pose for inferred static obstacles. This model is convenient since these optimal trajectories are already computed by GOFI. To reduce computation time, we pre-sample trajectories for each vehicle (both observed and inferred) at the beginning of each Monte Carlo tree search (MCTS) iteration. The ego-vehicle then plans to find maneuvers that avoid collisions with these trajectories.

## VI. EMPIRICAL EVALUATION

We next design and conduct an empirical analysis to answer the questions: (1) Does joint inference of goals and occluded factors increase inference accuracy?; (2) Does GOFI increase the safety of MCTS planning in the presence of occluded factors?

### A. Empirical Set-up

To answer our empirical questions, we conduct experiments in fixed-frame driving scenarios (shown in Fig. 2).<sup>2</sup> These scenarios were identified as realistic settings where the presence of occluded factors confound goal recognition and where recognition of occluded factors is necessary to avoid collisions. GOFI provides a general framework for joint inference and is instantiated with a particular method for PLANOPTIMAL. In our experimental implementation, we implement PLANOPTIMAL with A\* search over a finite set of pre-defined maneuvers. Implementation details for A\* follow Albrecht et al. [2].

**Baseline Methods:** We use the following methods as prediction and planning baselines. 1) **GR-Only:** performs inverse-planning goal recognition based on Albrecht et al. [2] without modelling occluded factors. This baseline is implemented by running Algorithm 1 with  $\mathcal{Z}$  only consisting of the state corresponding to no occluded factors. 2) **OF-Oracle:** the same as GR-Only except  $\mathcal{Z}$  only consists of the state corresponding to occluded factors present. This approach serves as an upper bound for inferring the true goals when occluded factors are present (as they are in our main experiments). 3) **Goal-Oracle:** performs occluded factor inference assuming the observed vehicle’s goal is

<sup>2</sup>Scenario videos can be found at: <https://www.five.ai/gofi>




Fig. 2: **Evaluation Scenarios.** In all scenarios the ego-vehicle is shown in blue; non-ego, observed vehicles are grey; and occluded factor vehicles are purple. The ego-vehicle’s goal is represented by a blue circle. Possible goals for the non-ego observed vehicle are shown by yellow/green circles, with optimal paths to each goal. In each case, the true goal of the non-ego observed vehicle is G1.

known, i.e., Algorithm 1 with  $\mathcal{G}^i = \{g^i\}$ . This approach is an upper bound for inferring occluded factors. 4) **MAP:** uses a maximum a posteriori combination of GOFI and MCTS. GOFI is used for goal and occluded factor inference but only the most likely occluded factor instantiation and inferred goal are considered for predicting the future trajectory of non-ego vehicles during MCTS. In all methods, we use a uniform prior on the possible goals and a prior of 0.1 on the occluded factor instantiation with an occluded vehicle present. We use a lower prior on an occluded vehicle’s presence since, in the absence of other evidence, we wish to avoid assuming existence of an occluded vehicle which could lead to overly cautious driving.

**Evaluation Scenarios:** In **Scenario 1**, the non-ego vehicle blocks the ego’s view of a pedestrian crossing the road before a junction. The non-ego vehicle slows down to allow the pedestrian to cross into the ego’s lane. This behaviour could also indicate preparation to turn. In **Scenario 2**, the ego-vehicle and observed non-ego travel to the end of their initial lane. The observed non-ego changes lane to avoid a stopped vehicle that the ego-vehicle cannot observe. It then switches back to the left lane after overtaking the stopped vehicle. In **Scenario 3**, a building obscures the ego’s view of an oncoming vehicle while the observed non-ego stops and waits. The wait could indicate that the non-ego’s goal is to turn. In **Scenario 4**, the ego-vehicle cannot observe a stopped car around the corner that causes the non-ego vehicle to slow down as it approaches the junction. The slow down could also be indicative of a vehicle preparing to turn (without signalling). Each scenario considers a single occluded factor and a single non-ego vehicle which has two possible goals. Inaccurate inference may result in a collision with the occluded factor (Scenarios 1-3) or the visible non-ego vehicle (Scenario 4). The initial position and velocity of each vehicle is randomised (except for the static occluded

vehicles in Scenarios 2 and 4).

Since our evaluation focuses on behaviour-based inference of occluded factors, in our main experiments we first test a setting where the ego-vehicle is always blind to occluded vehicles and pedestrians – they can only be inferred from the behaviour of the observed non-ego vehicle. In some of our scenarios, the ego-vehicle would eventually be able to see the occluded factor. Thus, in Section VI-C, we also evaluate safe driving in Scenarios 1, 2 and 3 when the ego-vehicle can eventually observe the occluded factors. These experiments show how behaviour-based inference of occluded factors complements perception and leads to safer driving.

## B. Inference Evaluation

To answer our first empirical question, we run GOFI and baselines to infer goals and occluded factors for the non-ego vehicle in each scenario. We compare GOFI and baselines in terms of their posterior probability on the true occluded factor and true goal for the non-ego vehicle (higher is better). These posterior probabilities are shown in Fig. 3.

From the middle plots in each subfigure of Fig. 3, we can observe that GOFI is always able to increase its posterior belief in the presence of an occluded vehicle as it observes more of the non-ego’s behaviour. The Goal-Oracle method is given the correct goal of the observed non-ego and thus only has to consider occluded factors. This makes its posterior an upper bound for GOFI. However, in some scenarios we observe that GOFI is near this upper bound.

The bottom plots in each subfigure of Fig. 3 give the posterior belief in the correct non-ego goal. In these plots, OF-Oracle is aware of the presence of the occluded factor and thus serves as an upper bound on obtainable posteriors. We note that in some scenarios there is sufficient ambiguity such that, even with this knowledge, the true goal is hard to infer. On the other hand, GR-Only does *not* model the possibility of occluded factors and thus obtains the lowest posterior beliefs on the true goal. In one case, GR-Only converges to belief in the wrong goal (Scenario 4). GOFI infers a higher posterior than GR-Only because it models the possibility that the non-ego’s actions are due to the presence of an occluded factor. It also never converges to belief in the incorrect goal.

Our scenarios only consider the case where an occluded factor is present. In Appendix II we include evaluation in Scenarios 2 and 4 in which the occluded factor is *not* present and the non-ego’s goal is switched to G2. In these settings we expect to see that GOFI converges to the prior belief on the presence of occluded factors and these additional results confirm this hypothesis.

## C. MCTS Integration Evaluation

Next, we evaluate the integration of GOFI with MCTS. Figure 4 shows the fraction of runs completed in each scenario without a collision. Of the non-oracle methods (GOFI, MAP, GR-Only), GOFI obtains the highest completion percentage in all scenarios. OF-Oracle knows an occluded factor exists and Goal-Oracle knows the true goal in all






Fig. 3: **Scenario Results.** *Top:* Trajectory of the observed non-ego vehicle. In Fig. 3b, we plot the vertical position of the non-ego over time since it better describes the non-ego’s behavior than velocity. *Middle:* Posterior probability for the belief in the presence of an occluded obstacle over time (higher is better). *Bottom:* Posterior probability for belief in the true goal over time. Beliefs are averaged over 20 repetitions (shaded area represents one standard error).

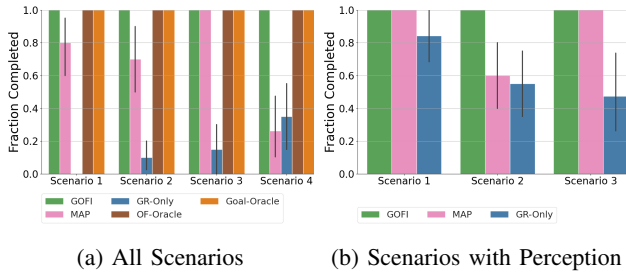


Fig. 4: Fraction of trials without a collision. Error bars give a bootstrap 95% confidence interval.

scenarios. This privileged knowledge allows these methods to upper bound the performance of GOFI but could not be realised in practice.

In terms of baselines without privileged information, the comparison to MAP shows the importance of our sampling-based integration of an inferred belief with MCTS. Since MAP plans with the most likely goal and occluded factor under the posterior, it ignores low-probability – but potentially costly – situations when selecting actions. GR-Only also experiences more collisions than GOFI as it cannot infer the presence of occluded vehicles. This comparison shows how using the behaviour of other vehicles to “sense” occluded factors yields safer planning.

Finally, while the completion rates shown in Fig. 4 are measured *without* perception of occluded factors, we also evaluate versions of Scenarios 1, 2, and 3 in which the ego-vehicle can perceive the occluded factor. In Scenario 1, the pedestrian is occluded if a line from the ego-vehicle to the pedestrian intersects a box around the non-ego vehicle. In Scenario 2, we base occlusion on the ego-vehicles position; once the ego-vehicle reaches the position where the non-ego vehicle began its overtake maneuver, we set the ego’s belief in the presence of the occluded vehicle to 1.0. In Scenario 3, we determine occlusion based on whether the ego-vehicle’s line of sight to the occluded vehicle is obscured by a building. This change improves the completion percentage of GR-Only and MAP (see Fig. 4b). However, due

to the fixed 1 Hz control frequency of the MCTS planner, these baselines still cannot always avoid collisions. GOFI complements perception with inference based on the non-ego’s behaviour and so is able to plan for the potential of an occluded factor sooner and always avoid collisions. While MAP also avoids all collisions in Scenarios 1 and 3, it fails to do so in Scenario 2.

## VII. CONCLUSION

In this paper we considered the problem of inferring the goals of an observed vehicle in the presence of occluded factors which may confound goal recognition. This problem arises in autonomous driving scenarios with occluded vehicles and pedestrians. We showed how unobserved occluded factors may confound goal recognition, introduced an interpretable inverse-planning algorithm for joint goal and occluded factor inference, and demonstrated how it can be integrated into MCTS action selection. In addition to more accurate goal recognition, our algorithm infers the presence of unobserved occluded factors based on the behaviour of observed vehicles. We performed an empirical evaluation of our framework in simulated autonomous driving scenarios that require occluded factor inference for safe driving. Our empirical results demonstrated the necessity of joint goal and occluded factor inference as goal recognition alone misidentified goals in the presence of occluded factors and resulted in more collisions in safety critical scenarios. Our Goal and Occluded Factor Inference algorithm can simultaneously infer goals and occluded factors, leading to fewer collisions in safety critical scenarios.

In the future, we aim to address the construction of the set of possible occluded factors. This problem requires identifying potential locations of occluded factors (i.e., occluded locations) as well as the set of behaviours consistent with identified locations. We will also consider how to select information gathering actions to resolve uncertainty about occluded factors.

## ACKNOWLEDGMENT

This research was in part supported by the Royal Society via an Industry Fellowship (S.A.), the Alan Turing Insti-

tute via project "Model Criticism in Multi-Agent Systems" (J.H.), and UK Research and Innovation via NPIF Innovation Placements (E.F.). J.H. was with Five AI and School of Informatics, University of Edinburgh while this research took place. A.R. and E.F. were interns at Five AI while this research took place. We would like to thank Majd Hawasly and others at Five AI for insightful comments on this work.

## APPENDIX I EXTENDED MCTS INTEGRATION DESCRIPTION

In this appendix we provide full pseudocode (Algorithm 2) for our integration of the beliefs obtained with GOFI into MCTS. MCTS runs a fixed number of iterations to build up a search tree in which nodes correspond to sequences of actions. Recall from Section V that we follow Albrecht et al. [2] and plan using an action-space of *macro-actions* where a macro-action is a sequence of pre-defined maneuvers such as turn, change-lane, or follow-lane. Maneuvers define trajectories consisting of a reference path and target speeds along the path, which are mapped to steering and acceleration controls via PID control or adaptive cruise control. Thus, even though planning is done using a high-level action-space, the MCTS iterations use low-level simulation of actions.

In the search-tree, each node,  $q$ , contains a value,  $Q(q, \mu)$ , for each macro-action  $\mu$  that gives the expected future reward, following choosing  $\mu$  in node  $q$ . Once these *action-values* are estimated, the optimal macro-action at time  $t$  can be selected as  $\arg \max_{\mu} Q(\text{root}, \mu)$ . Each iteration begins by sampling an instantiation of the occluded factors, non-ego goals, and trajectories for non-ego vehicles. The ego-vehicle then uses a given exploration strategy (we use UCB [5]) to select macro-action until a fixed search depth is reached or a termination condition is met. Upon termination, the final reward  $r$  is back-propagated through the search tree using the following update rule based on Q-learning [33]:

$$Q(q, \mu) \leftarrow Q(q, \mu) + \begin{cases} \delta^{-1}[r - Q(q, \mu)] & \text{if } q \text{ leaf node, else} \\ \delta^{-1}[\max_{\mu'} Q(q', \mu') - Q(q, \mu)] & \end{cases} \quad (3)$$

where  $\delta$  is the number of times maneuver  $\mu$  has been selected in node  $q$ .

## APPENDIX II ADDITIONAL EMPIRICAL RESULTS

In the main paper, we evaluated GOFI and baselines under a single instantiation of goals and occluded factors in each scenario. In this appendix we include results for Scenarios 2 and 4 with no occluded factor present and the visible non-ego's goal switched from G1 to G2. With these changes, GR-Oracle and OF-Oracle, that assume G1 and occluded factor presence respectively, use incorrect assumptions.

Fig. 5 shows the inferred probabilities for each variation. For occluded factor recognition, we see in both scenarios that Goal-Oracle is misled by its false assumption that the non-ego is heading to G1. The false assumption leads it to explain the non-ego's behaviour as indicative of the presence of an occluded factor (note that we end the plotted beliefs

---

### Algorithm 2 Monte Carlo Tree Search algorithm

---

**Returns:** optimal maneuver for the ego-vehicle  $\varepsilon$  in state  $\hat{s}(t)$

Perform  $k$  iterations:

- 1: Search node  $q.s \leftarrow \hat{s}t$  (root node)
  - 2: Search depth  $d \leftarrow 0$
  - 3: Sample occluded factor instantiation  $z \sim \Pr(z|\hat{s}_{1:t})$
  - 4: **for all**  $i \in \mathcal{I} \setminus \{\varepsilon\}$  **do**
  - 5:   Sample goal  $g^i \sim p(g^i | \hat{s}_{1:t}, z)$
  - 6:    $s_{1:T}^i \leftarrow \text{PLANOPTIMAL}(\hat{s}_t^i, g^i, z)$
  - 7: **while**  $d < d_{\max}$  **do**
  - 8:   Select maneuver  $\mu$  for  $\varepsilon$  applicable in  $q.s$
  - 9:    $s_{\tau, \ell} \leftarrow$  Simulate maneuver until it terminates, with other vehicles following their sampled trajectories  $s_{1:T}^i$
  - 10:    $r \leftarrow \emptyset$
  - 11:   **if** ego-vehicle collides during  $s_{\tau, \ell}$  **then**
  - 12:      $r \leftarrow c_{\text{coll}}$ .
  - 13:   **else if**  $s_{\ell}^{\varepsilon}$  achieves the ego-vehicle's goal  $g^{\varepsilon}$  **then**
  - 14:      $r \leftarrow c(s_{\ell:T})$
  - 15:   **else if**  $d = d_{\max} - 1$  **then**
  - 16:      $r \leftarrow r_{\text{term}}$ .
  - 17:   **if**  $r \neq \emptyset$  **then**
  - 18:     Use (3) to backprop  $r$  along search branches  $(q, \mu, q')$  traversed in the iteration
  - 19:     Start next iteration
  - 20:      $q'.s = s_{\ell}; q \leftarrow q'; d \leftarrow d + 1$
- Return maneuver  $\mu \in \arg \min_{\mu} Q(\text{root}, \mu)$
- 

at the point it is impossible for the non-ego to reach G1). In Scenario 2, GOFI initially increases its belief in the presence of an occluded factor after the lane-change. However, as the non-ego continues in the right lane, the belief returns to the prior. Once the non-ego turns, the belief returns to the prior since the non-ego's goal is sufficient for explaining the non-ego's observed behaviour. A similar trend is observed in Scenario 4 in response to the non-ego's slow-down and then eventual acceleration as it turns left. Note that, GOFI should *not* converge to a belief in the absence of an occluded factor given the observed non-ego behaviour as absence of evidence is not evidence of absence.

For goal recognition, since there is no occluded factor present, occluded factors do not have to be considered for identifying the non-ego's goal. Thus GR-Only has the most accurate inference for goal recognition. Since OF-Oracle (falsely) assumes the presence of an occluded factor, it explains the lane change as obstacle-avoidance and thus does not increase its belief in the true goal. In both scenarios, GOFI obtains a posterior belief that is between OF-Oracle and GR-Only. Recall from Fig. 3 in the main text that the relative performance of the baselines is flipped under a different setting of the non-ego goals and occluded factor presence. Taken together these results show that by jointly recognising the factors explaining the non-ego behaviours GOFI can converge to higher posteriors on the true factors.




Fig. 5: **Scenario Variation Results.** *Top:* Trajectory of the observed non-ego vehicle. *Middle:* Belief in the presence of an occluded obstacle over time (higher is better). *Bottom:* Belief in the true follow-road goal over time. Beliefs averaged over 20 repetitions (shaded area represents standard error).

## REFERENCES

- [1] Oladapo Afolabi, Katherine Driggs-Campbell, Roy Dong, Mykel J. Kochenderfer, and S. Shankar Sastry. People as sensors: Imputing maps from human actions. In *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 2342–2348. IEEE, 2018.
- [2] Stefano V. Albrecht, Cillian Brewitt, John Wilhelm, Balint Gyevnar, Francisco Eiras, Mihai Dobre, and Subramanian Ramamoorthy. Interpretable goal-based prediction and planning for autonomous driving. In *Proceedings of the IEEE Conference on Robotics and Automation (ICRA)*, 2021.
- [3] Stefano V. Albrecht and Subramanian Ramamoorthy. Exploiting causality for selective belief filtering in dynamic Bayesian networks. *Journal of Artificial Intelligence Research*, 55:1135–1178, 2016.
- [4] Stefano V. Albrecht and Peter Stone. Autonomous agents modelling other agents: A comprehensive survey and open problems. *Artificial Intelligence*, 258:66–95, 2018. DOI: 10.1016/j.artint.2018.01.002.
- [5] Peter Auer, Nicolo Cesa-Bianchi, and Paul Fischer. Finite-time analysis of the multiarmed bandit problem. *Machine Learning*, 47(2-3):235–256, 2002.
- [6] Chris L. Baker, Rebecca Saxe, and Joshua B. Tenenbaum. Action understanding as inverse planning. *Cognition*, 113(3):329–349, 2009.
- [7] Maxime Bouton, Alireza Nakhai, Kikuo Fujimura, and Mykel J. Kochenderfer. Scalable decision making with sensor occlusions for autonomous driving. In *2018 IEEE International Conference on Robotics and Automation (ICRA)*, pages 2076–2081. IEEE, 2018.
- [8] Sergio Casas, Wenjie Luo, and Raquel Urtasun. {IntentNet:} Learning to Predict Intention from Raw Sensor Data. In *Conference on Robot Learning*, pages 947–956, 2018.
- [9] Yuning Chai, Benjamin Sappm, Mayank Bansal, and Dragomir Anguelov. MultiPath: Multiple Probabilistic Anchor Trajectory Hypotheses for Behavior Prediction. In *Proceedings of the 3rd Conference on Robot Learning*, 2019.
- [10] Nachiket Deo and Mohan M. Trivedi. Convolutional social pooling for vehicle trajectory prediction. In *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, pages 1549–1548, Salt Lake City, Utah, USA, June 2018. IEEE.
- [11] 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 *International Conference on Machine Learning*, volume 119, pages 3145–3153. PMLR, November 2020.
- [12] Enric Galceran, Edwin Olson, and Ryan M. Eustice. Augmented vehicle tracking under occlusions for decision-making in autonomous driving. In *2015 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 3559–3565. IEEE, 2015.
- [13] Ezequiel González Debada, Adeline Ung, and Denis Gillet. Occlusion-aware motion planning at roundabouts. *IEEE Transactions on Intelligent Vehicles*, 2020.
- [14] ByeoungDo Kim, Chang Mook Kang, Jaekyum Kim Kim, Seung Hi Lee, Chung Choo Chung, and Jun Won Choi. Probabilistic vehicle trajectory prediction over occupancy grid map via recurrent neural network. In *2017 IEEE 20th International Conference on Intelligent Transportation Systems (ITSC)*, pages 399–404, Yokohama, Japan, October 2017. IEEE.
- [15] Levente Kocsis and Csaba Szepesvári. Bandit based Monte-Carlo planning. In *Proceedings of the 17th European Conference on Machine Learning*, pages 282–293. Springer, 2006.
- [16] Philip Koopman, Rob Hierons, Siddhartha Khastgir, John Clark, Michael Fisher, Rob Alexander, Kerstin Eder, Pete Thomas, Geoff Barrett, Philip Torr, et al. Certification of highly automated vehicles for use on uk roads: Creating an industry-wide framework for safety. Technical report, Five AI Ltd, 2019.
- [17] Karl Kurzer, Christoph Hörtnagl, and J Marius Zöllner. Parallelization of monte carlo tree search in continuous domains. *arXiv preprint arXiv:2003.13741*, 2020.
- [18] Jongmin Lee, Wonseok Jeon, Geon-Hyeong Kim, and Kee-Eung Kim. Monte-carlo tree search in continuous action spaces with value gradients. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 4561–4568, 2020.
- [19] Namhoon Lee, Wongun Choi, Paul Vernaza, Christopher B. Choy, Philip H. S. Torr, and Manmohan Chandraker. DESIRE: Distant future prediction in dynamic scenes with interacting agents. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 336–345, 2017.
- [20] Jiachen Li, Wei Zhan, and Masayoshi Tomizuka. Generic vehicle tracking framework capable of handling occlusions based on modified mixture particle filter. In *2018 IEEE Intelligent Vehicles Symposium (IV)*, pages 936–942. IEEE, 2018.
- [21] Luis Yoichi Morales, Akai Naoki, Yuki Yoshihara, and Hiroshi Murase. Towards predictive driving through blind intersections. In *2018 21st International Conference on Intelligent Transportation Systems (ITSC)*, pages 716–722. IEEE, 2018.
- [22] Henry Pulver, Francisco Eiras, Ludovico Carozza, Majd Hawasly, Stefano V. Albrecht, and Subramanian Ramamoorthy. PILOT: Efficient planning by imitation learning and optimisation for safe autonomous driving. *arXiv:2011.00509 [cs]*, November 2020.
- [23] Miquel Ramírez and Hector Geffner. Plan recognition as planning. In *Twenty-First International Joint Conference on Artificial Intelligence*, 2009.
- [24] Miquel Ramírez and Hector Geffner. Probabilistic plan recognition using off-the-shelf classical planners. In *24th AAAI Conference on Artificial Intelligence*, pages 1121–1126, 2010.
- [25] Miquel Ramírez and Hector Geffner. Goal recognition over pomdps: Inferring the intention of a pomdp agent. In *IJCAI*, pages 2009–2014. Citeseer, 2011.
- [26] Nicholas Rhinehart, Rowan McAllister, Kris Kitani, and Sergey Levine. PRECOG: Prediction conditioned on goals in visual multi-agent settings. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 2821–2830, 2019.
- [27] Shirin Sohrabi, Anton V. Riabov, and Octavian Udrea. Plan recognition as planning revisited. In *IJCAI*, pages 3258–3264, 2016.
- [28] Gita Sukthankar, Christopher Geib, Hung Hai Bui, David Pynadath, and Robert P Goldman. *Plan, activity, and intent recognition: Theory and practice*. Newnes, 2014.
- [29] Liting Sun, Wei Zhan, Ching-Yao Chan, and Masayoshi Tomizuka. Behavior planning of autonomous cars with social perception. In *2019 IEEE Intelligent Vehicles Symposium (IV)*, pages 207–213. IEEE, 2019.
- [30] Omer Sahin Tas and Christoph Stiller. Tackling existence probabilities of objects with motion planning for automated urban driving. *arXiv preprint arXiv:2002.01254*, 2020.
- [31] Scott Thiebes, Sebastian Lins, and Ali Sunyaev. Trustworthy artificial intelligence. *Electronic Markets*, October 2020.
- [32] Mor Vered and Gal A. Kaminka. Heuristic online goal recognition in continuous domains. *arXiv preprint arXiv:1709.09839*, 2017.
- [33] Christopher J. C. H. Watkins and Peter Dayan. Q-learning. *Machine learning*, 8(3-4):279–292, 1992.
- [34] Daniel Whitehouse, Edward J Powley, and Peter I Cowling. Determinization and information set monte carlo tree search for the card game dou di zhu. In *2011 IEEE Conference on Computational Intelligence and Games (CIG’11)*, pages 87–94. IEEE, 2011.
- [35] Zixu Zhang and Jaime F Fisac. Safe occlusion-aware autonomous driving via game-theoretic active perception. In *In proceedings of Robotics: Science and Systems (RSS)*, 2021.