When does a human replan? Exploring intent-based replanning in multi-objective path planning

Meher T. Shaikh, Michael A. Goodrich

Meher T. Shaikh, Michael A. Goodrich, "When does a human replan? Exploring intent-based replanning in multi-objective path planning," Proc. SPIE 10640, Unmanned Systems Technology XX, 106400G (3 May 2018); doi: 10.1117/12.2304546

Event: SPIE Defense + Security, 2018, Orlando, Florida, United States

When does a Human Replan? Exploring Intent-Based Replanning in Multi-Objective Path Planning

Meher T. Shaikh and Michael A. Goodrich

Brigham Young University, Provo, UT, USA

ABSTRACT

In goal-based tasks such as navigating a robot from location A to location B in a dynamic environment, human intent can mean to choose a specific trade-off between multiple competing objectives. For example, intent can mean to find a path that balances between "Go quickly" and "Go stealthily". Given human expectations about how a path balances such tradeoffs, the path should match the *human's intent* throughout the entire execution of the path even if the environment changes. If the path drifts from the human's intent because the environment changes, then a new robotic-path needs to be planned — referred to as *path-replanning*.

We discuss here three system-initiated triggers (prompts) for path-replanning. The objective is to create an interactive replanning system that yields paths that consistently match human intent. The triggers are to replan (a) at regular time intervals, (b) when the current robotic path deviates from the user intent, and (c) when a better path can be obtained from a different homotopy class. Further, we consider one user-generated replanning trigger that allows the user to stop the robot anytime to put the robot onto a new route. These four trigger variants seek to answer two fundamental critical questions: *When is a re-planned path acceptable to a human?* and *How should a planner involve a human in replanning?*.

Keywords: intent, intention, plans, BDI, reasoning, commitment, multi-objective path planing, tradeoffs, re-planning, graphical user interfaces, human-robot teams, human-robot interaction, human-robot collaboration

1. INTRODUCTION

The notion of *intent* has been conceptualized and defined by philosophers 1–9, psychologists 10–20, neuroscientists 21–24, and artificial intelligence researchers 25–31. Most of this literature expresses intent as a mental state that enables an agent to commit to achieve something in future. Many theories either consider or suggest that intentions are precursors to action or sequences of actions 1–3,5,6,11,14,20,32. Wikipedia uses a concise (albeit incomplete) summary of Bratman's notion of intent 2: *Intention is a mental state that represents a commitment to carrying out an action or actions in the future.*[*]

Much of the intent-based literature assumes a *rational agent*, which could be either human or a robot; this paper assumes that the human holds the intent and the robot executes intent. Task execution has two components: what it is to be achieved, that is, the desired outcome (the goal), and the means (trajectory) to achieve it. Accordingly, we assume intent includes (a) the agent's capabilities, (b) the interaction environment where trajectories are executed, and (c) the agent's commitment to a goal and trajectory over time. In this paper, a trajectory is a path taken to reach the goal. Paths are chosen based on constraints, objectives, and policies/strategies/plans that determine *how* intent is translated into action. This paper deals directly with the temporal aspect of intent, when a persistent commitment reaches its "expiration date" 33.

The primary contribution of this paper is a partial answer to the question: *when does a human replan* in dynamic environments such that the adverbial description of a task maintains intent while balancing multiple objectives.

[*]https://en.wikipedia.org/wiki/Intention

Unmanned Systems Technology XX, edited by Robert E. Karlsen, Douglas W. Gage,
Charles M. Shoemaker, Hoa G. Nguyen, Proc. of SPIE Vol. 10640, 106400G
 · doi: 10.1117/12.2304546

2. RELATED LITERATURE

The intent theory has been connected to many attitudes/pro-attitudes thereby bringing to the table its multiple notions. In addition to intentions being seen as a product of desires and beliefs that is strongly emphasized by Dennet 1, 34, 35 to name a few, intentions are also shown to be related to commitment, reasoning and partial plans 2, 31. They are associated with plans, actions, and time 2, 7, 33. Gibbs 36 portray intentions in the light of interactions. He reports that people hardly ever act independently, hence, in addition to intentions being private mental states, intentions are also emergent product of interactions. Intentions are also viewed as a kind of persistent goals, where persistence involves an agent's internal commitment to a course of events over time 33. The vast literature on intent 1–4, 8–13, 16–19, 21–29, 31, 33, 36 leads to following notions: *commitment*, *persistent* (that is, intentions are not to be abandoned atleast for some time), *beliefs*, *desires*, *rational processes*, *plans*, *partial plans*, *goals*, *action* and *time*.

We discuss intent here in light of human-robot teams or human-robot interaction (HRI). In HRI, it is usually the human who owns the intent 37 and communicates his intent either explicitly or implicitly to the robot. The robot assists the human to accomplish his intent. In this work, what human wants is linked to decision making under multiple conflicting objectives for robot navigation in a dynamic environment. Hence, intent is a tradeoff that dictates '*how*' the robot should navigate to the goal.

Success in human-robot tasks to a large extent depends on the success of communicating the intent to the robot. The techniques for communicating the human intent can be classified as (i) explicit and (ii) implicit. Explicit intent strategies include both verbal[38] such as natural language/speech commands, and non-verbal communication[39] such as eye-gaze, gesture and facial expression, as well as combination of these.[40, 41] Further, conventional user interfaces that use devices such as keyboard, mouse, joystick, hap-tic and touch interaction[42] had been around for more than half a century in order to commands robots. Note that, human inputs such as from knobs or from sliders or from natural language can be transitioned into vectors to represent an intent. Implicit intent, also known as *indirect intent* by,[43] include physiological signals such as ECG (Electrocardiogram), EMG (Electromyogram), EEG (Electroencephalogram), skin conductance, pupil-dilation[43–45] etc. Our path-replanning architecture uses the *adverb palette (AP)* discussed in Sec. 4.1 to communicate human intent to the robot.

3. OPERATIONAL DEFINITION OF INTENT

Intent in this paper is based on Bratman's 2 and Malle et al.'s 11 theories of intentions. Bratman approaches intentions by way of planning theory. Accordingly, intentions are *partial plans* brought about by deliberation and practical reasoning considering resources and coordination (both intrapersonal and interpersonal); plans which on a commitment get updated with time that eventually bring about the desired outcome. Note that partial plans does not mean plans are incomplete but rather that plans get updated in response to the interaction environment as the agent executes the initial partial plan. Updating and renewing a plan maintains intent.

This paper assumes that beliefs, desires, and intentions form the basis of an intentional action 1,2,11. *Desires* are what an agent wants or wishes for. When desires are combined with commitment, reasoning, and action, a subset of possible desired goals become intentions. A goal is a mental representation of a *desired outcome* that one wants to attain through action and *desired means*. Intention is therefore a commitment towards achieving the goal. While goals are outcomes that are measurable at the end of certain time, intentions also include the "journey" towards these goals satisfying desired attributes.

Belief is the knowledge that the agent carries about itself and the surrounding environment to bring about planning and action in order to attain the goal. For a path planning problem, the environment has information such as: 'where is the enemy?' 'how far is the goal state?', 'what are the alternative trajectories available to consider?' etc. Beliefs also include all the constraints and objectives that dictate or specify how the goal needs to be attained; these attributes determine the desired means to be met on the trajectory or path of achieving the goal.

Capabilities are critical for goal attainment. For human, capabilities are associated to practiced or acquired human skills and for a robot, capabilities would be its functionality dictated by the algorithms endowed in it. We use *capabilities*, *autonomy* and *algorithms* interchangeably for robots.

Initially, before the action begins the agent will have a prior knowledge about itself and the environment. Based on the prior knowledge, the agent partially plans the solution. Later on, as action proceeds, the agent updates its knowledge in coordination with itself and with the environment.

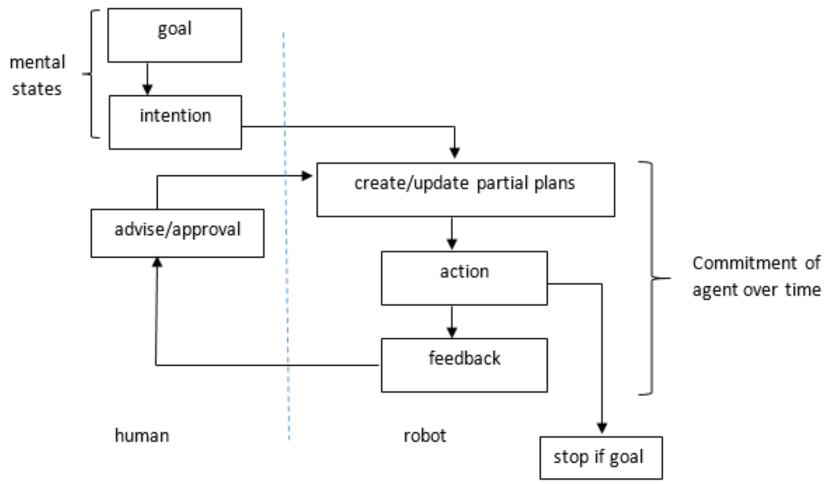

Figure 1: Intent framework for human-robot teams.

Fig. 1 illustrates how the different components above relate to each other for a problem that includes one human and one robot. Our operational definition of intent is:

> In an environment inhabited with agents of different *capabilities*, *desires* and *beliefs*, intent is a *commitment* of a rational agent to bring about a *desired outcome* in a *reasonable time* by shaping a sequence of environment states that (a) satisfies both a set of *constraints* and a set of *objectives*, and (b) executes a *plan/policy* toward the desired outcome..

4. REPLANNING ARCHITECTURE

An important element of intention is monitoring when and whether a particular intent is feasible or relevant, and when intent needs to be updated. For *a specific human intent*, if the real-time execution of the partial plan does not fulfill the human's intent, then alternative plans or course of actions need to be built online that would satisfy the intent — referred to as *replanning*. Similarly, intent can change based on emerging behaviors and constraints in an environment, especially one with multiple agents.[36] We define *triggers* as events that signal the human the need to reconsider the current intent solution along with a new plan. This calls for (a) some sort of replanning framework for intent and (b) some sort of user interface that enables intent monitoring and path-replanning.

The design of replanning architecture here is inspired by the planning considerations presented in Chapter 8 of 46 which says: "...if a planning model is to generate planning behaviors that somehow mimic those of a human planner, the model must attempt to replicate the various stages of planning,..." Accordingly, our replanning framework includes elements for (a) information exchange, (b) situation assessment, (c) course-of-action development, and (d) monitoring and replanning.

Our replanning system architecture has three entities: a GUI that we call the *adverb palette (AP)* , the robot, and path-planning/replanning algorithms.

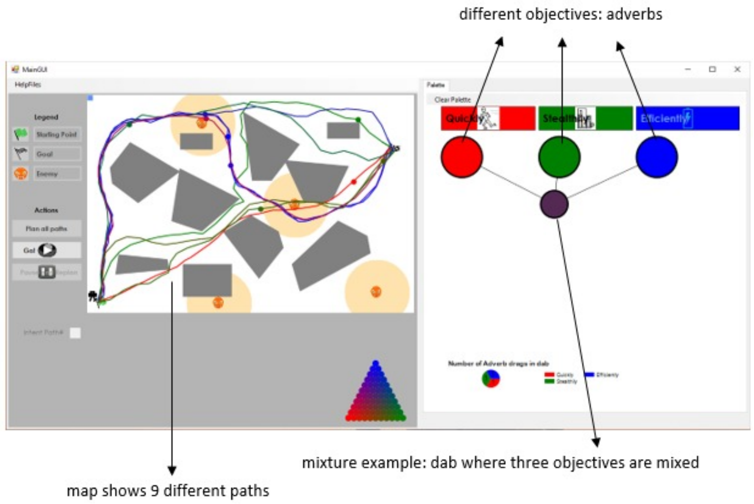

Figure 2: Map showing robot's potential paths. Robot: black entity, Enemy: orange haloed entities, Obstacles: gray blocks (trees or buildings)

4.1 Adverb Palette

The *adverb palette (AP)* is an interactive graphical interface that we described in 47, which was designed to express human intent; in this paper, we extend *AP* to support intent monitoring and replanning. The interface is tailored to a robot navigation task in a dynamic environment in which the robot needs to navigate from location A to location B under conflicting objectives. The environment is depicted by a map that shows the robot's current location, the goal where it has to reach, the enemy positions, and the obstacles. In this application, we have three objectives that are expressed as adverbs *quickly*, *stealthily*, and *efficiently*, each represented with a unique color on the user interface. The human expresses the intent using these adverbs that basically defines the objectives/constraints on how the robot should navigate to location B, which is the goal location. The adverb *quickly* is a command to the robot to take the shortest route based on Euclidean distance, the adverb *stealthily* is a command to the robot to take a route evading the enemy as much as possible, and the adverb *efficiently* is a command to the robot to take a route that minimizes fuel consumption. For a tradeoff, such as a path where both distance and stealth are intended, a mixture of *quickly* and *stealthily* needs to be created on the *AP*. The tradeoffs are thus represented as a mixture of colors. The tradeoff/color encodes the intent of the navigation task. We refer the reader to 47 for the basic working of the *AP*.

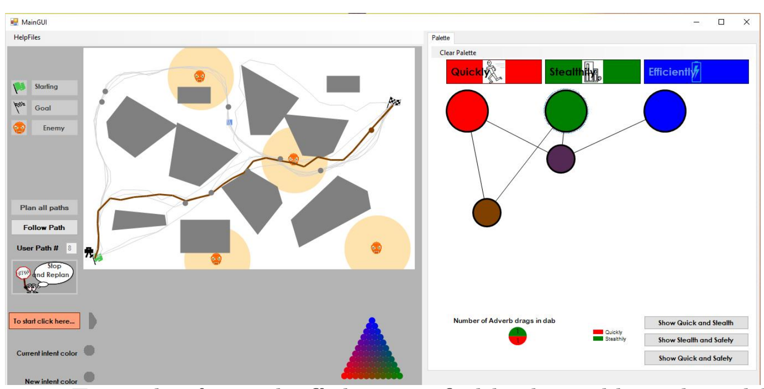

Figure 3: Example of a tradeoff that satisfied both quickly and stealthily.

At the outset of the navigation task, *AP* serves as a means to visualize multiple solutions (plans/options/paths) generated by the robot. Figure 2 shows the *adverb palette* with 9 candidate paths shown on the map. Each path starts from an initial location, indicated by the black robot in the lower left of the map, and ends at the goal location in the upper right of the map. The path is a series of waypoints through the map, which determines a trajectory that either optimizes a single objective or balances a mixture of objectives. Each trajectory from start to goal is constructed using different mixtures of the adverbial objectives. That is, each solution is a tradeoff between different objectives. The human communicates intent to the robot by selecting one of the tradeoff solutions using the right panel of the interface. For example, if the user clicks on the big red circle on right panel of Figure 2, then a quick path is desired; a red path on the map in Figure 2. On the other hand, if the user clicks on the big green circle, a stealthy path is selected for travel. Figure 3 shows a tradeoff between quickly and stealthily formed by mixing the two adverbs on the right panel resulting into the brown mixture. Thus, the human intent here is a tradeoff.

4.2 Planning, Execution, and Replanning

The robot is equipped with a tree-based planner, online fast marching tree*(*O-FMT**) 48 that is used to generate the initial set of plans. The planning and replanning problem is described in Section 5. The robot is an autonomous agent that has the capability to execute the path by following the path chosen by the user via *AP*.

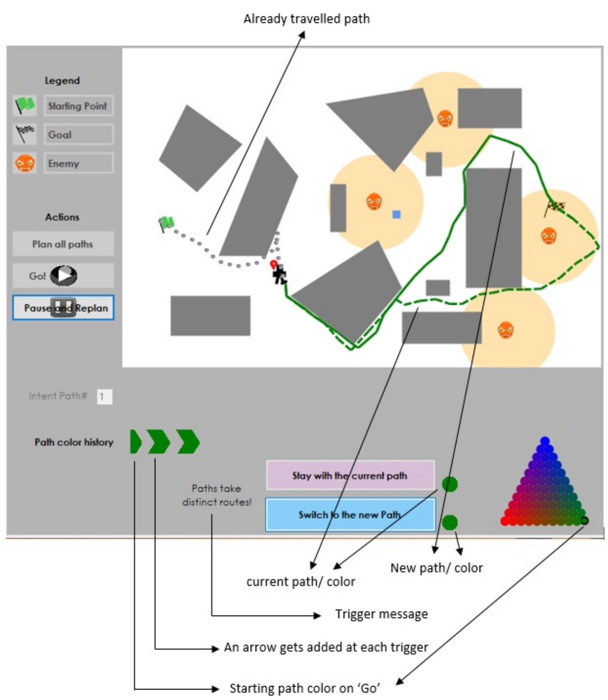

Figure 4: Robot's current and new path. Current path: dashed green, New path: solid green

During path execution, the robot's planner is capable of *updating* the current path and *proposing alternative* paths that may better match intent. It can generate a new path in less than 2 secs for a given intent based on the weighted combination of speed, stealth, and energy-efficiency objectives. Updating a path theoretically allows a robot to adapt its execution so that it matches intent as the world changes. Proposing alternative paths theoretically allows a robot to present alternatives that better match the human's intent or allow a human to change intent during execution. Figure 4 shows robot's new path/plan in a solid green trajectory, and the current path in a dashed green trajectory for an example stealthily intent set at the beginning. The updated plan assists replanning. Sec. 5.4 details the concept of how the new plan differs from the current one.

The robot can prompt the human to consider alternate paths at specific events called *triggers*. Potential triggers include replanning at (a) *regular time intervals*, (b) *when the current robotic path deviates from the user intent*, and (c) *when a better path can be obtained from a different homotopy class*. The different triggers provide an opportunity to the human to either approve or disapprove the new robotic path. We discuss each of these triggers in detail in Sec. 6. The robot is said to have successfully navigated to the goal (*location* B) if it maintains its intent throughout the entire navigation task or if the human is able to express a revised intent and the robot follows the revised intent.

At triggers, the human uses his judgment to either to remain on the current plan or change to the new plan, thereby collaborating towards a successful intended task execution. The new plan suggested by the robot may vary in four possible ways: (i) the old plan is no more according to intent, but the new plan is; (ii) the old and the new plan both follow intent but differ in trajectories; (iii) the old and the new plan differ only by a margin (dictated quantitatively), they both follow intent; (iv) the old plan and the new plan both do not follow intent. In addition to responding to the triggers generated by the robot, the human himself can pause the robot anytime and put the robot on a new plan which results into a *user initiated trigger*.

Human input is critical in automation 49. The AP serves as a platform by which the human monitors execution, becomes aware of triggers, and makes adjustments to paths or intent. The AP thus aids for intent exchange, situation assessment, course-of-action analysis and selection, and monitoring and replanning.

5. PATH PLANNING

The path-planning problem and planner descriptions presented in this section are adapted from 48, 50.

5.1 Problem

For robotic path planning, an environment at any time is a topological space $X \subset \mathbb{R}^d$, with an obstacle space X_{obs}, an initial state x_{init}, and a goal region x_{goal}. The obstacle-free space is denoted by $X_{free} = X \setminus X_{obs}$. Consider the set of J objectives determined by a vector cost function $\boldsymbol{c}(\cdot) = [c_1(\cdot), \ldots, c_J(\cdot)]^T$ defined by $\boldsymbol{c} : \mathbb{X} \to \mathbb{R}^J$. Note that $\boldsymbol{c}$ is defined for all points in X in free space. At a given time t, let X_{E} be the set of the locations of n enemies, $X_{\text{E}} = \{\mathbf{x}_{e_i} | \mathbf{x}_{e_i} \in X_{\text{free}}\}$, and let the location of the robot be denoted by $\mathbf{x}_{\text{rob}}$, and $\mathbf{x}_{\text{rob}}$ is not x_{goal}. The robot acts in this environment specification to create a trajectory towards the goal.

The term agent in this paper applies to the robot. The robot knows the goal location to attain and has the 'how' intent communicated to it from AP. The path-planning/path-replanning problem in this paper deals with a robot that navigates in a dynamic environment under multiple objectives, *quickly*, *stealthily*, and *efficiently*; thus, $J = 3$.

5.2 Terminology: What is a Path/Trajectory and What are its Costs?

A trajectory or a path is a continuous curve induced by an robot's algorithm parameterized by s, denoted by $\sigma : [0, s] \to X$. Note that, the trajectory satisfies (a) $\forall \tau \in [0, s]$, $\sigma(\tau) \in X_{free}$; (b) $\sigma(0) = x_{\text{init}}$, $\sigma(s) = x_{\text{goal}}$; (c) causes/influences a sequence of environments $X_1, \ldots, X_g$ where $i \in \{1, \ldots, g\}$, g is the number of elements in the sequence, the first element of the sequence, X_1, is adjacent to x_{init}, and finally the last element X_g is adjacent to x_{goal}.

Given what a trajectory or a path is, T is the set of all obstacle avoiding trajectories with an initial point as x_{init} and end point as x_{goal}.

At the start, before the robot starts moving, given a set of three objective functions, let $\Sigma = \{\sigma_p\}$ denote the set of Pareto optimal paths. Since we are doing path-planning on a two-dimensional plane, a path is a parameterized curve that exists in $\Re^2$. Thus, each path σ_p is a mapping from a parameter space to $\Re^2$. Without loss of generality, let the parameter space be the continuous interval $[0, 1]$. Thus,

$$\forall j \quad \sigma_p : [0, 1] \mapsto \Re^2. \tag{1}$$

For the kind of path-planning that we are doing, the path is constrained to begin at a starting location (x_0, y_0) (that is, x_{init}) and end at the goal location (x_f, y_f) (that is, x_{goal}), yielding the constraints on the path as follows:

$$\begin{aligned} \sigma(0) &= (x_0, y_0) \\ \sigma(1) &= (x_f, y_f). \end{aligned}$$

For a multiple objective problem expressed as cost functions, let J_i denote the i^{th} cost function. Suppose that we have three cost functions, $i \in \{1, 2, 3\}$. Each cost function assigns a real-valued cost to a path,

$$\forall i \quad J_i : \Sigma \mapsto \Re. \tag{2}$$

Thus, a cost function takes a path, σ, and assigns a real-value, $a \in \Re$, to it, $J_i(\sigma_p) = a$.

The tree-based planner path is made up of m vertices and weighted, directed edges. The edges vary in costs. Each weighted, directed edge is a cost to traverse from a parent vertex, o_k, to a child vertex, o_{k+1}; let $c(o_k, o_{k+1})$ denote the cost of this edge. A path σ_p is a sequence of edges through the tree, with the first vertex located at (x_0, y_0) and the last vertex located at (x_f, y_f). Because of varying cost edges, we don't have a uniform partition on the parameterization interval $[0, 1]$ so we will write the partition as $m + 1$ different points in [0,1], s_k, where $s_0 = 0$, $s_m = 1$ and $s_k < s_{k+1}$.

The cost of a path is the sum of the costs of the edges. Thus,

$$J_i(\sigma) = \sum_{k=0}^{m-1} c[\sigma(s_k), \sigma(s_{k+1})] \tag{3}$$

where $\sigma(s_k)$ equals the location of vertex k.

5.3 Normalization and Scaling: What is the Color of a Path?

Our prior work found that the color palette of the AP was a useful way for a human to express intent to a multi-objective path-planner.[47] Consequently, we need to assign a color to each path. There are three objectives and, for design and usability purposes, we assign colors from RGB space in a constrained way. The color is restricted to three components, one each from the RGB set of colors, but constrained such that the sum of the red component, green component, and blue component sum to one. Thus, we will create a color vector $\mathbf{h}(\sigma_p) = [R(\sigma_p), G(\sigma_p), B(\sigma_p)]$ satisfying

$$\begin{aligned} R(\sigma_p) &\in [0, 1] \\ G(\sigma_p) &\in [0, 1] \\ B(\sigma_p) &\in [0, 1] \\ \sum_{color \in \{R,G,B\}} color(\sigma_p) &= 1. \end{aligned}$$

We need to translate the cost triple $(J_1(\sigma_p), J_2(\sigma_p), J_3(\sigma_p))$ into a color vector. We do this by associating each color to a different cost function; without loss of generality we assign cost functions such that J_1 corresponds to red, J_2 to green, and J_3 to blue. For reasonable correspondence to colors, we create a normalized objective $o_i(\sigma_p)$ from $J_i(\sigma_p)$ for each path σ_p as follows:

$$o_i(\sigma_p) = \frac{J_i(\sigma_p) - \min_{\sigma' \in \Sigma} J_i(\sigma')}{\max_{\sigma' \in \Sigma} J_i(\sigma') - \min_{\sigma' \in \Sigma} J_i(\sigma')}. \tag{4}$$

The color of a path is defined as the vector, $\mathbf{h}(\sigma_p)$, that maximizes the cosine similarity between the objective vector.

$$\mathbf{o}(\sigma_p) = [o_1(\sigma_p), o_2(\sigma_p), o_3(\sigma_p)]^T$$

and the inverse color vector

$$\begin{aligned} \mathbf{h}' &= [R', G', B']^T \\ 1 &= R + G + B \end{aligned}$$

where

$$\begin{aligned} R' &= 1.01 - R \\ G' &= 1.01 - G \\ R' &= 1.01 - B \end{aligned}$$

yielding

$$\mathbf{h}(\sigma_p) = \arg\max_{\mathbf{h}'} \frac{\mathbf{h}' \cdot \mathbf{o}(\sigma_p)}{\| \mathbf{h}' \| \| \mathbf{o}(\sigma_p) \|} \tag{5}$$

Since for costs (Equation 4), lower values are better, we subtract each of the RGB component of color from 1.01 in the equation above so that the higher preferred color values get converted to lower values and thus can be matched with the corresponding lower costs and vice versa. Further, the color element is subtracted from 1.01 instead of 1 so as not to nullify the effect of an objective with a corresponding color component of 1 in $color \in R, G, B$ in the computation of cosine similarity in Equation 5.

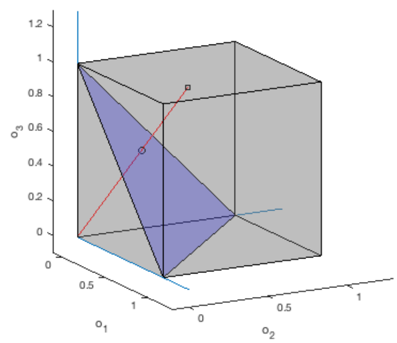

Figure 5: Objective space, $\mathbf{o}(\sigma_p)$, and color space, $\mathbf{h}$.

Figure 5 illustrates the different spaces involved in assigning a path a color. The axes in the figure represent the objectives o_1, o_2, and o_3. Since $o_i(\sigma_p) \in [0, 1]$, the unit cube shown in the figure bounds the ranges of the objectives for any possible path $\sigma_p \in \Sigma$. Colors are normalized such that they must sum to one, so the triangular simplex represents the set of possible colors.

The square box represents the objective vector $\mathbf{o}(\sigma_p)$ for a path σ_p, and the brown line segment indicates the vector emanating from the origin to the vector. The color of the path is given by the coordinates at which the red line segment intersects the triangular simplex, which occurs at the brown ellipse obtained by using Equation 5.

5.4 Replanning Trigger

Suppose that (a) a human has specified a desired path color and (b) path σ^h is the path from Σ that most closely matches that color. Suppose further that the robot has been following path σ^h for some period of time and has reached location $\sigma(s)$, where $s \in (0, 1)$; the open interval $(0, 1)$ indicates that the robot has been traveling for some positive time, meaning that $s \neq 0$, but hasn't reached the end of the path, meaning that $s \neq 1$. Even though s isn't technically a time, we can treat it as if it is a time unit, so suppose at time s^* something happens and the costs change. For simplicity, suppose that cost function J_i has changed. For example, suppose that objective i generates the edge cost in Equation 3 with high edge costs if the edge is close to an enemy, but the enemies move at time s. Should the robot change paths?

Figure 6a replicates Figure 5, but for only two objectives. The unit square represents the set of possible objective vectors; the small brown square is the end point of the objective vector, $\mathbf{o}$ for a particular path, σ_p; the brown line segment emanating from the origin is the objective vector $\mathbf{o}(\sigma)$; the diagonal blue line is the set of possible normalized colors; and the small brown ellipse represents the color for the path $\mathbf{h}(\sigma_p)$. The small brown square indicates the cost with respective to all objectives at time $s = 0$.

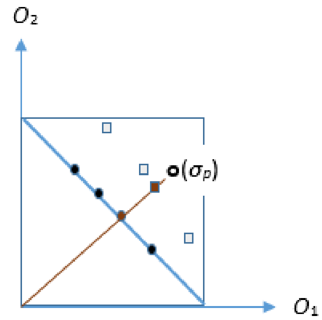

(a) $\mathbf{o}(\sigma_p)$, and color space, $\mathbf{h}$ for o_1 and o_2.

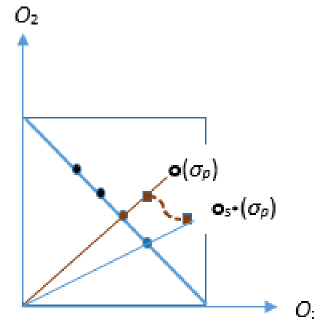

(b) Shifted path costs at time s^*.

Figure 6: Color and costs space for two objectives.

Figure 6b illustrates what happens when one of the cost functions changes at time s^*. Note that the path path σ^h hasn't changed, but the costs have changed in response to the changes in the environment. For example, say that the enemy has approached closer to this path at time s^* resulting into a higher stealthily cost. As a result the objective vector that includes the change in cost function has shifted down and to the right. Note the shift in costs indicated by a dashed curve in Figure 6b. Now, because the objective vector $\mathbf{o}^{s^*}(\sigma_p)$ has changed, the color associated with the path σ_p has changed from brown to light blue. Since we assumed that the human intent was indicated by the brown vector, the original path σ_p no longer matches the human intent. Should this be a trigger for replanning?

Thus, the costs of a path are associated with color and the human expresses intent by selecting a color. Once the robot starts moving the change in environment may lead to changes in objectives, which may correspond to a different color thereby indicating a deviation from intent. Sometimes the change in path costs may not induce a big color change, but other times the path costs may cause a large color change.

Note that when we replan, we don't care to compute the path from the start point, (x_0, y_0), to the end point, (x_f, y_f), anymore. Rather, we only care to compute a new path such that the replanned path is identical to σ^h up to time s^*; after time s^* the replanned path may differ from σ_h. In other words, the replanned path should shift to a better path from time s^* onwards. The problem is illustrated in Figure 7. The robot has followed the original orange path up until time s^*. At time s^*, it needs to decide whether to continue along the original orange path or switch to a new path that builds from the original orange path. The green path in the figure represents a new replanned path.

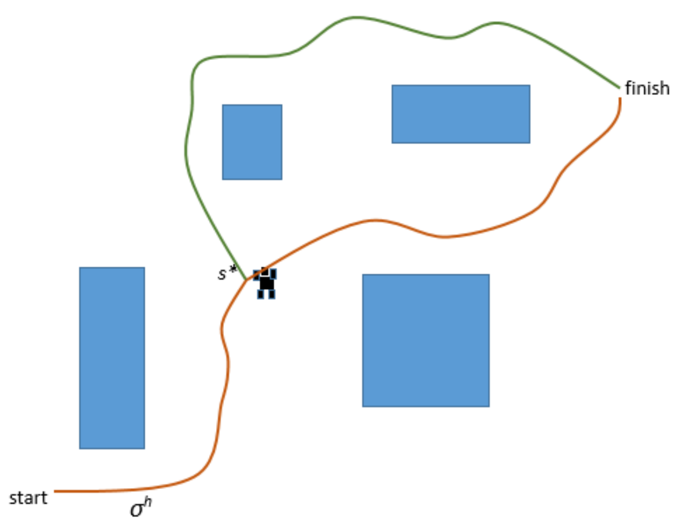

Figure 7: Branching from the original blue path.

6. REPLANNING TAXONOMY

Given the intent, path-planning, and replanning formalism, we can describe four different replanning triggers. For the remainder of the paper, we consider only two paths: the robot's current path and a path that is automatically replanned as the robot and enemies move in the world. The replanned path is a proposed change or suggestion to the human. Replanning triggers are possible ways that a human might use the current path and replanned path together to maintain intent.

On a trigger, the AP displays (a) the robot's *current location* with an "I'm here" status denoted with a robot with a red top icon, (b) the *current path* in a dashed pattern, (c) and the *replanned path* in a solid pattern. The path already travelled is shown as tiny circular footsteps on the map. On a trigger, the interface pops up buttons that allow the user to either 'Stay with the current path' or 'Switch to the new path'. The AP also displays *path color history* to help the human understand original intent and measure drift in intent in the lower left section of the left panel; the first triangular arrow of the history indicates the initial intent with which navigation had started. See Figure 4.

6.1 Time Trigger

Replanning at regular intervals is the simplest replan strategy. A *time trigger* occurs at deterministic time intervals to make the human aware of the current environment and the two paths: the current one and the replanned one. The two paths may or may not vary in different aspects such as the intent, homotopy (discussed later), and/or a combination of these. Figure 8 and Figure 9 show two occurrences of *time trigger* in an example stealthily navigation (the intent chosen by the user was *stealthily* on AP and is represented by green color). In our planned experiments, the robot generates a time trigger every n seconds unless a different trigger occurs.

Figure 8: Time trigger at point s_{t_1}.

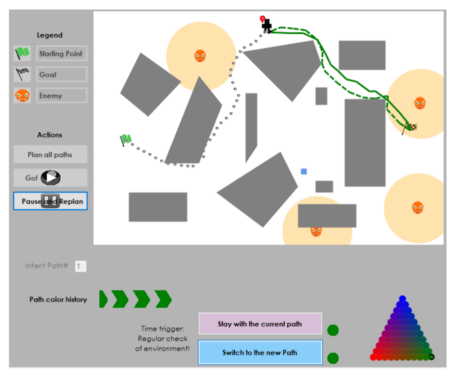

Figure 9: Time trigger at point s_{t_2}.

6.2 Change-in-Intent trigger

Section 5.4 showed how the costs of a path after time s^* may change due to changes in the environment during navigation. Let $c_{threshold}$ be the cosine similarity value given by Equation 5 for $\mathbf{h}$ associated with path $\mathbf{o}(\sigma_h)$, the original path. If the new costs of the path $\mathbf{o}(\sigma_h)$ change such that Equation 5 produces a value equal to or greater than $c_{threshold}$, then we say that the path (σ_h) maintains the original intent. However, if the new costs after time s^* change as illustrated in Figure 6b such that the Equation 5 yields a value below $c_{threshold}$ then the path (σ_h) is far enough away from $\mathbf{h}$ that the current path no longer matches the original intent.

Figure 10 shows an example of change in intent for a stealthily navigation task. The robot had started with a stealthily intent, green. At some time s, the current path color changes to red — the path became expensive because of the approaching enemy. The automatically replanned path better matches the original intent, so the human may want to switch paths.

6.3 Homotopy Trigger

Quoting from ,[51] path σ_1 is said to be *homotopic* to path σ_2 if σ_1 can be mapped to σ_2 without encroaching on any obstacle 52, 53. Otherwise, the two paths are said to belong to different homotopy class or said to be *non-homotopic*. We denote homotopic paths by $\sigma_1 \simeq \sigma_2$.

A *homotopy trigger* occurs when the robot's replanned path and original path are non-homotopic. In our planned experiments, the robot uses the algorithm 51 to check homotopy. The idea behind this trigger is that if the current and replanned paths go around obstacles in different ways, then the human may need to consider whether the path matches intent even if the path colors stay the same. Figure 11 shows an example of a homotopy trigger that occurrs for a navigation task meant for the robot to go *stealthily* as well as *quickly*, a brown color path.

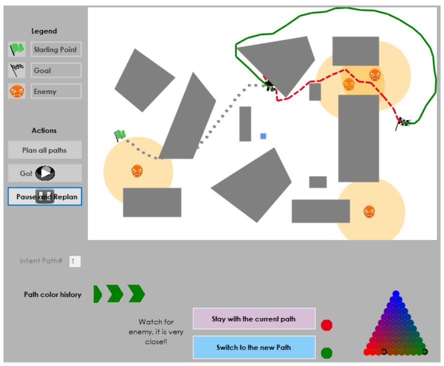

Figure 10: Change in intent trigger example.

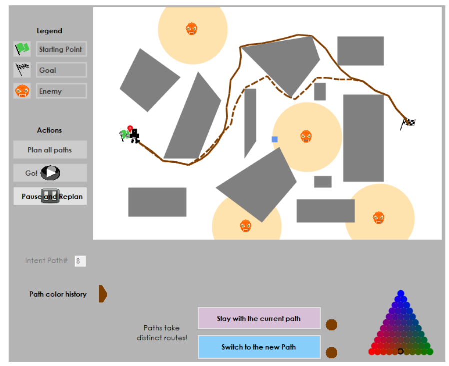

Figure 11: Homotopy trigger example.

In order to not irk a human from very frequent homotopy or change-in-intent triggers, in the planned experiments we will restrict the frequency of these triggers such that a time trigger would essentially separate any two homotopy or change-in-intent triggers. That said, the time trigger may show up a non-homotopic or intent-violated path if any as shown in Figure 8.

6.4 User-Initiated Trigger

In addition to the above three triggers generated by the robot, in the planned experiments the human can pause the robot anytime and put it on a replanned path. This is possible because during the walk the AP displays the robotic replanned path. The AP facilitates *user initiated trigger* with a 'Pause and Replan' button on the left panel.

7. SUMMARY AND FUTURE WORK

In this paper, we applied an operational definition of intent to human-robot teams where both the human and the robot work in collaboration towards a common goal. Using a replanning architecture, we applied the definition to robot navigation in dynamic environment and under multiple objectives. We proposed a replanning taxonomy to answer a critical question: when does a human replan such that the human-intent is preserved.

Our in-house simulations of replanning triggers show promise to have potential in maintaining human intent in HRI by engaging the best of the capabilities of the human and the robot thereby improving the chances towards successful goal attainment. Our hypothesis is that the triggers would help a human to judge and decide the paths at critical times that would eventually support the intended travel. In the very near future, we plan to conduct a user study to answer the question; does the adverbial description of task represent intent. Further, the user study will let us know about the critical triggers that would help maintain intent. To realize this goal, we would look for the subjective scores as to which triggers appealed to the users and which were found useful. We plan to record the navigation trigger sequence that would reflect the statistics of change in intent during travel.

ACKNOWLEDGMENTS

This work was funded by the United States Army Research Laboratory through contract W911NF-14-1-0633. All results, findings, and conclusions are the responsibility of the authors and do not necessarily reflect the opinions of the funding organization.

REFERENCES

[1] Dennett, D. C., [*The intentional stance*], MIT press (1989).

[2] Bratman, M. E., [*Intention, Plans, and Practical Reason*], Harvard University Press, Cambridge, MA, USA (1999).

[3] Anscombe, G. E. M., [*Intention*], Harvard University Press (1957).

[4] Bratman, M. E., "What is intention," *Intentions in communication* , 15–31 (1990).

[5] Goldman, A. I., "A causal theory of knowing," *The journal of Philosophy* **64**(12), 357–372 (1967).

[6] Davidson, D., "Intending," in [*Philosophy of history and action*], 41–60, Springer (1978).

[7] Mele, A. R. and Moser, P. K., "Intentional action," *Nous* **28**(1), 39–68 (1994).

[8] Searle, J. R., [*Rationality in action*], MIT press (2003).

[9] Searle, J. R., "Collective intentions and actions," *Intentions in communication* **401**, 401–02 (1990).

[10] Ajzen, I., "The theory of planned behavior," *Organizational behavior and human decision processes* **50**(2), 179–211 (1991).

[11] Malle, B. F., Moses, L. J., and Baldwin, D. A., [*Intentions and intentionality: Foundations of social cognition*], MIT press (2001).

[12] Malle, B. F. and Knobe, J., "The folk concept of intentionality," *Journal of Experimental Social Psychology* **33**(2), 101–121 (1997).

[13] Tomasello, M., [*The cultural origins of human cognition*], Harvard university press (2009).

[14] Gollwitzer, P. M., "Goal achievement: The role of intentions," *European review of social psychology* **4**(1), 141–185 (1993).

[15] Bruner, J. S., "Organization of early skilled action," *Child development* , 1–11 (1973).

[16] Custers, R., Eitam, B., and Bargh, J. A., "Conscious and unconscious processes in goal pursuit," (2012).

[17] Sheeran, P., "Intention—behavior relations: a conceptual and empirical review," *European review of social psychology* **12**(1), 1–36 (2002).

[18] Webb, T. L. and Sheeran, P., "How do implementation intentions promote goal attainment? a test of component processes," *Journal of Experimental Social Psychology* **43**(2), 295–302 (2007).

[19] Triandis, H. C., "Values, attitudes, and interpersonal behavior.," in [*Nebraska symposium on motivation*], University of Nebraska Press (1979).

[20] Baumeister, R. F., Masicampo, E., and Vohs, K. D., "Do conscious thoughts cause behavior?," *Annual review of psychology* **62**, 331–361 (2011).

[21] Jeannerod, M. and Jeannerod, M., [*The cognitive neuroscience of action*], vol. 1997, Blackwell Oxford (1997).

[22] Haggard, P., Clark, S., and Kalogeras, J., "Voluntary action and conscious awareness," *Nature neuroscience* **5**(4), 382 (2002).

[23] Haggard, P. and Clark, S., "Intentional action: Conscious experience and neural prediction," *Consciousness and cognition* **12**(4), 695–707 (2003).

[24] Montano, D. E. and Kasprzyk, D., "Theory of reasoned action, theory of planned behavior, and the integrated behavioral model," *Health behavior: Theory, research and practice* , 95–124 (2015).

[25] Rao, A. S. and Georgeff, M. P., "Modeling rational agents within a bdi-architecture.," *KR* **91**, 473–484 (1991).

[26] Georgeff, M., Pell, B., Pollack, M., Tambe, M., and Wooldridge, M., "The belief-desire-intention model of agency," in [*International Workshop on Agent Theories, Architectures, and Languages*], 1–10, Springer (1998).

[27] Holvoet, T. and Valckenaers, P., "Beliefs, desires and intentions through the environment," in [*Proceedings of the fifth international joint conference on Autonomous agents and multiagent systems*], 1052–1054, ACM (2006).

[28] Jennings, N. R., "Specification and implementation of a belief-desire-joint-intention architecture for collaborative problem solving," *International Journal of Intelligent and Cooperative Information Systems* **2**(03), 289–318 (1993).

[29] Levesque, H. J., Cohen, P. R., and Nunes, J. H., "On acting together," in [*AAAI*], **90**, 94–99 (1990).

[30] Cohen, P. R. and Levesque, H. J., "Teamwork," *Nous* **25**(4), 487–512 (1991).

[31] Cohen, P. R. and Levesque, H. J., "Intention is choice with commitment," *Artificial intelligence* **42**(2-3), 213–261 (1990).

[32] Nahmias, E., "Why we have free will," *Scientific American* **312**(1), 77–79 (2015).

[33] Cohen, P. R. and Levesque, H. J., "Persistence, intention, and commitment," in [*Reasoning About Actions & Plans*], 297–340, Elsevier (1987).

[34] Davis, W. A., "A causal theory of intending," *American Philosophical Quarterly* **21**(1), 43–54 (1984).

[35] Audi, R., "Intending," *The Journal of Philosophy* **70**(13), 387–403 (1973).

[36] Gibbs, R. W., "Intentions as emergent products of social interactions," *Intentions and intentionality: Foundations of social cognition* , 105–122 (2001).

[37] Bauer, A., Wollherr, D., and Buss, M., "Human–robot collaboration: a survey," *International Journal of Humanoid Robotics* **5**(01), 47–66 (2008).

[38] Meriçli, C., Klee, S. D., Paparian, J., and Veloso, M., "An interactive approach for situated task specification through verbal instructions," in [*Proceedings of the 2014 international conference on Autonomous agents and multi-agent systems*], 1069–1076, International Foundation for Autonomous Agents and Multiagent Systems (2014).

[39] Monajjemi, V. M., Wawerla, J., Vaughan, R., and Mori, G., "Hri in the sky: Creating and commanding teams of uavs with a vision-mediated gestural interface," in [*Intelligent Robots and Systems (IROS), 2013 IEEE/RSJ International Conference on*], 617–623, IEEE (2013).

[40] Trafton, J. G., Schultz, A. C., Perznowski, D., Bugajska, M. D., Adams, W., Cassimatis, N. L., and Brock, D. P., "Children and robots learning to play hide and seek," in [*Proceedings of the 1st ACM SIGCHI/SIGART conference on Human-robot interaction*], 242–249, ACM (2006).

[41] Rogalla, O., Ehrenmann, M., Zollner, R., Becher, R., and Dillmann, R., "Using gesture and speech control for commanding a robot assistant," in [*Robot and Human Interactive Communication, 2002. Proceedings. 11th IEEE International Workshop on*], 454–459, IEEE (2002).

[42] Demeester, E., Hüntemann, A., Vanhooydonck, D., Vanacker, G., Van Brussel, H., and Nuttin, M., "User-adapted plan recognition and user-adapted shared control: A bayesian approach to semi-autonomous wheelchair driving," *Autonomous Robots* **24**(2), 193–211 (2008).

[43] Croft, D., "Estimating intent for human-robot interaction," in [*IEEE International Conference on Advanced Robotics*], 810–815 (2003).

[44] Bien, Z. Z., Kim, J.-B., Kim, D.-J., Han, J.-S., and Do, J.-H., "Soft computing based emotion/intention reading for service robot," in [*AFSS International Conference on Fuzzy Systems*], 121–128, Springer (2002).

[45] Cabrera, M. E., Bogado, J. M., Fermin, L., Acuna, R., and Ralev, D., "Glove-based gesture recognition system," in [*Adaptive Mobile Robotics*], 747–753 (2012).

[46] Staff, N. R. C., Pew, R. W., and Mavor, A. S., [*Modeling human and organizational behavior: Application to military simulations*], National Academies Press (1998).

[47] Shaikh, M. T. and Goodrich, M. A., "Design and evaluation of adverb palette: A gui for selecting tradeoffs in multi-objective optimization problems," in [*Proceedings of the 2017 ACM/IEEE International Conference on Human-Robot Interaction*], 389–397, ACM (2017).

[48] Chandler, B. and Goodrich, M. A., "Online rrt* and online fmt*: Rapid replanning with dynamic cost,"

[49] Cummings, M. L., Clare, A., and Hart, C., "The role of human-automation consensus in multiple unmanned vehicle scheduling," *Human Factors* **52**(1), 17–27 (2010).

[50] Yi, D., Goodrich, M. A., and Seppi, K. D., "MORRF*: Sampling-based multi-objective motion planning," in [*Proceedings of the 24th International Conference on Artificial Intelligence*], 1733–1739, AAAI Press (2015).

[51] Yi, D., Goodrich, M. A., and Seppi, K. D., "Homotopy-aware rrt*: Toward human-robot topological path-planning," in [*Human-Robot Interaction (HRI), 2016 11th ACM/IEEE International Conference on*], 279–286, IEEE (2016).

[52] Jenkins, K. D., *The shortest path problem in the plane with obstacles: A graph modeling approach to producing finite search lists of homotopy classes*, PhD thesis, Monterey, California. Naval Postgraduate School (1991).

[53] Bhattacharya, S., Likhachev, M., and Kumar, V., "Topological constraints in search-based robot path planning," *Autonomous Robots* **33**(3), 273–290 (2012).