2023 IEEE International Conference on Robotics and Automation (ICRA 2023)
May 29 - June 2, 2023. London, UK

On Legible and Predictable Robot Navigation in Multi-Agent Environments

Jean-Luc Bastarache Christopher Nielsen Stephen L. Smith

***Abstract*— Legible motion is intent-expressive, which when employed during social robot navigation, allows others to quickly infer the intended avoidance strategy. Predictable motion matches an observer's expectation which, during navigation, allows others to confidently carryout the interaction. In this work, we present a navigation framework capable of reasoning on its legibility and predictability with respect to dynamic interactions, e.g., a passing side. Our approach generalizes the previously formalized notions of legibility and predictability by allowing dynamic goal regions in order to navigate in dynamic environments. This generalization also allows us to quantitatively evaluate the legibility and the predictability of trajectories with respect to navigation interactions. Our approach is shown to promote legible behavior in ambiguous scenarios and predictable behavior in unambiguous scenarios. In a multi-agent environment, this yields an increase in safety while remaining competitive in terms of goal-efficiency when compared to other robot navigation planners in multi-agent environments. The code of this work is made publicly available[1].**

I. INTRODUCTION

Without having to explicitly communicate their intentions, humans are able to seemingly effortlessly navigate amongst one another in a collision-free manner. A key constituent in human navigation is their ability to infer others' interaction intentions, e.g., to which side are they trying to pass, as shown in Fig. 1. In crowded environments, this inference allows them to cooperate in the interaction, enabling efficient navigation. Early robot navigation frameworks often overlooked or did not attempt to model these aspects, resulting in undesirable behaviors such as oscillations [1] or the freezing robot problem [2].

During locomotion, humans make use of numerous motion cues such as gaze [3] and head movement [4] to avoid collisions. Since mobile robots are not equipped with these motion cues, humans have much more difficulty in inferring their intentions. Studies have shown that humans are more conservative when avoiding moving inanimate objects [5] or objects with fixed limbs [6] than they are with humans.

Legible motion allows an observer to quickly and confidently infer an agent's intention [7]. Although largely studied for human robot interactions (HRI), legibility has become an important property to consider in the design of social robot navigation planners. Many navigation frameworks seeking to produce legible motion do so by considering the robot's legibility with respect to its global goal pose [8]–[10].

The authors are with the Department of Electrical and Computer Engineering, University of Waterloo, Waterloo, ON N2L 3G1, Canada {jbastarache,cnielsen,stephen.smith}@uwaterloo.ca

[1]https://github.com/jlbas/LPSNav

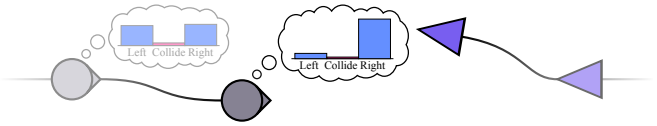

Fig. 1: Agent (gray) inferring a robot's (purple) avoidance strategy.

However, in dynamic environments, agents are unaware of others' global goals and adjust their trajectories based on the inferred avoidance strategy (e.g., the passing side).

Another important property of motion is that of predictability, which measures the degree to which motion matches an observer's expectation [7]. Predictability, as proposed in [7], is not well suited for motion planning applications since it is a function of entire trajectories.

Contributions: The contributions of this work are fourfold. First, we present an approach to explicitly model navigation interactions as dynamic goal regions. Second, we generalize legibility and predictability for static goal points as formalized in [7] to dynamic goal regions. Third, we propose an approximation of the predictability score for partial trajectories discussed in [11], which we use in our real-time planner. Lastly, we propose a navigation planner that is capable of optimizing over the legibility and the predictability of its motion with respect to navigation interactions.

Related Work: Recent research has sought to model social robot navigation as a cooperative collision avoidance task. In [12], the authors use topological braids to encode agents' joint behaviors. They show that their algorithm more rapidly decreases the uncertainty of the emerging avoidance strategy in the workspace. Deep reinforcement learning (DRL) approaches have demonstrated promising results [13], [14]. These approaches train policies that implicitly encode the agents' models and interaction intentions. However, complex learning-based approaches make it difficult to extract the social strategies being employed during navigation.

Early works claiming to generate legible motion often did so indirectly by targeting related properties [15]. In [16], the authors directly use legibility as an optimization criterion to generate legible motion, but only consider stationary goal points. Inspired by these ideas, researchers have begun to seek navigation strategies that allow an agent to quickly infer another agent's avoidance strategy. In [9], a cost function is formulated that takes into account social and context dependent costs. The authors show that their planner is able to generate legible motion with respect to the robot's underlying goal in the workspace. In [17], legible motion was shown to reduce the planning effort in a locomotion setting.

979-8-3503-2365-8/23/$31.00 ©2023 IEEE

In a multi-agent environment, *Social Momentum* makes use of topological braid theory to generate legible motion [18]. Our work is more closely related to this approach, since it explicitly reasons about the emerging collision avoidance strategy, rather than the global goal within the workspace.

II. PROBLEM FORMULATION

We consider a robot R moving in a planar workspace $\mathcal{W} \subseteq \mathbb{R}^2$ towards a stationary goal $\boldsymbol{g}_R \in \mathcal{W}$ and sharing the workspace with another dynamic agent A. An adaptation to the multi-agent case is provided in Section V. To only make use of what is readily observable in a social environment, it is assumed that the robot's goal is unknown to A and the robot has no means of explicitly communicating its intended goal. Through its onboard sensors, we assume that the robot has access to the position $\boldsymbol{p}_i(t) \in \mathcal{W}$ and velocity $\boldsymbol{v}_i(t) \in \mathbb{R}^2$, $i \in \{R, A\}$, for $\|\boldsymbol{p}_R(t) - \boldsymbol{p}_A(t)\| \leq d_{\text{sense}}$, where d_{sense} is the sensor range. The heading of the robot and agent at time t are, respectively, the angle that $\boldsymbol{v}_R(t)$ and $\boldsymbol{v}_A(t)$ make with respect to a fixed axis in the inertial frame.

The control space of the robot is a finite set of motion primitives $\mathcal{P}$ and each motion primitive $\rho_i \in \mathcal{P}$ has the same time duration $\delta t > 0$. A navigation plan for the robot is a sequence of primitive selections. The robot R is considered to be in a collision with the other agent A at time t if $\|\boldsymbol{p}_R(t) - \boldsymbol{p}_A(t)\| \leq r^C$ where $r^C > 0$ is a positive constant depending on the footprint of the robot and agent.

We seek to design a navigation strategy so the robot reaches its goal while simultaneously using the shape and speed of its trajectory to disambiguate its intentions to the other agent. To accomplish this, we take inspiration from the notions of legibility and predictability, in the sense of [7] and as reviewed in the next section, which have been validated in the HRI field [19]. The way in which legibility and predictability are to be optimized poses an important design decision. In this work, we design a navigation framework based on the following high-level principles:

1) A robot should disambiguate its intended navigation strategy by being legible.
2) Once legible, the robot should proceed predictably.
3) Motion should adhere to social norms (i.e., left or right passing conventions).

Notation: Let $\xi \colon [t_s, t_f] \subset \mathbb{R} \to \mathcal{W}$, $t_s < t_f$, denote a trajectory of the robot. The set of all trajectories is denoted by $\mathcal{T}$ and, for notational simplicity, we denote by $\xi_{\boldsymbol{a}\to\boldsymbol{b}}$ the trajectory along ξ from $\boldsymbol{p}_R(t_a)$ to $\boldsymbol{p}_R(t_b)$, where $t_s \leq t_a < t_b \leq t_f$, and $\boldsymbol{a} = \boldsymbol{p}_R(t_a)$ and $\boldsymbol{b} = \boldsymbol{p}_R(t_b)$.

III. LEGIBILITY AND PREDICTABILITY

We now review the notions of legibility and predictability.

Legible Motion: Let $\mathcal{G} \subset \mathcal{W}$ be a finite set and let $\mathbf{g}^\star \in \mathcal{G}$ denote the robot's goal. A partial trajectory $\xi_{\boldsymbol{s}\to\boldsymbol{t}}$ from the starting position to the current position is said to be legible when an observer can quickly and confidently infer $\boldsymbol{g}^\star$ [7]. This is modeled by the legibility inference function $f_L \colon \mathcal{T} \to \mathcal{G}$, mapping trajectories to goals, and $\xi_{\boldsymbol{s}\to\boldsymbol{t}}$ is legible if

$$f_L(\xi_{\boldsymbol{s}\to\boldsymbol{t}}) = \boldsymbol{g}^\star. \tag{1}$$

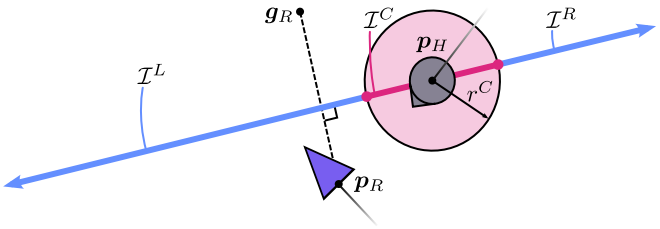

Fig. 2: Interaction line segmented into its collision segment $\mathcal{I}^C$ (magenta) and two rays, $\mathcal{I}^L$ and $\mathcal{I}^R$ (blue) representing passing on the left and right respectively.

The earliest t for which this is true is a measure of how legible the trajectory $\xi_{\boldsymbol{s}\to\boldsymbol{t}}$ is.

Predictable Motion: Given the goal $\boldsymbol{g}^\star$, the motion resulting from the entire trajectory $\xi_{\boldsymbol{s}\to\boldsymbol{g}^\star}$ from $\boldsymbol{s}$ to $\boldsymbol{g}^\star$ is said to be predictable when it matches the observer's inference [7]. This is modeled by the predictability inference function $f_P \colon \mathcal{G} \to \mathcal{T}$, mapping goals to trajectories, and $\xi_{\boldsymbol{s}\to\boldsymbol{g}^\star}$ is predictable if

$$f_P(\boldsymbol{g}^\star) = \xi_{\boldsymbol{s}\to\boldsymbol{g}^\star}. \tag{2}$$

The closeness of $\xi_{\boldsymbol{s}\to\boldsymbol{g}^\star}$ to the observer's inferred trajectory $f_P(\boldsymbol{g}^\star)$, given by a distance metric $d \colon \mathcal{T} \times \mathcal{T} \to \mathbb{R}_{\geq 0}$, is a measure of how predictable the trajectory $\xi_{\boldsymbol{s}\to\boldsymbol{t}}$ is.

IV. APPROACH

We present an approach that extends the notions of legibility and predictability from [7] to dynamic goal regions.

A. Defining an Interaction Using Dynamic Goal Regions

As the robot navigates towards its global goal $\boldsymbol{g}_R$, it will attain intermediate navigation goals, which we refer to as interaction goals. We consider three possible interaction goals: passing on the left, passing on the right and colliding.[2] The interaction goals are dynamic and infinite sets, distinguishing them from the static goal points in previous works [20]–[22].

Let $\mathcal{I}(t)$, hereinafter referred to as the interaction line, represent the line passing through the other agent's position $\boldsymbol{p}_A(t)$ oriented to be orthogonal to the vector pointing from the robot's position $\boldsymbol{p}_R(t)$ to its goal $\boldsymbol{g}_R(t)$ (see Fig. 2). Formally, at time t and with $\boldsymbol{p}_R(t) \neq \boldsymbol{g}_R$, define the unit vector

$$\boldsymbol{e}(t) \coloneqq \begin{bmatrix} 0 & -1 \\ 1 & 0 \end{bmatrix} \left(\frac{\boldsymbol{g}_R - \boldsymbol{p}_R(t)}{\|\boldsymbol{g}_R - \boldsymbol{p}_R(t)\|} \right) \tag{3}$$

and the real one-dimensional subspace $\mathcal{V}(t) \coloneqq \operatorname{span}\{\boldsymbol{e}(t)\}$, then the interaction line is the time-varying one-dimensional affine subspace

$$\mathcal{I}(t) \coloneqq \{\boldsymbol{v} + \boldsymbol{p}_A(t) : \boldsymbol{v} \in \mathcal{V}(t)\}. \tag{4}$$

The time-to-interaction (TTI) at time t is defined to be the infimum (possibly infinite) time it would take the robot R to reach the interaction line assuming the robot and other agent continue at their current velocities and $\mathcal{I}(t)$ translates with A. Robot R is said to be interacting with agent A if $\|\boldsymbol{p}_R(t) - \boldsymbol{p}_A(t)\| \leq d_{\text{sense}}$, the interaction line separates R

[2]The terminology for interactions is from the robot's point-of-view.

from its goal $\boldsymbol{g}_R$ and the TTI is less than or equal to a predefined maximum interaction time $t_{\mathcal{I}}^{\max} > 0$.

In order to define the interaction goals, the interaction line is segmented into passing and collision regions (see Fig. 2). The collision line segment, $\mathcal{I}^C$, is defined as the intersection between the interaction line and the closed disc centered at $\boldsymbol{p}_A$ of radius r^C

$$\mathcal{I}^C(t) \coloneqq \mathcal{I}(t) \cap \left\{\boldsymbol{p} \in \mathcal{W} : \|\boldsymbol{p} - \boldsymbol{p}_A(t)\| \leq r^C\right\}. \quad (5)$$

Each passing interaction, $\mathcal{I}^R$ and $\mathcal{I}^L$, is described by an open ray starting at the collision segment's endpoints, extending in the direction opposite to the other agent. The left passing side is defined as

$$\mathcal{I}^L(t) \coloneqq \left\{\boldsymbol{p}_A(t) + \alpha \boldsymbol{e}(t) : \alpha > r^C\right\} \quad (6)$$

and the right passing side is defined as

$$\mathcal{I}^R(t) \coloneqq \left\{\boldsymbol{p}_A(t) - \alpha \boldsymbol{e}(t) : \alpha > r^C\right\}. \quad (7)$$

These three interaction goals $\boldsymbol{\mathcal{G}} = \left\{\mathcal{I}^L, \mathcal{I}^C, \mathcal{I}^R\right\}$ represent dynamic goal regions fixed to the other dynamic agent. The passing interaction goals are denoted by $\boldsymbol{\mathcal{G}}^P \coloneqq \left\{\mathcal{I}^L, \mathcal{I}^R\right\}$. Orienting the interaction line to be orthogonal to $\boldsymbol{g}_R - \boldsymbol{p}_R(t)$ allows the robot to reason about a passing side irrespective of the other agent's heading. Note that the robot's global goal point, $\boldsymbol{g}_R \in \mathcal{W}$, is not an interaction goal, i.e. $\boldsymbol{g}_R \notin \boldsymbol{\mathcal{G}}$.

B. Robot and Observer Motion Models

We assume the observer expects the robot to be a rational agent seeking to move efficiently in the environment. As such, the observer's model of the robot's motion minimizes the cost functional

$$c[\xi] = (t_f - t_s)^2, \quad (8)$$

where $c\colon \mathcal{T} \to \mathbb{R}_{\geq 0}$ maps robot trajectories $\xi \in \mathcal{T}$ to the square of its duration.

Since the interaction goals defined in (5) to (7) are dynamic, evaluating (8) will require a trajectory prediction model. Given a prediction model chosen by the designer, the robot predicts the interaction goal from $\mathcal{I}(t)$ at the current time to the end of the interaction at time t_f with

$$\widehat{\mathcal{I}}_{t \to t_f} = \text{prediction}(\xi^A_{\boldsymbol{s} \to \boldsymbol{t}}, \mathcal{I}(t)), \quad (9)$$

where $\widehat{\mathcal{I}}_{t \to t_f} \in \boldsymbol{\mathcal{G}}$ is the predicted interaction in the closed interval $[t, t_f]$ and $\xi^A_{\boldsymbol{s} \to \boldsymbol{t}}$ is an observed segment of the other agent's trajectory.

C. Legibility and Predictability of Navigation Interactions

Given a robot trajectory $\xi_{\boldsymbol{s} \to \boldsymbol{t}}$ up to time t, we model the observer's legibility inference function (1) as returning the most likely interaction goal $\mathcal{I}(t)$ from the finite collection of possible goals $\boldsymbol{\mathcal{G}}(t)$:

$$f_L(\xi_{\boldsymbol{s} \to \boldsymbol{t}}) = \underset{\mathcal{I}(t) \in \boldsymbol{\mathcal{G}}(t)}{\operatorname{argmax}} P(\mathcal{I}(t) \mid \xi_{\boldsymbol{s} \to \boldsymbol{t}}). \quad (10)$$

We compute the above posterior probability following the derivations in [7]. This involves using Bayes' theorem to obtain the likelihood of $\mathcal{I}$, which we model using a Boltzmann policy whose partition function is approximated using Laplace's method as derived in [23]. Assuming a quadratic cost functional c, the posterior of interaction $\mathcal{I} \in \boldsymbol{\mathcal{G}}$ can be approximated by

$$P(\mathcal{I} \mid \xi_{\boldsymbol{s} \to \boldsymbol{t}}) \approx \frac{\exp\left(\beta\big(c[\xi^\star_{\boldsymbol{s} \to \widehat{\mathcal{I}}}] - c_{\widehat{\mathcal{I}}}[\xi_{\boldsymbol{s} \to \boldsymbol{t}}]\big)\right)}{\sum_{\bar{\mathcal{I}} \in \boldsymbol{\mathcal{G}}} P\left(\xi_{\boldsymbol{s} \to \boldsymbol{t}} \mid \bar{\mathcal{I}}\right) P\left(\bar{\mathcal{I}}\right)} P(\mathcal{I}), \quad (11)$$

where the denominator is a normalizer over goals $\mathcal{I} \in \boldsymbol{\mathcal{G}}$, $\xi^\star_{\boldsymbol{t} \to \widehat{\mathcal{I}}} = \operatorname{argmin}_{\xi_{\boldsymbol{t} \to \widehat{\mathcal{I}}} \in \mathcal{T}} c[\xi_{\boldsymbol{t} \to \widehat{\mathcal{I}}}]$ is the optimal cost to reach the predicted interaction goal (9), $c_{\widehat{\mathcal{I}}}$ is the cost to reach $\widehat{\mathcal{I}}$ through $\xi_{\boldsymbol{s} \to \boldsymbol{t}}$ and the optimal trajectory $\xi^\star_{\boldsymbol{t} \to \widehat{\mathcal{I}}}$, $P(\mathcal{I}) \in [0, 1]$ represents the prior on interaction $\mathcal{I}(t)$ with $\sum_{\mathcal{I} \in \boldsymbol{\mathcal{G}}} P(\mathcal{I}) = 1$ and $\beta_{\geq 0}$ acts as a rationality parameter [24]. To adhere to social norms, a larger prior could be assigned to the customary passing side.

Let us now assume the robot's interaction goal region at time t, $\mathcal{I}^\star(t)$ is known to the observer. We model the observer's predictability inference function (2) as the most likely trajectory $\xi_{\boldsymbol{t} \to \mathcal{I}^\star}$ from the set of possible trajectories $\mathcal{T}$ in the following sense:

$$f_P(\mathcal{I}^\star(t)) = \underset{\xi_{\boldsymbol{t} \to \mathcal{I}^\star(t)} \in \mathcal{T}}{\operatorname{argmax}} P(\xi_{\boldsymbol{t} \to \mathcal{I}^\star(t)} \mid \mathcal{I}^\star(t)). \quad (12)$$

At the start, i.e. $\xi(t_s) = s$, the observer's predictability inference (12) is the same as was proposed in [7]. Modeling f_P with (12) allows us to consider the inference of partial trajectories. As a result, the notion of predictability becomes well-suited for motion planning in dynamic environments. It should be noted that this is more closely related to the notion of t-predictability from [11]. However, rather than inferring a sequence of actions, we consider a trajectory.

To compute the trajectory inference (12), we model $P(\xi_{\boldsymbol{t} \to \mathcal{I}^\star})$ as a Boltzmann policy and again approximate its partition function using Laplace's method to obtain:

$$P(\xi_{\boldsymbol{t} \to \mathcal{I}^\star} \mid \mathcal{I}^\star) \approx \exp\left(\beta\big(c[\xi^\star_{\boldsymbol{t} \to \widehat{\mathcal{I}}^\star}] - c[\xi_{\boldsymbol{t} \to \mathcal{I}^\star}]\big)\right). \quad (13)$$

An illustrative example showing the evolution of the goal (11) and trajectory (13) conditionals during a swap scenario is given in Fig. 3.

D. Deriving the Optimal Costs for the CVM

We approximate the costs of the optimal trajectories appearing in (11) and (13) as the minimum time to reach each of the three interaction regions assuming the other agent follows the CVM. These are calculated by building on the idea of constant bearing control from [25] for a single integrator robot and by translating the interaction line at time t along with the predicted trajectory. To find the optimal costs to each interaction region, one need only compute the minimal times to reach the endpoints of the collision line segment, $\mathcal{I}^C$, and the interaction line (see Fig. 4).

We denote the positions of the passing on the right and left interaction rays' endpoints, respectively, by

$$\boldsymbol{p}_{\mathcal{I}^R}(t) \coloneqq \boldsymbol{p}_A(t) - r^C \boldsymbol{e}(t),\ \boldsymbol{p}_{\mathcal{I}^L}(t) \coloneqq \boldsymbol{p}_A(t) + r^C \boldsymbol{e}(t).$$

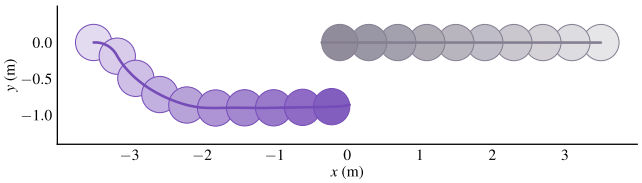

(a) Trajectories for an interaction between a robot (purple) and agent (gray).

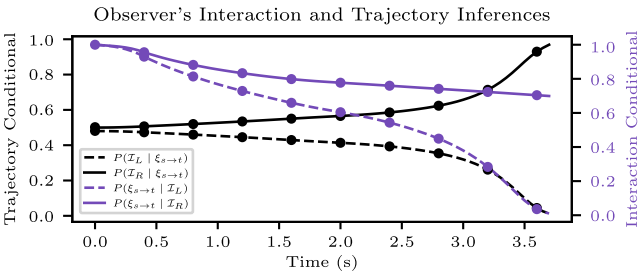

(b) Goal and trajectory inferences for passing on the right and on the left.

Fig. 3: Observer's inferences tracked along an interaction where the circles in (a) darken with time and correspond to the markers in (b).

Fixing the robot's speed to its maximum $v_R^{\max}$, the cost of the optimal trajectory to each endpoint is given by:

$$t^\star_{\boldsymbol{p}_\mathcal{I}}(t,\mathcal{I}) = \frac{\|\boldsymbol{p}_R(t) - \boldsymbol{p}_\mathcal{I}(t)\|}{v_R^{\max}\sqrt{1-\left(\frac{v_{\mathcal{I}_\perp}(t)}{v_R^{\max}}\right)^2} + v_{\mathcal{I}_\parallel}(t)}, \quad \mathcal{I} \in \boldsymbol{\mathcal{G}}^P(t), \tag{14}$$

where $v_{\mathcal{I}_\perp}$ and $v_{\mathcal{I}_\parallel}$ are respectively the perpendicular and parallel components of the robot's velocity relative to the vector pointing from $\boldsymbol{p}_R$ to $\boldsymbol{p}_A$. Computing the minimal time to reach the interaction line is given at each time t by:

$$t^\star_\mathcal{I}(t) = \frac{\|(\boldsymbol{p}_R(t) - \boldsymbol{p}_A(t))^\intercal\, \hat{\boldsymbol{q}}(t)\|}{v_R^{\max} + \boldsymbol{v}_A^\intercal(t)\hat{\boldsymbol{q}}(t)}, \tag{15}$$

where $\hat{\boldsymbol{q}} = (\boldsymbol{g}_R - \boldsymbol{p}_R)/\|\boldsymbol{g}_R - \boldsymbol{p}_R\|$. The three minimal times computed in (14) and (15) are assigned to the three interaction regions as follows:

$$(t^\star_{\mathcal{I}^R}, t^\star_{\mathcal{I}^C}, t^\star_{\mathcal{I}^L}) = \begin{cases} (t^\star_\mathcal{I}, t^\star_{\boldsymbol{p}_{\mathcal{I}^R}}, t^\star_{\boldsymbol{p}_{\mathcal{I}^L}}) & \text{if } t^\star_{\boldsymbol{p}_{\mathcal{I}^R}} < t^\star_\mathcal{I} < t^\star_{\boldsymbol{p}_{\mathcal{I}^L}} \\ (t^\star_{\boldsymbol{p}_{\mathcal{I}^R}}, t^\star_{\boldsymbol{p}_{\mathcal{I}^L}}, t^\star_\mathcal{I}) & \text{if } t^\star_{\boldsymbol{p}_{\mathcal{I}^L}} < t^\star_\mathcal{I} < t^\star_{\boldsymbol{p}_{\mathcal{I}^R}} \\ (t^\star_{\boldsymbol{p}_{\mathcal{I}^R}}, t^\star_\mathcal{I}, t^\star_{\boldsymbol{p}_{\mathcal{I}^L}}) & \text{otherwise.} \end{cases} \tag{16}$$

V. MOTION PLANNER

In this section, we describe how the robot reasons on its legibility and predictability with respect to interactions and we provide an adaptation for the multi-agent case.

A. Goal and Trajectory Conditionals of Motion Primitives

To trade-off between legibility and predictability, the robot queries the scores from its set of motion primitives $\mathcal{P}$. At each planning cycle, the trajectory (11) and interaction (13) conditionals are computed for each primitive with respect to each goal. To adapt more quickly in dynamic environments, we redefine t_s in (11) as $\max(t_s, t - t_p)$, so that the robot reasons about its more recent trajectory segments rather than its entire trajectory since the start of the interaction. The trajectory conditional for primitive ρ_i and interaction goal $\mathcal{I} \in \boldsymbol{\mathcal{G}}(t)$ becomes

$$P(\mathcal{I} \mid \xi_{\boldsymbol{s}\to\boldsymbol{\rho}_i}) \approx \frac{\exp\Big(-\beta\big(c[\xi^\star_{\boldsymbol{s}\to\widehat{\mathcal{I}}}] - c_{\widehat{\mathcal{I}}}[\xi_{\boldsymbol{s}\to\boldsymbol{\rho}_i}]\big)\Big)}{\sum_{\bar{\mathcal{I}}\in\boldsymbol{\mathcal{G}}} P\left(\xi_{\boldsymbol{s}\to\boldsymbol{\rho}_i} \mid \bar{\mathcal{I}}\right) P\left(\bar{\mathcal{I}}\right)} P(\mathcal{I}), \tag{17}$$

where $\xi_{\boldsymbol{s}\to\boldsymbol{\rho}_i} = \xi_{\boldsymbol{s}\to\boldsymbol{t}} + \xi_{\boldsymbol{t}\to\boldsymbol{\rho}_i}$ and $\boldsymbol{\rho}_i$ is the robot's position after completing primitive ρ_i.

In previous works [7], the robot's goal remains unchanged. In contrast, to adapt to the other agent's behavior, we make no such assumption and allow the robot to dynamically switch its intended passing side. During an interaction, the robot continuously updates $\mathcal{I}^\star$ to be the goal region which is most likely, i.e., we set $\mathcal{I}^\star = \operatorname{argmax}_{\mathcal{I}\in\boldsymbol{\mathcal{G}}} P(\mathcal{I} \mid \xi_{\boldsymbol{s}\to\boldsymbol{t}})$. The goal conditional for trajectory $\xi_{\boldsymbol{t}\to\boldsymbol{\rho}_i}$ becomes:

$$P(\xi_{\boldsymbol{t}\to\boldsymbol{\rho}_i} \mid \mathcal{I}^\star) \approx \exp\Big(\beta\big(c[\xi^\star_{\boldsymbol{t}\to\mathcal{I}^\star}] - c_{\mathcal{I}^\star}[\xi_{\boldsymbol{t}\to\boldsymbol{\rho}_i}]\big)\Big). \tag{18}$$

B. Optimizing Legibility and Predictability

In planning motion, the robot must balance between the objectives of legibility and predictability. To do this, we propose to scalarize the objectives and take a convex combination of these inferences

$$\rho^\star = \operatorname*{argmax}_{\rho_i\in\mathcal{P},\mathcal{I}\in\boldsymbol{\mathcal{G}}^P} (1-\lambda(t))P(\mathcal{I} \mid \xi_{\boldsymbol{s}\to\boldsymbol{\rho}_i}) + \lambda(t)P(\xi_{\boldsymbol{t}\to\mathcal{I}} \mid \mathcal{I}), \tag{19}$$

where $\lambda\colon [t_s, t_f] \to [0,1]$ is a function to be designed. When the robot is not interacting with another agent, as per Section IV-A, the robot simply optimize with respect to its global goal $\boldsymbol{g}_R \in \mathcal{W}$. Certain scenarios are more ambiguous than others; the authors in [16] state that scenarios where the legibility of the predictable trajectory is lower are more ambiguous. In order to promote legible motion when the robot's intentions are ambiguous and predictable motion otherwise, we set

$$\lambda(\mathcal{I}, \xi_{\boldsymbol{s}\to\boldsymbol{t}}) = \max\left(0, \min\left(1, \frac{\alpha(\mathcal{I}, \xi_{\boldsymbol{s}\to\boldsymbol{t}}) - a_L}{a_P - a_L}\right)\right) \tag{20}$$

where $\alpha(\mathcal{I}, \xi_{\boldsymbol{s}\to\boldsymbol{t}}) = |P(\mathcal{I}^L \mid \xi_{\boldsymbol{s}\to\boldsymbol{t}}) - P(\mathcal{I}^R \mid \xi_{\boldsymbol{s}\to\boldsymbol{t}})|$ and a_L and a_P are scalar parameters that determine the values

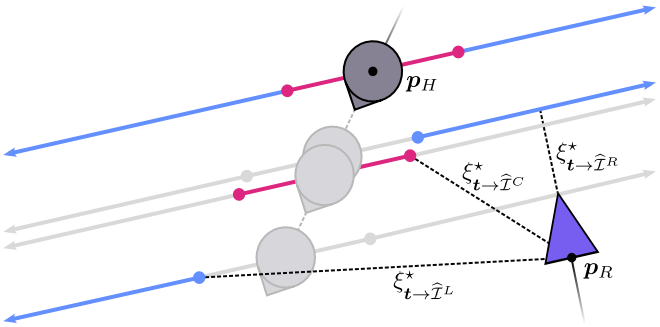

Fig. 4: Optimal trajectories for single integrator dynamics, $\xi^\star_{\boldsymbol{t}\to\widehat{\mathcal{I}}^R}$, $\xi^\star_{\boldsymbol{t}\to\widehat{\mathcal{I}}^C}$ and $\xi^\star_{\boldsymbol{t}\to\widehat{\mathcal{I}}^L}$, from the robot's current position to the predicted passing on the right, collision and passing on the left interaction regions respectively. This configuration corresponds to the first case in (16).

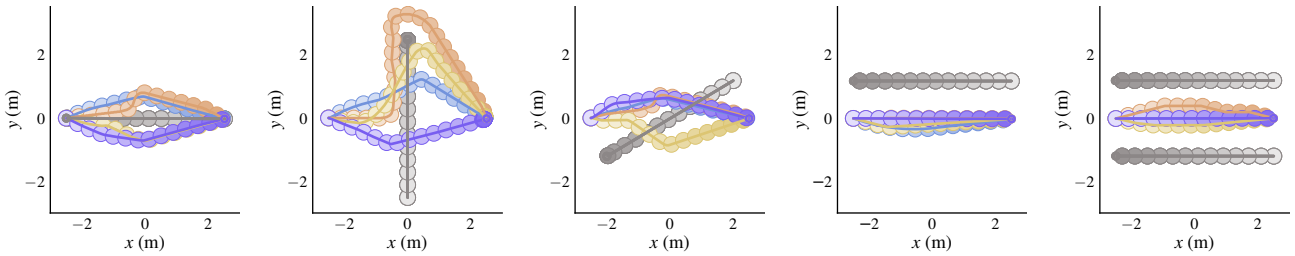

(a) Swap (b) T-junction (c) Obtuse (d) Pass (e) Split

Fig. 5: Overlay of time lapsed trajectories resulting from basic scenarios for LPSNav (purple), SM (blue), SA-CADRL (yellow) and GA3C-CADRL (orange) policies with an inattentive agent (gray). The circle sizesmatch the agents' radii and darken as the simulation progresses. The policies being evaluated have initial and final configurations of $(-3, 0)$m and $(3, 0)$m respectively.

of α at which the robot should strictly optimize legibility or predictability. Predictability is often required to convey intent [16]. In fact, strictly optimizing the legibility functional can lead to arbitrarily unpredictable motions [26]. We recommend setting $a_L < 0 < a_P$ to safeguard the robot from becoming too unpredictable. Therefore, when a passing side is much more likely than the other, λ approaches 1 (weighing primitives that match expectation more heavily), whereas if they are approximately equal, λ approaches 0 (weighing primitives that convey a passing side more heavily).

C. Multi-Agent Adaptation

We assign to each agent A in the set of interacting agents $\mathcal{A}$, its own interaction regions $\mathcal{I}_A(t) \in \boldsymbol{\mathcal{G}}_A(t)$ and lambda parameter $\lambda_A(t)$. To penalize robot motions that are illegible, we optimize over $\mathcal{A}$ by maximizing the minimum score:

$$\rho^\star = \underset{\rho_i \in \mathcal{P}}{\operatorname{argmax}} \left\{ \min_{A \in \mathcal{A}} \left\{ \max_{\mathcal{I}_A \in \boldsymbol{\mathcal{G}}_A} (1 - \lambda_A) P(\mathcal{I}_A \mid \xi_{\boldsymbol{s}_A \to \boldsymbol{\rho}_i}) + \lambda_A P(\xi_{\boldsymbol{t} \to \mathcal{I}_A} \mid \mathcal{I}_A) \right\} \right\}. \quad (21)$$

We also make r^C from (5) adapt to the density of the crowd.

VI. RESULTS

In this section, we evaluate the extent to which our approach (LPSNav) promotes legible behavior in ambiguous scenarios and predictable behavior in unambiguous scenarios. We also evaluate the multi-agent performance in randomly generated scenarios. We compare our framework to the following approaches:

- *ORCA* [27]: a collision-free navigation framework (assuming homogeneous agents) that minimizes the effort spent by minimally adjusting each agent's velocity,
- *SFM* [28]: a model that captures social interactions as a sum of forces resulting from the environment,
- *Social Momentum (SM)* [18]: a planning framework aimed at generating motion that clearly communicates an agent's intended collision avoidance strategy,
- *SA-CADRL* [29]: state-of-the-art socially aware DRL collision avoidance navigation framework,
- *GA3C-CADRL* [30]: adaptation of [29] to deal with an arbitrary number of agents.

A. Implementation Details

Agents are simulated as either inattentive, where they take the straight path to goal at their maximum speed, or as attentive, modeled using the ORCA framework [27], allowing cooperation in the interaction. The mobile robot is modeled as a second-order unicycle, with $\boldsymbol{u} = [u_a, u_\alpha]^\intercal$ the translational and angular acceleration inputs respectively where $|u_a| \leq 3\text{m/s}^2$ and $|u_\alpha| \leq 5\text{rad/s}^2$. The LPSNav agents are configured with $r^C \in [0.35, 0.65]$m, $t_s = 2$s, $a_L = -0.02$, $a_P = 0.5$ and $\beta = 1$.

B. Qualitative Results

Five basic scenarios with an inattentive agent (gray) are overlaid in Fig. 5. To emulate initially ambiguous and unambiguous scenarios, the first three scenarios are initialized on a collision course whereas the last two are not.

In the ambiguous scenarios, our policy and SM indicate their intention early. Our approach respects the passing side convention in the swap scenario and maintains a more conservative behavior in the t-junction by passing from behind. The DRL approaches exhibit a more aggressive swerve later in the interaction to avoid a collision, which specifically in the t-junction results in roundabout trajectories.

In the unambiguous scenarios, our policy chooses the straight path to goal, suggesting predictable behavior. The other policies compromise their goal-efficiency by needlessly seeking to increase their legibility, which in the split scenario reduces the legibility to the third agent.

C. Quantitative Results

To quantify the trade-off between legibility and goal-efficiency, we would ideally need access to the underlying inference being run by the observer on the robot's avoidance strategy. Since we cannot determine this in simulation, we use a combination of the minimal predicted distance (MPD) [31] and the extra distance traveled beyond the straight path as a proxy for this trade-off. The MPD is a continuous function of time, where at instant t, MPD(t) represents the minimum distance attained between the humans if they were to continue at their current velocities. In [31], they found that humans adapt their motion only if it is required, that is, when the MPD falls below a threshold of 1m.

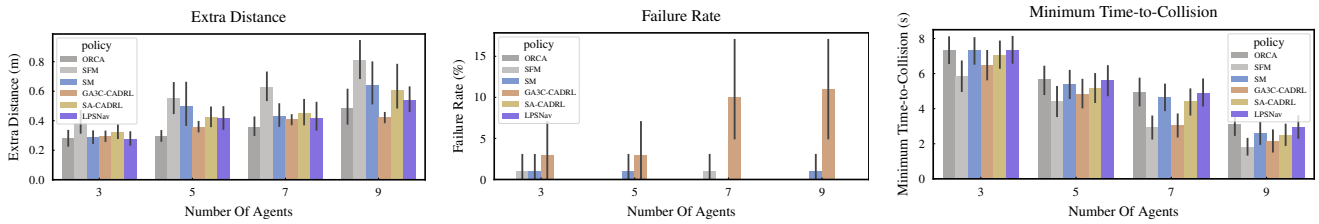

(a) Extra distance (b) Failure rate (c) Minimum time-to-collision

Fig. 6: Performance metrics averaged over 100 random scenarios for 5, 7 and 9 agents.

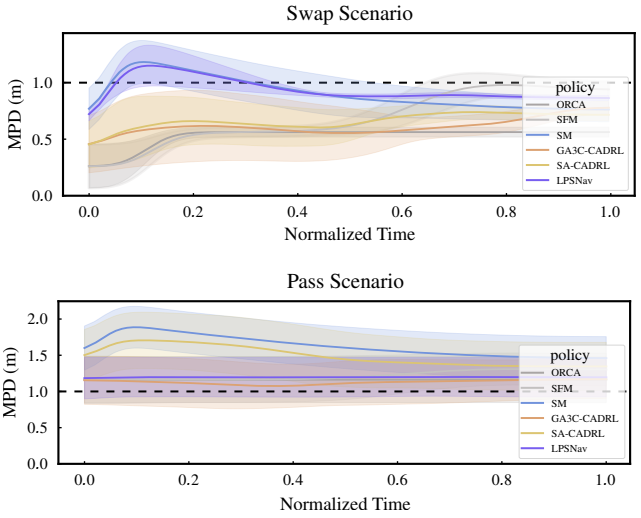

Fig. 7: The minimal predicted distance (MPD) aggregated over 100 random swap and pass scenarios.

TABLE I: The average extra distance and minimum distance to the other agent over 100 random swap and pass scenarios.

Policy	Extra Distance (m)		Minimum Distance (m)	
	Swap	Pass	Swap	Pass
ORCA	0.09	0.08	0.06	0.66
SFM	0.35	0.10	0.44	0.80
SM	0.20	0.14	0.26	0.96
CADRL	0.27	0.12	0.21	0.85
GA3C-CADRL	0.28	0.09	0.28	0.67
LPSNav	0.25	0.09	0.36	0.70

We track the MPD (Fig. 7) and measure the average extra distance and minimum distance to the other agent (Table I) across 100 random configurations with an attentive agent centered around the basic scenarios from (Fig. 5) and report the findings for the swap and pass scenarios.

In the ambiguous swap scenario, LPSNav and SM have the largest initial deviation to express their intent, as indicated by the rapid increase in MPD. The other policies exhibit subtle initial deviations and are forced to do a last minute avoidance maneuver, as indicated by the late peak in MPD. Although ORCA also has a subtle initial deviation, it does not swerve late and passes very close to the other agent. These findings are similarly observed in the other ambiguous scenarios.

In the unambiguous pass scenario, the MPD is initially acceptable according to [31]. LPSNav, ORCA and GA3C-CADRL are the most predictable as suggested by a smaller extra distance and relatively constant MPD value. However, GA3C-CADRL counterintuitively decreases its MPD, suggesting a decrease in legibility with respect to its passing side. SA-CADRL and SM unnecessarily increase the MPD. In the split scenario, SA-CADRL and GA3C-CADRL increase the MPD to one agent while decreasing it to the other, thus confirming the findings from the qualitative results.

We also compute the legibility and the predictability of the randomly generated basic scenarios using a discretized version of (11) and take an average over the 100 configurations. We performed the Mann-Whitney U test with a 95% confidence interval between LPSNav and each baseline. Across the ambiguous swap and obtuse scenarios, LPSNav shows a statistically significant improvement over each baseline in terms of legibility. Furthermore, LPSNav, along with ORCA are the most predictable across the unambiguous scenarios.

To evaluate the multi-agent performance (Fig. 6), we generate 100 random configurations by setting the starts and goals within an 8m $\times$ 8m area for 3, 5, 7 and 9 agents and by setting their maximum speed by randomly sampling $v_i^{\max} \sim \mathcal{N}(1.42, 0.26)\mathrm{m/s}$ [32]. Our approach has a competitive goal-efficiency and scales well with the number of agents, while remaining collision-free. Although ORCA is the most goal-efficient, it had the smallest minimum distance to the other agents, suggesting a more aggressive behavior. As a proxy for legibility and safety, we also report the minimum time-to-collision (TTC). By indicating its intent early, our policy maintains a high minimum TTC with a varying number of agents, thus promoting safe behavior.

VII. CONCLUSION

In this work, we presented an approach for representing navigation interactions as dynamic goals with which a motion planner can use to reason on its legibility and predictability with respect to a passing side. These properties of motion were used to promote legible behavior in ambiguous scenarios and predictable behavior otherwise. We also tested our framework's multi-agent performance, where it is competitive with state-of-the-art approaches in terms of goal-efficiency while remaining collision-free in randomly generated scenarios. Our work is limited in that it has only been tested in a simplified simulation environment. Future work should include a user study to validate the legibility and goal-efficiency trade-off in human-shared environments.

REFERENCES

[1] Y. Abe and M. Yoshiki, "Collision avoidance method for multiple autonomous mobile agents by implicit cooperation," in *IEEE/RSJ International Conference on Intelligent Robots and Systems*, vol. 3, 2001, pp. 1207–1212.

[2] P. Trautman and A. Krause, "Unfreezing the robot: Navigation in dense, interacting crowds," in *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2010, pp. 797–803.

[3] L. Nummenmaa, J. Hyönä, and J. K. Hietanen, "I'll walk this way: Eyes reveal the direction of locomotion and make passersby look and go the other way," *Psychological Science*, vol. 20, pp. 1454–1458, 2009.

[4] M. Hollands, A. Patla, and J. Vickers, ""look where you're going!": Gaze behaviour associated with maintaining and changing the direction of locomotion," *Experimental brain research*, vol. 143, pp. 221–30, 2002.

[5] W. Souza Silva, G. Aravind, S. Sangani, and A. Lamontagne, "Healthy young adults implement distinctive avoidance strategies while walking and circumventing virtual human vs. non-human obstacles in a virtual environment," *Gait & Posture*, vol. 61, pp. 294–300, 2018.

[6] S. D. Lynch, R. Kulpa, L. A. Meerhoff, J. Pettré, A. Crétual, and A.-H. Olivier, "Collision avoidance behavior between walkers: Global and local motion cues," *IEEE Transactions on Visualization and Computer Graphics*, vol. 24, pp. 2078–2088, 2018.

[7] A. D. Dragan, K. C. T. Lee, and S. S. Srinivasa, "Legibility and predictability of robot motion," in *ACM/IEEE International Conference on Human-Robot Interaction*, 2013, pp. 301–308.

[8] J. Guzzi, A. Giusti, L. M. Gambardella, G. Theraulaz, and G. A. Di Caro, "Human-friendly robot navigation in dynamic environments," in *IEEE International Conference on Robotics and Automation*, 2013, pp. 423–430.

[9] T. Kruse, P. Basili, S. Glasauer, and A. Kirsch, "Legible robot navigation in the proximity of moving humans," in *IEEE Workshop on Advanced Robotics and its Social Impacts*, 2012, pp. 83–88.

[10] C. Lichtenthäler and A. Kirsch, "Towards legible robot navigation - how to increase the intend expressiveness of robot navigation behavior," in *International Conference on Social Robotics*, 2013.

[11] J. F. Fisac, C. Liu, J. B. Hamrick, S. Sastry, J. K. Hedrick, T. L. Griffiths, and A. D. Dragan, *Generating Plans that Predict Themselves*. Springer International Publishing, 2020, vol. 13, pp. 144–159.

[12] C. I. Mavrogiannis and R. A. Knepper, "Multi-agent path topology in support of socially competent navigation planning," *The International Journal of Robotics Research*, vol. 38, pp. 338–356, 2019.

[13] M. Everett, Y. F. Chen, and J. P. How, "Motion planning among dynamic, decision-making agents with deep reinforcement learning," in *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2018, pp. 3052–3059.

[14] C. Chen, Y. Liu, S. Kreiss, and A. Alahi, "Crowd-robot interaction: Crowd-aware robot navigation with attention-based deep reinforcement learning," in *International Conference on Robotics and Automation*, 2019, pp. 6015–6022.

[15] M. Beetz, F. Stulp, P. Esden-Tempski, A. Fedrizzi, U. Klank, I. Kresse, A. Maldonado, and F. Ruiz-Ugalde, "Generality and legibility in mobile manipulation: Learning skills for routine tasks," *Autonomous Robots*, vol. 28, pp. 21–44, 2010.

[16] A. D. Dragan and S. S. Srinivasa, "Generating legible motion," in *Robotics: Science and Systems*, 2013.

[17] D. Carton, W. Olszowy, and D. Wollherr, "Measuring the effectiveness of readability for mobile robot locomotion," *International Journal of Social Robotics*, vol. 8, pp. 721–741, 2016.

[18] C. I. Mavrogiannis, W. B. Thomason, and R. A. Knepper, "Social momentum: A framework for legible navigation in dynamic multi-agent environments," in *IEEE International Conference on Human-Robot Interaction*, 2018, p. 361–369.

[19] A. D. Dragan, S. Bauman, J. Forlizzi, and S. S. Srinivasa, "Effects of robot motion on human-robot collaboration," in *ACM/IEEE International Conference on Human-Robot Interaction*, 2015, pp. 51–58.

[20] S. Pellegrinelli, H. Admoni, S. Javdani, and S. Srinivasa, "Human-robot shared workspace collaboration via hindsight optimization," in *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2016, pp. 831–838.

[21] B. Busch, J. Grizou, M. Lopes, and F. Stulp, "Learning legible motion from human–robot interactions," *International Journal of Social Robotics*, vol. 9, pp. 765–779, 2017.

[22] M. Faria, R. Silva, P. Alves-Oliveira, F. S. Melo, and A. Paiva, ""me and you together" movement impact in multi-user collaboration tasks," in *International Conference on Intelligent Robots and Systems*, 2017, pp. 2793–2798.

[23] A. Dragan and S. Srinivasa, "Formalizing assistive teleoperation," in *Robotics: Science and Systems*, 2012, pp. 73–80.

[24] C. Baker, J. Tenenbaum, and R. Saxe, "Goal inference as inverse planning," *Annual Meeting of the Cognitive Science Society*, vol. 29, pp. 779–784, 2007.

[25] S. D. Bopardikar, S. L. Smith, F. Bullo, and J. P. Hespanha, "Dynamic vehicle routing with moving demands - part i: Low speed demands and high arrival rates," in *American Control Conference*, 2009, pp. 1454–1459.

[26] E. Short, J. Hart, M. Vu, and B. Scassellati, "No fair!! an interaction with a cheating robot," in *ACM/IEEE International Conference on Human-Robot Interaction*, 2010, pp. 219–226.

[27] J. van den Berg, S. J. Guy, M. Lin, and D. Manocha, "Reciprocal n-body collision avoidance," in *Robotics Research*, 2011, pp. 3–19.

[28] D. Helbing and P. Molnár, "Social force model for pedestrian dynamics," *Physical Review E*, vol. 51, pp. 4282–4286, 1995.

[29] Y. F. Chen, M. Everett, M. Liu, and J. P. How, "Socially aware motion planning with deep reinforcement learning," in *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2017, pp. 1343–1350.

[30] M. Everett, Y. F. Chen, and J. P. How, "Collision avoidance in pedestrian-rich environments with deep reinforcement learning," *IEEE Access*, vol. 9, pp. 10 357–10 377, 2021.

[31] A.-H. Olivier, A. Marin, A. Crétual, and J. Pettré, "Minimal predicted distance: A common metric for collision avoidance during pairwise interactions between walkers," *Gait & Posture*, vol. 36, pp. 399–404, 2012.

[32] R. C. Browning, E. A. Baker, J. A. Herron, and R. Kram, "Effects of obesity and sex on the energetic cost and preferred speed of walking," *Journal of Applied Physiology*, vol. 100, pp. 390–398, 2006.